International Journal of Cloud Computing and Cyber Security (IJCCCS)
ISSN: XXXX-XXXX / Volume 1 Issue 2 November-December 2026
Paper Id: JCCCS -V1I2P106/ Page No: 73-86
Doi: XX.XXXXX/XXXXXXX/ IJCCCS -V1I2P106

Original Article

Reducing Energy Consumption in Continuous Cloud Threat Monitoring Through Adaptive Sampling

Wenke Lee[1] , Taesoo Kim[2], Frank Li[3]

[1,2,3] *Georgia Institute of Technology.*

Received Date: 28 November 2026 Revised Date: 02 December 2026 Accepted Date: 06 December 2026

Abstract: *Continuous threat monitoring of cloud-based systems is critical for recognizing potential cyber attacks; identifying unusual behavior; and safeguarding confidential data from dynamic, evolving cloud-based systems. However, ongoing gathering and analysis of security related telemetry can require significant amounts of computing power and energy, especially in larger-scale systems with large volumes of logging; network-event; authentication; and system-metrics data. This study presents a dynamic sampling methodology to provide both energy savings and adequate threat detection. The dynamic sampling framework adjusts the monitoring rate based on workload level; previous threat history; system sensitivity; and current indicators of possible threats. When the risk level is lower, the framework decreases the sampling rate to eliminate unnecessary data processing, storage, and computation. When anomalies occur or when the risk level increases, the monitoring system increases the sampling rate to enhance visibility and enable quick identification of incidents. Risk-based decision making, feedback-driven adaptation, and energy-aware resource allocation are used in the development of the framework to balance security with operational efficiency. Simulations of cloud workloads and representative security event data sets could be used to evaluate the dynamic sampling framework by measuring energy consumption; detection accuracy; false-positive rates; response times; and monitoring overhead. This research establishes a theoretical basis for sustainable cloud security monitoring and provides guidance on implementation issues such as sampling fairness; gaps in visibility; and resistance to adversarial manipulation.*

Keywords: *Cloud security, adaptive sampling, energy efficiency, threat monitoring, cloud computing, anomaly detection, security telemetry, sustainable computing, resource optimization, cybersecurity.*

I. INTRODUCTION

The growth of cloud computing has changed how businesses organize their data, create and run their programs, oversee their digital services and provide access to computer power. Cloud based systems offer scalability, adaptability, central control, and affordable options for accessing computer power. These benefits also present new security issues as cloud based systems work on multiple networks, virtual machines, containers, application programming interfaces, and linked services. Therefore, businesses must continually watch over their systems to find out about unauthorized entry, bad network traffic, people getting higher privileges than they should have, taking away business secrets and other types of attacks. As such, watching over systems continuously is a key element of modern day cloud security designs. Most traditional security monitoring systems gather and review a lot of security information at set times or continuously. While traditional security monitoring systems show what is happening in a system, they use up a lot of computer resources (CPU), take up a lot of space on hard drives, create a lot of network traffic, and require a lot of electricity. This issue gets even bigger when businesses use many cloud-based systems that include thousands of workloads and millions of security events. Because of this, researchers and administrators of cloud-based systems are finding more efficient ways to watch over their systems and protect them from attacks while using less electricity. Recent studies on using less electricity to detect intrusions and creating green edge-cloud security show that we need to find a balance between protecting our systems from attacks and conserving our resources.

A significant drawback to traditional cloud threat monitoring is that it relies on static data collection policies which do not adequately address the dynamic nature of threats to security. While continuous high frequency monitoring can be useful while there are potential threats present; continuing to use high levels of monitoring when there is little to no potential threats present can lead to wastefulness in terms of computational time and energy used. On the other hand, decreasing the level of monitoring based solely on the assumption that the threat level has decreased can cause gaps in visibility and therefore delay the ability to identify potential attacks. Adaptive sampling may provide a possible answer to this problem by allowing organizations to

This is an open access article under the CC BY-NC-ND license (https://creativecommons.org/licenses/by-nc-nd/2.0/)

dynamically adjust the frequency and amount of security related information they collect based upon how their systems are behaving, past experiences with threats, the characteristics of the work being performed, and what they believe the current threat level to be. In general, an adaptive monitoring system will typically operate at a much lower sampling frequency to minimize the cost associated with collecting, processing, transmitting, and storing large amounts of data. However, if anomalies such as unauthorized authentication attempts, unexpected network connections, unusually high resource consumption or any other indicator that indicates a threat exists, then the adaptive monitoring system will increase its sampling frequency and provide more computing power for those workloads that are exhibiting these types of behaviors. This type of strategy would require organizations to find the optimal balance between minimizing energy consumption while maintaining a rapid response capability to threats. Although some researchers have looked into adaptive monitoring, intelligent intrusion detection, and energy aware security operations, none have addressed how organizations should go about making sampling decisions in a continuous cloud threat monitoring environment taking into account detection delays, false negatives, operational complexity, and potential manipulation by adversaries.

This paper presents an adaptive sampling strategy to minimize the amount of energy consumed during continuous cloud threat monitoring while providing the necessary level of security visibility. An adaptive sampling strategy consists of four components: a risk estimation component; a workload awareness component; an escalation component; and a feedback component. The core concept behind adaptive sampling is that monitoring levels must change dynamically based on the risk associated with a particular resource or operational state, rather than being fixed for all resources or states. Risk indicators used by the proposed strategy include recent security alerts, authentication anomalies, network traffic patterns, workload criticality, and past frequencies of detected threats. Risk indicators are provided to a sampling controller that determines the most appropriate sampling rate for a given resource or state within defined operational parameters. In addition to optimizing energy usage through sampling rates, the proposed strategy includes minimum sampling rates for critical workloads, immediate escalation procedures for high-risk events, and mechanisms to protect against sudden drops in monitoring coverage. The purpose of this research is to create a formal methodology to evaluate energy consumption, threat detection accuracy, detection delay, false positive rate, and computational overhead under various sampling strategies. The primary contribution of this research is a risk-based approach to sustainable cloud security that views energy efficiency and threat detection as mutually supportive goals, not mutually exclusive operations. Controlled simulation studies, representative security data sets, and actual measurement of computing energy consumption can be utilized to examine each of these areas of interest collectively, in order to assist in developing scalable cloud monitoring strategies that allow for minimizing unnecessary computational overhead while maintaining adequate security effectiveness.

II. PROPOSED ADAPTIVE SAMPLING FRAMEWORK FOR ENERGY-EFFICIENT CLOUD THREAT MONITORING

The adaptive sampling framework proposed here offers a methodical way to decrease the energy used in continuous cloud threat monitoring, while still offering adequate protection against cyber threats. There are four primary parts of the framework: telemetry collection; risk assessment; adaptive sampling control; and performance evaluation. Telemetry collection involves obtaining security-related data from cloud infrastructure. Telemetry data includes but is not limited to: authentication log data; network traffic data; application event data; resource usage data for systems; and access control activity data. Rather than collecting all possible events at the same rate as one another, the framework classifies the monitored resources based upon how important they are to business operations, what their typical behaviors look like, and the amount of potential security risks associated with each. The telemetry data that was collected is provided to the risk assessment part of the framework. This part of the framework assesses the current state of affairs and determines when there have been deviations from established behavioral patterns. The architectural design of the framework allows for monitoring resources to be assigned according to the needs of the actual security situation, rather than simply having them operate on a predetermined schedule (i.e., fixed interval sampling). The goal of the framework is to reduce the number of times that events are recorded during periods where operations are stable. Reducing the recording rate during stable operational periods will reduce the number of processors being used, the amount of memory being utilized, the amount of data being transferred over the network, and the amount of data being written into storage.

The risk assessment phase of the monitoring methodology uses both past attack data and real-time indications of potential malicious activity to determine the proper level of monitoring intensity. Potential indications of possible malicious activity could include, but are not limited to; multiple failed attempts at authentication, unusual times of accessing resources, excessive amounts of data transferred over the network, abnormally large consumption of resources, and changes to user permissions. Using this information, each monitored workload is classified as either low, moderate, or high risk. Workloads classified as low

risk will have longer sampling periods, while moderate risk workloads will be sampled more frequently. Workloads classified as high risk will receive enhanced monitoring in order to increase the likelihood of detecting new attacks as soon as possible. Additionally, the methodology can take into account the sensitivity of a workload and the security policy of an organization, so that workloads providing critical functions retain adequate monitoring coverage regardless of how low the current risk estimate for that workload appears. A feedback loop continually assesses whether the current observed behavior matches the original estimate of risk for that workload, and makes adjustments to sampling strategies as necessary. This feedback loop is designed to minimize rapid fluctuations in sampling frequencies and provide a more stable environment for allocating resources. Finally, although sampling strategies are used to help manage the amount of data generated by the monitoring system, they should never be substituted for mandatory security logging or event-triggered alerts, which provide complete records of activity for purposes of demonstrating compliance, investigating incidents, or responding to security breaches.

The adaptive sampling controller will convert risk assessments into operational monitoring decisions. For instance, when a cloud-based application is running normally (i.e. experiencing normal traffic and consistent user authentication patterns), it will typically use a relatively low sampling frequency; however, if there is a sudden increase in suspicious requests, the controller can temporarily switch to a higher frequency monitoring decision. A simple controller could be represented by the equation \(f_t=f_{\min}+(f_{\max}-f_{\min})R_t\). In this equation, \(f_t\) is the chosen sampling frequency, \(f_{\min}\) and \(f_{\max}\) are the allowable frequencies for sampling, and \(R_t\) is a normalized risk score that ranges from zero to one. To turn this relationship into a practical production system, we must implement additional safety features, consider workload constraints, and use smoothing techniques to ensure reliable performance. Additionally, the controller must have minimum observation rate requirements, quick escalation rules, and only return to lower sampling frequencies once risk indicators have stabilized. Evaluations of the proposed framework must be compared to static-rate monitoring under the same workload, attacks, and hardware conditions. Metrics used in evaluations should include energy consumption, detection rate of threats, false positive rate of detection, detection latency, CPU utilization, memory usage, and telemetry data volumes. To establish whether any improvement observed is statistically significant, repeated evaluations must be performed using the same measurement techniques.

The effectiveness of the proposed framework relies upon a balance between energy efficiency, security visibility, and operational reliability. If the sampling rate is reduced too significantly, then brief periods of attack can go undetected; this is especially true if malicious activity takes place during those times that are not being observed. To address this issue, the framework must incorporate both an adaptive periodic sampling scheme and an event-driven alert mechanism along with targeted high-frequency collection and continuous monitoring of critical security signal(s). In addition, it would be beneficial to perform adversarial testing, since attackers will likely try to mimic normal user behavior or alter risk indicators to encourage lower monitoring rates. Other security protections may include: independently escalated trigger mechanisms; randomized sampling rates within acceptable parameters; protected audit logs; and periodic evaluation of sampling policy. In terms of energy management, the data collected should differentiate between the energy used for monitoring purposes and the base level of energy needed to maintain the operations of cloud-based applications. This differentiation will allow investigators to determine how much energy is saved due to the implementation of the adaptive sampling strategy. It should be noted that the purpose of the proposed framework is to create a basis for experimental research rather than to assume that there will be a specific degree of improved performance resulting from its adoption. Ultimately, the success of the adaptive sampling strategy will need to be demonstrated through replicated experimental studies showing that the adaptive sampling strategy results in reduced energy usage while maintaining sufficient levels of security detection performance across a variety of cloud-based application types.

III. ENERGY OPTIMIZATION AND INTELLIGENT RISK-BASED MONITORING IN CLOUD ENVIRONMENTS

A. Energy-Aware Telemetry Collection and Adaptive Sampling Strategies

The ability to collect energy-aware telemetry is a critical element in decreasing the operating costs of long-term continuous threat monitoring in the cloud. In many cases, traditional security monitoring systems collect logs, network events, authentication records, and system performance metrics on a consistent basis (e.g., every five minutes), regardless of how active the workload is or what type of threats currently exist. Although this provides consistency in visibility, it can also lead to excess processing and data transmission during periods when the system is behaving normally. An alternative approach to addressing this issue is an adaptive sampling strategy that adjusts the rate of telemetry collection based on workload behavior, the estimated likelihood of a threat occurring, and the level of sensitivity to the resource. For example, during times when there is little to no evidence of a potential threat, the monitoring system can lower its sampling rate for non-critical resources. By doing so, it will use less processor time, less memory space, less network bandwidth, and less storage space.

Conversely, if suspicious activity is identified, then the monitoring system can immediately raise its sampling rate to gather additional information about the suspicious activity. There are several factors that could be used by the sampling controller to determine the optimal collection rate. These include previous patterns of behavior, anomaly scores generated from security alerts, irregular authentication behaviors, variations in network traffic patterns, and the severity level of any generated security alerts. Additionally, energy-aware scheduling can be used to coordinate the collection of telemetry with the availability of computational resources. This can prevent unnecessary allocation of resources. While lowering sampling rates can reduce resource utilization, it should never result in the loss of critical security-related records, nor should it disrupt compliance with applicable regulations or impede forensic investigations. Therefore, critical workloads should always have minimum observation rates, and mandatory audit logs as well as high-priority security events should be available at all times when required. To enhance reliability, the framework can integrate both periodic sampling and event-driven collection. This would allow for immediate investigation of urgent events even if the normal sampling rate is operating at a lower frequency. Evaluation of performance should compare adaptive sampling with fixed-rate monitoring under identical workload conditions. Measurements of interest should include energy usage, telemetry volume collected, CPU utilization, detection accuracy, and response time. Overall, this strategy creates a systematic framework for achieving balance between energy efficiency and dependable visibility over distributed cloud infrastructure.

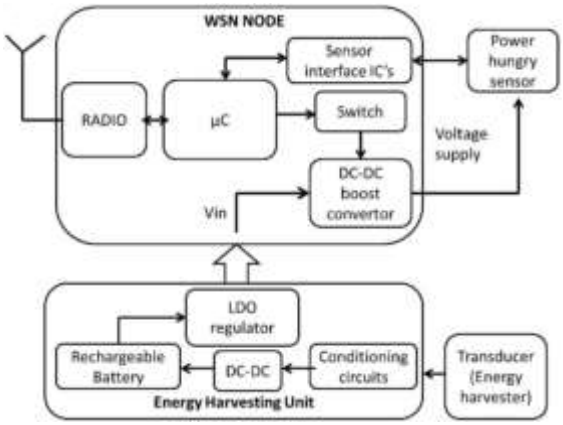


Figure 1: Generic Block Diagram of a WSN Platform

B. Intelligent Threat Detection, Feedback Control, and Performance Evaluation

The smart risk evaluation and feedback controller provide the ability to adaptively react to changing cloud security issues. Initially, the monitoring system creates behavioral baselines for each cloud workload by analyzing standard authentication, network communication, application activity, and resource usage patterns. Then, the system compares incoming observations with those baseline patterns to determine if there have been deviations from the normal behavior which would indicate unauthorized access, abnormal privileges changes, excessive traffic being generated, or an attempt to steal data. A risk evaluation component assigns a threat score to every monitored workload and sends that score to the adaptive sampling controller. As long as the threat score is low, the controller uses an energy efficient sampling configuration. However, when the threat score increases, the controller increases the frequency of sampling and provides more analysis capability for the workload that is experiencing high threat scores. Subsequent observations, security alerts, and investigation results allow the controller to modify its decision-making processes.

There are also hysteresis thresholds and minimum observation requirements that can limit unnecessary changes in sampling configurations. Additionally, independent emergency triggering mechanisms can be implemented to prevent attackers from using the reduced sampling frequency to their advantage. To validate the effectiveness of this methodology, researchers must conduct controlled experiments that use normal workloads, simulated attacks, and different levels of demand for cloud resources. Evaluation measures should include total energy consumed during monitoring, the percentage of threats identified, the false-positive rate, the average time to identify threats, telemetry reduction, and computational overhead. Researchers should use a clearly identified baseline to evaluate energy savings; the baseline should utilize similar hardware, workloads, and security situations when possible. Researchers should also determine if reductions in sampling affect certain types of workloads more than others or if they hide short-term attacks. Conducting multiple experiments and reporting clearly will help researchers verify if improvements are consistently reproducible or if they depend on specific experimental setups. Through the integration of

intelligent risk evaluation, adaptive feedback, and thorough performance evaluation, the proposed methodology aims to maintain sustainable cloud threat monitoring while providing adequate cybersecurity protection and operational reliability.

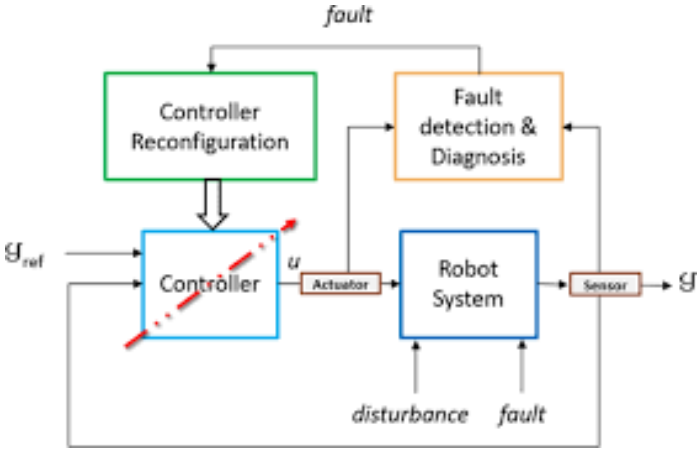


Figure 2: Fault Diagnosis and Controller Reconfiguration

IV. IMPLEMENTATION AND OPERATIONAL EVALUATION OF ADAPTIVE CLOUD THREAT MONITORING

A. Risk-Based Sampling Frequency Management

The ability to manage the frequency at which risk-based samples are taken will allow cloud-based monitoring systems to adjust their allocation of computational resources in response to evolving security conditions. Risk-based sampling frequency management is a method of managing the frequency at which samples are collected based upon the classification of monitored workloads. Indicators that may be used to classify monitored workloads include but are not limited to; authentication failures, unusual network traffic, privilege modifications, previous attacks, and application behaviors. Depending on the perceived risk level associated with the workload, the sampling controller will determine the proper sampling rate. Routine monitoring of low-risk workloads will occur less frequently than routine monitoring of moderate-risk workloads. Moderate-risk workloads will be monitored at a standard rate. High-risk workloads will be sampled more frequently than standard rate monitoring and analyzed further to identify possible threats quickly. Additionally, the sampling controller must assess workload criticality, compliance requirements and minimal monitoring standards prior to adjusting sampling intervals. Adjusting sampling intervals based on these factors will prevent energy optimization from compromising the safety of sensitive applications. Also, escalation thresholds that have been established prior to the adjustment of sampling intervals can automatically increase the frequency of sampling if suspicious events occur. Because the framework continuously assesses risk indicators, it is better suited to operate in dynamic cloud environments than a fixed frequency monitoring approach. It is intended to minimize redundant telemetry data collection while maintaining adequate visibility into potential malicious actions.

B. Cloud Resource Utilization and Energy Consumption Reduction

Energy consumption for continuous security monitoring is determined by how much energy is being used up by your cloud resources. Your system will use more energy as you collect, send over the wire, analyze, and save a large amount of data (telemetry). You will use up more processor cycles, memory space, network traffic, and storage when collecting/sending/analyzing/saving large amounts of data. This means that if you're continually collecting/sending/analyzing/saving large amounts of data, you'll be using up more energy. Your energy use will go down if you stop doing this. Our suggested method will reduce the number of times you collect/send/analyze/save data. It will do this through adaptive sampling. When you're having low-risk periods in terms of security threats and you have a good idea of what's going to happen next, then your agents don't need to check in as often. Your agents can then focus on other things instead of checking in on a regular basis. As a result, you can spend less money on running your agents. Adaptive sampling also allows you to allocate your computing power to other areas of your virtual machines/containers/cloud applications. In addition, adaptive sampling will help you coordinate your monitoring activities with the processing capabilities of your hardware so that you aren't fighting for resources. You shouldn't assume that you'll save energy just because you aren't collecting as many events. To prove that you're saving energy, you need to measure how much energy you're using. You should compare your total energy usage for monitoring against a baseline measurement that was taken under similar operating conditions as before but with a constant rate of event collection. Also, remember to include the overhead of the controller and any extra processing that takes place during a

security incident in your comparison. Some examples of measurements you could take include CPU utilization, amount of telemetry collected, time it takes to process the data, and power used to run your systems. These types of measurements will allow you to see if your approach is actually saving you energy. These measurements will also help you decide whether your energy savings come at too high a cost in terms of operations.

C. Threat Detection Accuracy and Monitoring Reliability

When integrating adaptive sampling in cloud security operations, a key factor that must be met is to provide high levels of threat detection accuracy. Although a low-frequency sample can help to save energy, there are several negative consequences including an increased risk of missing short-duration attacks or delaying the time it takes to detect suspicious activity. Therefore, the proposed framework is designed to utilize a combination of risk-based sampling with event-triggered alerting, minimum observation periods and fast response mechanisms. Risk-based sampling will enable the system to focus its attention on threats that pose the greatest risk to the organization. Additionally, if the system is capable of detecting high-priority security signals (e.g., unauthorized privilege changes and critical access violations), it should immediately initiate an investigation into the matter. Furthermore, the framework must continuously record all relevant auditing data regardless of whether or not optional sampling reductions have been implemented. The reliability of the framework can be tested by applying it to representative security datasets and conducting controlled attack simulations. These simulations would involve simulating unauthorized access attempts, unusual network communications and suspicious resource usage. Some of the parameters used to evaluate the effectiveness of the framework include detection rates, false positive rates, false negative rates and mean time to detection. Researchers should then assess how well the results obtained via the proposed framework compare to those obtained via traditional fixed-rate monitoring techniques. Repeated experiments under varied workload conditions can demonstrate whether the framework continues to operate reliably in terms of its ability to identify potential threats. By providing such assurances, researchers can confirm that energy-saving measures do not compromise the essential security requirements of the organization.

D. Performance Assessment and Future Scalability

Performance assessment determines if adaptive sampling offers an improvement on traditional cloud monitoring techniques. An evaluation of performance should involve evaluating performance across various workload types; evaluate performance under different levels of traffic; evaluate performance with a variety of potential threats; and evaluate performance with a range of different sampling settings. Researchers can evaluate performance among three different approaches to cloud monitoring: (1) fixed rate cloud monitoring (fixed-rate); (2) risk-based adaptive sampling; and (3) event-driven cloud monitoring. Each approach should be evaluated under similar experimental conditions. In addition, the experiments should collect data regarding energy consumption, threat detection accuracy, reduction in telemetry, latency of response, CPU usage, and memory usage. It will be required to assess all aspects of performance together because some strategies for reducing energy consumption may introduce significant delay in detecting threats. Also, the scalability of each approach should be assessed by adding more workloads being monitored, more security events occurring, and more cloud resources available. As the complexity of the cloud environment increases, the framework should continue to provide reliable decisions and maintain low computational overhead. Finally, the ability of each approach to withstand adversary manipulation of risk scores resulting in the reduction of monitoring intensity below acceptable levels should be evaluated. All results obtained from these evaluations will help to improve the strategies used for determining which samples to collect, estimating risk associated with each sample, and allocating resources appropriately. To validate that each approach is effective and viable for use in large-scale energy efficient cloud security environments, it is imperative that the research methodology employed in evaluating the approaches be documented clearly and the results generated be published openly.

V. COMPARATIVE ANALYSIS OF ADAPTIVE SAMPLING AND CONVENTIONAL CLOUD THREAT MONITORING

The continuous threat monitoring in the cloud has to find a balance among the effectiveness of security, its reliability of operation and how much energy it consumes. Most of the time, traditional methods of monitoring either produce security telemetry on a regular schedule (for example every hour) or process all events that have been generated in the cloud; this produces consistent visibility over what is happening in the system but may also lead to excessive use of computer resources when there is little activity going on. On the other hand, the adaptive sampling method we propose produces monitoring data according to the level of estimated threat, the type of work being performed by the workload, and the amount of resources that are sensitive to degradation. This method aims to minimize unneeded processing of telemetry data, while at the same time providing enough observations to identify suspicious behaviors. The following table provides a conceptual comparison of the conventional monitoring methods and our adaptive sampling method. The comparison illustrates expected operational characteristics of each method instead of experimental results. Therefore, actual performance improvements depend upon

several factors including the architecture of the cloud, the way the workload is distributed, the characteristics of the attack, the sampling policy, and the quality of the implementation. Therefore, energy consumption, accuracy in detecting attacks, latency in responding to attacks, and the computational overhead should be evaluated through controlled experiments prior to making any conclusions regarding the practical application of our adaptive sampling method.

Table: Adaptive Sampling for Cloud Threat Monitoring

Evaluation criterion	Conventional monitoring	Proposed adaptive sampling
Sampling frequency	Fixed or continuously active	Dynamically adjusted according to risk
Energy consumption	May remain high during low-risk periods	Aims to reduce unnecessary monitoring expenditure
Telemetry collection	Uniform or continuous collection	Risk-sensitive collection frequency
Threat detection	Depends on monitoring coverage and detection algorithms	Intensifies monitoring when suspicious activity emerges
CPU utilization	May increase with continuous processing	Potentially reduced during stable periods
Storage requirements	May generate substantial telemetry volumes	Aims to reduce routine telemetry volume
Detection latency	Depends on collection frequency and processing capability	May improve during escalated monitoring but can increase at lower sampling rates
False-positive management	Depends on detection rules and analytical models	Can incorporate risk assessment and feedback mechanisms
Critical workload protection	Determined by predefined monitoring policies	Maintains minimum sampling requirements for sensitive resources
Scalability	Resource requirements may increase with telemetry volume	Risk-based allocation may improve resource utilization
Operational complexity	Often simpler to configure	Requires risk estimation, feedback control, and policy management
Performance evaluation	Measures security and resource utilization	Measures security, energy consumption, and adaptive control overhead

Comparative analysis suggests that adaptive sampling provides an effective way to improve the efficiency of cloud threat monitoring, especially during periods where the level of activity on each workload varies significantly, as well as the risk associated with each workload. The main benefit of this type of sampling is that it allows the organization to allocate monitoring resources based on current operational conditions instead of sampling at the same rate regardless of operational conditions. However, the use of this type of sampling creates additional complexity since risk assessment, sampling decision making, and feedback loops all require computing resources and proper setup. Furthermore, if attacks occur in-between sampling intervals or if the attack pattern looks like normal behavior then there may be short-term visibility gaps. Therefore, the proposed framework should include event-driven alerting, minimum sampling thresholds, independent emergency trigger mechanisms and ongoing recording of mandatory security events. In addition, experimental testing should evaluate both types of sampling under similar workloads and attacks. Testing should also measure energy consumption, threat detection rates, false positive rates, detection latency and processing overhead. This testing will allow organizations to assess if the increased complexity is worth the potential energy savings and if security performance remains acceptable in varying types of cloud environments. Ultimately, adaptive sampling should be used in conjunction with existing security controls rather than replacing essential monitoring functions; thereby allowing organizations to implement more sustainable and responsive cloud security operations.

VI. CHALLENGES AND PRACTICAL CONSIDERATIONS IN ADAPTIVE CLOUD THREAT MONITORING

Adaptive sampling is used as part of a continuous cloud threat monitoring strategy. However, there are numerous technical and operational challenges associated with using adaptive sampling as a means to optimize energy use in continuous cloud threat monitoring. One such challenge is establishing the appropriate sampling rate for various types of applications that operate in the same cloud environment. Different applications running in the cloud present varying degrees of risk to the organization and have varying levels of operational sensitivity. Therefore, establishing a single sampling rate for all of the applications in the cloud may

result in too little monitoring for high-risk applications and/or excessive monitoring for low-risk applications. To overcome this limitation, the proposed framework should classify workloads based on their risk level; establish risk-based threshold values; evaluate past behavior; and provide minimum sampling requirements. Another challenge related to implementing adaptive sampling is accurately assessing energy consumption attributable to monitoring activity. Processor utilization, memory usage, network traffic, and storage-related activities all contribute differently to total energy consumption. Therefore, researchers need to create a baseline for energy consumption during normal operations, and then compare energy consumption during monitoring activities against the baseline. Additionally, researchers should consider the computational overhead created by conducting risk assessments and making adaptive decisions when evaluating the feasibility of implementing adaptive sampling as a means to reduce energy consumption in continuous cloud threat monitoring.

Another factor to consider is the potential loss of visibility into a threat if you Sampl with a less frequent sampling schedule. Many types of cyber attacks happen in very short amounts of time, therefore, if you only Sample at long intervals, you will be unable to see those types of attacks. Attackers will also try to act like they are normal users, change the amount of work being done on the computer, or use the fact that they know when the sampling is scheduled to avoid detection. To reduce the risk of these things happening, the framework needs to use both random periodic sampling and event based sampling as well as independent security alerts, and high frequency sampling as needed during suspicious activity. Random sampling based on fixed operating parameters can also help make it harder for the attacker to predict when you will be sampling. Even though routine sampling may vary, all of the critical audit records and security events needed for regulatory compliance should continue to be available. Feedback loops should constantly evaluate how effective the monitoring is based upon the results of detections, how the user's workload is changing, and security incidents. Finally, access control and integrity controls are essential to prevent unauthorized changes to sampling policies. These controls allow for consistent monitoring coverage while at the same time allowing the system to process fewer telemetry data during periods where there is little variation in activity.

When implementing adaptive sampling on large-scale cloud architectures, scalability, inter-operability, and operational management should be considered. Large-scale cloud architectures typically involve a mixture of virtual machines (VMs), containerized applications, distributed systems, multiple cloud providers, and central security event management systems. Therefore, the proposed adaptive sampling framework should allow for standardized telemetry formats, modular integrations and policy management of heterogeneous resources. The performance of the adaptive sampling framework should be assessed based upon its ability to maintain consistent energy reduction as the size of the monitored environment increases; this assessment should include evaluation under varying levels of workload intensity, attacks, and architectural configurations. Metrics used to evaluate the adaptive sampling framework could include: energy consumption, telemetry reduction, detection accuracy, false positive rate, response latency, and overhead associated with the controller. Additionally, researchers should assess the adaptive sampling framework's recovery characteristics in instances where monitoring components fail, risk scores lose reliability, or communication between cloud service providers is lost. Fallback mechanisms may be developed to restore a predetermined level of monitoring when decisions made by the adaptive sampling framework are no longer reliable. To successfully implement adaptive sampling as a practical component of sustainable cloud security, researchers need to develop an adaptive sampling architecture that includes security safeguards, and provides for repeatable evaluations, that demonstrate measurable reductions in energy usage while providing the same level of visibility, reliability, and responsiveness needed for effective detection of threats.

VII. ADVANCED OPTIMIZATION TECHNIQUES FOR ADAPTIVE CLOUD THREAT MONITORING

A. Machine Learning-Based Threat Risk Prediction

Risk prediction based on machine learning can enhance adaptive sampling through finding commonality in the behavior of emerging security risks within cloud computing environments. The proposed framework can use historical security log data, authentication data, network traffic characteristics, and system performance metrics to create behavioral models of each workload. Supervised learning models can categorize previously identified behaviors, whereas unsupervised anomaly detection methods can recognize deviations from established behaviors. Risk estimates generated by these models can assist the sampling controller in deciding whether to increase, decrease or maintain current levels of observation. For example, if there have been numerous failed authentication attempts in conjunction with unusual network traffic patterns, this would indicate a higher level of risk and therefore justify increased telemetry collections. On the other hand, if a workload continues to exhibit normal behaviors for an extended period of time, then it may be able to be monitored less frequently as long as minimum security requirements are met. It is important to evaluate the effectiveness of the model using representative data sets and relevant performance metrics such as precision, recall, false positive rates, and detection latency. Additionally, the framework must

account for model drift, missing observations, and adversarial manipulation. By combining predictive analysis with rule-based safety mechanisms, more accurate sampling decisions can be made while reducing unnecessary computational costs.

B. Dynamic Resource Allocation and Monitoring Coordination

Dynamic allocation of resources provides a way for a cloud monitoring system to dynamically adjust the amount of computing power it provides to process telemetry from various applications, depending on the volume of telemetry generated by each application and its security needs. In large cloud systems, there are often applications that generate vastly different amounts of telemetry and have different security needs. Providing the same amount of computing power to every monitoring application will likely result in inefficient use of the available computing power and increased energy consumption. This framework can coordinate sampling rate, priority for processing telemetry, and analytical resource allocation based on workload criticality, current system usage, and potential threat level. If a workload becomes highly vulnerable to attack, more processing power can be provided to support security event processing; whereas, if a workload remains relatively safe, less routine collection can be performed. When coordinating resource allocation, consideration should be given to available CPU cycles, memory usage, network bandwidth and the cost of moving telemetry data among distributed services. Aggressive resource reallocation can cause delay or impact application performance if the monitoring tasks compete with essential workloads. Therefore, appropriate resource limits, scheduling priorities, and minimal service guarantees must be implemented. Evaluation of the framework should include measuring actual power consumed, processing overhead, and response time to security threats across various load levels. This coordination can improve the overall efficiency of monitoring operations, without implying that reducing telemetry will necessarily produce proportional energy savings.

C. Feedback-Driven Adaptation and Continuous Improvement

The feedback based adjustment enables the monitoring system to revise its sampling strategies based upon how well it detects security issues and performs operationally. Instead of solely using fixed sampling rates, the controller can regularly evaluate detection performance, false alarm rates, energy use, and workload patterns to decide if the existing configuration is still suitable. If suspicious events are identified, then the monitoring rate may be temporarily increased until the system determines that the threat level has decreased. As long as there are no further indications of instability, the controller may slowly reduce the sampling rate back down to a lower level as permitted by policy. Hysteresis limits help avoid rapid transitions from one sampling rate to another, while separate emergency trigger mechanisms provide protection from excessive reductions in monitoring coverage. To continuously improve the system, researchers must examine missed detections, investigate unusual behaviors exhibited by the system, and update risk models when workload patterns change. Additionally, researchers should incorporate several metrics into risk assessments in order to avoid having attackers exploit the feedback mechanism. Furthermore, researchers should protect the integrity of audit records that contain information used in risk assessments. Finally, researchers can assess the success of the adaptation process through repeated simulation studies incorporating different levels of traffic volume, duration of attacks, and availability of resources. By comparing the results of these simulations with those of fixed-rate monitoring, researchers will be able to determine whether the feedback-driven control approach results in increased energy efficiency without significantly compromising detection performance or operational reliability.

VIII. EXPERIMENTAL EVALUATION AND PERFORMANCE MEASUREMENT OF THE PROPOSED FRAMEWORK

Experimental evaluations are necessary in order to determine if the proposed adaptive sampling strategy will allow organizations to decrease power usage and still maintain effective detection capabilities for threats to cloud computing environments. In order to assess the performance of adaptive sampling against conventional static rate monitoring systems, both types of monitoring strategies must be applied to the same type of cloud workload (i.e., VMs, containers, auth), the same security settings (i.e., IDS/IPS, firewalls, etc.), and the same threat model (i.e., DOS, brute force, etc.). To create a controlled experimental environment, researchers can create virtual machines, deploy containerized applications, implement authentication systems, use traffic generators to generate network activity, and implement centralized security monitoring components. Researchers can then introduce normal operational activities and simulated security incidents into the environment to assess how the system responds to different risk levels. During the course of the experiment, researchers should measure energy consumption, telemetry volume, threat detection accuracy, detection latency, CPU usage, and memory overhead. If direct energy measurement is not feasible, researchers can employ validated power measurement techniques and differentiate between estimated and measured values. The following table identifies potential evaluation metrics, why they would be measured, and which way researchers expect them to improve based upon their hypotheses. However, these are merely hypotheses; researchers will need to demonstrate that their expectations hold true by performing repeated trials on identical hardware platforms with well-documented sampling configurations.

Table: Evaluation Metrics for Cloud Threat Monitoring

Evaluation metric	Measurement method	Expected outcome
Energy consumption	Measure monitoring-related power usage over time	Reduced energy expenditure
Telemetry volume	Count collected events and transmitted data	Lower routine data volume
Threat detection rate	Calculate detected attacks relative to total attacks	Maintain an acceptable detection rate
False-positive rate	Measure benign events incorrectly classified as threats	Avoid unnecessary increases
Detection latency	Measure time between attack initiation and detection	Maintain acceptable response time
CPU utilization	Record processor usage during monitoring	Lower average overhead during stable periods
Memory utilization	Measure memory consumed by monitoring components	Limit unnecessary resource usage
Sampling adaptation time	Measure the time required to change monitoring intensity	Rapid escalation under elevated risk
Controller overhead	Measure resources used by risk assessment and control	Keep adaptation costs low
Scalability	Increase workloads and telemetry generation rates	Maintain stable performance as workloads grow

The findings of the experiments conducted on the subject will allow us to assess whether there is a significant reduction in energy usage without causing an unacceptable degradation in the level of security provided. We will measure the amount of energy used as a percentage of the total amount of energy used under baseline conditions. In addition, we will compare the number of detections, false positives, and detection times of each of the experiments using the same attack scenario, so that we can find any possible trade-offs. Our experiments will involve short duration attacks; attacks with low frequency malicious activity; variations in workload intensity; and attempts to alter risk indicators. Multiple experiments may help determine if the observed effects are consistent, and we will perform statistical analysis on the data collected in order to quantify the degree of variability in the data. We must also take into account the energy consumed by our adaptive control systems; additional processing required during periods of increased risk; and storage and/or network costs related to providing feedback to our adaptive controls. We must present our results clearly and distinguish between the results obtained through measurement and those predicted by theory. If our adaptive sampling system reduces the volume of telemetry data collected but increases the rate at which detections go unnoticed, then we must modify our sampling strategies prior to implementing them in actual production environments. In general terms, we will consider our proposed framework successful only if we obtain experimental evidence demonstrating that our adaptive sampling system achieves significant energy efficiencies while maintaining an acceptable level of threat detection accuracy, operational reliability, and response performance within representative cloud environments.

IX. DEPLOYMENT STRATEGIES AND SECURITY GOVERNANCE FOR ADAPTIVE CLOUD THREAT MONITORING

A. Practical Deployment and Integration with Cloud Security Infrastructure

Deployment of the proposed adaptive sampling framework will require careful consideration of how it relates to currently deployed cloud security infrastructure so as not to affect the functioning of existing monitoring operations. The framework is designed to function as an additional level of control over the sampling frequency applied by telemetry collection agents prior to sending data to central security analytics platforms. To do so, the framework uses workload characteristics, security indicators, and resource utilization to determine the optimal sampling frequency. Possible mechanisms for integrating the framework into current cloud environments include; virtual machines, container orchestration platforms, cloud-native logging services, intrusion detection systems, and security information and event management solutions. Initially, during the deployment phase, the framework should run in observation mode, which means it will collect data regarding the workload behavior, but not make automatic decisions to change sampling frequencies. This will provide administrators with a way to create baseline measurements, identify critical resources, and assess the risk of implementing adaptive monitoring. Once the baseline measurements have been established, administrators can begin to apply sampling adjustments to selected non-critical workloads. Sampling policies should define minimum observation frequencies, escalation thresholds, and maximum permitted

intervals. Sampling frequencies should never be reduced below those required to capture critical authentication events, changes to privileged access, or mandatory audit records. Additionally, integration testing should verify compatibility with existing alerting mechanisms, incident response processes, and data retention requirements. A phased approach to implementing adaptive sampling will reduce disruption to operations, and allow for the correction of configuration issues prior to larger scale deployment.

A successful deployment of adaptive sampling requires an effective coordination process among the adaptive controller, cloud resource management systems and security operations personnel. The adaptive controller will provide communication of sampling decision making via authenticated interfaces and keep a record of why sampling intensity changes. Administrators will have access to view risk classification, review sampling adjustment and override automated decisions as required. In multi-cloud environments, consistent policy definition can facilitate better management while providing flexibility for each workload to adhere to environment-specific requirements. Deployment testing should include service disruptions, controller failure, unexpected traffic spikes and workload behavioral changes. When risk assessment is unavailable or unreliable, the system should default back to a predefined safe monitoring configuration rather than automatically reduce visibility. Performance dashboards can show energy usage, telemetry volume, sampling rates, detection results and resource utilization to facilitate operational oversight. Performance metrics provide organizations with data to assess if adaptive sampling offers enough benefit to justify implementing and maintaining it. Therefore, a practical deployment should include controlled rollout, inter-operability testing, transparent decision-making processes and clearly defined recovery procedures to create a solid basis for energy-efficient cloud threat monitoring.

B. Security Governance, Compliance, and Responsible Monitoring Optimization

Organizations need to develop security governance policies that are designed to provide the necessary guidance for implementing adaptive sampling practices, which are consistent with the organization's overall security strategy. While there is potential to reduce energy consumption by reducing the amount of data collected (telemetry) during sampling, organizations must be careful to avoid making sampling decisions that would result in lost auditing capabilities, reduced ability to conduct incident investigations, or failure to comply with relevant regulatory requirements. Therefore, the proposed framework should differentiate between optional monitoring observations and those security-related logs that are mandated to be collected and/or retained based on organizational policies and/or regulatory mandates. In addition, organizations should avoid disabling mandatory logging based solely on the fact that a workload has been identified as being low-risk. Organizations should define those workloads that are considered to be "critical," set forth minimum levels of monitoring for each workload, define the circumstances under which sampling adjustments may be made, and outline the criteria for escalating issues as soon as possible. Access to the adaptive controller should be limited through proper authentication, authorization, and change management processes. Organizations should document all policy changes, sampling decisions, and major risk assessments, and periodically review these documents to ensure that the monitoring strategy remains adequate and relevant to meet the organization's needs. These controls will help prevent organizations from using energy optimization as an excuse for failing to adequately monitor their systems. Finally, responsible governance includes assessing whether sampling decisions result in unevenly distributed protection among different application types or among different user groups.

Because cloud environments, threat models, and business objectives continuously evolve; the continuous monitoring of security controls is required. To maintain effectiveness, security personnel must periodically assess the efficacy of their detection capabilities, identified incidents that were not captured, false positive rates, energy consumption, and the effectiveness of their sampling policies. When gaps in monitoring are discovered, the framework must provide a means to rapidly update policy and return to higher frequency sampling. Additionally, independent security assessments may be used to determine if an attacker would be capable of altering the risk indicator used to trigger sampling; if an attacker could predict the sampling interval and use that knowledge to avoid detection; or if an attacker could disrupt the feedback loop that informs sampling decisions. Organizations should also document the limitations of adaptive sampling and create clear guidelines for responding to incidents where incomplete observations delayed detection. Energy efficiency should be evaluated along with security effectiveness, service reliability, and operational transparency as opposed to being viewed as the primary goal of optimization. When automated decision-making affects how intensely security personnel monitor activity, administrators must retain the ability to oversee the system's behavior and receive understandable explanations regarding its operation. Organizations should conduct periodic testing using simulated attack scenarios to ensure that emergency escalations and fallback mechanisms operate properly. By incorporating governance, compliance reviews, audits, and continuous improvement into their adaptive monitoring strategy, organizations can mitigate the risks associated with adaptive monitoring while striving for more sustainable cloud operations.

Ultimately, the proposed methodology must demonstrate evidence that decreasing spending on monitoring is consistent with maintaining the level of protection expected within the cloud security environment.

X. FUTURE RESEARCH DIRECTIONS AND RECOMMENDATIONS FOR ADAPTIVE CLOUD THREAT MONITORING

Continuous cloud threat monitoring will be enhanced by developing energy efficient adaptive sampling. Future studies should emphasize increasing accuracy of risk assessment, reducing cost associated with continuous cloud monitoring, and providing robust threat detection in dynamically changing cloud environments. These recommendations represent crucial areas for further developing the proposed framework and enabling it to be implemented effectively.

- **Advanced Machine Learning:** Future research can explore how to use machine learning models to identify new threats based on past security events, the behavior of workloads, and network traffic patterns. If predictive models are developed they could allow the monitoring controller to increase the rate at which samples are collected prior to a potential incident occurring. Researchers will need to assess the model's accuracy, computational resource needs, and ability to resist attacks aimed at manipulating the model.Further investigation should be conducted into algorithms that can simultaneously optimize both the sampling rate, the required resources (e.g., CPU, RAM), and the energy consumed. Multi-objective optimization methods can simultaneously consider energy use, detection accuracy, response latency, and available resources. Additionally, experimental studies can assess if these types of algorithms perform better than constant-rate monitoring in various types of cloud workloads.
- **Integration with Cloud-Native Security Platforms:** Adaptive sampling may be used in conjunction with container orchestrators, cloud log managers, intrusion detection systems (IDS), and central security analytics tools. Integration of adaptive sampling with these types of platforms will allow for the coordinated monitoring of virtual machines, containers, applications and distributed services. Future adaptive sampling implementations should focus on interoperability, scalable policy management and communication reliability among monitoring components.
- **Enhancing the Ability to Detect Attacks of Short Duration:** The reduction of sampling rates can provide an opportunity for an attacker to conduct short-lived attacks in the time frame of two subsequent measurements. Therefore, future studies should focus on creating a combination of different types of monitoring methods, such as combining adaptive periodic sampling with event based alerts, targeted high-frequency collection, and independent emergency detection mechanisms. Additionally, these methods need to be evaluated against various lengths of attacks and different kinds of network traffic.
- **Standardize the process of measuring energy consumption for each type of telemetry collection, security analysis, data transmission and adaptive control:** Researchers need to develop standardized processes to measure the amount of energy used to collect telemetry, conduct security analyses, transmit data, and use adaptive controls. Where applicable, researchers should make direct measurements of the amount of energy being consumed during those activities using identical hardware and workload conditions. This will allow researchers to determine whether observed energy savings are a result of decreased events or reduced processing capacity.
- **Improving security management and reliability:** Future frameworks will include minimal monitoring requirements, secure audit logs, automated recovery systems, and transparent sampling choices. Research will examine how these protective measures function during failure of controllers, abnormal workload changes, and attempts to manipulate risk scores. Compliance rules will continue to be separate from optional reductions in sampling.
- **Testing in More Cloud Environments:** To test the scalability and generalizability of adaptive sampling, additional studies need to evaluate public, private, hybrid and multi-cloud environments. The researchers should test their techniques using workloads that vary in size, types, and other characteristics. Researchers also need to assess how much energy is consumed by each type of cloud environment, what percentage of attacks are detected, what the false-positive rate is for each attack type, how long it takes to detect an attack, and what the controller overhead is for each trial. By testing in more cloud environments and evaluating these metrics over multiple trials, researchers can demonstrate whether their findings can be generalized to a broader range of cloud environments.

XI. CONCLUSION

Cloud computing is becoming more prevalent; therefore, there is a greater requirement to continuously monitor for threats to the security of the cloud environment. Cloud environments produce large quantities of security telemetry via authentication processes, network communications, application processes, system metrics, and access control events. Continuous monitoring helps identify potential threats to the security of the cloud environment as well as unauthorized access to resources within the cloud environment. However, the production of large volumes of telemetry results in increased computational overhead, storage

requirements, network usage, and energy consumption. This paper, titled Reducing Energy Consumption in Continuous Cloud Threat Monitoring Through Adaptive Sampling, proposes an adaptive sampling framework that addresses these issues by altering the level of monitoring based upon workload activity, perceived threat level, and resource availability. The primary goal of this work is to lower the cost associated with monitoring for threats to the security of the cloud environment while providing adequate security visibility so that threats to the security of the cloud environment may be detected effectively.

The suggested model combines telemetry collection, risk evaluation, adaptive sampling control, threat identification and optimization based on feedback. The model distinguishes between those times when the system is operating normally (low risk) and those times when the system needs to be observed more closely due to potential risks. During the time of low risk, the model can lower the frequency at which it routinely samples for certain workloads; therefore, reducing unnecessary processing of telemetry data and the use of the associated resources. As soon as a possible security breach occurs, such as an unusual authentication attempt, unusual network activity, excessive resource usage, etc., the controller can raise the frequency of sampling and allocate additional analysis tools. An alternative to traditional monitoring with a set rate is provided by this risk-aware approach. Traditional monitoring applies the same collection policy regardless of what is happening with the security environment. To ensure that energy optimization does not create unacceptable gaps in security, there will still be minimum monitoring requirements, event-driven alerts, independent trigger mechanisms for escalations, and secure audit records. Therefore, adaptive sampling should be used in conjunction with current security controls rather than replacing mandatory logging or existing incident detection methods.

Another key benefit of our methodology is how we evaluate both the energy efficiency of the proposed solution and its security performance. While reducing telemetry data alone does not prove that a monitoring system has realized significant energy savings, we must assess the overall monitoring energy consumption, threat detection rates, false positive rates, detection latencies, CPU usage, memory requirements and the added cost of the adaptive controller. When comparing fixed rate and adaptive monitoring solutions, researchers must employ similar hardware configurations, representative workloads, identical attack scenarios and multiple experimental trials. This type of assessment will demonstrate whether any energy reductions are accompanied by unacceptable increases in missed detections or response times. Our proposed framework has not been evaluated in practice; therefore, we must evaluate the success of the framework based on controlled experiments instead of assuming improvements in performance. We will need to report transparently on our experimental design and experimental results to establish whether the framework is applicable to various types of cloud infrastructure.

Adaptive sampling may have some advantages, but it poses several problems; these include: risk estimation (of short duration) attacks (adversarial) manipulation (of workloads) and compliance with regulations. In addition to investigating machine learning-based risk predictions, energy-aware resource allocation, feedback-driven controls, and integration with cloud native security platforms, future implementations of the adaptive sampling process should also evaluate the performance of adaptive policies in public, private, hybrid, and multicloud environments. Additionally, robust fallback mechanisms are required to return to the pre-defined level of monitoring when risk assessment is no longer reliable or the adaptive controller fails. Finally, organizations should implement human oversight, auditability, and policy enforcement to verify that sampling decisions are in alignment with organizational security obligations. By addressing each of the aforementioned issues, researchers can improve the dependability and practicality of the adaptive sampling framework.

Adaptive sampling is an area of study that could help reduce the unnecessary amount of power consumed during continuous cloud threat monitoring. The adaptive sampling framework was designed to find a compromise between the cost of operations (i.e., operating expenses), the use of computational resources (i.e., hardware) and protecting against cyber attacks. To succeed, adaptive sampling must be able to show documented evidence of energy savings, while still providing the level of threat detection needed to protect the organization's assets. Adaptive sampling is expected to provide organizations with a way to create a more sustainable, scalable and responsive cloud security infrastructure. Additionally, adaptive sampling supports the overall goal of creating cloud monitoring systems that protect organizations' assets, but do so responsibly with respect to the use of computational resources.

XII. REFERENCES

[1] Ahmad, I., Namal, S., Ylianttila, M. and Gurtov, A. (2016) 'Security in software defined networks: A survey', IEEE Communications Surveys & Tutorials, 18(1), pp. 231–250.

[2] Alpcan, T. and Başar, T. (2010) Network Security: A Decision and Game-Theoretic Approach. Cambridge: Cambridge University Press.

[3] Armbrust, M. et al. (2010) 'A view of cloud computing', Communications of the ACM, 53(4), pp. 50–58.

[4] Buczak, A.L. and Guven, E. (2016) 'A survey of data mining and machine learning methods for cyber security intrusion detection', IEEE Communications Surveys & Tutorials, 18(2), pp. 1153–1176.

[5] Buyya, R., Yeo, C.S., Venugopal, S., Broberg, J. and Brandic, I. (2009) 'Cloud computing and emerging IT platforms: Vision, hype, and reality for delivering computing as the 5th utility', Future Generation Computer Systems, 25(6), pp. 599–616.

[6] Chandola, V., Banerjee, A. and Kumar, V. (2009) 'Anomaly detection: A survey', ACM Computing Surveys, 41(3), Article 15.

[7] Cisco Systems (2020) Cisco Annual Internet Report (2018–2023) White Paper. San Jose, CA: Cisco Systems.

[8] Dastjerdi, A.V. and Buyya, R. (2016) 'Fog computing: Helping the Internet of Things realize its potential', Computer, 49(8), pp. 112–116.

[9] Demertzis, K., Iliadis, L. and Spartalis, S. (2017) 'A spiking one-class anomaly detection framework for cyber-security on industrial control systems', Computers & Security, 74, pp. 269–284.

[10] Dinh, T.N., Li, X. and Thai, M.T. (2014) 'Network clustering and community detection', ACM Computing Surveys, 45(4), Article 54.

[11] Elahi, M., Khorasani, M. and Bakhshi, A. (2013) 'Energy-efficient cloud computing: A review of techniques and approaches', International Journal of Computer Applications, 82(11), pp. 1–7.

[12] Fadlullah, Z.M., Tang, F., Mao, B., Kato, N., Akashi, O., Inoue, T. and Mizutani, K. (2017) 'State-of-the-art deep learning: Evolving machine intelligence toward tomorrow's intelligent network traffic control systems', IEEE Communications Surveys & Tutorials, 19(4), pp. 2432–2455.

[13] García-Teodoro, P., Díaz-Verdejo, J., Maciá-Fernández, G. and Vázquez, E. (2009) 'Anomaly-based network intrusion detection: Techniques, systems and challenges', Computers & Security, 28(1–2), pp. 18–28.

[14] Gonzalez, N., Miers, C., Redígolo, F., Simplicio, M., Carvalho, T., Näslund, M. and Pourzandi, M. (2012) 'A quantitative analysis of current security concerns and solutions for cloud computing', Journal of Cloud Computing: Advances, Systems and Applications, 1, Article 11.

[15] Han, J., Kamber, M. and Pei, J. (2012) Data Mining: Concepts and Techniques. 3rd edn. Waltham, MA: Morgan Kaufmann.

[16] Harchol-Balter, M. (2013) Performance Modeling and Design of Computer Systems: Queueing Theory in Action. Cambridge: Cambridge University Press.

[17] Iqbal, S., Gupta, M. and Kumar, P. (2014) 'Energy efficiency in cloud computing: A review', International Journal of Computer Applications, 94(2), pp. 1–6.

[18] Jain, R. (1991) The Art of Computer Systems Performance Analysis: Techniques for Experimental Design, Measurement, Simulation, and Modeling. New York: Wiley.

[19] Kent, K. and Souppaya, M. (2006) Guide to Computer Security Log Management. NIST Special Publication 800-92. Gaithersburg, MD: National Institute of Standards and Technology.

[20] Krügel, C., Toth, T. and Kirda, E. (2002) 'Service specific anomaly detection for network intrusion detection', in Proceedings of the ACM Symposium on Applied Computing. New York: ACM, pp. 201–208.

[21] Mell, P. and Grance, T. (2011) The NIST Definition of Cloud Computing. NIST Special Publication 800-145. Gaithersburg, MD: National Institute of Standards and Technology.

[22] Mell, P. and Grance, T. (2011) Guidelines on Security and Privacy in Public Cloud Computing. NIST Special Publication 800-144. Gaithersburg, MD: National Institute of Standards and Technology.

[23] Mishra, P., Pilli, E.S., Varadharajan, V. and Tupakula, U. (2017) 'Intrusion detection techniques in cloud environment: A survey', Journal of Network and Computer Applications, 77, pp. 18–47.

[24] NIST (2012) Computer Security Incident Handling Guide. NIST Special Publication 800-61 Revision 2. Gaithersburg, MD: National Institute of Standards and Technology.

[25] Patcha, A. and Park, J.-M. (2007) 'An overview of anomaly detection techniques: Existing solutions and latest technological trends', Computer Networks, 51(12), pp. 3448–3470.

[26] Ristenpart, T., Tromer, E., Shacham, H. and Savage, S. (2009) 'Hey, you, get off of my cloud: Exploring information leakage in third-party compute clouds', in Proceedings of the 16th ACM Conference on Computer and Communications Security. New York: ACM, pp. 199–212.

[27] Sommer, R. and Paxson, V. (2010) 'Outside the closed world: On using machine learning for network intrusion detection', in Proceedings of the IEEE Symposium on Security and Privacy. Los Alamitos, CA: IEEE, pp. 305–316.

[28] Subashini, S. and Kavitha, V. (2011) 'A survey on security issues in service delivery models of cloud computing', Journal of Network and Computer Applications, 34(1), pp. 1–11.

[29] Varghese, B. and Buyya, R. (2018) 'Next generation cloud computing: New trends and research directions', Future Generation Computer Systems, 79, pp. 849–861.

[30] Zissis, D. and Lekkas, D. (2012) 'Addressing cloud computing security issues', Future Generation Computer Systems, 28(3), pp. 583–592.