IJMRSETM

e-ISSN: 2395 - 7639

INTERNATIONAL JOURNAL OF MULTIDISCIPLINARY RESEARCH

IN SCIENCE, ENGINEERING, TECHNOLOGY AND MANAGEMENT

INTERNATIONAL STANDARD SERIAL NUMBER INDIA

ISSN: 2395-7639

International Journal of Multidisciplinary Research in Science, Engineering, Technology & Management (IJMRSETM)

(A Monthly, Peer Reviewed Online Journal)

Visit: www.ijmrsetm.com

Volume 5, Issue 11, November 2018

Security Analysis of Hadoop-Based Big Data Systems: Threats and Strategies

Harsha Vardhan Reddy Goli

Software Developer, Innovative Intelligent Solutions LLC, Tennessee, USA

ABSTRACT: The rapid growth of big data has led to the widespread adoption of distributed computing frameworks such as Apache Hadoop, which offers scalable and efficient processing of massive datasets. While Hadoop has revolutionized data analytics across industries, its distributed and complex architecture introduces significant security vulnerabilities. As organizations increasingly rely on Hadoop-based infrastructures for sensitive data processing, ensuring the confidentiality, integrity, and availability of these systems becomes paramount.

This paper presents a comprehensive security analysis of Hadoop-based big data environments, focusing on the identification of prevalent threats and the evaluation of current mitigation strategies. We begin by outlining the fundamental components of the Hadoop ecosystem, including HDFS, MapReduce, and YARN, followed by an examination of intrinsic and external security risks such as weak authentication, data leakage, denial-of-service (DoS) attacks, and inadequate access controls.

Using a threat modeling approach and real-world case references, we categorize the most critical vulnerabilities and assess the effectiveness of native and third-party security mechanisms, including Kerberos authentication, Apache Ranger, Apache Knox, and encryption frameworks. Our analysis also highlights the trade-offs between system performance and security overhead.
The study contributes to the field by synthesizing existing research, evaluating practical defense strategies, and identifying gaps in current security implementations. We conclude with recommendations for enhancing Hadoop security through integrated policy enforcement, real-time monitoring, and future research directions in privacy-preserving analytics and machine learning-based threat detection.

KEYWORDS : Big Data, Hadoop Security, Data Privacy, Threat Analysis, Cybersecurity Strategies, Distributed Systems

I. INTRODUCTION

In the digital era, organizations across sectors are generating and collecting data at an unprecedented scale. This surge has given rise to the concept of **big data**, characterized by its volume, velocity, and variety. To derive meaningful insights from these vast datasets, enterprises rely on distributed computing frameworks capable of processing data efficiently and cost-effectively. Among these, **Apache Hadoop** has emerged as a leading open-source platform, widely adopted for its scalability, fault tolerance, and support for parallel processing.

Large datasets may be stored and analyzed across clusters of commodity hardware thanks to Hadoop's ecosystem, which includes elements like MapReduce, Yet Another Resource Negotiator (YARN), and the Hadoop Distributed File System (HDFS). Hadoop's extensibility and architectural flexibility have made it a vital component of contemporary data infrastructure, but they also present difficult security issues. Hadoop's distributed and modular architecture makes it vulnerable to a variety of security risks, such as insider assaults, denial-of-service (DoS) attacks, and unauthorized data access, in contrast to conventional centralized systems.

The motivation for this paper stems from the growing reliance on Hadoop for mission-critical and sensitive data operations, which demands a robust understanding of its **security posture**. Despite the availability of several security tools and configurations, many deployments remain **vulnerable** due to misconfigurations, lack of awareness, or inadequate integration of security features.

ISSN: 2395-7639

International Journal of Multidisciplinary Research in Science, Engineering, Technology & Management (IJMRSETM)

(A Monthly, Peer Reviewed Online Journal)

Visit: www.ijmrsetm.com

Volume 5, Issue 11, November 2018

This paper aims to **analyze the security vulnerabilities** inherent in Hadoop-based big data systems, identify the **most prominent threats**, and evaluate the effectiveness of existing **countermeasures and frameworks**. The scope encompasses both native Hadoop mechanisms and third-party solutions, with a focus on authentication, access control, data protection, and threat detection.

II. BACKGROUND AND RELATED WORK

2.1 Overview of Hadoop Architecture

A popular open-source platform for processing massive datasets in a distributed and scalable manner is Apache Hadoop. It consists of multiple essential elements:

- Large files are divided into blocks and dispersed among cluster nodes by the Hadoop Distributed File System (HDFS), a distributed storage layer. With a NameNode handling metadata and DataNodes holding real data blocks, it uses a master-slave architecture.
- **MapReduce**: A programming model and processing engine for parallel computation of large data volumes. It consists of two main phases—**Map**, which processes and filters data, and **Reduce**, which aggregates the intermediate results.
- **Yet Another Resource Negotiator (YARN):** Serves as the layer for work scheduling and resource management. On a Hadoop cluster, it permits the simultaneous operation of several data processing engines (such as MapReduce and Apache Spark).
- **Additional Ecosystem Tools**: Apache Hive (data warehousing), Pig (data flow scripting), HBase (NoSQL storage), and others extend the functionality of the core Hadoop platform.

While Hadoop provides scalable processing, its design originally prioritized performance and scalability over security, making it vulnerable in untrusted environments.

2.2 Previous Research on Big Data Security

A growing body of research addresses the unique security challenges posed by big data systems. Early studies highlighted concerns related to **data confidentiality**, **multi-tenancy**, and **policy enforcement** in distributed environments. Researchers have explored securing data-at-rest and data-in-transit using cryptographic methods, and have developed **access control models** tailored to distributed file systems. Notably, work by Gantz & Reinsel (IDC, 2011) and Zikopoulos et al. (2012) emphasized the importance of embedding security into big data platforms from the ground up.

2.3 Known Vulnerabilities in Hadoop Systems

Several vulnerabilities in Hadoop have been identified over the years, including:

- **Weak default security settings** that allow unauthenticated access to core services.
- **Lack of encryption** for RPC and HTTP communications within HDFS and YARN.
- **Insecure data exposure** through web interfaces and REST APIs.
- **Improper configuration of Kerberos** leading to token theft or bypass.
- **Absence of fine-grained access control** and audit capabilities in early versions.

These weaknesses make Hadoop clusters attractive targets for attackers seeking to exfiltrate sensitive information or disrupt data processing pipelines.

2.4 Literature Survey on Threat Models and Security Mechanisms

Threat modeling for big data platforms typically includes **STRIDE** (Spoofing, Tampering, Repudiation, Information disclosure, Denial of service, Elevation of privilege) and **CIA triad** (Confidentiality, Integrity, Availability). Tools like **Apache Ranger**, **Sentry**, and **Knox** have been proposed and evaluated in the literature for their capabilities to enforce authentication, authorization, auditing, and gateway-level protection.

ISSN: 2395-7639

International Journal of Multidisciplinary Research in Science, Engineering, Technology & Management (IJMRSETM)

(A Monthly, Peer Reviewed Online Journal)

Visit: www.ijmrsetm.com

Volume 5, Issue 11, November 2018

Researchers such as Ghosh et al. (2013) and Zhang et al. (2014) proposed enhanced access control mechanisms, while others focused on privacy-preserving analytics using differential privacy and data masking techniques. However, real-world deployments often lack comprehensive adoption of these solutions.

2.5 Gap Analysis

Despite significant advancements, several gaps persist in securing Hadoop environments:

- Security implementations are often **optional and manually configured**, leading to inconsistencies.
- Lack of **real-time threat detection** and incident response integration within the Hadoop ecosystem.
- Trade-offs between **performance and security overhead** are not well addressed in many studies.
- Limited research on **holistic frameworks** that integrate

III. SECURITY THREAT LANDSCAPE IN HADOOP SYSTEMS

The architecture of Hadoop, while enabling distributed processing and scalable storage, presents multiple attack surfaces. Due to its modular design and reliance on various open-source components, Hadoop is vulnerable to a wide range of security threats originating both from within and outside the organizational perimeter.

3.1 Internal vs. External Threats

External threats typically originate from attackers outside the organizational network who exploit vulnerabilities to gain unauthorized access, disrupt operations, or steal data. These include techniques such as denial-of-service (DoS) attacks, network eavesdropping, and brute-force attempts on unsecured interfaces.

Internal threats, on the other hand, come from authorized users or insiders who misuse their access, intentionally or unintentionally. Insider threats are especially dangerous because of the privileged access insiders may possess and the lack of comprehensive monitoring in many Hadoop deployments.

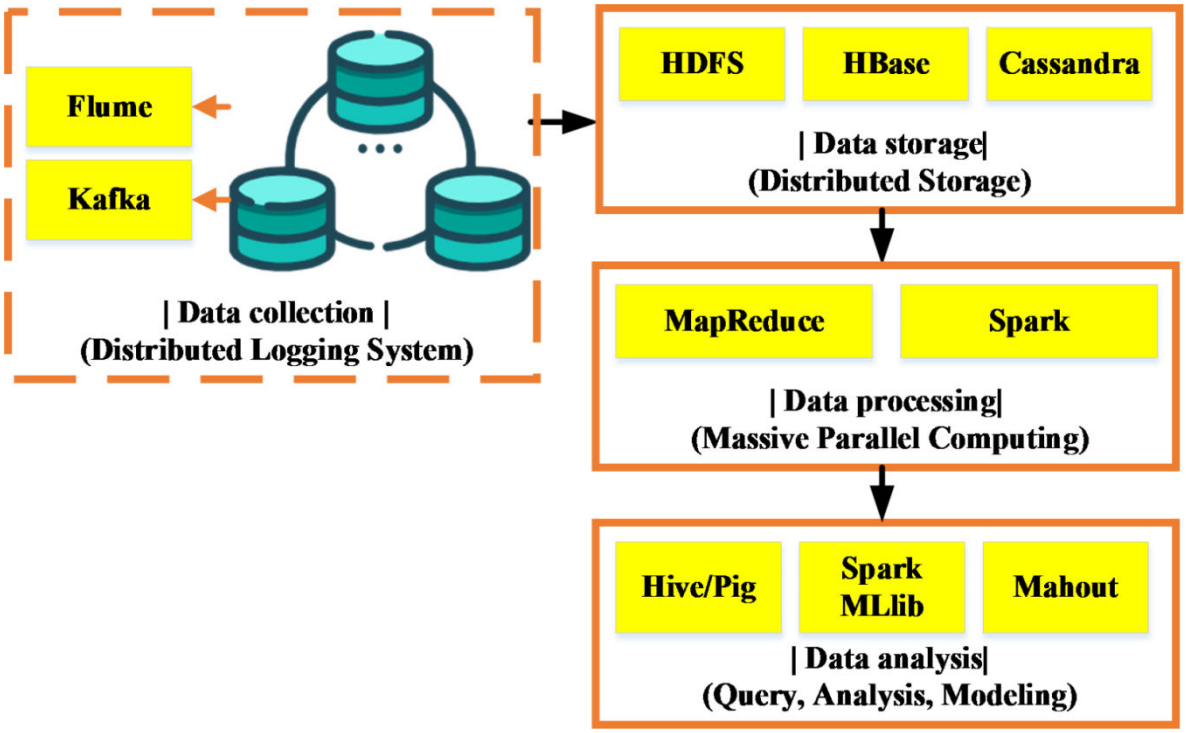

ISSN: 2395-7639

International Journal of Multidisciplinary Research in Science, Engineering, Technology & Management (IJMRSETM)

(A Monthly, Peer Reviewed Online Journal)

Visit: www.ijmrsetm.com

Volume 5, Issue 11, November 2018

3.2 Classification of Threats

3.2.1 Data Confidentiality Breaches

Data confidentiality refers to the protection of sensitive data from unauthorized disclosure. In Hadoop systems, breaches often occur due to:

- Unencrypted data at rest or in transit.
- Misconfigured HDFS permissions.
- Insecure APIs and web interfaces.
- Lack of network isolation between nodes.

Attackers can sniff traffic between nodes or exploit exposed services to access confidential data, such as customer records or financial logs.

3.2.2 Integrity Threats

Data integrity ensures that information is not altered in an unauthorized manner. Threats to integrity in Hadoop environments may include:

- Unauthorized modification of datasets or log files.
- Tampering with MapReduce jobs.
- Malware injected through third-party libraries.

Such attacks can cause faulty analytics, corrupt machine learning models, or lead to misinformed business decisions.

3.2.3 Availability Issues (e.g., DoS Attacks)

Availability ensures uninterrupted access to data and services. Hadoop is susceptible to availability-based attacks such as:

- Distributed Denial-of-Service (DDoS) attacks targeting NameNode or ResourceManager services.
- Exhaustion of cluster resources via malformed or overly intensive MapReduce jobs.
- Exploiting job queues to flood the system and block legitimate processing tasks.

Given the central role of the NameNode in HDFS, its failure can render the entire cluster unusable.

3.2.4 Insider Threats

Employees or privileged users may intentionally leak data, bypass security policies, or make configuration changes that weaken the system. Common insider threats include:

- Sharing authentication tokens or Kerberos tickets.
- Accessing unauthorized directories or jobs.
- Modifying audit logs to hide unauthorized actions.

Because Hadoop lacks native real-time anomaly detection, insider threats can go unnoticed for extended periods.

3.2.5 Weak Authentication and Access Control

Authentication and access control are foundational to secure system operation. In many deployments, these mechanisms are weak due to:

- Improper Kerberos configuration or absence thereof.
- Inconsistent or non-existent role-based access control.
- Use of plaintext passwords in configuration files.
- Open ports and unsecured services.

ISSN: 2395-7639

International Journal of Multidisciplinary Research in Science, Engineering, Technology & Management (IJMRSETM)

(A Monthly, Peer Reviewed Online Journal)

Visit: www.ijmrsetm.com

Volume 5, Issue 11, November 2018

Unauthorized users can exploit these gaps to gain access to sensitive operations or data.

3.3 Case Studies and Real-World Incidents

- In 2017, numerous Hadoop clusters were exposed on the public internet due to misconfigurations, resulting in ransomware attacks that encrypted data and demanded payment in cryptocurrency.
- A security audit of a Hadoop deployment in a Fortune 500 company revealed over 300 DataNodes were accessible without authentication, allowing for unrestricted data download and job submission.
- In a well-documented incident, attackers exploited open web UIs to inject malicious MapReduce jobs that launched cryptomining operations within a corporate Hadoop cluster.

IV. VULNERABILITY ANALYSIS

In Hadoop, several security vulnerabilities can compromise the integrity and confidentiality of data within a cluster. **Insecure communication** is one of the most significant concerns, as Hadoop often relies on **RPC (Remote Procedure Calls)** and **HTTP** protocols for communication between nodes. Without proper encryption, sensitive data transmitted across these channels can be intercepted and manipulated by attackers. This lack of **TLS/SSL encryption** leaves the system open to man-in-the-middle (MITM) attacks, where malicious actors can exploit unsecured data exchanges.

Weak authentication mechanisms, particularly **Kerberos misconfiguration**, are another common vulnerability. Kerberos plays a vital role in authenticating users and services, but if not properly configured, it can lead to **token leaks** or unauthorized access. Inadequate key management or weak token handling can allow attackers to impersonate legitimate users or services, thereby bypassing security controls.

Additionally, **poor authorization mechanisms** in Hadoop can lead to unauthorized access to sensitive data. Misconfigured **HDFS permissions** or improperly set role-based access control (RBAC) settings may allow users or applications to access data or perform actions they shouldn't be able to, potentially exposing confidential information.
Another vulnerability comes from **unpatched components** and **plugin vulnerabilities**. Many Hadoop components, such as **Hive**, **Pig**, or **HBase**, often rely on third-party libraries and plugins. If these components are not regularly updated or patched, they can become easy targets for exploitation due to known security flaws that attackers can leverage to gain unauthorized access or execute arbitrary code.

Finally, **multitenancy and data leakage risks** are a concern in shared environments where multiple users or applications access the same Hadoop cluster. Without proper isolation mechanisms, one tenant's data can leak into another's, either inadvertently or through deliberate attacks. This is particularly dangerous when sensitive data from one department or organization is accessible by others, leading to potential compliance violations and data breaches.
To mitigate these vulnerabilities, organizations must adopt a comprehensive security strategy that includes encryption, robust authentication, strict access control policies, regular updates, and careful management of multitenant environments.

V. SECURITY STRATEGIES AND COUNTERMEASURES

Strong security measures became more and more necessary as Hadoop was embraced by business settings. Hadoop had few built-in security features at first, but as time went on, a number of tactics and defenses were developed to safeguard private information and guarantee legal compliance. Kerberos authentication is the foundation of Hadoop's built-in security, offering a safe means of confirming user identities throughout the cluster. Furthermore, HDFS permissions limit unlawful user actions within the distributed file system by enforcing access control at the file and directory levels.
The Hadoop ecosystem has robust add-on technologies to meet greater enterprise-grade and granular security needs. With the help of Apache Ranger's unified security administration, Hadoop components such as HDFS, Hive, HBase, and others can have fine-grained access control and auditing. While Apache Sentry facilitates role-based authorization for data access in Hive and Impala systems, Apache Knox serves as a gateway, protecting REST APIs and streamlining perimeter security.
In addition to application-level security, network-level safeguards are essential for Hadoop cluster protection. These include virtual private networks (VPNs) to separate Hadoop infrastructure from public networks, firewalls to restrict

ISSN: 2395-7639

International Journal of Multidisciplinary Research in Science, Engineering, Technology & Management (IJMRSETM)

(A Monthly, Peer Reviewed Online Journal)

Visit: www.ijmrsetm.com

Volume 5, Issue 11, November 2018

access to vital ports and services, and the use of TLS/SSL encryption for secure communication between nodes. At the data level, encryption both in transit and at rest guarantees that information is shielded from unwanted access while it is being transmitted over a network and stored on a disk.

Hadoop also supports robust **access control models**, such as **Role-Based Access Control (RBAC)** and Policies based on user roles or certain characteristics, such as data sensitivity or time of access, can be enforced with the use of attribute-based access control, or ABAC. Lastly, monitoring user activity, identifying irregularities, and guaranteeing adherence to data governance rules all depend on auditing and logging systems. These logs, which are an essential component of the overall Hadoop security architecture, offer insight into access patterns, alterations, and possible security issues. When combined, these tactics offer a thorough foundation for safeguarding Hadoop systems in intricate, data-intensive settings.

VI. COMPARATIVE EVALUATION OF SECURITY SOLUTIONS

Security Solution	Effectiveness	Overhead	Scalability	Usability	Implementation Complexity
Kerberos Authentication	High	Moderate	High	Moderate	Moderate
HDFS Permissions	High	Low	High	High	Low
Apache Ranger	High	Moderate to High	High	Moderate	High
Apache Knox	High	Low to Moderate	High	High	Moderate
Apache Sentry	High	Low to Moderate	High	Moderate	Moderate
TLS/SSL (Network Encryption)	High	Moderate	High	High	Moderate

Summary of Trade-offs and Considerations:

- **Effectiveness:** Tools like **Kerberos**, **Apache Ranger**, and **Apache Sentry** provide strong security measures with centralized control over access, proving highly effective in managing large datasets and user identities.
- **Overhead:** Solutions like **Kerberos** and **Apache Ranger** introduce some overhead, especially in high-traffic environments, but the trade-off is generally considered worthwhile for strong security. Tools like **HDFS Permissions** and **TLS/SSL** have lower overhead but may introduce slight performance impacts.
- **Scalability:** Most security solutions scale well with Hadoop clusters, but solutions like **Apache Ranger**, **Apache Sentry**, and **Kerberos** are specifically designed to work efficiently with large, distributed environments.
- **Usability:** Tools like **HDFS Permissions** and **TLS/SSL** are relatively easy to set up, while others, such as **Apache Ranger** and **Kerberos**, require a deeper understanding of the Hadoop ecosystem and may involve more complex configuration.
- **Security-Performance Trade-offs:** High levels of security, such as with **encryption** and **Kerberos authentication**, often introduce performance penalties, especially in high-throughput environments. Solutions

ISSN: 2395-7639

International Journal of Multidisciplinary Research in Science, Engineering, Technology & Management (IJMRSETM)

(A Monthly, Peer Reviewed Online Journal)

Visit: www.ijmrsetm.com

Volume 5, Issue 11, November 2018

like **Apache Knox** and **Apache Ranger** balance security and performance, although the need for fine-grained access control may occasionally slow down data access.

- **Implementation Complexity:** Tools like **Apache Knox** and **TLS/SSL** have moderate to high implementation complexity due to integration and network configuration requirements. On the other hand, **HDFS Permissions** and **RBAC** provide relatively simple, yet effective, solutions that are easier to implement.

VII. CHALLENGES AND OPEN RESEARCH ISSUES

Despite the extensive security strategies and countermeasures available for Hadoop-based systems, several challenges remain that hinder the adoption of robust security frameworks in production environments. These challenges are compounded by the dynamic nature of big data workloads, the evolving threat landscape, and the growing complexity of Hadoop deployments.

7.1 Performance Impact of Security Layers

One of the primary challenges in securing Hadoop systems is the **performance overhead** introduced by security mechanisms. Encryption, authentication, and fine-grained access control often come at the cost of **increased latency** and **resource consumption**, particularly in large-scale data processing. For example, the implementation of Kerberos authentication and data encryption can significantly slow down data access and job execution, leading to performance degradation. While these security features are essential for protecting sensitive data, striking the right balance between security and system performance remains a critical issue for enterprise adoption.

7.2 Dynamic Policy Enforcement

Hadoop environments often involve dynamic workloads and evolving data access requirements. In such settings, enforcing **security policies** in real-time can be challenging. The difficulty lies in continuously monitoring access patterns, detecting anomalies, and adjusting policies to accommodate new users, services, or business requirements. The absence of automated, **adaptive security policy enforcement** systems that can respond to changes in workload and user behavior leaves Hadoop systems vulnerable to misconfigurations and policy lapses.

7.3 Security in Multi-Cloud Hadoop Deployments

The **multi-cloud** paradigm, where organizations deploy Hadoop clusters across multiple cloud providers, presents additional security complexities. Ensuring consistent security policies, data integrity, and access controls across diverse cloud environments introduces **interoperability issues**. Different cloud providers offer varying levels of security, and managing sensitive data across these platforms without exposing it to unauthorized access requires careful planning. Additionally, the **lack of standardized security frameworks** for multi-cloud deployments complicates the task of securing Hadoop clusters in a distributed, hybrid infrastructure.

7.4 Machine Learning for Anomaly Detection

While traditional security mechanisms like Kerberos and firewalls are effective at enforcing access control, they often fail to detect **new or sophisticated attacks** in real-time. **Machine learning** offers significant promise for anomaly detection, enabling the identification of unusual patterns or behaviors indicative of potential security breaches. However, **training effective models** for anomaly detection in Hadoop environments is challenging due to the massive scale of data and the variability of workloads. Moreover, incorporating machine learning techniques without introducing performance bottlenecks requires further research in the optimization of algorithms and models.

7.5 Privacy-Preserving Analytics

Data analytics privacy is becoming more and more crucial as businesses continue to use data to inform their decisions. Techniques that protect privacy, such homomorphic encryption and differential privacy, are becoming more popular as viable ways to reduce the possibility of data leaks during analytics. While homomorphic encryption enables calculations on encrypted data without disclosing sensitive information, differential privacy guarantees that the inclusion of any one data point does not materially alter the results of data analytics. However, when used in large-scale Hadoop systems, both methods have difficulties with complexity and processing overhead. Future research must focus on finding effective ways to incorporate these privacy-preserving strategies into Hadoop clusters.

ISSN: 2395-7639

International Journal of Multidisciplinary Research in Science, Engineering, Technology & Management (IJMRSETM)

(A Monthly, Peer Reviewed Online Journal)

Visit: www.ijmrsetm.com

Volume 5, Issue 11, November 2018

VIII. CONCLUSION

In conclusion, this paper provides a comprehensive security analysis of Hadoop-based big data systems, examining the key threats, vulnerabilities, and countermeasures that impact their operation. Hadoop's distributed architecture and reliance on open-source components make it inherently susceptible to security risks such as data breaches, integrity threats, and availability issues. However, through the integration of Hadoop-native features like Kerberos authentication and HDFS permissions, along with third-party tools like Apache Ranger, Knox, and Sentry, enterprises can mitigate many of these risks. At the network and data levels, implementing encryption and access control models such as RBAC and ABAC further enhances system security. Despite these protective measures, challenges remain, particularly with performance impacts, dynamic policy enforcement, multi-cloud deployments, and privacy-preserving analytics. The growing complexity of big data environments calls for continued research into machine learning for anomaly detection and more efficient privacy techniques like differential privacy and homomorphic encryption. For enterprises using Hadoop, securing big data environments requires a layered approach that balances security, performance, and usability. As the Hadoop ecosystem evolves, adopting more automated, adaptable security frameworks and addressing the open research issues will be critical to building resilient and secure big data infrastructures.

REFERENCES

1. Mahmoud, H., Hegazy, A., & Khafagy, M. H. (2018, February). An approach for big data security based on Hadoop distributed file system. In 2018 International Conference on Innovative Trends in Computer Engineering (ITCE) (pp. 109-114). IEEE.
2. Gupta, M., Patwa, F., & Sandhu, R. (2018, March). An attribute-based access control model for secure big data processing in hadoop ecosystem. In Proceedings of the Third ACM Workshop on Attribute-Based Access Control (pp. 13-24).
3. Singh, U., Solanki, N. K., Varma, M. K., & Sevak, T. (2017, October). A review on big data protection of Hadoop. In 2017 2nd International Conference on Communication and Electronics Systems (ICCES) (pp. 943-950). IEEE.
4. Rathore, M. M., Paul, A., Ahmad, A., Anisetti, M., & Jeon, G. (2017). Hadoop-based intelligent care system (HICS) analytical approach for big data in IoT. ACM Transactions on Internet Technology (TOIT), 18(1), 1-24.
5. Shetty, M. M., & Manjaiah, D. H. (2016, October). Data security in Hadoop distributed file system. In 2016 International Conference on Emerging Technological Trends (ICETT) (pp. 1-5). IEEE.
6. Yao, Q., Tian, Y., Li, P. F., Tian, L. L., Qian, Y. M., & Li, J. S. (2015). Design and development of a medical big data processing system based on Hadoop. Journal of medical systems, 39, 1-11.
7. Jam, M. R., Khanli, L. M., Javan, M. S., & Akbari, M. K. (2014, October). A survey on security of Hadoop. In 2014 4th International Conference on Computer and knowledge Engineering (ICCKE) (pp. 716-721). IEEE.
8. Chhabra, G. S., Singh, V., & Singh, M. (2018). Hadoop-based analytic framework for cyber forensics. International Journal of Communication Systems, 31(15), e3772.
9. Dou, Z., Khalil, I., Khreishah, A., & Al-Fuqaha, A. (2017). Robust insider attacks countermeasure for Hadoop: Design and implementation. IEEE Systems Journal, 12(2), 1874-1885.
10. Erraissi, A., Belangour, A., & Tragha, A. (2017). A Comparative Study of Hadoop-based Big Data Architectures. Int. J. Web Appl., 9(4), 129-137.
11. Tripathi, S., Gupta, B., Almomani, A., Mishra, A., & Veluru, S. (2013). Hadoop based defense solution to handle distributed denial of service (ddos) attacks. Journal of Information Security, 4(3), 150-164.
12. Khan, S., Shakil, K. A., & Alam, M. (2018). Cloud-based big data analytics—a survey of current research and future directions. Big Data Analytics: Proceedings of CSI 2015, 595-604.

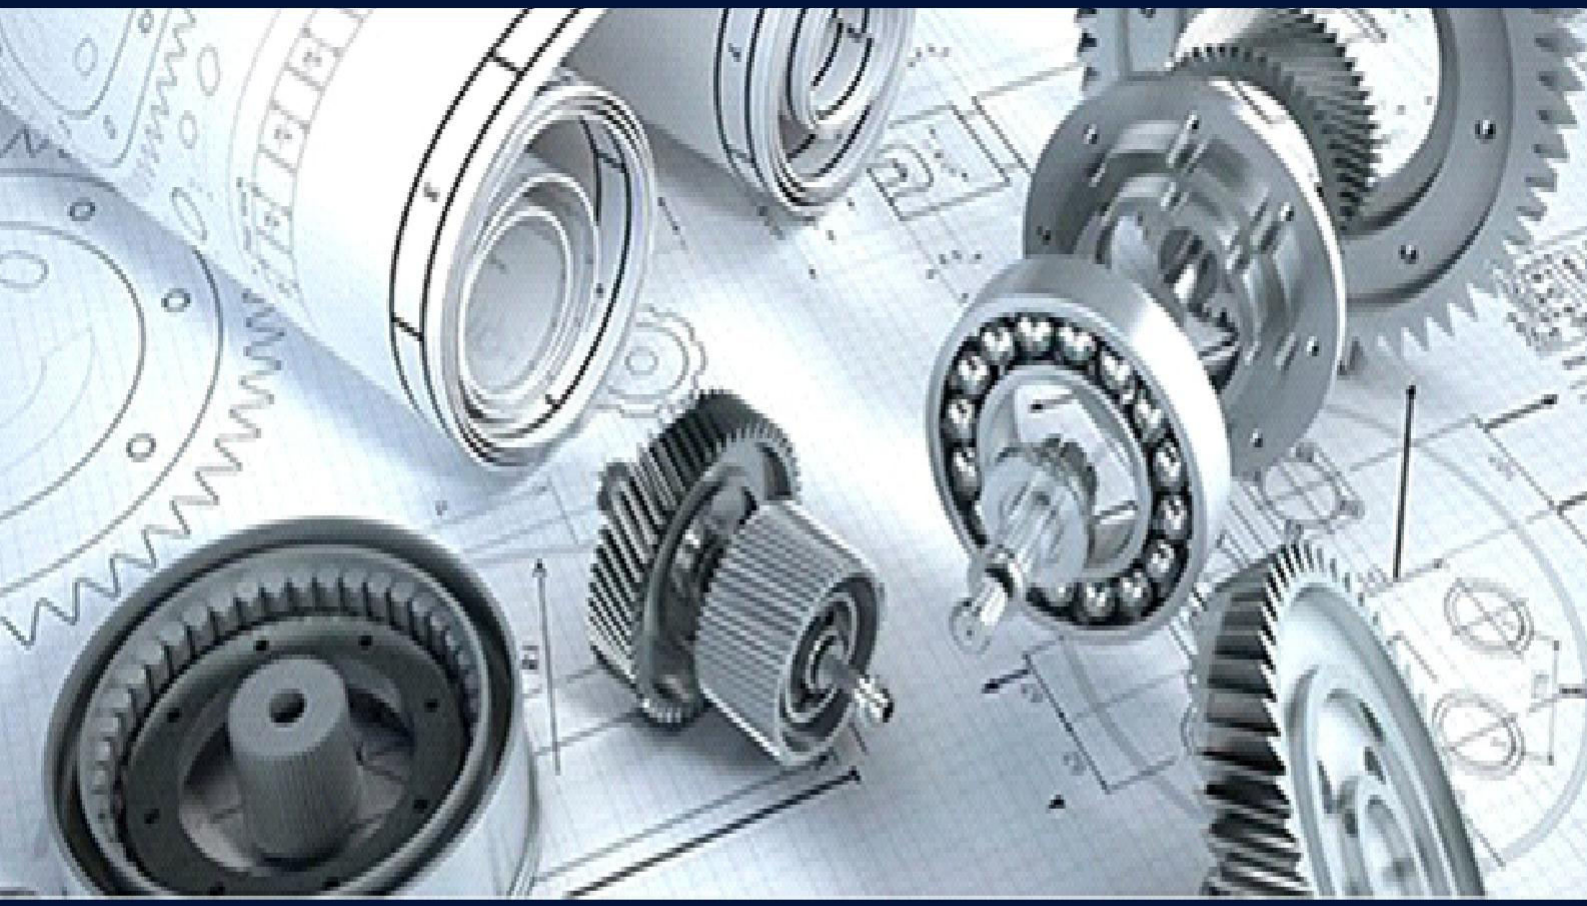

INTERNATIONAL JOURNAL OF MULTIDISCIPLINARY RESEARCH

IN SCIENCE, ENGINEERING, TECHNOLOGY AND MANAGEMENT

+91 99405 72462 +91 63819 07438 ijmrsetm@gmail.com