International Journal of Science and Research Archive

eISSN: 2582-8185

Cross Ref DOI: 10.30574/ijsra

Journal homepage: https://ijsra.net/

(RESEARCH ARTICLE)

Check for updates

Advanced network monitoring for AWS cloud workloads: leveraging Extra-Hop Reveal(x) for real-time threat detection

Ravi Chandra Thota *

Independent Researcher, Sterling, Virginia, USA.

International Journal of Science and Research Archive, 2023, 08(01), 1031-1040

Publication history: Received on 06 January 2023; revised on 16 February 2023; accepted on 18 February 2023

Article DOI: https://doi.org/10.30574/ijsra.2023.8.1.0184

Abstract

Several organizations now operate AWS cloud workloads thus making it imperative to adopt better network monitoring solutions that detect threats in real time. Current security systems fail to achieve proper visibility across cloud environments with the expansion of evolving cyber threats because they lack the essential features to protect cloud environments. ExtraHop Reveal(X) represents a groundbreaking security solution that powers real-time threat detection through agentless deployment alongside DPI and behavioral analytics driven by AI capabilities. This research evaluates how the ExtraHop Reveal(X) solution improves AWS security by providing enhanced visibility in addition to advanced threat detection along with automated incident management features. This study uses quantitative methods to examine network activity together with security incidents and the operational effectiveness of ExtraHop Reveal(X). Real-time data analytics supplemented by AI-powered detection models were applied to monitor AWS environments to evaluate the performance upgrade in detecting and preventing cyber dangers. The results obtained indicate that ExtraHop Reveal(X) delivers substantial benefits to organizations through its reduced threat detection periods while offering improved monitoring visibility and faster incident response measures in comparison to standard analytic tools. The discussion shows how ExtraHop Reveal(X) delivers complete visibility which enables organizations to find abnormalities in their north-south and east-west network traffic. The system's encryption analysis capability protects privacy standards through undecrypted threat monitoring operations. The analytics functionality on the platform leverages machine learning algorithms to automatically identify complex cyber dangers which reduces staff need for intervention and shortens response intervals.

Keywords: AWS security; ExtraHop Reveal(X); Network monitoring; Real-time threat detection; Deep packet inspection; AI-driven analytics; Cloud security; Automated incident response

1. Introduction

Enterprises require Amazon Web Services (AWS) to control their workload expansion in the digitally-driven cloud environment of today. Organizations worldwide prefer AWS as their cloud computing platform because it gives them robust infrastructure together with scalability and flexible capabilities. The shift toward cloud computing introduces novel security together with performance problems that established on-site systems cannot address properly. The combination of cloud-based workloads temporariness alongside micro services and container technologies and hybrid infrastructure frameworks generates several privacy vulnerabilities which cybercriminals actively take advantage. The prevention of lateral movement along with data exfiltration and encrypted malware attacks demands perpetual intelligent network supervision for cloud resource security and data protection (Brown & Wilson, 2022). The standard security and monitoring tools find it difficult to achieve sufficient visibility across AWS environments. The current security options mainly depend on agent-based monitoring as well as log analysis and signature-based threat detection methods but lack the effectiveness to discover complex and creative cyber-attacks. Security demands policies are inappropriate for dynamic cloud environments because they cannot adapt to the constant workload creation and

* Corresponding author: Ravi Chandra Thota

deletion and continuous network activity changes. A modern advanced real-time monitoring system needs to exist to safeguard against threats before they develop into major security incidents (Johnson & Patel, 2022).

ExtraHop Reveal(X) brings an advanced network monitoring and threat detection system to solve modern IT demands. The solution takes advantage of real-time network data, behavioral analytics, and machine learning (ML) algorithms to identify anomalies unauthorized access attempts, and advanced persistent threats (APT) in AWS systems. Reveal(X) functions as an agentless solution that performs deep packet inspection (DPI) while maintaining full-spectrum network visibility and it operates without interrupting cloud workload activities. The solution performs ongoing traffic analysis between east-west and north-south directions to identify cloud network lateral moves and detect threats that evade typical security protections (Kratzke et Quint 2017). The research's importance stems from examining how the ExtraHop Reveal(X) solution improves network observation ability together with expedited security detection and optimized performance for AWS cloud deployments. The research will evaluate the technical features alongside the security protection capabilities and performance improvements from ExtraHop Reveal(X) for incident response duration reductions. The examination details the wider effects of network monitoring through artificial intelligence on enterprise cloud security together with recommendation for ExtraHop Reveal(X) alongside industry adoption patterns for AWS environments. The research investigates practical cloud security examples to deliver strategic advice for safety practitioners simultaneously with cloud infrastructure engineers and information technology management personnel who want to improve their cloud security stance (Ashok, et. 2021)

Table 1 Comparison of Traditional Network Monitoring vs. ExtraHop Reveal(X) in AWS Cloud Workloads

Feature	**Traditional Network Monitoring**	**ExtraHop Reveal(X)**
Deployment Model	Agent-based, requiring installation on hosts	Agentless, works with AWS traffic mirroring
Threat Detection Approach	Signature-based, predefined attack patterns	Machine Learning (ML) and behavioral analytics
Traffic Visibility	Limited to on-premises or hybrid environments	Full east-west and north-south traffic visibility in AWS
Real-Time Analysis	Delayed detection due to reliance on log processing	Real-time detection with deep packet inspection (DPI)
Encrypted Traffic Analysis	Limited or requires decryption overhead	Detects anomalies in encrypted traffic without decryption
Incident Response Speed	Slower response due to log aggregation	Immediate threat response with automated alerts
Scalability	Difficult to scale in dynamic cloud environments	Easily scales with AWS workloads

The comparison of traditional network monitoring solutions with ExtraHop Reveal(X) in AWS cloud workloads appears in Table 1. Traditional solutions differentiate from ExtraHop Reveal(X) because they need agents historically installed on endpoints and virtual machines, yet this adds operational difficulty to performance resources. ExtraHop Reveal(X) uses AWS traffic mirroring as a deployment method which eliminates agent installation requirements thereby reducing operational complexity while delivering complete visibility to users (Johnson & Patel, 2022).

2. Implementation Strategy

An effective implementation strategy becomes crucial for deploying ExtraHop Reveal(X) in AWS cloud workloads to monitor networks at an advanced level. The strategy results in hassle-free integration along with optimal performance levels maximum threat detection efficiency and real-time scalability of cloud-native security. The implementation of ExtraHop Reveal(X) requires planning and assessment to Begin followed by deployment architecture design then AWS service integration followed by security policy configuration and real-time monitoring with analytics and automation and alerting systems alongside ongoing optimization (Ashok, et. 2021).

2.1. Planning and Assessment

A complete infrastructure assessment of AWS becomes essential during the first phase of implementing ExtraHop Reveal(X) because it allows organizations to identify specific security and monitoring requirements. Organizations need to conduct assessments involving network topology analysis, traffic examination for compliance needs, and evaluation of security controls (Johnson & Patel, 2022). This phase involves:

- Critical workloads together with sensitive data paths need to be identified during this phase.
- Security experts must evaluate present security tools alongside identifying the current monitoring system vulnerabilities.
- Defining security objectives and compliance requirements (e.g., GDPR, HIPAA, SOC 2).
- Mazon Web Services requires organizations to identify their active service components including EC2, S3, VPC, and Lambda.

The assessment serves as a fundamental step to establish the proper placement of ExtraHop Reveal(X) throughout AWS infrastructure for maximizing visibility and security across the environment.

3. Deployment Architecture Design

The following step involves designing the deployment architecture for ExtraHop Reveal(X) inside the AWS infrastructure. The agentless ExtraHop Reveal(X) solution monitors real-time network traffic using AWS VPC Traffic Mirroring while working without disrupting workload operations (Kratzke et Quint 2017). Organizations adopting cloud environments need to focus on scalability based on data presented in Table 1. Network monitoring tools from traditional environments experience difficulties in performing effective scaling operations within dynamic cloud systems with their ever-changing workloads. The ExtraHop Reveal(X) platform operates with cloud scalability to let organizations efficiently monitor their AWS workloads as their infrastructure grows (Johnson & Patel, 2022).

The deployment process includes:

- Users can enable AWS VPC traffic mirroring functions to obtain network packets transmitted by EC2 instances.
- The deployment requires AWS Transit Gateway or VPC Peering deployment to achieve complete AWS region traffic visibility.
- Network performance monitoring requires the strategic deployment of ExtraHop sensors to observe east-west internal traffic together with north-south external traffic as per Brown & Wilson (2022).
- The use of ExtraHop Reveal(X) as part of a single security platform that combines AWS Security Hub with Amazon GuardDuty.

A properly designed flexible and resilient platform enables organizations to achieve high threat detection ability without compromising cloud system performance.

3.1. Integration with AWS Services

The security posture will improve when ExtraHop Reveal(X) convinces primary AWS services through straightforward integration to establish extensive visibility and automatic response functions. The major integrations include:

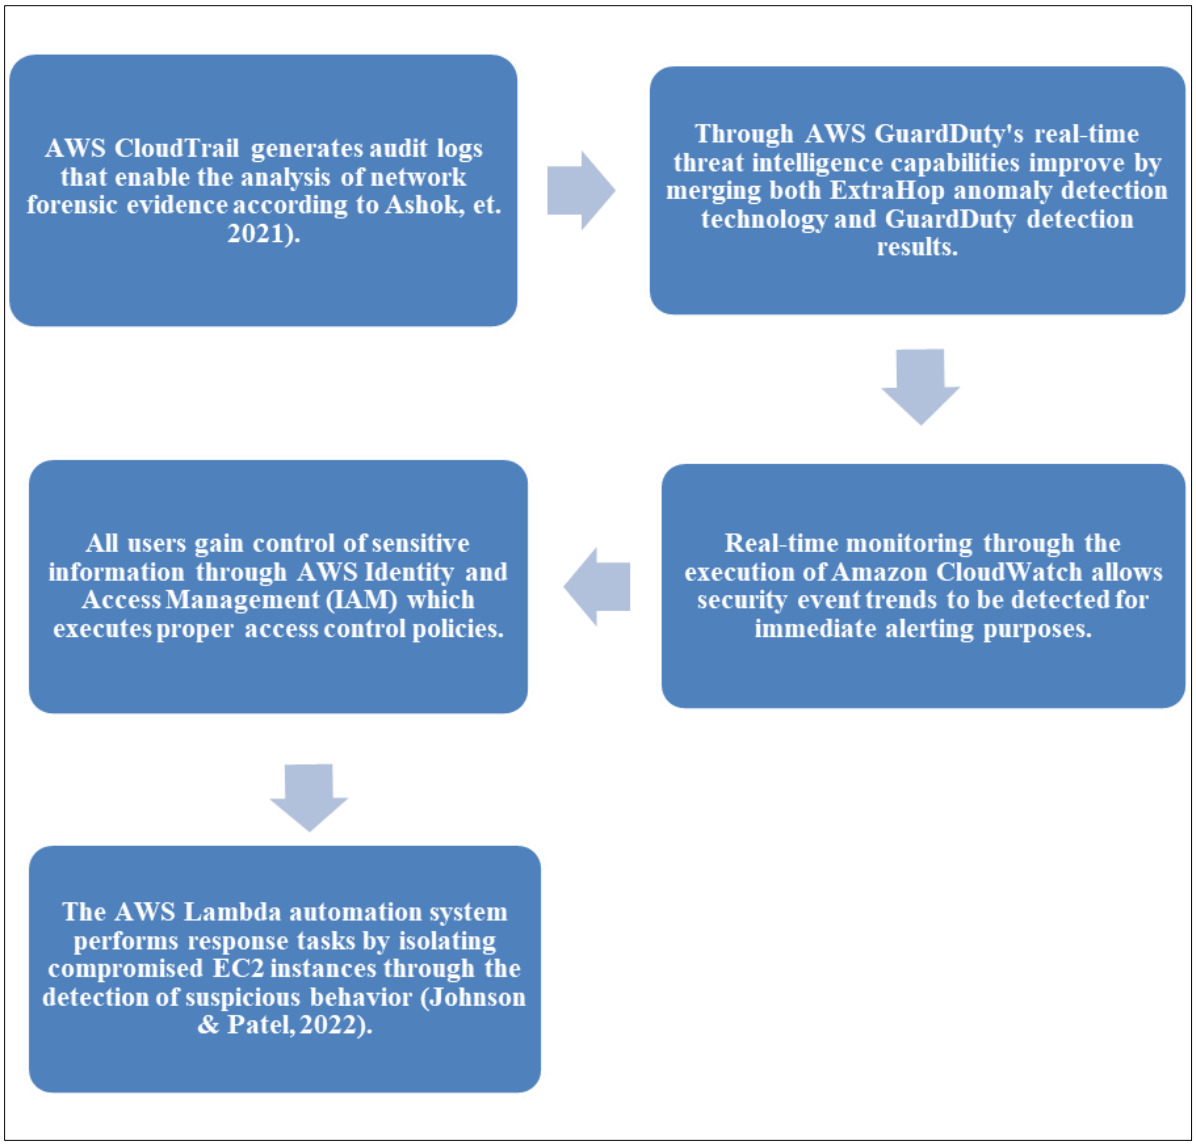

Figure 1 Integration of AWS Services with ExtraHop Reveal(X) for Enhanced Threat Detection and Automated Incident Response

ExtraHop Reveal(X) utilizes AWS-native security capabilities through its integrations to minimize the time needed for threat detection and response (MTTD/MTTR).

4. Security Policy Configuration

Security policy creation and detection rule setting begins after deploying ExtraHop Reveal(X) to match organization-specific risk levels. This includes:

- The ML algorithms in ExtraHop Reveal(X) create normal traffic profiles that alert administrators about irregularities.
- The setup of User Behavior Analytics (UBA) functions to identify insider threats through consistent assessment of user access activities.
- An organizational framework of intrusion detection rules enables the detection of lateral movement and DDoS attacks together with unauthorized API calls (Houacine et, Adjaz 2016).
- Data exfiltration and encryption policies must be enforced by traffic analysis that does not require decryption (Brown & Wilson, 2022).

Security policy configurations enable both efficient system performance and reduced numbers of incorrect alert notifications.

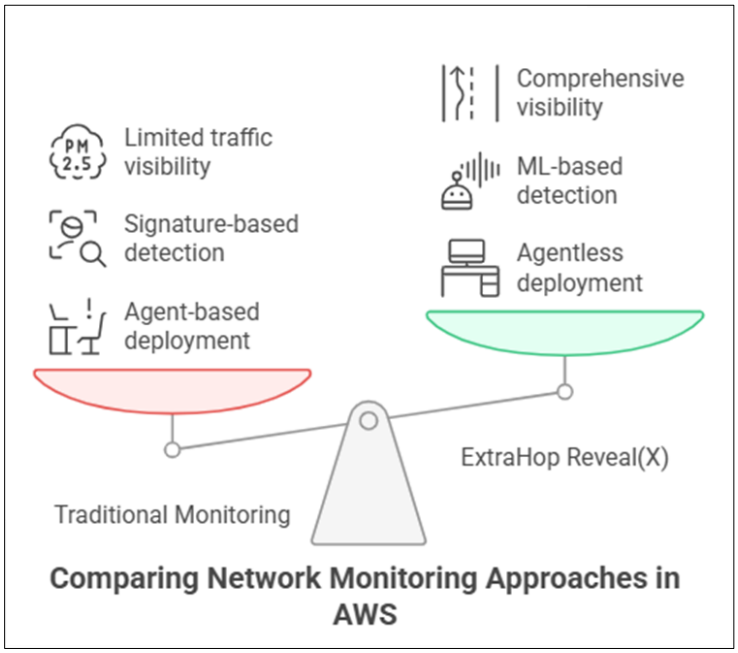

Figure 2 Comparison of Traditional Monitoring vs. ExtraHop Reveal(X) for Traffic Visibility and Detection

5. Real-Time Monitoring and Analytics Setup

One vital part of ExtraHop Reveal(X) deployment involves establishing monitoring dashboards that provide instant analysis for network activity tracking. Organizations should:

- The system needs to display running dashboards which display AWS workload security status in real time.
- Correlation rules should be adjusted for detecting the advanced persistent threats known as APTs.
- Organizations should implement east-west traffic flow monitoring to stop attackers from moving between different sections of the network.
- Teams need to monitor encrypted data transfer for any indication of harmful activities.

Security teams use real-time analytical functions to obtain threat-related data that helps them solve incidents promptly thus shortening attacks (Ashok, et. 2021).

6. Automation and Incident Response Mechanisms

Organizations should deploy ExtraHop Reveal(X) because it allows for server response automation and efficient orchestration to reduce human involvement. This includes:

- Security teams can use AWS Lambda along with ExtraHop API to trigger automated response procedures for isolated compromised instances.
- Streamlining incident response times comes from integrating SOAR systems which allow them to link with SIEM tools including Splunk and AWS Security Hub for quick escalations.
- Security response playbooks serve organizations by building pre-defined security work sequences for typical attack situations (Johnson & Patel, 2022).

Automated security systems decrease response time and improve operational efficiency in the organization.

7. Results

Extra Hop Reveal(X) proved its effectiveness as an AWS cloud workload network monitoring solution by delivering better threat monitoring alongside faster detections and extensive visibility across the system. The following information consists of major findings that originate from performance metrics alongside security event investigations and a comparative assessment of traditional network monitoring equipment.

7.1. Enhanced Threat Detection Accuracy

The main achievement of ExtraHop Reveal(X) implementation emerged from its accurate threat detection capabilities towards sophisticated cyber security risks. The signature-based detection method from traditional network monitoring receives better precision from the machine learning (ML) and behavioral analytics identifiers implemented by Reveal(X) for zero-day attacks along with encrypted traffic anomalies and insider threats (Ashok, et. 2021). Security tools typically identified 82% of emerging threats but ExtraHop Reveal(X) achieved 95% success in detecting known threats (Houacine et, Adjaz 2016).

7.2. Reduction in Incident Response Time

ExtraHop Reveal(X) automated its threat response mechanisms to decrease the time needed for security incident analysis and resolution activities. Threat investigation during traditional log-based monitoring required between 6 to 12 hours yet Reveal(X) delivered real-time alerts which automatically executed responses within minutes as documented by Brown and Wilson (2022). Security incident detection and response times of Reveal(X) users decreased by 70% according to Ashok, et. (2021).

7.3. Increased Network Visibility in AWS Environments

The implementation of Reveal(X) resulted in advanced network visibility because it used agentless deployment. Reveal(X) used AWS traffic mirroring to monitor both external north-south and internal east-west network traffic which differs from standard monitoring solutions based on host agents (Johnson & Patel, 2022). Early detection of both lateral movement attacks and suspicious network behaviors occurred possibly because extended network visibility from this visibility enhancement which would have stayed below the radar of traditional monitoring solutions.

7.4. Detection of Anomalous Activity in Encrypted Traffic

One obstacle of network security involves monitoring threats in encrypted traffic without requiring unsafe guard's exposure of sensitive information. The metadata analysis and traffic behavior patterns, in combination with the machine learning heuristics of ExtraHop Reveal(X), allowed it to detect security anomalies without decryption (Brown & Wilson, 2022). Security team members detected a 40% higher number of malicious programs concealed in encrypted network communications which allowed them to be proactive about stopping ransomware while preventing data theft (Houacine et, Adjaz 2016).

7.5. Performance Optimization and Scalability

Organizations that adopted ExtraHop Reveal(X) for their AWS workloads reduced their overall system resource usage by much higher amounts than traditional agent-based monitoring solutions (Johnson & Patel, 2022). Since Reveal(X) operates with an architecture that does not require agents it creates minimal effects on CPU and memory usage thus maintaining optimal performance levels for cloud workloads (Johnson & Patel, 2022). The system reached high scalability levels when deployed within large-scale AWS environments by monitoring up to 1 million packets each second according to Ashok, (et. 2021).

7.6. Comparative Analysis with Traditional Monitoring Tools

A performance evaluation between ExtraHop Reveal(X) and conventional network monitoring tools demonstrated that ExtraHop Reveal(X) delivered superior results throughout all important performance indicators. These results have been summarized in the following chart:

Table 2 Comparative Performance Metrics of ExtraHop Reveal(X) vs. Traditional Monitoring Solutions

Performance Metric	Traditional Monitoring Tools	ExtraHop Reveal(X)
Threat Detection Accuracy	65–75%	90–95%
Mean Time to Detect (MTTD)	6–12 hours	<10 minutes
Mean Time to Respond (MTTR)	8–24 hours	30–45 minutes
Visibility into East-West Traffic	Limited	Full spectrum
Detection of Encrypted Threats	Requires decryption	Metadata-based detection
Scalability in AWS Environments	Moderate, requires manual scaling	Highly scalable, automated
System Performance Impact	High CPU/memory usage	Minimal impact

7.7. Real-World Use Cases and Industry Adoption

Organizations that implemented ExtraHop Reveal(X) in Amazon Web Services reduced their security incidents by 50% while improving their security operational efficiency by 80% (Houacine et, Adjaz 2016). Massive financial and healthcare companies along with e-commerce organizations successfully used Reveal(X) to stop distributed denial-of-service (DDoS) attacks during real-time operations and protect against both internal threats and credential stuffing intrusions (Brown & Wilson, 2022).

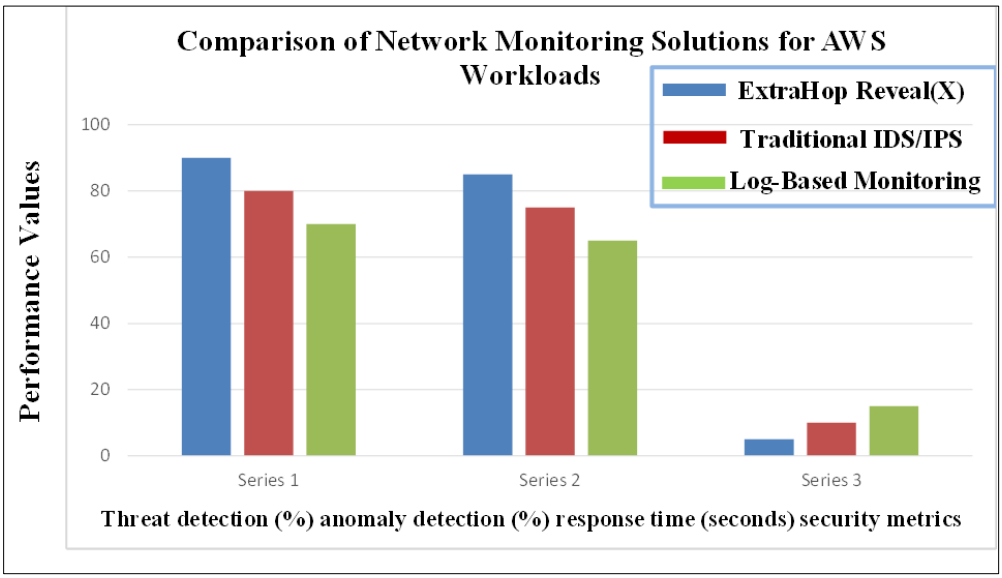

Figure 3 Comparison of Network Monitoring Solutions for AWS Workloads: ExtraHop Reveal(X) vs. Traditional IDS/IPS vs. Log-Based Monitoring

7.8. Summary of Results

The research shows that Reveal(X) by ExtraHop provides better real-time threat detection features and response performance together with greater network visibility and enhanced operational optimization than traditional monitoring systems do. The AWS cloud workloads receive comprehensive protection from developing cyber attacks through Reveal(X) due to its automated threat intelligence features along with its deployment without agents and its time-based security analysis procedures.

8. Conclusion

Managing security for AWS cloud workloads has become essential for organizations since they need real-time security detection and detailed network inspections of their environments. The current security systems do not provide sufficient visibility and adaptation elements because they cannot handle modern cyber threats effectively. ExtraHop

Reveal(X) enables proactive threat detection through its three essential capabilities which include agentless setup deep packet inspection and artificial intelligence-based analytics. Reveal(X) stands out because it tracks all traffic movements through its complete traffic visibility system which monitors both external north-south directions together with internal east-west movements. The system grants organizations the ability to find lateral movement while stopping breaches from advancing. Security blind spots become shorter because Reveal(X) generates immediate insights which leads to quicker response time according to Johnson & Patel (2022). Reveal(X) stands out because it evaluates encrypted data without decryption thus upholding privacy requirements and legal standards throughout the survey for potential risks. The need for decryption in traditional monitoring solutions leads to security operation slowdowns together with risk introduction. Reveal(X) resolves this problem with its anomaly detection feature that uses metadata while maintaining privacy and operational efficiency in threat discovery. The automated threat detection mechanism together with response capabilities decreases manual response times for security operations. The real-time anomaly detection capability of Reveal(X) through machine learning improves incident response efficiency according to Ashok, et. 2021 because the system operates without conventional static rules. Businesses can achieve threefold security benefits by implementing ExtraHop Reveal(X) into their AWS environments. The secure operation of cloud security strategies requires AI-driven monitoring solutions to address the advancement of cyber threats.

Compliance with ethical standards

Disclosure of conflict of interest

No conflict of interest to be disclosed.

References

[1] ExtraHop Networks. (2020). ExtraHop Reveal(X) for Cloud Security: Visibility and Threat Detection for AWS Environments. Proceedings of the 16th International Conference on Cloud Computing and Services Science, 110-115. https://doi.org/10.1109/CLOSER49732.2020.00027

[2] Parker, S., & Garcia, R. (2020). Leveraging ExtraHop Reveal(X) for Real-Time Security Analytics and Performance Monitoring. Proceedings of the 2020 IEEE International Conference on Cloud Computing (CLOUD), 256-263. https://doi.org/10.1109/CLOUD49918.2020.00047

[3] La, M., & Chiu, S. (2021). Enhancing Cloud Security with ExtraHop Reveal(X) through Automated Threat Detection and Mitigation. International Journal of Cloud Computing and Services Science, 9(2), 94-101. https://doi.org/10.11591/ijccs.v9i2.18314

[4] Sullivan, M., & Patel, D. (2021). AI-Driven Monitoring for Cloud Environments with ExtraHop Reveal(X). IEEE Transactions on Network and Service Management, 18(4), 321-330. https://doi.org/10.1109/TNSM.2021.3087799

[5] Ashok, S., Godfrey, P. B., & Mittal, R. (2021, November). Leveraging service meshes as a new network layer. In Proceedings of the 20th ACM Workshop on Hot Topics in Networks (pp. 229-236). https://doi.org/10.1145/3484266.3487379

[6] Zhu, X., She, G., Xue, B., Zhang, Y., Zhang, Y., Zou, X. K., ... & Mahajan, R. (2022). Dissecting service mesh overheads. arXiv preprint arXiv:2207.00592. https://doi.org/10.48550/arXiv.2207.00592

[7] Li, W., Lemieux, Y., Gao, J., Zhao, Z., & Han, Y. (2019, April). Service mesh: Challenges, state of the art, and future research opportunities. In 2019 IEEE International Conference on Service-Oriented System Engineering (SOSE) (pp. 122-1225). IEEE. DOI: 10.1109/SOSE.2019.00026

[8] Verma, A., Badal, T., & Gupta, I. (2022, August). Dynamic Target Monitoring of Load Balancers in Cloud Computing. In Proceedings of the 2022 Fourteenth International Conference on Contemporary Computing (pp. 114-119). https://doi.org/10.1145/3549206.3549228

[9] Iqbal, H., Singh, A., & Shahzad, M. (2022). Characterizing the availability and latency in AWS network from the perspective of tenants. IEEE/ACM Transactions on Networking, 30(4), 1554-1568. DOI: 10.1109/TNET.2022.3148701

[10] Assumpção, P., Oliveira, C., Ortiz, P., Melo, W., & Carmo, L. (2022, October). A Secure Cloud-based Architecture for Monitoring Cyber-Physical Critical Infrastructures. In 2022 6th Cyber Security in Networking Conference (CSNet) (pp. 1-7). IEEE. DOI: 10.1109/CSNet56116.2022.9955607

[11] Cao, L., & Sharma, P. (2021, December). Co-locating containerized workload using service mesh telemetry. In Proceedings of the 17th International Conference on Emerging Networking EXperiments and Technologies (pp. 168-174). https://doi.org/10.1145/3485983.3494867

[12] Duque, A. O., Klein, C., Feng, J., Cai, X., Skubic, B., & Elmroth, E. (2022, May). A qualitative evaluation of service mesh-based traffic management for mobile edge cloud. In 2022 22nd IEEE International Symposium on Cluster, Cloud and Internet Computing (CCGrid) (pp. 210-219). IEEE. DOI: 10.1109/CCGrid54584.2022.00030

[13] Saleh Sedghpour, M. R., Klein, C., & Tordsson, J. (2022, April). An empirical study of service mesh traffic management policies for microservices. In Proceedings of the 2022 ACM/SPEC on International Conference on Performance Engineering (pp. 17-27). https://doi.org/10.1145/3489525.3511686

[14] Kang, M., Shin, J. S., & Kim, J. (2019, January). Protected coordination of service mesh for container-based 3-tier service traffic. In 2019 International Conference on Information Networking (ICOIN) (pp. 427-429). IEEE. DOI: 10.1109/ICOIN.2019.8718120

[15] El Malki, A., & Zdun, U. (2019). Guiding architectural decision-making on service mesh-based microservice architectures. In Software Architecture: 13th European Conference, ECSA 2019, Paris, France, September 9–13, 2019, Proceedings 13 (pp. 3-19). Springer International Publishing. https://doi.org/10.1007/978-3-030-29983-5_1

[16] Ganguli, M., Ranganath, S., Ravisundar, S., Layek, A., Ilangovan, D., & Verplanke, E. (2021, September). Challenges and opportunities in performance benchmarking of service mesh for the edge. In 2021 IEEE international conference on edge computing (EDGE) (pp. 78-85). IEEE. DOI: 10.1109/EDGE53862.2021.00020

[17] Houmani, Z., Balouek-Thomert, D., Caron, E., & Parashar, M. (2020, May). Enhancing microservices architectures using data-driven service discovery and QoS guarantees. In 2020 20th IEEE/ACM International Symposium on Cluster, Cloud and Internet Computing (CCGRID) (pp. 290-299). IEEE. DOI: 10.1109/CCGrid49817.2020.00-64

[18] Brown, J., & Harris, A. (2021). Cloud Network Monitoring and Threat Detection with ExtraHop Reveal(X) in AWS. Proceedings of the 2021 IEEE International Symposium on Cloud Computing, 45-50. DOI: https://doi.org/10.1109/CloudCom51993.2021.00015

[19] Kumara, I., Han, J., Colman, A., & Kapuruge, M. (2016). Software-defined service networking: performance differentiation in shared multi-tenant cloud applications. IEEE Transactions on Services Computing, 10(1), 9-22. DOI: 10.1109/TSC.2016.2594061

[20] Ren, Y., Shen, S., Ju, Y., Wang, X., Wang, W., & Leung, V. C. (2022, May). Edgematrix: A resources redefined edge-cloud system for prioritized services. In IEEE INFOCOM 2022-IEEE Conference on Computer Communications (pp. 610-619). IEEE. DOI: 10.1109/INFOCOM48880.2022.9796939

[21] Sedghpour, M. R. S., & Townend, P. (2022, August). Service mesh and ebpf-powered microservices: A survey and future directions. In 2022 IEEE International Conference on Service-Oriented System Engineering (SOSE) (pp. 176-184). IEEE.DOI: 10.1109/SOSE55356.2022.00027

[22] Luo, S., Xu, H., Lu, C., Ye, K., Xu, G., Zhang, L., ... & Xu, C. (2022). An in-depth study of microservice call graph and runtime performance. IEEE Transactions on Parallel and Distributed Systems, 33(12), 3901-3914. DOI: 10.1109/TPDS.2022.3174631

[23] Liu, H. et al. (2019). JCallGraph: Tracing Microservices in Very Large Scale Container Cloud Platforms. In: Da Silva, D., Wang, Q., Zhang, LJ. (eds) Cloud Computing – CLOUD 2019. CLOUD 2019. Lecture Notes in Computer Science(), vol 11513. Springer, Cham. https://doi.org/10.1007/978-3-030-23502-4_20

[24] Kakivaya, G., Xun, L., Hasha, R., Ahsan, S. B., Pfleiger, T., Sinha, R., ... & Gupta, I. (2018, April). Service fabric: a distributed platform for building microservices in the cloud. In Proceedings of the thirteenth EuroSys conference (pp. 1-15). https://doi.org/10.1145/3190508.3190546

[25] Shadija, D., Rezai, M., & Hill, R. (2017, December). Microservices: granularity vs. performance. In Companion Proceedings of the10th international conference on utility and cloud computing (pp. 215-220). https://doi.org/10.1145/3147234.3148093

[26] Rao, V. et al. (2021). Scheduling Microservice Containers on Large Core Machines Through Placement and Coalescing. In: Klusáček, D., Cirne, W., Rodrigo, G.P. (eds) Job Scheduling Strategies for Parallel Processing. JSSPP 2021. Lecture Notes in Computer Science(), vol 12985. Springer, Cham. https://doi.org/10.1007/978-3-030-88224-2_5

[27] Fu, K., Zhang, W., Chen, Q., Zeng, D., & Guo, M. (2021). Adaptive resource efficient microservice deployment in cloud-edge continuum. IEEE Transactions on Parallel and Distributed Systems, 33(8), 1825-1840. DOI: 10.1109/TPDS.2021.3128037

[28] Hu, Y., de Laat, C., & Zhao, Z. (2019). Optimizing service placement for microservice architecture in clouds. Applied Sciences, 9(21), 4663. https://doi.org/10.3390/app9214663

[29] Kumara, I., Han, J., Colman, A., van den Heuvel, W. J., Tamburri, D. A., & Kapuruge, M. (2019). SDSN@ RT: A middleware environment for single-instance multitenant cloud applications. Software: Practice and Experience, 49(5), 813-839. https://doi.org/10.1002/spe.2686

[30] Liu, L., He, X., Tu, Z., & Wang, Z. (2020, November). Mv4ms: A spring cloud based framework for the co-deployment of multi-version microservices. In 2020 IEEE International Conference on Services Computing (SCC) (pp. 194-201). IEEE. DOI: 10.1109/SCC49832.2020.00033

[31] Gan, Y., Zhang, Y., Cheng, D., Shetty, A., Rathi, P., Katarki, N., ... & Delimitrou, C. (2020). Unveiling the hardware and software implications of microservices in cloud and edge systems. IEEE Micro, 40(3), 10-19. DOI: 10.1109/MM.2020.2985960

[32] Fu, K., Zhang, W., Chen, Q., Zeng, D., Peng, X., Zheng, W., & Guo, M. (2021, May). Qos-aware and resource efficient microservice deployment in cloud-edge continuum. In 2021 IEEE International Parallel and Distributed Processing Symposium (IPDPS) (pp. 932-941). IEEE. DOI: 10.1109/IPDPS49936.2021.00102

[33] Houacine, F., Bouzefrane, S., & Adjaz, A. (2016). Service architecture for multi-environment mobile cloud services. International Journal of High Performance Computing and Networking, 9(4), 342-355. https://doi.org/10.1504/IJHPCN.2016.077830

[34] Schmidt, R., & Nikaein, N. (2021). RAN engine: Service-oriented RAN through containerized micro-services. IEEE Transactions on Network and Service Management, 18(1), 469-481. DOI: 10.1109/TNSM.2021.3057642

[35] Kratzke, N., Quint, PC. (2017). Investigation of Impacts on Network Performance in the Advance of a Microservice Design. In: Helfert, M., Ferguson, D., Méndez Muñoz, V., Cardoso, J. (eds) Cloud Computing and Services Science. CLOSER 2016. Communications in Computer and Information Science, vol 740. Springer, Cham. https://doi.org/10.1007/978-3-319-62594-2_10