

An elastic scaling method for cloud security

Áine MacDermott, Qi Shi, Madjid Merabti & Kashif Kifayat

PROTECT: Research Centre for Critical Infrastructure Computer Technology and Protection

Liverpool John Moores University, Liverpool

a.mac-dermott@2008.ljmu.ac.uk and {q.shi,m.merabti,k.kifayat}@ljmu.ac.uk

Abstract

Cloud computing is being adopted in critical sectors such as transport, energy and finance. This makes cloud computing services critical in themselves. When cyber attacks and cyber disruptions happen, millions of users are affected. A cyber disruption in this context means a temporary or permanent loss of service, with impact on users of the cloud service who rely on its continuity. Intrusion detection and prevention methods are being developed to protect this sensitive information being stored, and the services being deployed. There needs to be an assurance that the confidentiality, integrity and availability of the data and resources are maintained. This paper presents a background to the critical infrastructure and cloud computing progression, and an overview to the cloud security conundrum. Analysis of existing intrusion detection methods is provided, in addition to our observation and proposed elastic scaling method for cloud security.

1. Introduction

As more sectors adopt cloud services in their computing environment, the trend will also reach IT services operating critical infrastructure. Critical infrastructures are physical or mechanical processes mostly controlled electronically by systems, usually called supervisory control and data acquisition (SCADA) or process control systems (PCSs), composed of computers interconnected by networks. Examples include transportation, energy, telecommunications, and water. When critical infrastructures were first implemented, the security and protection of their management system was not a primary concern [1]. In recent years, critical infrastructure have become increasingly dependent on ICT; more interconnected and linked to the Internet. Consequently, this makes these systems more vulnerable and increases the risk of cyber-attack [2].

SCADA systems have evolved over the years from being monolithic, to distributed, to networked. Research has shown that cloud computing will eventually reach the IT services that are operating critical infrastructures [3–6]. There is a similarity between critical infrastructure and cloud computing, as they are primarily large distributed data sets and may possess the same underlying issues. The emergence of the cloud computing paradigm could be beneficial for the operation and performance of these complex infrastructures.

With the technical development and market growth in cloud computing, organisations that provide, operate and maintain IT systems for critical infrastructure are making the decision as to when they should make the computing paradigm shift. Cloud services can offer efficient access to large IT infrastructures that benefit from the economy of scale. It would be highly desirable to maintain irrecoverable and valuable data obtained from critical infrastructure within secure cloud infrastructures [5]. However, the reality of today's advanced malware and targeted attacks is that 100% protection is not realistic. Reducing attack vectors and marginalising the impact of an attack is the practical approach.

The layout of this paper is as follows: Section 2 provides background on critical infrastructure and their cloud computing progression, as well as cloud attributes and security concerns. Section 3 details existing approaches to detecting intrusions in the cloud environment. In Section 4 we present our observation and elastic scaling method for this protection problem. In Section 5 we evaluate our approach, and in Section 6 discuss some thoughts on this area. Our conclusions and future work are highlighted in Section 7.

2. Background

2.1. Critical infrastructure overview

Critical infrastructures, such as the power grid and water distribution facilities, include a high

number of devices over a large geographical area. These infrastructures face significant threats as the growth in the use of industrial control systems such as SCADA systems, and their integration to networks in order to coordinate and manage the actions of these devices. The importance of protecting these infrastructures has been particularly highlighted by the increase in advanced persistent threats (APTs), such as 'Stuxnet' and 'Duqu', which were designed to target these control systems and disrupt their functionality [7]. Effective protection of SCADA systems is therefore crucial, as these are important components of critical infrastructures, and it is apparent that existing methods do not meet the security requirements of such interconnected infrastructures [4].

The evolution of SCADA systems has also raised concerns about cyber-related vulnerabilities. The SCADA industry is transitioning from a legacy environment, in which systems were isolated from the Internet and focused on reliability instead of security, to a modern environment where networks are being leveraged to help improve efficiency.

Traditionally, these infrastructures were inherently secure systems, as they were largely based on dedicated communication links. Nowadays, modern infrastructures make use of IT technologies, where wireless sensor networks (WSNs) with open access have become an integral part of virtually any critical infrastructure. Since IT infrastructures have become an integral part of almost all organisations, cloud computing will have a significant impact on them.

Critical infrastructure currently makes use of the benefits offered by general IT services, so benefiting from the intricate cloud computing paradigm is expected. Embracing the cloud environment is a natural extension of remote access as it removes the requirement for the user to be in the same location as the infrastructure. Remote access to critical infrastructure is already common practice, i.e. remote access to intelligent electronic devices (IEDs) or user interfaces in a substation for maintenance. While this could provide improved performance, concerns over protecting sensitive data and services in this environment remain [8].

Previous work of ours MacDermott et al. [4] has detailed a way in which critical infrastructure could utilise the cloud environment for improved performance and analysis of the automation processes.

2.2. Cloud computing overview

Cloud computing is a style of computing where elastic IT related capabilities are provided as an optimised, cost-effective and on-demand utility [9]. This can be considered as an evolution of e-business as cloud computing helps enterprises create and

deliver IT solutions in a more flexible and cost-effective way. Cloud providers usually build up large scale data centres and provide cloud users with computational resources in three service models:

- Software as a Service (SaaS)
- Platform as a Service (PaaS)
- Infrastructure as a Service (IaaS)

Figure 1 details these service models and gives examples of each [10]:

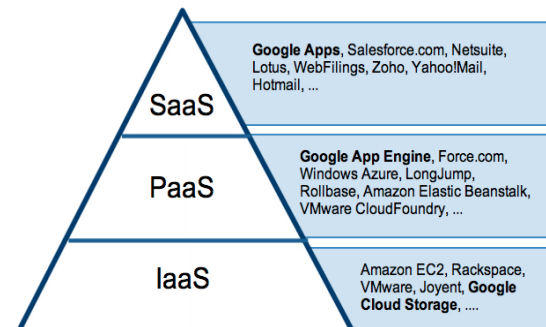


FIGURE 1: SERVICE MODELS AND EXAMPLES

There are security issues at each level of the cloud computing paradigm. These levels are application level, virtual level and physical level. The application level comprises of Software as a Service (SaaS), in which enterprises host and operate their applications over the Internet so that the customers can access it [11]. One benefit of this model is that customers do not need to buy software licences or any additional equipment for hosting the application.

The virtual level includes Platform as a Service (PaaS) and Infrastructure as a Service (IaaS). PaaS provides a platform for building and running customer applications. Enterprises can build applications without installing any tools on their local systems and can deploy them without many difficulties. IaaS provides a convenient option for organisations by transferring the IT infrastructure to the cloud provider. It is the responsibility of the cloud provider to tackle issues of management, such as configuring servers, routers, firewalls to name a few [11].

The physical level refers to the infrastructure upon which clouds are deployed. Security requirements and threats associated with each of these services are depicted in Table 1.

Cloud deployment models include public, private, community and hybrid:

- A public cloud is available to the general public or large industry group, owned by an organisation selling cloud services. A third party provides infrastructure, platform and software. The management, operational and

TABLE 1: SECURITY REQUIREMENTS AND THREATS ASSOCIATED WITH EACH SERVICE LEVEL.

Level	Service level	Security requirements	Threats
Application level	Software as a Service (SaaS)	<ul style="list-style-type: none"> ▪ Access control ▪ Communication protection ▪ Data protection from exposure ▪ Privacy in multitenant environment ▪ Service availability ▪ Software security 	<ul style="list-style-type: none"> ▪ Data interruption ▪ Exposure in network ▪ Interception ▪ Modification of data at rest and in transit ▪ Privacy breach ▪ Session hijacking ▪ Traffic flow analysis
Virtual level	Platform as a Service (PaaS), Infrastructure as a Service (IaaS)	<ul style="list-style-type: none"> • Access control • Application security • Cloud management control security • Communication security • Data security • Secure images • Virtual cloud protection 	<ul style="list-style-type: none"> • Connection flooding • Distributed denial of service (DDoS) attack • Defacement • Disrupting communications • Exposure in network • Impersonation • Programming flaws • Software modification • Software interruption • Session hijacking • Traffic flow analysis
Physical level	Physical data centre	<ul style="list-style-type: none"> • Hardware security • Hardware reliability • Legal not abusive use of cloud computing • Network protection • Network resources protection 	<ul style="list-style-type: none"> • Network attacks • Connection flooding • Distributed denial of service (DDoS) attack • Hardware interruption • Hardware theft • Hardware modification • Misuse of infrastructure • Natural disasters

security requirements are provisioned and shared between users and providers with a service level agreement (SLA).

- A private cloud operates for a single organisation. The infrastructure can be located in the organisational unit or in a third party unit's data centre. Private clouds grant complete control over how data is managed and what security measures are in place.
- A community cloud is shared by several organisations, supporting a specific community. The infrastructure is placed in more than one organisation in the community or third party's data centre. Management and operational tasks are split between data centre owner, organisations and third party.
- Hybrid clouds are the combination of more than one cloud deployment model, as previously described. All the infrastructure, platform and software are portable and can switch between the deployment models in the hybrid architecture [13].

In cloud environments, network perimeters will no longer exist from the cloud user's perspective, which renders traditional security protection methods, such as firewalls, inapplicable to cloud applications. The ability to clearly identify, authenticate, authorise and monitor who or what is accessing the assets of an organisation is essential.

3. Detecting intrusions in the cloud environment

Despite security issues delaying its adoption, cloud computing has already become an unstoppable force; thus, security mechanisms to ensure its secure adoption are an immediate need [14]. The distributed and open structure of cloud computing and services becomes an attractive target for potential cyber-attacks by intruders. The traditional intrusion detection and prevention systems (IDS/IPS) are largely inefficient to be deployed in cloud computing environments due to their openness and specific essence [15]. In addition, the deployment of intrusion detection and prevention systems varies per solution and is something that is not cohesive in its approach.

In the cloud environment, where massive amounts of data are generated due to high network access rates, an IDS must be robust against noise data and false positives. Since cloud infrastructure have

enormous network traffic, traditional IDSs are not efficient to handle such a large data flow. Due to the large data sets, classification techniques require a huge amount of memory and CPU usage.

Hamad and Al-Hoby [16] implemented the Cloud Intrusion Detection Service (CIDS), which can be deployed by cloud providers to enable clients to subscribe with the IDS in a service-based manner, i.e. "Security-as-a-Service". It is a re-engineered version of Snort, which is an open-source network IDS/IPS. The model outperforms currently used solutions for service-based IDS but at the same time provides minimal overhead to the case of traditional IDS deployment for single network protection.

Dhage et al. [17], convey that when there is only one IDS in the entire network, the load on it increases as the number of host's increases. It is difficult to keep track of different kinds of attacks or intrusions, which are acting on each of the host present in the network. An architecture in which mini IDS instances are deployed between each user of the cloud and the cloud service provider is proposed. As a result, the load on each IDS instance will be less than that on a single IDS and for this reason, the small IDS instance will be able to do its work in a more efficient way.

The work of Lee [18], proposes a multi-level IDS and log management method based on consumer behaviour for applying IDS effectively to the cloud system. They assign a risk level to users' behaviour based on analysis of their behaviour over time. By applying differentiated levels of security strength to users based on the degree of anomaly increases the effective usage of resources. Their method proposes the classification of generated logs by anomaly level. This is so that the system administrator analyses logs of the most suspected users first.

Lo et al. [19] present a cooperative intrusion detection system framework for cloud computing networks. They deploy an IDS in each cloud region, and each entity cooperates with each other through the exchange of alerts to reduce the impact of DoS attacks. A Snort based IDS is implemented and the three main modules are plugged into the system: block, communicate, defense. A co-operate agent is used to receive alerts from other IDSs, and they are analysed using a majority vote in order to determine the accuracy of results. If deemed a legitimate alert, the blocking rule is implemented. By co-operative operation among these agents, early detection and prevention technique is implemented. Therefore, IDSs deployed in cloud computing regions except the victim one could prevent this kind of attack.

Randles and Lamb [20], focus on tackling distributed load balancing for cloud computing and present a comparative study of algorithms considered. It is expressed that as the system increases in size and complexity, the rule sets become unwieldy. This means that it may not be

possible to maintain a viable monitoring and response cycle to manage the computational workload. For example, the execution of one rule may cause an event, triggering another rule set or set of rules, dependent on the current state. Methods are sought that promote load balancing on the global cloud scale via actions and interactions at the component level; however, a combination of algorithms seems clear.

In Mahmood and Agrawal [21], the focus is on 'Principal Component Analysis Neural network Algorithm' (PCANNA) which is used to reduce the number of computing resources, both memory and CPU time required to detect an attack. Feature reduction is used to remove useless information from the original high dimensional database of cloud traffic data. A back propagation algorithm is applied on reduced cloud traffic data for classification. Their contribution shows that dimensional reduction techniques help compact similar alerts and correlate alerts coming from heterogeneous platforms on several sites to detect intrusions that are more complex.

Alsafi et al. [22] propose an integrated intrusion handling model for cloud computing, which combines anomaly and signature detection. Their focus is on stopping an attack, rather than detecting it. Actions their proposed method should take include terminating the user session that is being used during the attack, block access to the target from the offending user account, IP address, or other attacker attribute. The integrated model uses signature matching with normal traffic profiling to enhance attack detection. They propose to deploy their IDS in the virtual machine (VM) itself as well as the virtual network in order to monitor the activities within the system.

Cloud defence strategy needs to be distributed so that it can detect and prevent the attacks that originate within the cloud itself and from the users using the cloud technology from different geographic locations through the Internet. As the popularity of the services provided in the cloud environment grows rapidly, the exploitation of possible vulnerabilities grows at the same pace.

4. Our observation

In the service-oriented architecture of the cloud, collaboration means data are coming from many different sources so existing IDS techniques will not be able to process data of this scale. Our survey of related work identified current weaknesses with existing approaches:

- Overload with a high volume of traffic
- Fail to scale to satisfy high speed networks
- Loss of accuracy

- Inaccurate profile of usage
- Require human intervention, which can slow down response time.
- Simply flags suspect behaviour
- High false alarm rate
- Ineffective log management.
- Cannot detect novel attacks.

Distributed systems need to maintain a balance between communication overheads and the addition of process power, as resources can become constrained. Distributed IDS detect attacks by analysing large sets of traffic. This traffic is often analysed by taking a sample, and a large percentage of attacks can be detected quite quickly, whereas novel attacks are often missed.

Since cloud computing supports a distributed service oriented paradigm, multi-domain and multi-users administrative infrastructure, it is more prone to security threats and vulnerabilities, such as data breaches, data loss, service hijacking, denial of services (DoS) attacks, malicious insiders to name a few [9]. The Cloud Security Alliance report “Top Threats to Cloud Computing,” published in March 2010 identified the following threats in their initial document:

- Abuse and misuse of cloud computing
- Insecure application programming interfaces
- Malicious insiders
- Shared technology vulnerabilities
- Data loss
- Data leakage
- Account, service, and session hijacking

The updated report published in 2013, entitled “The Notorious Nine – Cloud Computing Top Threats in 2013” [23] includes distributed denial of service (DDoS) as a key threat, not originally considered in 2010. By forcing the victim cloud service to consume inordinate amounts of finite system resources such as processor power, memory, disk space or network bandwidth, the attacker (or attackers, as is the case in DDoS attacks) causes an intolerable system slowdown and leaves all the legitimate service users confused.

It can be inferred that as the use of cloud in organisations grows, so will the rate of denial of service attacks. These attacks against the cloud are launched to deny service availability to end users. While DDoS attacks tend to generate a lot of fear and media attention, they are by no means the only form of DoS attack. Asymmetric application-level

DoS attacks take advantage of vulnerabilities in web servers, databases, or other cloud resources, allowing a malicious individual to take out an application using a small attack payload – in some cases less than 100 bytes long [23].

Unavailability of services due to cloud outages can cause monetary loss to cloud providers and operational loss to cloud users. Hosting infrastructure services, and storing sensitive data in the cloud environment brings with it security and resilience requirements that existing cloud services are not well placed to address. IDS mechanisms require an extensive use of hardware, especially CPU and memory, and may cause unintentional resource exhaustion or a bottleneck. We undertook a comparison of current protection approaches for the cloud environment, to ensure that our approach would have a minimal impact on the clouds infrastructure and the services operating within.

4.1. Analysis

Traditional network monitoring schemes are not scalable to high-speed networks. Where many solutions sample traffic, we believe sampling should fall within a prescribed error tolerance level. For detecting intrusions in the cloud environment, sampling costs are of paramount importance.

A solution to the problems we have identified would have to encompass the following attributes:

- Automate detection
- Scalable
- Elastic
- Traffic filter/gateway entity
- Improved efficiency over current methods
- Nominal profile that updates as parameters adjust

Our proposed approach, incorporating these attributes will use the resources of the cloud environment to sample at a higher resolution so it has improved effectiveness. Our proposed traffic filter entity is effectively a gateway. It channels information from different cloud services and analyses them to determine an attack, e.g. DDoS. The measurements required to obtain a comprehensive view on the status of the cloud lead to the generation of a very large volume of data coming from multiple distributed locations. Hence, a scalable monitoring system should be able to efficiently collect, transfer and analyse such volume of data without impairing the operations of the cloud.

The traffic filter entity needs to determine, firstly, if an attack is occurring. Many methods to detect intrusions have predefined thresholds or behaviours of their traits. In some cases, these could simply be

high surges of traffic and are legitimate. Our approach allows the traffic filter/gateway entity to scale out, using more resources of the cloud to balance the load of traffic and analyse in a higher resolution. When we start to see an attack, we scale out, i.e. clone/spawn the traffic filter service. It would be highly desirable to use the resources of the cloud to protect the cloud.

There would be an elastic traffic filter in each cloud provider domain. The removal of the human element to analyse 'flagged' actions would improve the efficiency of the proposed approach, as this action would be automated. We need to design an algorithm to deal with distributed cloud attacks. Static sampling algorithms, such as Simple Random Sampling were originally considered as an efficient approach, however, they tend to oversample at peak periods when efficiency and timeliness are crucial, thus not ensuring the accuracy of estimation. The biggest challenge in employing a sampling algorithm on a given network is scalability.

Combining stateful and stateless signatures for attack detection, differs from related work that mainly employs one or two signature approaches. The main benefit from the combination is that not all malicious packets have to be inspected in order to ascertain the presence of an attack. This improves detection efficiency and makes attack detection feasible within the routing infrastructure.

4.2. Case Study

A, B, C, and D represent services in a cloud environment. The traffic filter (TF) is a scalable gateway, through which communication between services in the cloud environment passes through as a normal occurrence. An exemplary scenario illustrating our proposed solution is illustrated in Figure 2:




FIGURE 2: SERVICES IN A CLOUD ENVIRONMENT COMMUNICATING THROUGH THE TRAFFIC FILTER

A-C are colluding to attack the cloud service. D is the target, and TF is the scalable traffic filter gateway.

Service D (Victim) sends a request to Service A, B or C. The Services attempt to send malicious traffic

to the user/collusion. This could even be a flash crowd type attack. In this scenario, all the communication passes through the traffic filter, so there is a nominal profile of actions and behaviour stored.

Analysis takes place in the following steps:

Analysis Stage 1: Step 1: Anomaly detected.

Analysis Stage 2: Step 2: Scale out/Spawn TF entity.

Step 3: Mitigate/Drop packets.

There is no single point of failure as the traffic filter is elastic so it scales the analysis of traffic and balances this with the new entity it has spawned, as illustrated in Figure 3:




FIGURE 3: TRAFFIC FILTER SCALING ITS RESOURCES TO BALANCE THE LOAD FOR ANALYSIS

A flow diagram depicting the actions the traffic filter would take is presented in Figure 4:

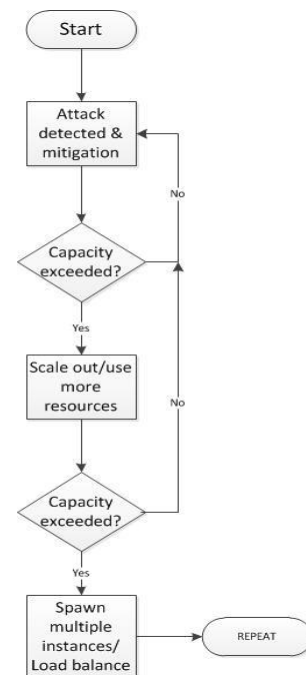


FIGURE 4: FLOW DIAGRAM FOR TRAFFIC FILTER ACTIONS

Giving the traffic filter the ability to scale out/spawn its resources in order to balance the load of analysing the large volume of data means that it can use the resources of the cloud to analyse at a higher resolution. Figure 4 is a high level view of the actions that would be taken in our proposed solution. Ensuring that the analysis of traffic does not cause the entity to exceed its capacity is important, as this often happens with existing approaches that overload with a high volume of traffic.

Much of the research in this area drops packets when they believe the traffic may be malicious, but this is often inaccurate. Additionally, some work sends a report or alert to the administrator or third party that an attack has just occurred. The ability to effectively analyse the large volume of data, through balancing the load means that the results will be more efficient, and appropriate actions can be taken. New algorithms optimised for detecting cloud attacks in an effective manner are needed. Additionally, having a solution with the ability to adapt to varying computational and network loads in order to not be invasive is needed also.

5. Evaluation of approach

It was originally considered a positive solution to detecting intrusions in the cloud environment. However, further research into implementing the traffic filter entity proved worrisome. When we were finalising designs, we concluded that if the communication between different services needs to go through the traffic filter, in practice it is hard to implement.

It would be desirable for communication to go through the traffic filter/gateway in order to gain a consistent understanding of the occurring communication. In contrast, design requirements for intrusion detection systems should not affect current operations in any way. Otherwise, if the traffic filter entity needed to be modified and upgraded, the services would too.

We concluded that the traffic filter could play some role for data collection, but not for a main player between different services. The use of agents to collect data could achieve the same operational objectives we desire, as they would not use operations to interfere with the current operation of services in the cloud. The use of independent agent entities that travel to and from services in the cloud environment, gather traffic logs and perform analysis could be applicable. We are currently analysing the feasibility of enhancing elements of the existing cloud infrastructure, to provide the functionality we proposed for the traffic filter attributes.

6. Discussion

One problem for this research environment is that there is no unified detection and prevention approach, or a globally accepted metric or standard to evaluate against. It is clear that an IDS alone cannot protect the cloud environment from attack. As such, there has been an increase in IDPS for this environment.

Desired characteristics for IDPS include optimised performance, minimum error and maximum protection [15]. The ability to adapt to changes in user behavior and system behaviour over time is also anticipated. An IDPS should be part of normal services and not affect the operation of the cloud environment in any way.

The structure of an IDPS is based upon two types: individual and collaborative. An individual arrangement of IDPS is achieved by physically integrating it within a firewall. A collaborative IDPS consists of multiple IDPSs over a large network where each one communicates with each other. Each IDPS has two main functional components: detection element and correlation handler. Detection elements consist of several detection components, which monitor their own sub-network or host individually and generate low-level alerts. The correlation handler transforms the low level alerts into a high level report of an attack [15].

The current lack of collaboration among different components within a cloud provider or among different providers for detection or prevention of attacks is an area we aim to focus on. Cloud service providers have the scale and resources to address and prevent cyber-attacks in a more professional way than most other organisations [24]. We believe it is feasible to allow cloud providers to collaborate, as it would be beneficial for them if they could.

Many solutions can only detect specific attacks, not unknown ones, and this is deterring the utilisation of the environment. A hybrid IDPS is needed for protecting the cloud environment from attack with optimised performance and protection with minimum error [15]. A hybrid approach combines two or more network intrusion detection techniques; signature based detection, anomaly based detection, and soft computing techniques. Using a hybrid approach can improve the accuracy of the IDS when compared to individual approaches.

Attacks and failures are inevitable; therefore, it is important to develop approaches to understanding the cloud environment under attack. Investigation into the appropriate 'points' in the cloud to deploy monitoring and attack detection functionality is imperative.

The four areas considered for deployment are in the VM, in the hypervisor or host system, in the virtual network, or in the traditional network.

- In the VM: Deploying a solution in the VM allows you to monitor the activity within the system, and detect and alert on issues that may rise.
- In the hypervisor or the host system: Deploying a solution in the hypervisor allows you not only to monitor the hypervisor, but anything travelling between VMs on that hypervisor. It is a more central location for intrusion detection, but there may be performance issues or dropping of some information if the amount is too large.
- In the virtual network/VLAN: Deploying a solution to monitor the virtual network allows you to monitor the network traffic between VMs on the host, as well as traffic between the VMs and host. This 'network' traffic never hits the traditional network.
- In the traditional network: Deploying a solution here allows you to monitor, detect, and alert on traffic that passes over the traditional network infrastructure. However, this is quite problematic as we may miss virtual traffic as it is encrypted.

We believe the optimal deployment location is on the virtual network/VLAN. To communicate between VMs, they talk over a virtual network. This would be a suitable place for an IDPS as communicating occurs through this point, and it would be easier to build a nominal profile of activities and behaviours. The use of a module that uses signature analysis of captured attack statistics, but also utilises a behaviour module to determine if the detected occurrence is actually an attack. This could in turn improve efficiency over current methods that only utilise one method.

7. Conclusions and future work

Critical infrastructure vendors will inevitably take advantage of the benefits offered by the cloud computing paradigm, but while this may offer improved performance and scalability, the associated security threats deter this progression. This paper has shown our plans to build upon our initial research into protecting services in the cloud environment, through our proposed elastic scaling method for cloud security. Our future work includes enhancing the attributes of our method and measuring how an attack against the cloud's infrastructure would affect performance. New algorithms optimised for detecting cloud attacks in an efficient manner are needed, and this is something we will explore further.

8. References

- [1] Á. MacDermott, Q. Shi, M. Merabti, and K. Kifayat, "Intrusion Detection for Critical Infrastructure Protection," in *Proceedings of 13th Annual Postgraduate Symposium on Convergence of Telecommunications, Networking and Broadcasting (PGNet 2012)*, 2012.
- [2] W. Hurst, M. Merabti, and P. Fergus, "Behavioural Observation for Critical Infrastructure Security Support," in *The Seventh IEEE European Modelling Symposium (EMS2013)*, 2013.
- [3] D. Wallom, M. Turilli, A. Martin, A. Raun, G. Taylor, N. Hargreaves, and A. McMoran, "myTrustedCloud: Trusted Cloud Infrastructure for Security-critical Computation and Data Managment," *2011 IEEE Third International Conference on Cloud Computing Technology and Science (CloudCom)*, pp. 247–254, Nov. 2011.
- [4] Á. MacDermott, Q. Shi, M. Merabti, and K. Kifayat, "Protecting Critical Infrastructure Services in the Cloud Environment," in *Proceedings of the 12th European Conference on Information Warfare and Security (ECIW)*, 2013, pp. 336–343.
- [5] OTE, "Discussion on the Challenges for the Development of a Context for : Secure Cloud computing for Critical infrastructure IT," Greece, 2012.
- [6] M. Sch, R. Bless, F. Pallas, J. Horneber, and P. Smith, "An Architectural Model for Deploying Critical Infrastructure Services in the Cloud," in *IEEE Cloud Com 2013*, 2013.
- [7] T. Miyachi, H. Narita, H. Yamada, and H. Furuta, "Myth and reality on control system security revealed by Stuxnet," in *2011 Proceedings of SICE Annual Conference (SICE)*, 2011, pp. 1537–1540.
- [8] M. T. Khorshed, A. B. M. S. Ali, and S. A. Wasimi, "A survey on gaps, threat remediation challenges and some thoughts for proactive attack detection in cloud computing," *Future Generation Computer Systems*, vol. 28, no. 6, pp. 833–851, Jun. 2012.
- [9] Á. MacDermott, Q. Shi, M. Merabti, and K. Kifayat, "Considering an elastic scaling model for cloud security," in *The 8th International Conference for Internet Technology and Secured Transactions (ICITST-2013)*, 2013.
- [10] W. Chun, "What is cloud computing?," Google Developers Academy, 2012. [Online]. Available: <https://developers.google.com/appengine/training/intro/whatiscc>.
- [11] K. Annapureddy, "Security Challenges in Hybrid Cloud Infrastructures," in Aalto University, T-110.5290 Seminar on Network Security, 2010.
- [12] A. Bakshi and Y. B. Dujodwala, "Securing Cloud from DDOS Attacks Using Intrusion Detection System in Virtual Machine," *2010 Second International Conference*

on Communication Software and Networks (ICCSN'10), pp. 260–264, 2010.

Telecommunications, Networking and Broadcasting (PGNet 2013), 2013.

[13] KPMG, “The Cloud Changing the Business Ecosystem,” 2011.

[14] S. Taghavi Zargar, H. Takabi, and J. Joshi, “DCDIDP: A Distributed, Collaborative, and Data-driven Intrusion Detection and Prevention Framework for Cloud Computing Environments,” *Proceedings of the 7th International Conference on Collaborative Computing: Networking, Applications and Worksharing (CollaborateCom)*, pp. 332–341, 2011.

[15] A. Patel, M. Taghavi, K. Bakhtiyari, and J. Celestino Júnior, “An intrusion detection and prevention system in cloud computing: A systematic review,” *Journal of Network and Computer Applications*., vol. 36, no. 1, pp. 25–41, Jan. 2013.

[16] H. Hamad and M. Al-Hoby, “Managing Intrusion Detection as a Service in Cloud Networks,” *International Journal of Computer Applications*., vol. 41, no. 1, pp. 35–40, Mar. 2012.

[17] S. N. Dhage and B. B. Meshram, “Intrusion detection system in cloud computing environment,” *International Journal of Cloud Computing*., vol. 1, no. 2/3, p. 261, 2012.

[18] J. Lee, M. Park, and J. Eom, “Multi-level Intrusion Detection System and log management in Cloud Computing,” *2011 IEEE 13th International Conference on Advanced Communication Technology*, no. 1, pp. 552–555, 2011.

[19] C.C. Lo, C.C. Huang, and J. Ku, “A Cooperative Intrusion Detection System Framework for Cloud Computing Networks,” in *2010 39th International Conference on Parallel Processing Workshops*, 2010, pp. 280–284.

[20] M. Randles, D. Lamb, and A. Taleb-Bendiab, “A Comparative Study into Distributed Load Balancing Algorithms for Cloud Computing,” in *2010 IEEE 24th International Conference on Advanced Information Networking and Applications Workshops*, 2010, pp. 551–556.

[21] Z. Mahmood and C. Agrawal, “Intrusion Detection in Cloud Computing environment using Neural Network,” *International Journal of Research in Computer Engineering & Electronics*., vol. 1, no. 1, pp. 1–4, 2012.

[22] H. M. Alsafi, W. M. Abdullallah, and A. K. Pathan, “IDPS : An Integrated Intrusion Handling Model for Cloud Computing Environment,” *International Journal of Computing & Information Technology (IJCIT)* vol. 4., no.1 (2012): pp. 1-16.

[23] Cloud Security Alliance, “The Notorious Nine Cloud Computing Top Threats in 2013,” 2013, Cloud Security Alliance, 2013.

[24] Á. MacDermott, Q. Shi, M. Merabti, and K. Kifayat, “Detecting Intrusions in the Cloud Environment,” in *14th Annual Postgraduate Symposium on Convergence of*