Nanotechnology Perceptions
ISSN 1660-6795
www.nano-ntp.com

Critical study of AWS Security Tools and Features for Hadoop Deployments: Review and Future Perspectives

Khalil Nabab Pinjari[1], Prasadu Peddi[2], Yogesh Kumar Sharma[3]

[1]Research Scholar, Shri JJT University, Churela, Jhunjhunu, Rajasthan, India, pinjarikhalil.hadoop@gmail.com
[2]Research Guide, Department of Computer Science & Engineering, Shri JJT University, Churella, Jhunjhunu, Rajasthan, India, peddiprasad37@gmail.com
[3]Professor, Department of Computer Science and Engineering, Koneru Lakshmaiah Education Foundation, Green Field, Vaddeswaram, Guntur, Andhra Pradesh, India, dr.sharmayogeshkumar@gmail.com

As organizations increasingly adopt Hadoop for managing and analyzing vast datasets, ensuring robust security for these deployments becomes critical. Amazon Web Services (AWS) provides a comprehensive suite of security tools and features to safeguard Hadoop environments against potential threats. This abstract explores the key AWS security capabilities tailored for Hadoop, including identity and access management (IAM), network security, encryption, and compliance. It delves into services like AWS Key Management Service (KMS) for data encryption, AWS Identity and Access Management for fine-grained control, and Amazon Virtual Private Cloud (VPC) for secure network configurations. Additionally, the role of AWS CloudTrail and Amazon GuardDuty in monitoring and detecting security anomalies is examined. By leveraging AWS-native security features and best practices, organizations can achieve a secure, scalable, and compliant Hadoop deployment, enabling them to harness big data's potential without compromising on safety.
Keywords: Hadoop Security; AWS Security Tools; Big Data Security; Cloud Security; Encryption and Compliance; Future Security Trends.

1. Introduction

In an era defined by data-driven decision-making, the rapid expansion of big data technologies has become pivotal to enterprises and industries worldwide. Hadoop, a widely-adopted open-source framework, enables the efficient storage, processing, and analysis of vast amounts of

data. However, with great power comes great responsibility; the security of such data-rich environments becomes paramount, especially when deployed in a cloud ecosystem like Amazon Web Services (AWS). AWS, a leading cloud services provider, offers robust tools and features to address the intricate security challenges faced by big data implementations.

The convergence of Hadoop's distributed computing capabilities with AWS's scalable and resilient cloud infrastructure has enabled organizations to achieve unprecedented agility and innovation. However, this integration introduces unique security concerns that require meticulous planning and execution. Big data deployments are inherently complex, dealing with sensitive and voluminous data that must be protected against breaches, unauthorized access, and other potential vulnerabilities. Deploying a secure Hadoop framework on AWS necessitates leveraging a comprehensive security strategy encompassing both native AWS features and Hadoop's built-in safeguards.

Figure 1: Domains associated with AWS

This paper explores the foundational principles of deploying a secure AWS security framework tailored to a Hadoop ecosystem. By addressing key aspects such as identity and access management, data encryption, network security, and compliance monitoring, this framework ensures a fortified environment. These principles not only safeguard data integrity and confidentiality but also promote operational excellence and adherence to industry standards.

The explosive growth of big data technologies has revolutionized the way organizations process, analyze, and derive insights from data. With applications ranging from healthcare and finance to social media and e-commerce, big data has become an indispensable tool for modern decision-making. However, this rapid adoption of big data systems has also introduced significant security challenges. The scale, variety, and velocity of big data make it a prime target for cyber threats, necessitating the development of robust security frameworks tailored

to the unique demands of big data environments.

1. Complexity of Big Data Environments

Big data systems are inherently complex, typically involving a wide array of technologies, including distributed storage systems like Hadoop and Spark, data processing frameworks, and real-time analytics platforms. This complexity increases the attack surface, providing multiple entry points for malicious actors. Traditional security mechanisms, which are often designed for smaller, centralized systems, struggle to address the distributed and dynamic nature of big data environments.

For instance, data is often stored and processed across multiple nodes and geographical locations, making it challenging to ensure consistent security policies. Moreover, as big data systems integrate heterogeneous data sources—from structured databases to unstructured data streams—the task of safeguarding data integrity and confidentiality becomes even more daunting. A well-defined security framework can help standardize and streamline security measures across these disparate components, ensuring a unified and proactive approach to threat mitigation.

2. Data Privacy and Regulatory Compliance

The increasing emphasis on data-driven decision-making has amplified concerns about data privacy. With stringent regulations such as the General Data Protection Regulation (GDPR), the California Consumer Privacy Act (CCPA), and other regional data protection laws, organizations must ensure that their big data practices comply with legal requirements. Failure to do so can result in severe financial penalties and reputational damage.

Big data systems often process sensitive information, including personally identifiable information (PII), financial records, and healthcare data. Without a robust security framework, organizations risk unauthorized access, data breaches, and misuse of this sensitive information. An effective security framework incorporates mechanisms for data anonymization, encryption, and access control, ensuring that sensitive data remains protected throughout its lifecycle. Additionally, such a framework can facilitate audit trails and reporting capabilities, which are essential for demonstrating compliance during regulatory assessments.

3. Threats to Data Integrity and Availability

Data integrity and availability are critical for the reliability of big data analytics. Cyber threats such as ransomware attacks, data poisoning, and Distributed Denial of Service (DDoS) attacks can compromise the quality and accessibility of data, leading to flawed analytics and disrupted operations. In big data environments, the cascading effects of such attacks can be particularly devastating, given the interdependence of data pipelines and analytics workflows.

For example, a ransomware attack on a distributed file system can render large volumes of data inaccessible, paralyzing decision-making processes. Similarly, data poisoning—where malicious actors inject false or misleading data—can skew analytics and lead to incorrect conclusions. To counter these threats, a security framework must include real-time monitoring, anomaly detection, and incident response mechanisms. These measures can help identify and mitigate threats before they escalate, ensuring the resilience of big data operations.

4. Challenges in Identity and Access Management (IAM)

Managing identities and access privileges in big data environments is a complex task, given the sheer volume of users, devices, and applications involved. Unauthorized access remains one of the leading causes of data breaches, underscoring the need for robust IAM practices.

A security framework for big data must implement fine-grained access controls, enabling organizations to enforce the principle of least privilege. Role-based access control (RBAC) and attribute-based access control (ABAC) are particularly effective in managing permissions in diverse big data ecosystems. Additionally, the framework should support multi-factor authentication (MFA) and continuous identity verification to enhance security. By incorporating these measures, organizations can minimize the risk of insider threats and unauthorized access to sensitive data.

5. Securing Data in Transit and at Rest

Data in big data systems is constantly in motion, flowing between data sources, processing nodes, and storage systems. This dynamic nature of data movement makes it vulnerable to interception and unauthorized access. Ensuring the security of data both in transit and at rest is therefore a cornerstone of any big data security framework.

Encryption is a fundamental technique for protecting data in transit and at rest. Advanced encryption standards (AES) and secure communication protocols such as Transport Layer Security (TLS) can safeguard data against eavesdropping and tampering. However, encryption alone is not sufficient. The security framework must also address key management challenges, ensuring that encryption keys are stored and managed securely. Furthermore, organizations should adopt data masking and tokenization techniques to add additional layers of protection, particularly for sensitive information.

6. The Role of Artificial Intelligence and Machine Learning

Given the scale and complexity of big data environments, traditional security measures often fall short in detecting and responding to threats in real time. Artificial intelligence (AI) and machine learning (ML) can play a pivotal role in enhancing big data security frameworks. By analyzing vast volumes of data, AI-driven tools can identify patterns indicative of potential security breaches, enabling proactive threat detection and mitigation.

For instance, anomaly detection algorithms can flag unusual activity, such as unauthorized data access or abnormal network traffic, which might indicate a cyber attack. Similarly, ML models can continuously learn and adapt to evolving threat landscapes, providing organizations with a dynamic and intelligent defense mechanism. Integrating AI and ML capabilities into the security framework can significantly enhance its effectiveness and efficiency.

7. The Need for a Holistic Approach

Addressing the security challenges of big data requires a holistic approach that goes beyond technical solutions. A comprehensive security framework must encompass organizational policies, employee training, and regular security audits. By fostering a culture of security awareness, organizations can reduce the risk of human error, which remains a major contributor to data breaches.

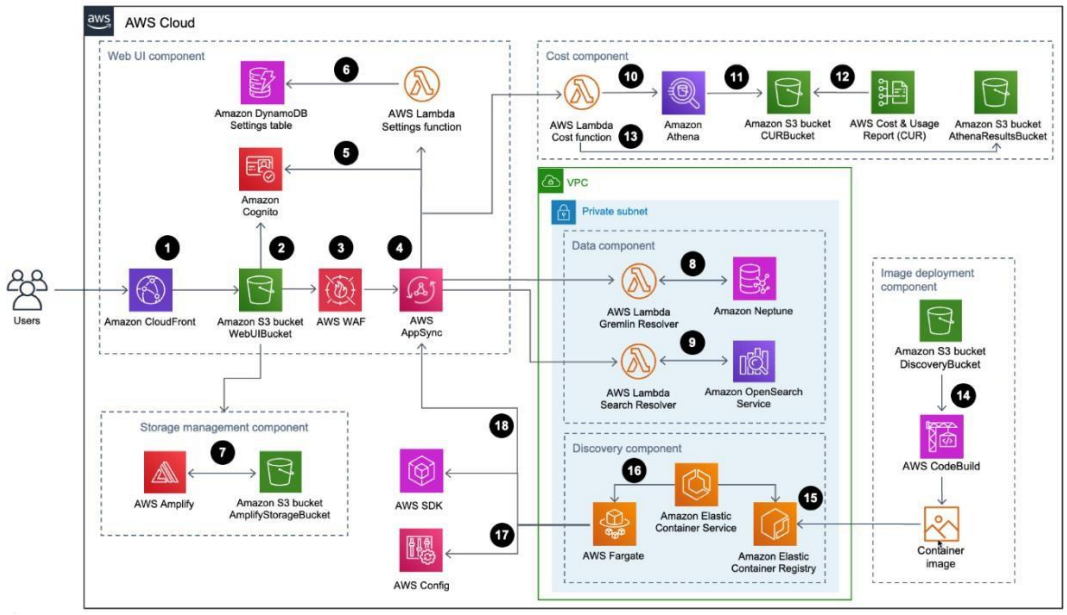

Figure 2 AWS Cloud

Moreover, collaboration between stakeholders—including IT teams, data scientists, and business leaders—is essential for developing and maintaining an effective security framework. This collaborative approach ensures that security measures align with organizational goals and adapt to evolving business needs. The growing reliance on big data technologies demands a proactive approach to security. A well-designed security framework can help organizations navigate the complex security landscape of big data, safeguarding sensitive information, ensuring compliance, and maintaining the integrity and availability of data. By addressing the unique challenges of big data environments, such a framework not only mitigates risks but also enables organizations to harness the full potential of big data with confidence.

The paper is organized as follows: Section 1 provides an introduction to Hadoop deployments and the importance of security in cloud-based environments. Section 2 reviews the current AWS security tools and features relevant to Hadoop, including identity management, encryption, and monitoring. Section 3 critically examines the strengths and limitations of these tools, with a focus on scalability and compliance. Section 4 explores case studies and real-world applications to highlight practical implications. Section 5 discusses future perspectives, emerging trends, and potential enhancements in AWS security for big data ecosystems. Finally, Section 6 concludes the study, summarizing key insights and recommendations.

2. Key Challenges in Securing Hadoop on AWS

Incident response is further complicated by the distributed nature of Hadoop. Isolating compromised nodes or addressing breaches in real-time necessitates automated response mechanisms, which can be challenging to configure and maintain. Additionally, ensuring that incident response plans align with organizational and regulatory requirements is an ongoing

effort. Securing Hadoop on AWS involves addressing a myriad of challenges, from data encryption and access control to compliance and cost management. Each layer of security, whether it pertains to network configurations, user permissions, or vulnerability management, requires meticulous planning and execution. By leveraging AWS-native tools alongside Hadoop's security features, organizations can create a resilient and secure environment. However, the dynamic nature of security threats and the complexity of distributed systems necessitate continuous vigilance, regular audits, and proactive security enhancements.

1. Data Protection and Encryption

Hadoop clusters often store and process vast amounts of sensitive data. On AWS, while services like S3 and EBS provide encryption options, securing data at rest and in transit remains a critical challenge. Organizations must ensure that all Hadoop data nodes and related services implement encryption mechanisms, such as TLS for data in transit and KMS (Key Management Service) for data at rest. Misconfigurations in these encryption setups can lead to data breaches, exposing private or sensitive information.

Additionally, the distributed nature of Hadoop complicates encryption. Data replication across nodes must adhere to encryption standards, ensuring that replicas remain secure. AWS's shared responsibility model places the onus on users to configure these settings correctly, making robust encryption a vital area of focus.

2. Access Control and Identity Management

Effective access control in a Hadoop environment on AWS involves integrating the right identity and access management (IAM) strategies. However, managing roles and permissions for diverse users and applications accessing Hadoop clusters can become cumbersome. The complexity grows when combining Hadoop's internal security features (like Ranger or Knox) with AWS IAM policies.

AWS's temporary credentials and federated access provide scalability and flexibility, but misconfigured roles or overly permissive policies can create vulnerabilities. A significant challenge is ensuring granular access controls for different Hadoop components, such as Hive, HDFS, and YARN, without introducing excessive administrative overhead or breaking compliance requirements.

3. Network Security

The distributed architecture of Hadoop requires constant communication between nodes. In AWS, securing this communication involves setting up Virtual Private Clouds (VPCs), security groups, and network ACLs. Ensuring that only authorized traffic flows between Hadoop nodes, as well as between the cluster and external systems, is essential. However, incorrect configurations can expose Hadoop clusters to unauthorized access or man-in-the-middle attacks.

Public-facing interfaces, like web UIs for Hadoop components, add another layer of risk. Restricting access to such interfaces using AWS security groups and ensuring proper authentication mechanisms like Kerberos or LDAP are in place is crucial. Despite these measures, monitoring and auditing network traffic for potential anomalies remain persistent challenges.

4. Compliance and Regulatory Requirements

Organizations operating in regulated industries must ensure their Hadoop deployments on AWS adhere to standards like GDPR, HIPAA, or PCI DSS. Achieving compliance involves maintaining detailed audit logs, data lineage tracking, and enforcing data access policies.

Hadoop's native logging and auditing capabilities, combined with AWS tools like CloudTrail and CloudWatch, can help meet these requirements. However, integrating these systems to provide comprehensive compliance reporting can be complex. Any misstep in configuring compliance-related settings could result in significant legal or financial repercussions.

5. Security Patching and Vulnerability Management

Managing security patches across a distributed Hadoop cluster is inherently challenging. On AWS, organizations must ensure that instances running Hadoop components are updated promptly to address vulnerabilities. Delays in patching nodes can expose clusters to exploits targeting known weaknesses.

While AWS provides tools like Systems Manager to automate updates, the risk of downtime or incompatibility with existing configurations complicates the process. Furthermore, ensuring that Hadoop's open-source dependencies are updated regularly requires constant vigilance and a deep understanding of the software stack.

6. Insider Threats and Misuse

The threat of insider attacks is particularly concerning for Hadoop clusters on AWS due to the broad access that administrators and developers might have. Misuse of administrative privileges or inadvertent errors in configuring security settings can lead to data leakage or breaches.

Mitigating this risk involves implementing strict role-based access controls (RBAC), regularly auditing access logs, and leveraging AWS services like CloudTrail to monitor and alert on unusual activity. However, balancing ease of use with stringent security measures remains a delicate task.

7. Integration of Native and Third-Party Security Tools

Hadoop's ecosystem includes tools like Ranger for fine-grained access control and Knox for gateway security. Integrating these tools with AWS services like IAM, KMS, and GuardDuty often requires custom configurations and additional expertise.

Third-party tools may offer advanced features but can introduce compatibility and maintenance challenges. For example, configuring an external SIEM (Security Information and Event Management) solution to ingest logs from Hadoop and AWS resources requires careful planning and integration testing.

8. Cost Management of Security Features

AWS provides a range of security services, from encryption to threat detection, but these come with associated costs. Organizations must balance the expense of implementing comprehensive security measures against their budget constraints. Overprovisioning security tools or mismanaging resources can lead to unnecessary expenses.

Optimizing the cost of securing Hadoop clusters requires a strategic approach, including the judicious use of spot instances, cost-effective encryption options, and periodic cost reviews. However, cost-saving measures should not compromise the security of sensitive data or critical operations.

9. Monitoring and Incident Response

Detecting and responding to security incidents in Hadoop clusters hosted on AWS demands robust monitoring and alerting systems. AWS services like GuardDuty and Security Hub can provide insights into potential threats, but correlating these alerts with Hadoop-specific logs requires integration and expertise. The security challenges associated with deploying Hadoop on AWS can be broadly categorized into several areas:

1. Data Privacy and Protection: Hadoop clusters often handle sensitive information such as personal identifiers, financial records, and intellectual property. Unauthorized access to this data can result in significant reputational and financial damage.

2. Access Control: Ensuring that only authorized users and applications have access to the Hadoop ecosystem is critical. Mismanaged permissions can lead to accidental or intentional data breaches.

3. Network Security: The communication between Hadoop nodes and external systems must be secured to prevent eavesdropping, data tampering, or man-in-the-middle attacks.

4. Compliance and Auditability: Organizations must adhere to regulatory requirements such as GDPR, HIPAA, and PCI DSS, necessitating the implementation of stringent security controls and audit mechanisms.

5. Elasticity and Scalability: The dynamic scaling of resources on AWS poses challenges in maintaining consistent security policies across expanding and contracting infrastructure.

3. AWS Security Tools and Features for Hadoop Deployments

In modern big data ecosystems, Hadoop remains a cornerstone for managing and analyzing vast volumes of data. However, the security challenges associated with Hadoop deployments can be significant, especially when these systems are hosted in cloud environments like Amazon Web Services (AWS). AWS offers a comprehensive suite of security tools and features that can help secure Hadoop deployments, ensuring data integrity, confidentiality, and compliance.

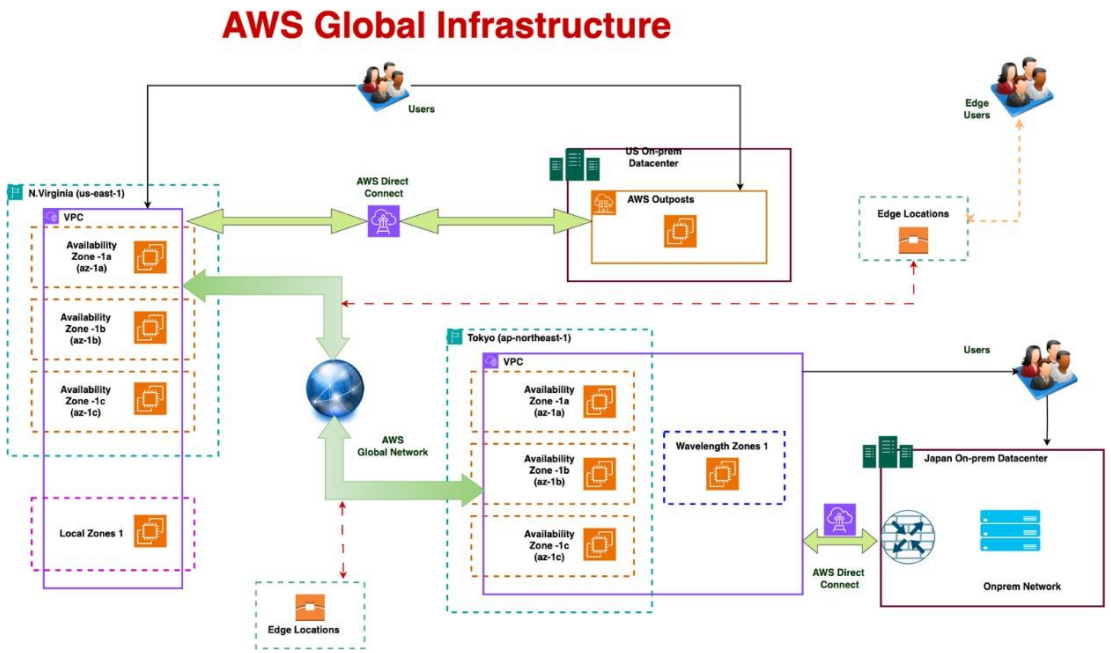

Figure 3: AWS Global Infrastructure

This paper explores these tools and their relevance to Hadoop deployments.

1. Identity and Access Management (IAM)

AWS Identity and Access Management (IAM) is a foundational security service that enables fine-grained control over user permissions and access to AWS resources. For Hadoop deployments, IAM can be leveraged to:

- Control Access to Hadoop Clusters: Define roles, policies, and permissions to ensure only authorized users and services can access Hadoop clusters.
- Integration with Kerberos: Hadoop supports Kerberos authentication, and IAM roles can be configured to provide secure key distribution and access controls.
- Temporary Credentials: Using AWS Security Token Service (STS), temporary and limited-time access to Hadoop resources can be granted, minimizing long-term exposure to credentials.

2. Amazon S3 Security Features

Amazon Simple Storage Service (S3) is commonly used as a data lake for Hadoop deployments. AWS provides robust security features for S3, ensuring that data stored and processed by Hadoop is secure:

- Bucket Policies and Access Control Lists (ACLs): Enforce strict access rules to regulate who can read, write, or delete data in S3 buckets.
- Server-Side Encryption: Enable encryption at rest using AWS Key Management Service (KMS) managed keys or customer-provided keys.

- Object Locking: Implement versioning and retention policies to prevent unauthorized deletion or modification of critical data.
- Logging and Monitoring: Enable S3 access logs and integrate them with services like AWS CloudTrail to monitor access patterns and detect anomalies.

3. Security Groups and Network Access Controls

Securing Hadoop's network infrastructure is critical to preventing unauthorized access and ensuring the integrity of data in transit. AWS provides the following tools:

- Security Groups: Act as virtual firewalls for EC2 instances, allowing fine-grained control over inbound and outbound traffic to Hadoop nodes.
- Network Access Control Lists (NACLs): Provide an additional layer of subnet-level traffic filtering.
- VPC Endpoints: Establish private connectivity to S3, DynamoDB, or other AWS services, bypassing the need for public internet exposure.
- Elastic Load Balancer (ELB): Distribute traffic securely across Hadoop cluster nodes, with integrated support for TLS termination.

4. Encryption Tools for Data at Rest and in Transit

Encryption is a vital component of securing Hadoop deployments. AWS provides multiple encryption mechanisms:

- AWS KMS: Centralized key management to encrypt data stored in S3, EBS, or RDS, commonly used in Hadoop environments.
- Hadoop Native Encryption: Configure Hadoop to work with AWS KMS for encrypting HDFS data at rest.
- Transport Layer Security (TLS): Secure data in transit between Hadoop nodes and client applications by enabling TLS encryption.
- Certificate Management: Use AWS Certificate Manager (ACM) to provision, manage, and deploy SSL/TLS certificates for Hadoop services.

5. AWS Security Monitoring and Threat Detection

Continuous monitoring and threat detection are essential to maintain the security posture of Hadoop deployments. AWS offers several services to aid in this:

- AWS CloudTrail: Track and log all API activity related to Hadoop resources, ensuring accountability and facilitating forensic analysis.
- Amazon GuardDuty: Detect anomalies and potential threats by analyzing logs from S3, VPC Flow Logs, and CloudTrail.
- AWS Config: Monitor configuration changes to Hadoop-associated AWS resources, ensuring compliance with best practices and security policies.

- AWS Security Hub: Consolidate security findings from various AWS services, providing a centralized view of potential risks.

6. Compliance and Governance Tools

Hadoop deployments often need to adhere to regulatory requirements. AWS provides tools to simplify compliance and governance:

- AWS Artifact: Access compliance reports and audit certifications relevant to your Hadoop deployment.
- AWS Organizations: Enforce policies across multiple AWS accounts managing Hadoop clusters.
- Service Control Policies (SCPs): Define and enforce organization-wide permission boundaries for Hadoop-related AWS accounts.
- Amazon Macie: Identify and protect sensitive data in S3 buckets, commonly used with Hadoop data lakes.

7. Best Practices for Securing Hadoop on AWS

To maximize the security of Hadoop deployments on AWS, consider the following best practices:

- Use Multi-Factor Authentication (MFA): Add an additional layer of security for administrative access to AWS and Hadoop resources.
- Implement Least Privilege Access: Grant users and services only the permissions they need to perform their tasks.
- Regularly Patch and Update Systems: Keep Hadoop clusters and underlying AWS resources updated to mitigate vulnerabilities.
- Enable Logging and Analytics: Use AWS's extensive logging and analytics tools to monitor and respond to potential threats in real time.
- Automate Security Policies: Use AWS tools like AWS Config Rules and Lambda to automate the enforcement of security policies.

AWS provides a robust set of security tools and features that are well-suited for securing Hadoop deployments. By leveraging these tools, organizations can ensure their Hadoop-based big data systems remain secure, compliant, and operationally efficient. Properly implementing these features not only protects data but also enhances trust and reliability in the overall system.

AWS offers a rich set of features that, when properly configured, provide robust security for Hadoop implementations. Some of the core AWS services and tools include:

- Identity and Access Management (IAM): Centralized control over user and application permissions, enabling fine-grained access to AWS resources.
- Virtual Private Cloud (VPC): Securely isolates Hadoop clusters within private networks, ensuring restricted access.

- AWS Key Management Service (KMS): Simplifies the encryption and management of cryptographic keys for securing data at rest and in transit.
- CloudTrail and CloudWatch: Facilitate real-time monitoring and logging of user activity, providing insights into potential security anomalies.
- AWS Shield and WAF: Protect Hadoop deployments from Distributed Denial-of-Service (DDoS) attacks and other web-based threats.

By strategically integrating these tools with Hadoop's native security features—such as Kerberos authentication and Hadoop Distributed File System (HDFS) encryption—organizations can create a multi-layered security framework. This approach not only enhances the overall security posture but also ensures resilience against evolving cyber threats.

4. Objectives of the Security Framework

The primary goal of deploying an AWS security framework for Hadoop is to safeguard the confidentiality, integrity, and availability of data and systems. Specific objectives include:

1. Risk Mitigation: Identifying and addressing potential security vulnerabilities before they can be exploited.
2. Regulatory Compliance: Ensuring adherence to legal and industry standards for data protection and privacy.
3. Operational Efficiency: Streamlining security operations to minimize overhead and optimize resource utilization.
4. Scalability and Flexibility: Enabling secure scaling of Hadoop clusters to meet dynamic workload demands.
5. Incident Response: Establishing a robust framework for detecting, analyzing, and responding to security incidents.

In this paper, we delve deeper into the methodologies and best practices for implementing a secure Hadoop deployment on AWS, examining real-world scenarios and case studies to illustrate the practical application of these principles.

References

[1]. Adam, Omer, Young Choon Lee, and Albert Y. Zomaya. "Stochastic Resource Provisioning for Containerized Multi-Tier Web Services in Clouds." IEEE Transactions on Parallel and Distributed Systems 28, no. 7 (July 1, 2017): 2060–73. https://doi.org/10.1109/TPDS.2016.2639009.

[2]. Adam, Omer Y., Young Choon Lee, and Albert Y. Zomaya. "Constructing Performance-Predictable Clusters with Performance-Varying Resources of Clouds." IEEE Transactions on Computers 65, no. 9 (September 1, 2016): 2709–24. https://doi.org/10.1109/TC.2015.2510648.

[3]. Akshatha, P.S., and S.M. Dilip Kumar. "MQTT and Blockchain Sharding: An Approach to User-Controlled Data Access with Improved Security and Efficiency." Blockchain: Research and Applications 4, no. 4 (December 2023): 100158. https://doi.org/10.1016/j.bcra.2023.100158.

[4]. Al-Dhuraibi, Yahya, Fawaz Paraiso, Nabil Djarallah, and Philippe Merle. "Elasticity in Cloud

Computing: State of the Art and Research Challenges." IEEE Transactions on Services Computing 11, no. 2 (March 1, 2018): 430–47. https://doi.org/10.1109/TSC.2017.2711009.

[5]. Al-Dulaimy, Auday, Javid Taheri, Andreas Kassler, M. Reza HoseinyFarahabady, Shuiguang Deng, and Albert Zomaya. "MultiScaler: A Multi-Loop Auto-Scaling Approach for Cloud-Based Applications." IEEE Transactions on Cloud Computing 10, no. 4 (October 1, 2022): 2769–86. https://doi.org/10.1109/TCC.2020.3031676.

[6]. Ali, Sijjad, Shuaib Ahmed Wadho, Aun Yichiet, Ming Lee Gan, and Chen Kang Lee. "Advancing Cloud Security: Unveiling the Protective Potential of Homomorphic Secret Sharing in Secure Cloud Computing." Egyptian Informatics Journal 27 (September 2024): 100519. https://doi.org/10.1016/j.eij.2024.100519.

[7]. Amekraz, Zohra, and Moulay Youssef Hadi. "Higher Order Statistics Based Method for Workload Prediction in the Cloud Using ARMA Model." In 2018 International Conference on Intelligent Systems and Computer Vision (ISCV), 1–5. Fez: IEEE, 2018. https://doi.org/10.1109/ISACV.2018.8354078.

[8]. Balaji, Mahesh, Ch. Aswani Kumar, and G. Subrahmanya V.R.K. Rao. "Predictive Cloud Resource Management Framework for Enterprise Workloads." Journal of King Saud University - Computer and Information Sciences 30, no. 3 (July 2018): 404–15. https://doi.org/10.1016/j.jksuci.2016.10.005.

[9]. Barcelona-Pons, Daniel, and Pedro García-López. "Benchmarking Parallelism in FaaS Platforms." Future Generation Computer Systems 124 (November 2021): 268–84. https://doi.org/10.1016/j.future.2021.06.005.

[10]. Belal, Mohamad Mulham, and Divya Meena Sundaram. "Comprehensive Review on Intelligent Security Defences in Cloud: Taxonomy, Security Issues, ML/DL Techniques, Challenges and Future Trends." Journal of King Saud University - Computer and Information Sciences 34, no. 10 (November 2022): 9102–31. https://doi.org/10.1016/j.jksuci.2022.08.035.

[11]. Bello, Yahuza, Alaa Awad Abdellatif, Mhd Saria Allahham, Ahmed Refaey Hussein, Aiman Erbad, Amr Mohamed, and Mohsen Guizani. "B5G: Predictive Container Auto-Scaling for Cellular Evolved Packet Core." IEEE Access 9 (2021): 158204–14. https://doi.org/10.1109/ACCESS.2021.3126048.

[12]. Bhaskaran, Harini Shree, Miriam Gordon, and Suresh Neethirajan. "Development of a Cloud-Based IoT System for Livestock Health Monitoring Using AWS and Python." Smart Agricultural Technology 9 (December 2024): 100524. https://doi.org/10.1016/j.atech.2024.100524.

[13]. Breternitz, Mauricio, Keith Lowery, Anton Charnoff, Patryk Kaminski, and Leonardo Piga. "Cloud Workload Analysis with SWAT." In 2012 IEEE 24th International Symposium on Computer Architecture and High Performance Computing, 92–99. New York, NY, USA: IEEE, 2012. https://doi.org/10.1109/SBAC-PAD.2012.13.

[14]. Lilhore, Umesh Kumar, Sarita Simaiya, Musaed Alhussein, Surjeet Dalal, Khursheed Aurangzeb, and Amir Hussain. "An Attention-Driven Hybrid Deep Neural Network for Enhanced Heart Disease Classification." Expert Systems (2024): e13791.

[15]. Bundela, Rajmani, Namrata Dhanda, and Rajat Verma. "Load Balanced Web Server on AWS Cloud." In 2022 International Conference on Computing, Communication, and Intelligent Systems (ICCCIS), 114–18. Greater Noida, India: IEEE, 2022. https://doi.org/10.1109/ICCCIS56430.2022.10037657.

[16]. Calatrava, Amanda, Hernán Asorey, Jan Astalos, Alberto Azevedo, Francesco Benincasa, Ignacio Blanquer, Martin Bobak, et al. "A Survey of the European Open Science Cloud Services for Expanding the Capacity and Capabilities of Multidisciplinary Scientific Applications." Computer Science Review 49 (August 2023): 100571. https://doi.org/10.1016/j.cosrev.2023.100571.

[17]. Yadav, Sudha, Harkesh Sehrawat, Vivek Jaglan, Yudhvir Singh, Surjeet Dalal, and Dac-Nhuong

Le. "Developing Model-Agnostic Meta-Learning Enabled Lightbgm Model Asthma Level Prediction in Smart Healthcare Modeling." Scalable Computing: Practice and Experience 25, no. 6 (2024): 4872-4885.

[18]. Calheiros, Rodrigo N., Rajiv Ranjan, and Rajkumar Buyya. "Virtual Machine Provisioning Based on Analytical Performance and QoS in Cloud Computing Environments." In 2011 International Conference on Parallel Processing, 295–304. Taipei, Taiwan: IEEE, 2011. https://doi.org/10.1109/ICPP.2011.17.

[19]. Saini, Himani, Gopal Singh, Sandeep Dalal, Umesh Kumar Lilhore, Sarita Simaiya, and Surjeet Dalal. "Enhancing cloud network security with a trust-based service mechanism using k-anonymity and statistical machine learning approach." Peer-to-Peer Networking and Applications (2024): 1-26.

[20]. Centofanti, Carlo, Walter Tiberti, Andrea Marotta, Fabio Graziosi, and Dajana Cassioli. "Taming Latency at the Edge: A User-Aware Service Placement Approach." Computer Networks 247 (June 2024): 110444. https://doi.org/10.1016/j.comnet.2024.110444.

[21]. Dalal, Surjeet, Umesh Kumar Lilhore, Bijeta Seth, Magdalena Radulescu, and Sofiane Hamrioui. "A Hybrid Model for Short-Term Energy Load Prediction Based on Transfer Learning with LightGBM for Smart Grids in Smart Energy Systems." Journal of Urban Technology (2024): 1-27.

[22]. Chen, Jiajun, Chi Wan Sung, and Terence H. Chan. "Heterogeneity Shifts the Storage-Computation Tradeoff in Secure Multi-Cloud Systems." IEEE Transactions on Information Theory 69, no. 2 (February 2023): 1015–36. https://doi.org/10.1109/TIT.2022.3206868.

[23]. Chen, Yunliang, Lizhe Wang, Xiaodao Chen, Rajiv Ranjan, Albert Y. Zomaya, Yuchen Zhou, and Shiyan Hu. "Stochastic Workload Scheduling for Uncoordinated Datacenter Clouds with Multiple QoS Constraints." IEEE Transactions on Cloud Computing 8, no. 4 (October 1, 2020): 1284–95. https://doi.org/10.1109/TCC.2016.2586048.

[24]. Dalal, Surjeet, Umesh Kumar Lilhore, Sarita Simaiya, Magdalena Radulescu, and Lucian Belascu. "Improving efficiency and sustainability via supply chain optimization through CNNs and BiLSTM." Technological Forecasting and Social Change 209 (2024): 123841.

[25]. Edeh, Michael Onyema, Surjeet Dalal, Musaed Alhussein, Khursheed Aurangzeb, Bijeta Seth, and Kuldeep Kumar. "A novel deep learning model for predicting marine pollution for sustainable ocean management." PeerJ Computer Science 10 (2024): e2482.