International Journal of Science and Engineering
ISSN (Online): 2454 - 2016
Volume 10 Issue 02 June 2024
DOI: 10.53555/ephijse.v10i2.301

INTELLIGENT COST OPTIMIZATION FOR EMR WORKLOADS USING CUSTOM APIS

Yasodhara Varma[1*], Manivannan Kothandaraman[2]

[1*]*Vice President at JPMorgan Chase & Co*
[2]*Vice President, Senior Lead Software Engineer, JP Morgan Chase & Co*
Corresponding Author:

Abstract
Given the rising demand for big data processing and the current dynamic economic environment, which presents challenges for businesses, especially in terms of managing cloud computing costs has become a large issue. Although cloud infrastructure provides scalability and adaptability, poor management of it could cause significant expenses. Large data set processing inside distributed computing systems largely relies on Amazon Elastic MapReduce (EMR). Although Electronic Medical Records in the healthcare sector and other businesses managing large amounts of data might find EMS suitable, its dynamic and scalable capabilities could potentially lead to cost inefficiencies. Inappropriate scaling or too generous resource allocation can lead to resource waste and higher running expenses. This white paper examines cutting-edge cost-optimal methods designed especially for E MR workload management. Electronic Medical Records (EMRs) are confidential patient records requiring efficient, safe, reliable processing—qualities lacking in which case significant charges could arise. The study stresses on a simple approach the usage of tailored application programming interfaces (APIs). These APIs let one automate important chores such dynamic job scheduling, real-time instance selection, autonomous scaling, and ongoing cost monitoring. By means of automation, businesses may guarantee that computing resources are distributed precisely where and when needed, therefore avoiding the inefficiencies connected with set configurations. Dynamic work scheduling distributes tasks depending on real-time data, therefore optimizing resource use all day. Companies can identify the most reasonably priced and task-appropriate computer instances by means of selective instance selection, therefore avoiding a homogeneous approach that might result in inefficiencies or insufficient performance. By allowing systems to dynamically change resources in response to various demands, autonomous scaling guarantees performance while eliminating needless resource allocation and so maximizes efficiency.

Keywords: *Intelligent Cost Optimization, EMR Workloads, Custom APIs, Cloud Computing, Data Processing, Dynamic Job Scheduling, Selective Instance Selection, Automatic Scaling, Real-Time Cost Control, Resource Efficiency, Automation, Smart Decision-Making*

1. Introduction

Amazon EMR and other cloud-based data processing systems help organizations properly handle large amounts of data. Still, businesses routinely struggle with:

- Unnecessary cloud spending brought on by too distributed resources.
- Underused clusters pay idle expenses without producing any useful computing.
- Less than ideal job scheduling leads to higher costs and longer processing times.
- Lack of real-time spending visibility makes proactive expenditure control difficult.

Customized APIs that integrate with EMR help organizations to streamline cost governance processes, improve cluster lifecycle management, and adjust settings depending on workload demands, so tackling these challenges.For organizations working on large-scale data processing projects, controlling cloud costs is still a constant challenge. Big data processing makes considerable use of Amazon EMR (Elastic MapReduce); nevertheless, if not correctly calibrated, its dynamic and scalable properties may lead to inefficient resource allocation and increased expenses. Resource over-provisioning, underutilized clusters, poor work scheduling, and inadequate real-time cost visibility are common problems that organizations run upon. Organizations could suffer significant operational inefficiencies and cost loads in the lack of preemptive initiatives.

Using customized APIs to automate scaling, improve resource use, and save unnecessary expenses, this white paper looks at intelligent cost optimization solutions for EMR workloads.Dynamic job scheduling, smart instance selection, auto-scaling improvements, and real-time cost control can help organizations save a lot of money while yet preserving dependability and performance. Artificial intelligence powered automation will help businesses to stop cost overruns and proactively adapt to match changing workloads.

This paper highlights advanced methodologies in cost forecasting based on predictive analytics, hybrid cloud solutions for cost control, and sustainability-oriented optimization approaches matching modern environmental goals. We highlight useful cases in which organizations have successfully implemented intelligent EHR optimization, hence improving operational efficiency and resulting in measurable cost savings. This article underlines the importance of AI-driven insights for better decision-making in EMR cost governance together with a review of cost-cutting techniques. By implementing automated resource management and applying cloud-native technologies, businesses may assist to guarantee that their EHR loads are reliable, scalable, and reasonably priced. This proactive approach enables businesses to maximize their cloud expenditures and keep outstanding performance in a world of dynamic big data.

1.1 Key Cost Optimization Strategies for EMR

1.1.1 Dynamic Cluster Scaling with Custom APIs

Custom APIs can increase EMR's scalability outside of set auto-scaling constraints by:

- Dynamic change of node counts and real-time tracking of workload changes.
- Actively changing cluster dimensions via predictive scaling guided by past performance.
- Deliberately disabling underutilized nodes rather than dependent on AWS auto-termination rules.

Offering better and more flexible scaling capabilities than traditional cloud-based auto-scaling solutions, custom APIs are revolutionizing organizational management of EMR (Elastic MapReduce) workloads. AWS often does not sufficiently maximize demand forecasting and real-time resource allocation, even if it offers basic auto-scaling depending on predefined parameters. Custom APIs fill in for this gap by allowing proactive resource deactivation, predictive scaling informed by past data, and dynamic node changes. Unlike normal scaling rules that turn on depending on set criteria, customizable APIs continuously monitor task information and performance indicators to recognize changes in workload. This real-time study guides clusters to decrease in load reduction and expand in demand surges. Customized APIs, for example, can quickly find underutilization and deactivate a node, therefore lowering unnecessary costs and resource depletion instead of waiting for a node to be inactive for a set duration. Customized APIs usually serve in their capacity to project future resource requirements depending on past trends.

Should data show increasing Monday morning EMR demand, the API can aggressively grow the cluster to prevent congestion and guarantee continuous service. This proactive strategy reduces generally associated delays connected with reactive scaling techniques and increases general system responsiveness. Custom APIs project execution times and look at job metadata to boost accuracy and scalability. Understanding the numerous criteria of every role helps the system to allocate resources correctly. While longer or more important chores can take first priority, jobs meant to be finished fast could be assigned less resources. This exact management guarantees best use of resources and helps to minimize overprovisioning. Artificial intelligence disturbs this ecology greatly. Fast study of big-scale performance data lets artificial intelligence find trends and autonomously change resource allocation, so improving and flexible cluster capacities. Acting as a discriminating supervisor, the artificial intelligence component guarantees flawless operations even with changing employment demand. This approach improves environmentally friendly computing, reduces costs, and improves performance. Modern sustainability seeks that efficient use of cloud resources lowers waste output and energy usage. Companies have more flexibility free from tight scaling restrictions. Custom APIs let companies create resource management plans fit for their particular operations. This is really beneficial in industries requiring both high-performance computers and effective cost control. Custom API interfaces provide a strong basis for scalable and sustainable computing since they permit EMR clusters to more effectively and economically control complex data processing. By means of real-

time scaling, predictive changes, and artificial intelligence-driven automation, custom APIs finally significantly improve EMR cluster management. These characteristics allow businesses to better and more strategically control data volumes as well as eliminate waste and reduce costs—all of which support continuous, high-performance operations.

1.2 Intelligent Instant Choice

Many organizations use default EC2 instance types for EMR uses, which could cause ineffective spending. Using a custom API helps one to:

- Choose dynamically depending on work complexity, memory constraints, and cost-effectiveness different instance kinds.
- Use contingency plans to enable workloads flow between spot and on-demand events as needed.
- Terminate instances when processing endpoints are reached to apply instance lifespan optimization.

By connecting with AWS pricing APIs, organizations may automate reasonably priced instances choosing to maximize savings. Artificial intelligence and machine learning techniques could assess prior performance data and provide informed recommendations on appropriate instance kinds, therefore guaranteeing an optimal balance between cost and performance.

1.3 Smart Job Scheduling to Cut Empty Costs

Cloud inefficiency is largely influenced by idle cluster time. Custom APIs can improve efficiency of scheduling by:

- Grouping related chores to cut waste of resources.
- Suspending clusters automatically during low demand to reduce needless running costs.
- Using dependant-aware scheduling will ensure that jobs run consecutively without overlapping resource conflict.
- Using serverless and container-based architectures—like Kubernetes with E MR—organizations may maximize scheduling and reduce wasted computing resources.

1.4 Real-Time Cost Monitoring & Alerts

Organizations usually only notice too high expenses when they receive their monthly AWS invoice. A proactive approach consists of:

- Implementing real-time cost tracking systems with finely tuned E-MR expenditure analysis.
- Setting automated alerts when spending more than allowed.
- Creating cost analyses that reveal which jobs and circumstances most affect expenses helps one understand which ones most affect which others.

Combining AWS offerings such as Cost Explorer and Cloud Watch provides companies real-time data to support rapid configuring changes. Combining predictive data with team future expenditure forecasts would enable better resource allocation and development of strategies. Strong machine learning techniques leveraging prior performance data can predict workload peaks and initiate proactive modifications, hence enabling dynamic cluster scaling. This method guarantees exactly what is needed, hence minimizing delay and preventing over-provisioning. Effective resource allocation can be somewhat facilitated by a load-aware cluster management system. By assigning exactly the necessary resources for certain operations, the approach lowers idle time and increases processing efficiency instead of scaling complete node arrays. A tiered scale helps to cut expenses and increases performance. Custom APIs improve availability and cost-efficiency by means of dynamic transitions between instance types depending on real-time workload demands and AWS price data using technologies like Spot Instance Advisor and EC2 Auto Scaling. Good parallel processing and resource economy guaranteed by effective task scheduling help to lower conflicts. Moreover, putting processing nodes near the data source reduces data transportation costs, improving performance—especially for transformational or high-volume applications. When utilized above permissible limits, APIs can set alarms, therefore allowing rigorous cost control and controlled scale-back. These concepts taken together offer an intelligent EHR design maximizing performance, affordability, and scalability by leveraging self-regulating intelligence. Companies implementing proactive, intelligent cluster management have shown up to 40% total cost savings.

2. Case Study: Minimizing Expenses in a Context Rich in Data

Rising EMR costs were experienced by a financial services company handling petabytes of transaction data. By means of customized APIs for intelligent instance selection, flexible scaling, and immediate cost control, they achieved:

- Using suitable instance types results in a 35% drop in EC2 costs.
- 50% less idle cluster time by way of automatic shutdown methods.
- Enhanced operational efficiency helps to enable speedy completion of activities free from needless provisioning.

Maximum cost savings came from additional enhancements to job scheduling and dependent-driven resource allocation. The company turned to artificial intelligence-based cost forecasting to improve its EMR cost control strategy.

2.1 Future paths and growing trends

As cloud computing and big data processing evolve, rising elements will increasingly influence cost efficiency in EMR systems:

Increasingly, businesses are utilizing serverless methods to independently control resource allocation and cut pointless expenses. Real-time EMR cost prediction and improvement led by artificial intelligence-based analytics will progressively help to optimize machine learning. By means of hybrid and multi-cloud architectures, organizations are dividing their cloud strategies, therefore distributing workloads among various providers and lowering reliance on a single vendor and hence increasing cost efficiency. Reducing carbon footprints allows one to maximize cloud resources, hence improving sustainability and harmonizing with corporate sustainability objectives.

2.2 Custom API Based Optimizing of Instance Selection

Using real-time workload metrics, previous task performance data, and AWS pricing information via a customized API would help one select the most suitable EC2 instance for EMR load. This automatic support:naturally selecting instance types based on job size, complexity, and duration. Automatically selecting instance forms based on work size, complexity, and duration. Using AWS spot instances with a backup plan to run tasks at less cost while still keeping availability. Running instance lifecycle management wherein automatically terminating instances upon work completion.

Moreover, the API can change the choice depending on:

Designed for memory-intensive applications like machine learning model training, these are instances maximized for memory use. Designed for computing, these instances satisfy high-performance processing demandsDesigned for data-intensive jobs requiring high I/O throughput, they are storage-oriented solutions. Automatically choosing instance kinds based on project size, complexity, and duration.

2.3 Maximizing Spot Instances for Economic Performance

Spot events provide notable cost savings when weighed against on-demand events.

Still, they run the danger of being kicked off by AWS should capacity be required elsewhere. A carefully written bespoke API can reduce this risk by:

Seeing real-time spot market pricing and choosing the most cost-effective solutions.

building fallback systems to transfer work to events scheduled for on-demand consumption when spot events cancel out. Analyzing work completion times using predictive analytics helps to match them with pertinent spot events. This approach provides operational dependability and reduces EMS load associated costs. Over events booked on demand, spot events provide clear cost savings. Should capacity grow elsewhere, AWS still runs the risk of deleting them. A carefully designed bespoke API can lower this risk by tracking real-time spot market pricing and selecting the most reasonably priced solutions; moreover, it can construct fallback systems to reallocate work to events prepared for on-demand consumption should spot events be canceled. Predictive analytics analysis of work completion timeframes facilitates matching of them with relevant spot occurrences. This method provides operational dependability and lowers EMS load related expenditures. Including artificial intelligence models also helps to predict disruptions, hence allowing seamless transitions and less disturbance.

2.4 AI-activated Instance Suggestions

Analyzing past use patterns and suggesting appropriate configurations helps artificial intelligence (AI) and machine learning (ML) to enhance instance choosing. Incorporating artificial intelligence-driven recommendations can help businesses choose the best instance types for particular workload categories. Change cluster configurations in line with expected job criteria. Reduce hand-off by allowing completely automated, reasonably priced instance building. These developments lower costs and improve the dependability and efficiency of EMR activities. Including artificial intelligence generated recommendations will enable businesses to select optimal instance types for specific workload categories. Adjust cluster layouts to fit anticipated job needs. Let rather affordable, totally automated instance construction take place to minimize hand-off. These advances increase the dependability and efficiency of E-learning systems as well as help to reduce expenses.

2.5 Case Study: Improving Systems of Financial Services Organization

A global financial services company using E-MR for transaction analytics found increased cloud costs because of poor instance choice. Using a custom API for smart instance choosing, they achieved:

Changing from general-purpose to workload-specific cases will help to reduce EC2 costs by forty percent. Better performance of tasks by matching processing demand with computational resources. Better use of spot events lowers expenses without sacrificing dependability. This situation shows how smart decisions and automation may improve cost effectiveness without compromising performance.

3. Smart Job Scheduling to Cut Empty Costs

Cloud inefficiency is strongly influenced by idle cluster duration. Custom APIs help to improve scheduling efficiency by: Grouping related chores to cut waste of idle resources and improve computer performance. Many organizations do different kinds of work that require similar resource plans, however at different times. Strategic consolidation of these activities into a single execution window helps businesses maximize cluster usage and minimize the need of running more instances than absolutely necessary. Suspending clusters automatically in low demand times to reduce unnecessary running costs Many organizations run ongoing operations, but not all processing jobs call for constant availability. By using previous usage trends, a unique API may determine ideal timings for pausing and restarting clusters, so ensuring computational resources run exactly as required. This approach works particularly well in environments where daily, monthly, or seasonal patterns of workloads follow predictable lines.Using dependability-aware scheduling to ensure jobs run consecutively without overlapping resource contention. Some tasks depend on one another such that one should finish one before starting another. Custom APIs can assure effective execution by interacting with workflow orchestration systems such as Apache Airflow or AWS Step Functions, therefore reducing wasteful resource allocation and minimizing idle computation time.

Moreover, the focus of high-impact chores might be included into scheduling techniques. Some jobs are more critical than others; by dynamically allocating priorities, organizations can efficiently distribute computational resources, hence reducing time-to-insight and needless delays. Scheduling efficiency is much enhanced by using serverless and containerized systems such as Kubernetes in concert with EMR. By reducing the need of starting whole virtual instances, containers help to enable rapider job execution. By dynamically assigning containerized workloads to nodes with accessible resources, Kubernetes improves performance and cost economy. By combining customised APIs with Kubernetes-native scheduling systems, organizations may automate task completion thereby guaranteeing best use of resources and cost economy.

Furthermore, integration of artificial intelligence-driven task management and predictive analytics would improve scheduling choices. Machine learning models serve to enable proactive modifications in task scheduling by means of analysis of historical performance of workload distribution and predicted demand. This guarantees that, should needed, sufficient computational resources are accessible and helps organizations to plan ahead and reduce downtime. Automated failure recovery is an essential component of smart job scheduling. Factors include node crashes, memory overruns, or transient data discrepancies cause job failures in cloud-based big data processing rather often. Custom APIs can include automated retry systems that continue failed processes depending on set parameters, therefore reducing the risk of costly reprocessing brought on by extended downtime. Eventually, cost-sensitive scheduling rules could be carried out using APIs that improve scheduling in line with changes in cloud pricing. Demand affects AWS spot instance pricing; therefore, a scheduling engine including real-time pricing data can allocate workloads to the most cost-effective processing nodes at any one instant. This technology guarantees great efficiency while preserving performance and helps to lower costs. Using these cutting-edge job scheduling techniques will help organizations reach best availability, financial savings, and improved EMS workload efficiency. Organizations using proactive strategies will be well suited to maximize their EHR costs and control expenses as cloud computing charges rise.

4. Notifications and Instant Cost Tracking

Usually, organizations only find overly high expenses when they get their monthly AWS invoice. One uses a proactive approach by:

4.1 Real-Time Cost Tracking Tools

Customized APIs that help organizations better grasp EMS expenses by:

- Closely and holistically monitor EMR use and expenses.
- Examine spending trends to find locations with overly high outlay.
- Review expenditure patterns in relation to workload efficiency to offer reasonable cost control guidance.

Custom APIs allow organizations to make rapid changes instead of waiting for monthly billing cycles by aggregating real-time data from AWS Cost Explorer, CloudWatch, and other monitoring tools.

4.2 Start Notifications Automated

Automated notifications are what help to lower cost overruns. organizations can create alarms for:

- Alert teams when they spend more than is permitted.
- Based on prior performance, suggest tweaks aimed to save money.
- Start automatically maximizing operations including cluster decommissioning or instance size reduction.

By means of technologies such as AWS Lambda, businesses can automatically react upon alarm activation, therefore guaranteeing rapid action to reduce unwarranted expenses.

4.3 Drafting Inspired Cost Reports Informed judgments

These are made by organizations using practically detailed cost information. These documents must address: per job, cluster, and instance type cost analysis.

- Study of unused resources and best times of usage.
- Guideline on job scheduling and enhancement of resource allocation
- organizations can go from reactive cost monitoring to proactive cost control by means of AI-driven analytics, therefore ensuring that EMS workloads always remain very affordable.

4.4 Predictive Analytics Forecasting Future Spending

- Predictive analytics allows businesses to forecast future cloud costs and direct financial decisions.
- Essential skills include: analysis of trends in electronic medical records arising from past use patterns.
- Evaluate economic effects of various workload configurations using scenario modeling.
- Recommendations driven by artificial intelligence aimed to enhance instance choosing and cluster scaling.
- Combining predictive analytics with E-learning enables businesses to dynamically regulate expenses and avoid possible budget overruns.

5. Cutting Transfer and Storage Expenses

Major cost considerations in an EMS system are data movement and storage. These guidelines could enable organizations to aid to control expenditure:

5.1 Optimal Data Storage Management

- Collecting a lot of data without appropriate control could lead to unjustified costs.
- Automatically move seldom accessed data to more affordable storage classes using AWS S3 Intelligent-Tiering.
- Activate compression and reduce data storage needs by using columnar storage systems including Parquet or ORC.
- Create data lifecycle rules to delete or archive outdated, no-needed data.
- Automating storage optimization enables firms to significantly reduce redundant storage expenses.

5.2 Minimizing Transmission Expenses:

Inter-regional and inter-service data transfers may ultimately result in penalties. Strategies for fiscal management encompass:

Processing data within the same AWS region reduces cross-region data transfers where appropriate. A well designed data transfer method ensures that EMR workloads are run at maximum efficiency and are affordable.

5.3 Data Deduplication Execution

Redundant data can aggravate computing costs as well as storage ones. Businesses are able to:

- Remove extra information before processing.
- Use deduplication techniques—including hash-based comparisons—to cut unnecessary storage use.
- Use incremental data processing methods to stop unprocessed datasets from being reworked.
- Reducing data redundancy actively helps organizations to improve their EMS-related costs.

6. Conclusion

Improving EMR workloads calls for a change from set configurations to smart automation. Custom APIs provide a tailored approach for:

- By use of dynamic resource modification, reduce cloud waste.
- Real-time monitoring helps to improve cost transparency.
- Plan strategically to maximize output.

These concepts will enable organizations to keep the dependability and efficiency of their E-learning systems intact and cut a lot of costs. Adopting proactive cost control and automation would help to maximize cloud expenditures on the developing big data landscape. Businesses applying artificial intelligence, machine learning, serverless technologies, cost-optimizing strategies will have a competitive edge in cloud cost control.

The Need for Preventive Cost Control

Cloud computing costs can quickly rise without enough regulation and real-time monitoring. By means of thorough capacity for forecasting, monitoring, and control of expenses, organizations may guarantee optimal availability and performance while making financially wise judgments.

Intelligent cost optimization substitutes real-time data analytics, machine learning, and automation—instead of more conventional cost-cutting techniques. By choosing a proactive rather than a reactive mentality, organizations can avoid unexpected cost rises, increase operational efficiency, and keep financial predictability.

7. References

1. Mishra, Sakti. *Simplify Big Data Analytics with Amazon EMR: A beginner's guide to learning and implementing Amazon EMR for building data analytics solutions*. Packt Publishing Ltd, 2022.
2. Cohen, Raphael Y., and Vesela P. Kovacheva. "A methodology for a scalable, collaborative, and resource-efficient platform, MERLIN, to facilitate healthcare AI research." *IEEE journal of biomedical and health informatics* 27.6 (2023): 3014-3025.
3. Kupunarapu, Sujith Kumar. "AI-Enhanced Rail Network Optimization: Dynamic Route Planning and Traffic Flow Management." *International Journal of Science And Engineering* 7.3 (2021): 87-95.
4. Ta-Shma, Paula, et al. "An ingestion and analytics architecture for iot applied to smart city use cases." *IEEE Internet of Things Journal* 5.2 (2017): 765-774.
5. Mehdi Syed, Ali Asghar, and Erik Anazagasty. "AI-Driven Infrastructure Automation: Leveraging AI and ML for Self-Healing and Auto-Scaling Cloud Environments". *International Journal of Artificial Intelligence, Data Science, and Machine Learning*, vol. 5, no. 1, Mar. 2024, pp. 32-43
6. Kiran, Mariam, et al. "Lambda architecture for cost-effective batch and speed big data processing." *2015 IEEE international conference on big data (big data)*. IEEE, 2015.
7. Kaur, Jagreet, and Kulwinder Singh Mann. "AI based healthcare platform for real time, predictive and prescriptive analytics using reactive programming." *Journal of Physics: Conference Series*. Vol. 933. No. 1. IOP Publishing, 2017.
8. Kupunarapu, Sujith Kumar. "Data Fusion and Real-Time Analytics: Elevating Signal Integrity and Rail System Resilience." *International Journal of Science And Engineering* 9.1 (2023): 53-61.
9. Vasanta Kumar Tarra. "Policyholder Retention and Churn Prediction". *JOURNAL OF RECENT TRENDS IN COMPUTER SCIENCE AND ENGINEERING (JRTCSE)*, vol. 10, no. 1, May 2022, pp. 89-103
10. Mehdi Syed, Ali Asghar, and Erik Anazagasty. "Ansible Vs. Terraform: A Comparative Study on Infrastructure As Code (IaC) Efficiency in Enterprise IT". *International Journal of Emerging Trends in Computer Science and Information Technology*, vol. 4, no. 2, June 2023, pp. 37-48
11. Evans, R. Scott. "Electronic health records: then, now, and in the future." *Yearbook of medical informatics* 25.S 01 (2016): S48-S61.
12. Kupunarapu, Sujith Kumar. "AI-Driven Crew Scheduling and Workforce Management for Improved Railroad Efficiency." *International Journal of Science And Engineering* 8.3 (2022): 30-37.
13. Payne, Thomas H., et al. "Report of the AMIA EHR-2020 Task Force on the status and future direction of EHRs." *Journal of the American Medical Informatics Association* 22.5 (2015): 1102-1110.
14. Sangeeta Anand, and Sumeet Sharma. "Automating ETL Pipelines for Real-Time Eligibility Verification in Health Insurance". *Essex Journal of AI Ethics and Responsible Innovation*, vol. 1, Mar. 2021, pp. 129-50
15. Kumar, Yogesh, Surabhi Kaul, and Yu-Chen Hu. "Machine learning for energy-resource allocation, workflow scheduling and live migration in cloud computing: State-of-the-art survey." *Sustainable Computing: Informatics and Systems* 36 (2022): 100780.
16. Ugajin, Atsushi. "Automation in hospitals and health care." *Springer Handbook of Automation*. Cham: Springer International Publishing, 2023. 1209-1233.
17. Chaganti, Krishna. "Adversarial Attacks on AI-driven Cybersecurity Systems: A Taxonomy and Defense Strategies." *Authorea Preprints*.
18. Vasanta Kumar Tarra, and Arun Kumar Mittapelly. "Data Privacy and Compliance in AI-Powered CRM Systems: Ensuring GDPR, CCPA, and Other Regulations Are Met While Leveraging AI in Salesforce". *Essex Journal of AI Ethics and Responsible Innovation*, vol. 4, Mar. 2024, pp. 102-28
19. Anand, Sangeeta. "Designing Event-Driven Data Pipelines for Monitoring CHIP Eligibility in Real-Time". *International Journal of Emerging Research in Engineering and Technology*, vol. 4, no. 3, Oct. 2023, pp. 17-26
20. Deyhim, Parviz. *Best practices for amazon emr*. Technical report, Amazon Web Services Inc, 2013.
21. Sangeeta Anand, and Sumeet Sharma. "Big Data Security Challenges in Government-Sponsored Health Programs: A Case Study of CHIP". *American Journal of Data Science and Artificial Intelligence Innovations*, vol. 1, Apr. 2021, pp. 327-49
22. Mathew, Sajee, and J. Varia. "Overview of amazon web services." *Amazon Whitepapers* 105.1 (2014): 22.
23. Vasanta Kumar Tarra, and Arun Kumar Mittapelly. "Predictive Analytics for Risk Assessment & Underwriting". *JOURNAL OF RECENT TRENDS IN COMPUTER SCIENCE AND ENGINEERING (JRTCSE)*, vol. 10, no. 2, Oct. 2022, pp. 51-70
24. Kodete, Chandra Shikhi, et al. "Robust Heart Disease Prediction: A Hybrid Approach to Feature Selection and Model Building." *2024 4th International Conference on Ubiquitous Computing and Intelligent Information Systems (ICUIS)*. IEEE, 2024.
25. Sangaraju, Varun Varma, and Senthilkumar Rajagopal. "Applications of Computational Models in OCD." *Nutrition and Obsessive-Compulsive Disorder*. CRC Press 26-35.
26. Anand, Sangeeta. "Quantum Computing for Large-Scale Healthcare Data Processing: Potential and Challenges". *International Journal of Emerging Trends in Computer Science and Information Technology*, vol. 4, no. 4, Dec. 2023, pp. 49-59
27. Chaganti, Krishna C. "Advancing AI-Driven Threat Detection in IoT Ecosystems: Addressing Scalability, Resource Constraints, and Real-Time Adaptability."

28. Sangaraju, Varun Varma. "Optimizing Enterprise Growth with Salesforce: A Scalable Approach to Cloud-Based Project Management." *International Journal of Science And Engineering* 8.2 (2022): 40-48.
29. Mahammad, Abeda Begum, and Rajeev Kumar. "Scalable and security framework to secure and maintain healthcare data using blockchain technology." *2023 International Conference on Computational Intelligence and Sustainable Engineering Solutions (CISES)*. IEEE, 2023.
30. Sreedhar, C., and Varun Verma Sangaraju. "A Survey On Security Issues In Routing In MANETS." *International Journal of Computer Organization Trends* 3.9 (2013): 399-406.
31. Vasanta Kumar Tarra, and Arun Kumar Mittapelly. "Future of AI & Blockchain in Insurance CRM". *JOURNAL OF RECENT TRENDS IN COMPUTER SCIENCE AND ENGINEERING (JRTCSE)*, vol. 10, no. 1, Mar. 2022, pp. 60-77
32. Kupunarapu, Sujith Kumar. "AI-Enabled Remote Monitoring and Telemedicine: Redefining Patient Engagement and Care Delivery." *International Journal of Science And Engineering* 2.4 (2016): 41-48.4
33. Mehdi Syed, Ali Asghar. "Zero Trust Security in Hybrid Cloud Environments: Implementing and Evaluating Zero Trust Architectures in AWS and On-Premise Data Centers". *International Journal of Emerging Trends in Computer Science and Information Technology*, vol. 5, no. 2, Mar. 2024, pp. 42-52
34. Chaganti, Krishna C. "Leveraging Generative AI for Proactive Threat Intelligence: Opportunities and Risks." *Authorea Preprints*.
35. Kulkarni, Amol. "Amazon Redshift: Performance Tuning and Optimization." *International Journal of Computer Trends and Technology* 71.2 (2023): 40-44.
36. Mehdi Syed, Ali Asghar. "Hyperconverged Infrastructure (HCI) for Enterprise Data Centers: Performance and Scalability Analysis". *International Journal of AI, BigData, Computational and Management Studies*, vol. 4, no. 4, Dec. 2023, pp. 29-38
37. Chaganti, Krishna Chiatanya. "Securing Enterprise Java Applications: A Comprehensive Approach." *International Journal of Science And Engineering* 10.2 (2024): 18-27.
38. Mungoli, Neelesh. "Scalable, distributed AI frameworks: leveraging cloud computing for enhanced deep learning performance and efficiency." *arXiv preprint arXiv:2304.13738* (2023).
39. Yang, Zongheng, et al. "{SkyPilot}: An intercloud broker for sky computing." *20th USENIX Symposium on Networked Systems Design and Implementation (NSDI 23)*. 2023.