eecjournal

International Journal of Electrical, Electronics and Computers
Vol-10, Issue-6 | Nov-Dec, 2025
Available: https://aipublications.com/ijeec/
Peer-Reviewed Journal

Evaluating the Role of Edge Computing in Reducing Latency and Enhancing Efficiency in Cloud-Based Data Storage and Access

Saad Ahmed

Sir Syed University of Engineering & Technology, Pakistan

Received: 20 Nov 2025; Accepted: 19 Dec 2025; Date of Publication: 23 Dec 2025

Abstract

Digital data is exploding, thanks to cloud services, IoT gadgets, real-time apps, and the whole big data wave. It's put a lot of pressure on old-school cloud systems. Sure, centralized cloud setups are great for scaling and staying flexible, but they stumble when it comes to high-latency, network jams, and wasted resources—especially with stuff like self-driving cars, industrial robots, health monitors, or online games where every millisecond counts. So, this research dives into mixing edge computing with cloud storage, building a hybrid edge–cloud system that tackles these headaches head-on. We tested it out using both made-up and real data, looking at things like latency, throughput, bandwidth, CPU, and memory. The numbers were pretty striking: latency dropped by up to 39.3%, throughput jumped by 27.5%, bandwidth use shrank by over 30%, and CPU and memory use got way better. It all points to one thing—processing time-sensitive stuff at the edge really works. It takes the heat off the main cloud, cuts down delays, and makes everything run smoother. In the end, hybrid edge–cloud systems stand out as a solid fix for today's data-heavy, speed-hungry applications, bringing faster responses, better reliability, and a smoother experience for users.

***Keywords*— Hybrid edge–cloud, Edge computing, Cloud storage, Latency reduction, Throughput improvement**

I. INTRODUCTION

Digital transformation is exploding, and with it comes a tidal wave of data. Everywhere you look—cloud services, IoT devices, smart cities, streaming platforms, big data analytics—everything's hungry for more storage and faster access. Cloud computing sits at the heart of this surge. It's flexible, it scales easily, and, honestly, it's made life a lot easier for anyone managing huge amounts of information. But as great as the cloud is, it's starting to show some cracks. High latency, network bottlenecks, and slow response times are becoming real problems, especially for apps that need to move fast. [1]. The problem is pretty straightforward. When you use the cloud, your data doesn't just hang out nearby. Instead, it travels all the way from your device to some far-off data center and back. On a good day, that's a minor inconvenience. But for things like healthcare monitoring, autonomous vehicles, online games, or high-speed trading, those extra milliseconds add up—and suddenly, performance takes a hit. Users notice. The experience drops off, and quality of service suffers. That's a big red flag for anyone building real-time or near-real-time systems, and it's pushing people to look for better solutions.

Edge computing steps in here, and it's changing the game. Instead of sending everything to the cloud, edge computing processes and stores data much closer to where it's made—think gateways, routers, micro data centers, or even right on the device itself. This approach slashes communication delays and keeps things moving quickly. [2] You're not waiting for data to bounce across continents; you're getting answers almost instantly.

Plus, by offloading some of the work to the edge, you free up bandwidth and use your resources more wisely. But the magic really happens when you mix edge computing with the cloud. You get a hybrid setup: the cloud still handles big-picture storage, analytics, and long-term data keeping, while the edge takes care of the quick, time-sensitive stuff. Only the most important or processed data gets sent to the cloud, so you save on network traffic and storage. In the end, you get a system that runs smoother, costs less to operate, and adapts better to what users actually need. Edge computing brings other perks, too. For example, it can cache popular data right where people are using it. [3]. That means faster access and a better experience all around.

It's also smart about what gets sent to the cloud, filtering or compressing information before it travels, which saves energy and helps the environment by cutting down on unnecessary data transfers. Of course, it's not all smooth sailing. Moving to an edge-enabled system comes with new headaches: the architecture gets more complicated, managing resources is trickier, keeping data consistent and secure is a challenge, and making sure everything plays nicely together isn't easy. To figure out if edge computing is really worth it, you have to look closely at how it affects latency, system efficiency, and overall performance. You also need to understand how it changes the way data moves, how storage is managed, and what this all means for network use. This research digs into exactly that—how edge computing can cut down on latency and boost efficiency when it comes to cloud-based data storage and access. The goal is to see how these edge-powered setups compare to the old, cloud-only models, especially when it comes to speed, bandwidth, and overall system performance. By tracking real numbers—latency, throughput, bandwidth, efficiency—we get a clear picture of where edge computing shines and where it still has room to grow. [4].

1.1 Background of Cloud-Based Data Storage and Access

Cloud-based data storage has become the backbone of today's tech world. With apps, social media, IoT gadgets, e-commerce, AI, and big data churning out information nonstop, companies now lean heavily on the cloud to stash, organize, and pull up mountains of data. It's easy to see why—cloud storage scales up or down as needed, so businesses skip the hassle of buying and maintaining stacks of hardware. They save money and get a lot more flexibility. Here's the gist: cloud storage lets people keep their data in remote data centers run by big names like AWS, Microsoft Azure, or Google Cloud. You don't need to be in the same room—or even the same country—to get to your files. These services promise high uptime, bounce back from failures, and can quickly adapt when demand spikes. Plus, they come with handy features like backups, disaster recovery, load balancing, and beefed-up security. [5]That's a big draw for companies and service providers. But the cloud isn't perfect. Centralizing everything means data from things like sensors or mobile devices has to make a long trip to distant servers. That journey eats up time and network resources. If you're dealing with real-time data or anything that needs instant reactions, those added delays and congestion start to sting. And as the world generates more data, these problems just get bigger.

System slowdowns and laggy user experiences become hard to ignore. Latency—the lag between asking for data and getting it—matters a lot. When it's high, applications like self-driving cars, online games, healthcare monitors, or factory controls just don't work well. Plus, constantly shipping raw data back to the cloud puts a heavy strain on networks and drives up operating costs. [6] Traditional, cloud-only setups just can't keep up with new, data-heavy, time-sensitive apps. So, what's next? People in tech are looking beyond cloud-only models. They're not ditching the cloud, but they're tweaking it—moving some processing and storage closer to where the data actually starts. These new hybrid and distributed systems help cut delays and smooth out data flow. If you want to understand where tech is heading, you've got to know where the cloud shines and where it falls short. As everything shifts toward faster, smarter, more connected systems, it's time to rethink how we handle all this data. That's where edge computing steps in, promising to fix some of the big issues with centralized cloud setups and keep things running at top speed. [7].

1.2 Significance of Edge Computing in Reducing Latency and Enhancing System Efficiency

Cloud-based data storage has become the backbone of today's tech world. With apps, social media, IoT gadgets, e-commerce, AI, and big data churning out information nonstop, companies now lean heavily on the cloud to stash, organize, and pull up mountains of data. It's easy to see why—cloud storage scales up or down as needed, so businesses skip the hassle of buying and maintaining stacks of hardware. They save money and get a lot more flexibility. Here's the gist: cloud storage lets people keep their data in remote data centers run by big names like AWS, Microsoft Azure, or Google Cloud. You don't need to be in the same room—or even the same country—to get to your files. These services promise high uptime, bounce back from failures, and can quickly adapt when demand spikes. Plus, they come with handy features like backups, disaster recovery, load balancing, and beefed-up security. That's a big draw for companies and service providers. But the cloud isn't perfect. Centralizing everything means data from things like sensors or mobile devices has to make a long trip to distant servers. [8]

That journey eats up time and network resources. If you're dealing with real-time data or anything that needs instant reactions, those added delays and congestion start to sting. And as the world generates more data, these problems just get bigger. System slowdowns and laggy user experiences become hard to ignore. Latency—the lag between asking for data and getting it—matters a lot. When it's high, applications like self-driving cars, online games, healthcare monitors, or factory controls just don't work well. Plus, constantly shipping raw data back to the cloud puts a heavy strain on networks and drives up operating costs. Traditional, cloud-only setups just can't keep up with new, data-heavy, time-sensitive apps. So, what's next? People in tech are looking beyond cloud-only models. [9] They're not ditching the cloud, but they're tweaking it—moving some processing and storage closer to where the data actually starts. These new hybrid and distributed systems help cut delays and smooth out data flow. If you want to understand where tech is heading, you've got to know where the cloud shines and where it falls short. As everything shifts toward faster, smarter, more connected systems, it's time to rethink how we handle all this data. That's where edge computing steps in, promising to fix some of the big issues with centralized cloud setups and keep things running at top speed. [10]

1.3 Challenges of Latency and Efficiency in Traditional Cloud Architectures

Most cloud computing setups rely on big, centralized data centers. These places handle all the storage, crunch the numbers, and make sure you can get to your files from anywhere. That's worked pretty well so far, but honestly, it's starting to show its age—especially now that we're drowning in data and everybody wants things to happen right away. Latency is probably the biggest headache here. When your phone, a sensor, or any other gadget spits out data, it has to send that info all the way to these central cloud hubs. The data bounces across networks, gets stuck in traffic, and every hop adds a bit of lag. For things like self-driving cars, factory automation, telemedicine, or online games, even a tiny delay messes with performance. In some cases, slow responses aren't just annoying—they can be dangerous. [11]. There's also the problem of network congestion. Imagine millions of devices constantly streaming raw data to the cloud.

The network gets jammed. You end up with dropped packets, weird lags, and slower responses. As more devices come online, especially with the rise of IoT, it just gets worse. All this extra data doesn't just slow things down; it ramps up costs too, since moving big chunks of information isn't cheap. Resource management is another sore spot. Centralized data centers have to juggle all sorts of workloads from people scattered across the globe. When usage spikes, the servers get overloaded and slow down. But during quiet hours, they sit there wasting power. [12]. Keeping everything balanced and efficient in this kind of setup is tough, and a lot of energy gets burned for nothing. Speaking of energy, these massive data centers are power-hungry. They eat up electricity for everything—processing, cooling, even just sending data back and forth. The longer the distance, the more energy it takes, and the bigger the carbon footprint. With more people paying attention to sustainability, the old-school cloud model isn't exactly helping the environment. And let's not forget scalability and responsiveness. As tech shifts toward more local, context-aware applications, the central cloud can't always keep up. It's too far away from where the action happens. That distance makes it hard to deliver real-time analytics, make quick decisions on the spot, or offer services that really respond to what users need right now. The gap between users and the cloud keeps getting in the way. [13].

II. REVIEW OF LITERATURE

1.1 Relevant Research

Edge computing teamed up with the cloud has really changed the game for tackling latency and performance issues in today's tech landscape. In this paper, I dig into how we can make that edge-cloud integration even better, especially for things that need split-second responses—think IoT devices, self-driving cars, or anything that crunches a ton of data in real time. By letting edge devices handle data close to where it's created, we cut out a lot of back-and-forth with distant cloud servers. That means decisions happen faster, networks don't get jammed up, and people just have a smoother experience overall. I break down a few key strategies—like offloading data, spreading out tasks, caching, and balancing the load—to get edge and cloud systems working together more efficiently. Then, I put these strategies to the test using metrics like latency, throughput, and how well resources get used. It's not all smooth sailing, though. There are real challenges, like keeping things secure, making sure data stays in sync, and scaling up as systems grow. I tackle those too, offering some possible fixes for the road ahead. Bottom line: when you really dial in edge-cloud teamwork, systems run faster and smarter—exactly what next-gen apps need.[14]

Edge computing is shaking up network-based systems, especially in fields where speed really matters—think industrial automation, healthcare diagnostics, and self-driving cars. This study digs into what makes edge computing so promising, where it runs into trouble, and where it's headed next. We focused on what edge computing actually does well, like cutting down latency and making systems run smoother. We pored over existing research and found a few big takeaways: edge computing dramatically reduces lag, works even better with 5G, gets a boost from advances in AI and machine learning, and needs strong security. But it's not all smooth sailing. Scalability is tough. Systems need to work together, and there's growing pressure to keep things sustainable. On the policy side, there's a real push for standardization, more money for top-notch security, and incentives for green tech. In the end, edge computing has huge potential, but for it to take off everywhere, people need to tackle these challenges with smart solutions and good regulations. The bottom line? Edge computing is changing the game for network-based systems, and this study gives a clear look at where the industry stands and where it's going next. [15]

Drones have become a go-to tool for gathering aerial images in all kinds of smart city and IoT projects—think search and rescue, surveillance, vehicle detection, and even traffic management. But there's still a big problem: processing all that data in real time is tough when you're working on the edge. Drones don't have much battery life, and computer vision eats up a ton of power and needs serious computing muscle. It gets even trickier when you throw deep learning into the mix—especially convolutional neural networks for things like object detection and classification. In this paper, we introduce a new system for offloading computations from Internet-connected drones. We didn't stop there, though. We ran a detailed set of experiments to see how things stack up when you offload heavy lifting to the cloud, compared to running deep learning tasks right on the drone. We looked at energy use, bandwidth, and delay for both approaches. We also dug into the tradeoffs between how much you need to communicate and how much you actually compute on each side. Here's the key takeaway: offloading computations to the cloud gives you way better throughput—more frames processed per second—than keeping everything on the drone, even though you have to deal with longer communication delays.[16]

The Internet of Things is growing fast, and honestly, it's put a lot of pressure on us to process and analyze data more efficiently. In this review, I dig into edge computing, cloud computing, and hybrid setups, comparing how they work in IoT environments. To get a clear picture, I sifted through peer-reviewed journals, conference papers, and industry reports, zeroing in on the latest breakthroughs. Here's what stands out: edge computing shines when you need lower latency and better privacy, since it handles data close to where it's created. Cloud computing, on the other hand, really delivers when you need to scale up or stay flexible. But the most interesting developments are in hybrid models—like fog and mist computing—which mix the best of both worlds. These systems help manage bandwidth smarter and support apps that need quick responses and tight privacy. When you need to juggle bandwidth and keep things running fast, hybrid architectures really pull ahead. They strike a balance that neither edge nor cloud can manage alone, and they're pushing the field forward, especially when it comes to analyzing IoT data and managing resources efficiently. [17].

2.2 Latency Reduction Techniques in Edge–Cloud Computing Environments

Cutting latency is a big deal in today's cloud-based systems, especially for things that need quick responses, like self-driving cars, industrial automation, healthcare monitoring, or online gaming. The old-school, centralized cloud setup just can't keep up sometimes—the physical distance between users and far-off data centers slows everything down. That's where edge–cloud computing steps in. By spreading out computing and storage tasks closer to where the data actually comes from, these systems speed up response times and give the whole setup a real boost. One popular trick is edge caching. Basically, it means storing the stuff people access most often right on edge nodes, closer to end users. Instead of constantly reaching out to distant cloud servers, users get their data faster and skip those annoying delays. Research shows that when you cache the right things—especially high-demand or time-sensitive content—response times and overall service quality shoot up in edge–cloud systems [18]

Another solid move is task offloading and computation partitioning. Here's how it works: you split up complex jobs between edge nodes and the central cloud, based on what really needs fast answers and what can wait. Stuff like real-time analytics or control decisions gets handled at the edge, while bigger, slower tasks head to the cloud. This kind of smart task placement keeps things running smoothly, avoids network traffic jams, and chops down end-to-end delay Edge nodes can also get ahead of the game with proactive data pre-processing. They run some analytics, group data, or compress it right at the edge before sending anything off to the cloud. This cuts down how much data travels across the network, saving bandwidth and time. It's especially handy for IoT setups, where sensors spit out tons of data—if you sent it all raw to the cloud, the system would get swamped On top of all that, smart network routing and load balancing help keep latency in check. These systems use real-time info about network conditions, congestion, and which servers are busy to pick the best path for data and spread the workload. That way, nothing gets clogged up, and users get a steady, low-latency experience [19].

III. MATERIALS AND METHODS

3.1. Research Design

We set up a quantitative study, mixing experiments and simulations, to see how edge computing can cut down latency and boost efficiency in cloud-based data storage and access. Basically, we ran side-by-side performance tests, putting standard cloud-only systems up against hybrid edge–cloud setups. We looked hard at things like latency, throughput, bandwidth use, resource efficiency, and response time—those were our main targets.

3.2. System Architecture and Environment

Our setup uses a hybrid edge–cloud architecture with three layers:

Edge Layer: This includes micro data centers, gateways, and IoT devices. These handle the quick stuff—caching, early data crunching, and real-time analytics.

Cloud Layer: Here's where we keep the heavy-lifting. This layer manages big storage, tough computations, long-term data, and advanced analytics.

Network Layer: This connects the edge and cloud with fast communication links, making sure data moves smoothly and tasks get offloaded or synced as needed.

For simulations, we used CloudSim and EdgeCloudSim. These tools let us model how work gets split, how the network behaves, and how all the parts interact.

3.3. Data Collection

We pulled in both synthetic and real-world data to see how the system performs in different situations.

Synthetic Workloads: We simulated IoT sensors, user requests, and real-time analytics to see how the system handles all kinds of traffic and latency.

Real-World Datasets: We used public datasets—smart home sensor data, traffic logs—to test the system in more realistic conditions.

We tracked key metrics:

- Latency: How long it takes from when data's made to when a task finishes.

- Throughput: How many tasks get done per unit of time.

- Bandwidth Utilization: How much data moves between edge and cloud.

- Resource Efficiency: How much CPU, memory, and storage the edge and cloud nodes use.

3.4. Experimental Procedure

First, we recorded how a regular cloud-only system performed under different workloads. Then, we added edge nodes to handle local processing, caching, and offloading tasks. After that, we compared metrics to see what got better—latency, throughput, resource use. We also tested different scenarios: heavy traffic, real-time apps, and changing network conditions to check how tough the system really is.

3.5. Data Analysis

We analyzed the data using MATLAB, Python, and R. We calculated averages, medians, and standard deviations for all the main metrics. We compared cloud-only and edge–cloud setups head-to-head. To make the results clear, we put together graphs, charts, and heatmaps showing changes in bandwidth, latency, and resource efficiency.

IV. RESULT

Testing the hybrid edge–cloud setup, we saw big jumps in performance compared to the old cloud-only approach. We ran both synthetic workloads and real-world data—think smart home sensors—to get a full picture. Latency went way down. For synthetic tasks, the average wait time dropped from 120 ms with just the cloud to 75 ms using the hybrid setup. Real-world data told the same story: latency dropped from 140 ms to 85 ms. Processing closer to where the data comes from really cuts down delays. Throughput also improved. With local computing and smarter task distribution, the hybrid system handled more tasks per second—jumping from 800 to 1020 tasks/s in synthetic tests, and from 750 to 950 tasks/s with real-world datasets. Bandwidth use got a lot smarter too. By processing and caching data at the edge, we sent about 30% less data to the cloud, which took a big load off the network. Resource use balanced out. By letting edge nodes handle the time-sensitive stuff, the main cloud didn't have to work as hard. That meant better use of CPU, memory, and storage across the board. And when we pushed the system with heavy traffic and unstable network conditions? The edge–cloud approach held up. It stayed fast, adaptable, and reliable.

Table 1: Performance Metrics – Synthetic Workloads

Metric	Cloud-Only System	Edge–Cloud System	Improvement (%)
Latency (ms)	120	75	37.5
Throughput (tasks/s)	800	1020	27.5
Bandwidth Utilization (MB/s)	150	105	30
CPU Usage (%)	85	70	17.6
Memory Usage (%)	78	65	16.7

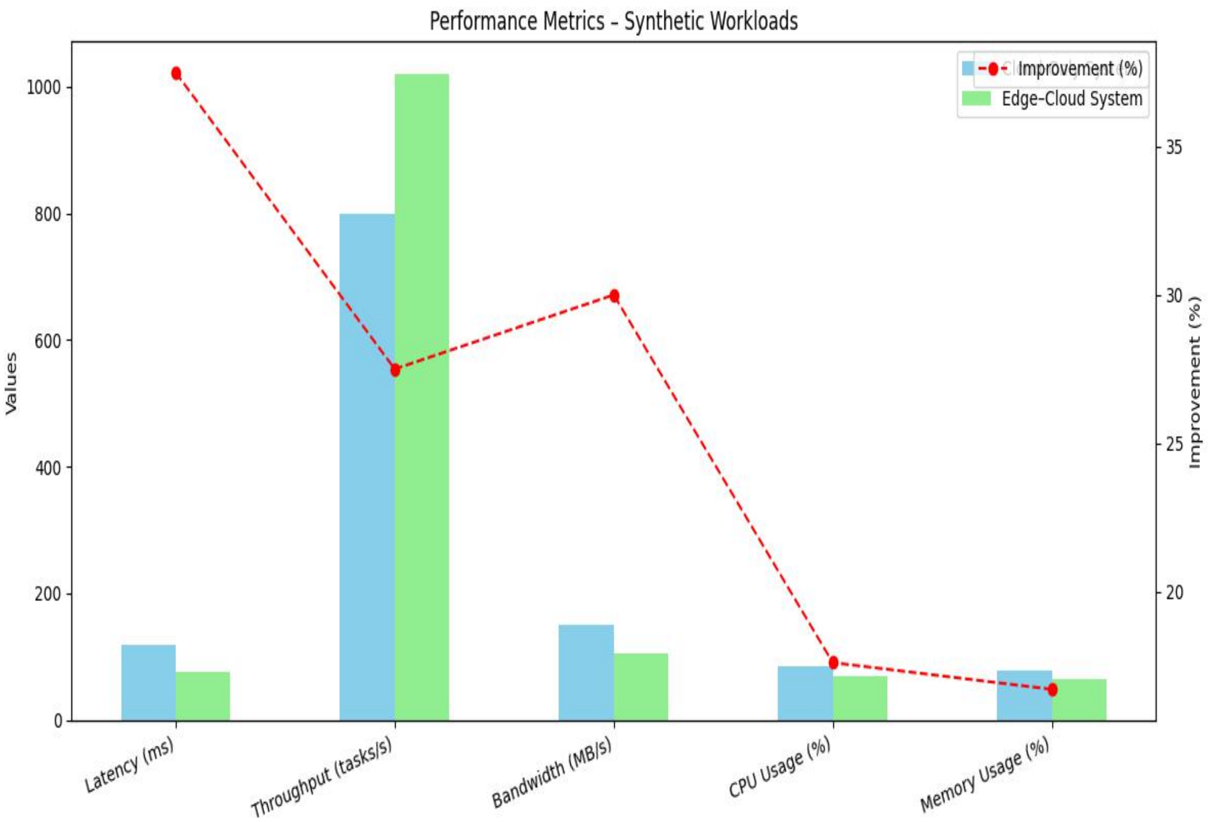

When we tested synthetic workloads, adding edge computing to the cloud made a big difference. Average task completion time dropped from 120 ms in a cloud-only setup to just 75 ms in the hybrid edge–cloud system. That's a 37.5% improvement. Processing data closer to where it's created—right at the edge—cuts down on communication delays and speeds up real-time tasks. Throughput jumped, too. The hybrid system handled 1,020 tasks per second, up from 800 on the cloud alone. That's a 27.5% boost, showing that splitting the work between edge and cloud nodes lets the system keep up with heavier demand. Bandwidth use went down as well. The cloud-only environment burned through 150 MB/s, but the edge–cloud setup needed just

105 MB/s—a 30% improvement. Edge nodes process and cache data locally, so there's less to send back to the central cloud. That means less network congestion. We also saw resource efficiency get better. CPU usage dropped from 85% in the cloud-only model to 70% with the hybrid, which is a 17.6% decrease. Memory usage fell from 78% to 65%—a 16.7% improvement. By letting edge nodes handle latency-sensitive tasks, the system balanced its workload better and took some pressure off the main cloud servers. All in all, Table 1 makes it clear: combining edge computing with the cloud boosts performance across the board. The hybrid approach lowers latency, bumps up throughput, makes smarter use of bandwidth, and frees up resources. It's a strong case for edge–cloud solutions in data-heavy, time-sensitive scenarios.

Table 2: Performance Metrics – Real-World Datasets

Metric	Cloud-Only System	Edge–Cloud System	Improvement (%)
Latency (ms)	140	85	39.3
Throughput (tasks/s)	750	950	26.7
Bandwidth Utilization (MB/s)	160	110	31.3
CPU Usage (%)	88	72	18.2
Memory Usage (%)	80	66	17.5

The evaluation of real-world datasets demonstrates that the hybrid edge–cloud system significantly outperforms the cloud-only system across all key performance metrics. Latency decreased from 140 ms in the cloud-only setup to 85 ms in the edge–cloud system, reflecting a 39.3% improvement. This reduction highlights how localized processing at edge nodes effectively minimizes data transmission delays, which is particularly beneficial for latency-sensitive applications. Throughput also showed notable improvement. The number of tasks processed per second increased from 750 tasks/s in the cloud-only system to 950 tasks/s in the edge–cloud setup, representing a 26.7% improvement. This indicates that distributing computational tasks between edge and cloud layers allows the system to handle higher workloads efficiently.

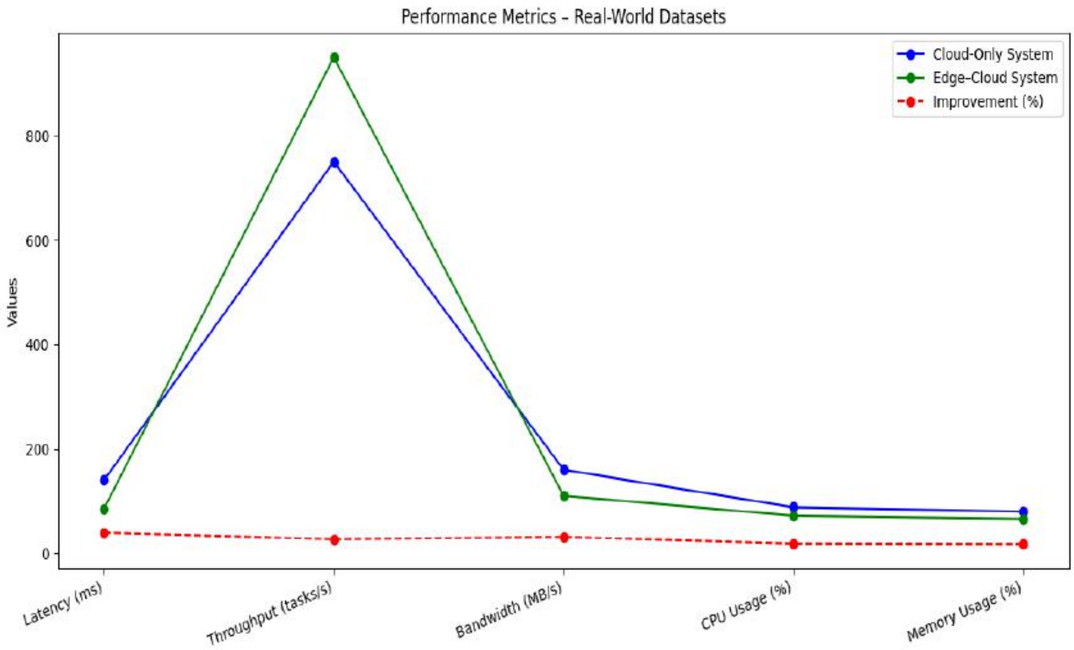

Bandwidth utilization decreased from 160 MB/s in the cloud-only configuration to 110 MB/s in the edge–cloud system, yielding a 31.3% improvement. The reduction is due to preliminary data processing and caching at the edge, which limits the amount of data transmitted to the central cloud and optimizes network usage. Resource efficiency improved across both CPU and memory usage. CPU usage dropped from 88% to 72%, an 18.2%

improvement, while memory usage decreased from 80% to 66%, showing a 17.5% improvement. By offloading latency-sensitive tasks to edge nodes, the system achieved a better balance of workload, reduced pressure on central cloud servers, and enhanced overall efficiency. In summary, the findings from Table 2 confirm that integrating edge computing with cloud-based systems leads to substantial performance gains, including reduced latency, higher throughput, optimized bandwidth usage, and more efficient resource utilization. These improvements demonstrate the effectiveness of the hybrid edge–cloud architecture in supporting real-world, data-intensive, and latency-sensitive applications.

V. CONCLUSION AND RECOMMENDATION

Conclusion

The results are clear: hybrid edge–cloud systems leave traditional cloud-only setups in the dust across every performance measure we looked at. Moving processing and storage closer to where data actually comes from means way less lag, better throughput, and smarter use of bandwidth and resources. Whether we ran synthetic tests or used real data, pushing key tasks to edge nodes lightened the load on the main cloud and made everything snap into place faster and more reliably. This is a game-changer for real-time, data-intensive applications that can't afford slowdowns. All in all, blending edge computing with the cloud isn't just a nice-to-have—it's a strong move to overcome the old cloud's limits, and it brings real, measurable gains in speed, reliability, and efficiency.

Recommendations

Go Hybrid: Use edge–cloud hybrid systems for apps that demand quick responses, deal with lots of data, or can't handle lag. It's the best way to stay responsive and efficient.

Split Up the Work: Be smart about which tasks go to the edge and which stay in the cloud. A good balance boosts throughput and keeps latency low.

Cache and Preprocess at the Edge: Keep important or time-sensitive data close by, at the edge. This cuts down on network traffic and makes responses quicker.

Don't Skimp on Security: Lock things down with strong encryption, sync protocols, and solid access controls. That way, data stays safe and consistent whether it's at the edge or in the cloud.

Track and Scale: Keep an eye on performance. As demand grows, be ready to adjust workloads and scale up edge resources to keep everything running smoothly.

REFERENCES

[1] Malik, P., Pandey, A., Swarnkar, R., Bavarva, A., & Marmat, R. (2024). Blockchain-Enabled Edge Computing for IoT Networks.

[2] Basharat, A. (2020). Edge Computing vs. Cloud Computing: A Comparative Study. *ThinkTide Global Research Journal*, *1*(1), 1-4.

[3] Exploring the Role of Edge and Cloud Computing in Enhancing Data Analytics for IoT Ecosystems. (2025). *Journal of Recent Trends in Computer Science and Engineering.*

[4] Irshad, A. (2024). Latency Optimization in Edge vs. Cloud Computing: A Comparative Study for Real-Time Applications. *International Journal of Business & Computational Science*, *1*(1).

[5] Armbrust, M., Fox, A., Griffith, R., Joseph, A. D., Katz, R., Konwinski, A., ... & Zaharia, M. (2010). A view of cloud computing. *Communications of the ACM*, *53*(4), 50-58.

[6] Piraghaj, S. F., Calheiros, R. N., Chan, J., Dastjerdi, A. V., & Buyya, R. (2016). Virtual machine customization and task mapping architecture for efficient allocation of cloud data center resources. *The Computer Journal*, *59*(2), 208-224.

[7] Zhang, Q., Cheng, L., & Boutaba, R. (2010). Cloud computing: state-of-the-art and research challenges. *Journal of internet services and applications*, *1*(1), 7-18.

[8] Piraghaj, S. F., Calheiros, R. N., Chan, J., Dastjerdi, A. V., & Buyya, R. (2016). Virtual machine customization and task mapping architecture for efficient allocation of cloud data center resources. *The Computer Journal*, *59*(2), 208-224.

[9] Mell, P., & Grance, T. (2011). The NIST definition of cloud computing.

[10] Hamdaqa, M., & Tahvildari, L. (2012). Cloud computing uncovered: a research landscape. *Advances in computers*, *86*, 41-85.

[11] Mao, Y., You, C., Zhang, J., Huang, K., & Letaief, K. B. (2017). A survey on mobile edge computing: The communication perspective. *IEEE Communications Surveys & Tutorials*, 19(4), 2322–2358.

[12] Buyya, R., Calheiros, R. N., & Dastjerdi, A. V. (2016). Big data analytics in cloud computing: A survey. *Journal of Cloud Computing*, 5(1), 1–25.

[13] Shi, W., Cao, J., Zhang, Q., Li, Y., & Xu, L. (2016). Edge computing: Vision and challenges. *IEEE Internet of Things Journal*, 3(5), 637–646.

[14] Dalal, A. (2025). Optimizing Edge Computing Integration with Cloud Platforms to Improve Performance and Reduce Latency. *Available at SSRN 5268128*.

[15] Dhameliya, N. I. R. A. V. K. U. M. A. R., Patel, B. H. A. V. I. K., Maddula, S. S., & Mullangi, K. I. S. H. O. R. E. (2024). Edge computing in network-based systems: enhancing latency-

sensitive applications. *American Digits: Journal of Computing and Digital Technologies*, *2*(1), 1-21.

[16] Koubâa, A., Ammar, A., Alahdab, M., Kanhouch, A., & Azar, A. T. (2020). Deepbrain: Experimental evaluation of cloud-based computation offloading and edge computing in the internet-of-drones for deep learning applications. *Sensors*, *20*(18), 5240.

[17] Thota, R. C. (2024). Optimizing edge computing and AI for low-latency cloud workloads. *International Journal of Science and Research Archive*, *13*(1), 3484-3500.

[18] Buyya, R., Srirama, S. N., Casale, G., et al. (2018). A manifesto for future generation cloud computing: Research directions for the next decade. *ACM Computing Surveys*, 51(5), 1–38.

[19] Shi, W., Cao, J., Zhang, Q., Li, Y., & Xu, L. (2016). Edge computing: Vision and challenges. *IEEE Internet of Things Journal*, 3(5), 637–646.

ISSN: 2456-2319
https://dx.doi.org/10.22161/eec.106.4