



LJMU Research Online

MacDermott, Á, Shi, Q, Merabti, M and Kifayat, K

Hosting critical infrastructure services in the cloud environment considerations

<http://researchonline.ljmu.ac.uk/id/eprint/6934/>

Article

Citation (please note it is advisable to refer to the publisher's version if you intend to cite from this work)

MacDermott, Á, Shi, Q, Merabti, M and Kifayat, K (2015) Hosting critical infrastructure services in the cloud environment considerations. International Journal of Critical Infrastructures, 11 (4). pp. 365-381. ISSN 1475-3219

LJMU has developed [LJMU Research Online](#) for users to access the research output of the University more effectively. Copyright © and Moral Rights for the papers on this site are retained by the individual authors and/or other copyright owners. Users may download and/or print one copy of any article(s) in LJMU Research Online to facilitate their private study or for non-commercial research. You may not engage in further distribution of the material or use it for any profit-making activities or any commercial gain.

The version presented here may differ from the published version or from the version of the record. Please see the repository URL above for details on accessing the published version and note that access may require a subscription.

For more information please contact researchonline@ljmu.ac.uk

<http://researchonline.ljmu.ac.uk/>

Hosting critical infrastructure services in the cloud environment considerations

Áine MacDermott*, Qi Shi, Madjid Merabti,
and Kashif Kifayat

PROTECT: Research Centre for Critical Infrastructure Computer
Technology and Protection,
School of Computing and Mathematical Sciences,
Liverpool John Moores University,
Liverpool, L3 3AF, UK
Email: a.mac-dermott@2008.ljmu.ac.uk
Email: q.shi@ljmu.ac.uk
Email: m.merabti@ljmu.ac.uk
Email: k.kifayat@ljmu.ac.uk
*Corresponding author

Abstract: Critical infrastructure technology vendors will inevitably take advantage of the benefits offered by the cloud computing paradigm. While this may offer improved performance and scalability, the associated security threats impede this progression. Hosting critical infrastructure services in the cloud environment may seem inane to some, but currently remote access to the control system over the internet is commonplace. This shares the same characteristics as cloud computing, i.e., on-demand access and resource pooling. There is a wealth of data used within critical infrastructure. There needs to be an assurance that the confidentiality, integrity and availability of this data remains. Authenticity and non-repudiation are also important security requirements for critical infrastructure systems. This paper provides an overview of critical infrastructure and the cloud computing relationship, whilst detailing security concerns and existing protection methods. Discussion on the direction of the area is presented, as is a survey of current protection methods and their weaknesses. Finally, we present our observation and our current research into hosting critical infrastructure services in the cloud environment, and the considerations for detecting cloud attacks.

Keywords: critical infrastructure; intrusion detection; cloud computing; cloud security; critical infrastructure protection; critical sectors.

Reference to this paper should be made as follows: MacDermott, Á., Shi, Q., Merabti, M. and Kifayat, K. (xxxx) 'Hosting critical infrastructure services in the cloud environment considerations', *Int. J. Critical Infrastructures*, Vol. X, No. Y, pp.xxx-xxx.

Biographical notes: Áine MacDermott is a PhD research student studying at Liverpool John Moores University in the School of Computing and Mathematical Sciences, where she received her BSc (Hons) in Computer Forensics in 2011. Her current research towards her PhD is focusing on protecting critical infrastructure services in the cloud environment. Her research interests include critical infrastructure protection, computer network security, and digital forensics.

Qi Shi is a Professor in Computer Security in the School of Computing and Mathematical Sciences at Liverpool John Moores University in the UK. He received his PhD in Computing from the Dalian University of Technology, P.R. China. Prior to joining the Liverpool John Moores University, he worked as a Research Associate for the Department of Computer Science at the University of York in the UK. His research interests include security protocol design, ubiquitous computing security, formal models, sensor network security, computer forensics, and intrusion detection. He has published extensively, and is supervising several projects, in these research areas. He also actively participates in other academic activities including conference program committee and journal editorial board memberships.

Madjid Merabti is a Professor, and Director of School, in the School of Computing and Mathematical Sciences at Liverpool John Moores University in the UK. He is a graduate of Lancaster University in the UK. His research interests include computer security, critical infrastructure protection, distributed multimedia systems, computer networks, and mobile computing. He is widely published in these areas and has a number of government and industry supported research projects. He is a programme chair for many international conferences and serves as an editor and an editorial board member for a number of international journals.

Kashif Kifayat received his PhD in Computer Science from Liverpool John Moores University. He worked as a Research Fellow in Network Security for two years in the School of Computing and Mathematical Science at Liverpool John Moores University, and he is currently working as a Lecturer. His research interests include network security, security of complex and scalable systems, and security in wireless sensor networks.

1 Introduction

Critical infrastructures include sectors such as energy resources, finance, food and water distribution, health, manufacturing, government services and in some cases people. The consequences of an attack on one of these could result in loss of life, economic damage or a devastating effect on the operation of government services and military defence. In recent years, critical infrastructures have become increasingly more interlinked and are often connected to the internet; consequently, this makes these systems more vulnerable and exposed to the threat of cyber attacks (Hurst, 2011). In this paper, we refer to critical infrastructure as those whose disruption could have a high socioeconomic impact. The critical infrastructure systems that support major industries, such as manufacturing, transportation, and energy, are highly dependent on information communication systems for their command and control. While a high dependence on industrial control systems, i.e., supervisory control and data acquisition (SCADA) still exists, critical infrastructure systems are migrating to new communication technologies. As a result, common communications protocols and open architecture standards are replacing the diverse and disparate proprietary mechanics of industrial control systems. This replacement can have both positive and negative impacts (Department of Homeland Security, 2009).

Utilising cloud computing within an environment that historically has not had any internet connectivity would appear trivial to some, however research has shown that cloud computing will reach the information and communication technology (ICT)

services that are operating critical infrastructure (Dekker, 2013; OTE, 2012; Khajeh-Hosseini et al., 2010). Operators of critical infrastructures, in particular the ICT that supports gas and electricity utilities and government services, are considering using the cloud to provision their high assurance services. This is reflected in a recent white paper produced by the European Network and Information Security Agency (ENISA) (Dekker, 2013), which provides specific guidelines in this area.

Deploying high assurance services in the cloud increases cyber security concerns, as successful attacks could lead to outages of key services that our society depends on, and disclosure of sensitive personal information. To address these concerns, a range of security measures must be put in place, such as cryptographic storage and network firewalls. Industrial control systems have evolved from being monolithic, to distributed, to networked. Remote access is common practice, i.e., remote access to intelligent electronic devices (IEDs) or user interfaces in a substation for maintenance. In addition, historians within the infrastructure can take operational and production data from the SCADA environment and publish it to web assessable portals to be viewed by corporate users. Automation has become an indispensable part of service provision and has increased exponentially, as demand for digital services and interconnectivity has increased. The reliance on these systems has resulted in ICT playing a key role in the provisioning of services that critical infrastructures deliver to the general population. Cloud computing can be conveyed as the next logical progression as the cloud paradigm is already being used for crucial assets within the critical infrastructure industry.

The aim of this paper is to provide an outline of the cloud computing and critical infrastructure connexion, conveying the evident utilisation by critical infrastructure technology vendors, and highlighting the security concerns and subsequent disinclination to host services in this environment. Requirements for cloud computing for critical infrastructure will be defined; in addition, a use case will express some benefits of adopting this approach. We will expand on our observation of this area and detail our research into the use of predictive algorithms, such as CUSUM, for detecting attacks against the cloud, i.e., distributed denial of service (DDoS) for unavailability of data, which, if it was infrastructure data would cause high operational and financial loss. Transfer of data is time critical within these systems, and availability of services is imperative, thus disruptions could have high socioeconomic impact.

The remainder of this paper is organised as follows. Section 2 provides background on critical infrastructure and cloud computing, detailing their relationship, and associated security concerns. Section 3 illustrates cloud computing requirements for critical infrastructure and describes a use case based on our previous work. In Section 4, we compare numerous protection methods for cloud computing and highlight their weaknesses; we discuss algorithms for detecting cloud attacks, and offer insight on this area. Finally, we present our conclusions and future work in Section 5.

2 Background

2.1 Critical infrastructure

Critical infrastructures deliver critical services that are keys to economy, safety of people, the community, and the functioning of government. They are key service providers greatly relied upon by governments and the general population. There are other

infrastructures society relies upon, but they are not considered in the scope of this paper. Many critical infrastructure areas have become heavy ICT users with automation playing a key role in production. Use of ICT has also begun to expand in areas such as agriculture, food, and water where control systems and the use of sensor equipment are helping to facilitate production and become more adaptive to the growing demands being placed on them.

Clearly, the use of ICT is becoming more pervasive in all areas of critical infrastructure (Hurst et al., 2013). These infrastructures face significant threats due to the growth in the use of SCADA systems and increasingly integrating networks. Although the complex infrastructure provides great capabilities for operation, control, business and analysis, it also increases the security risks due to cyber-related vulnerabilities (Ten et al., 2010). 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.

There are similarities 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. Adoption of cloud technologies allows critical infrastructure to benefit from dynamic resource allocation for managing unpredictable load peaks.

2.2 *Cloud computing*

The National Institute of Standards and Technology (NIST) defines five essential characteristics of cloud computing (Mell and Grance, 2011):

- On-demand self-service: A user should be able to acquire or release resources without requiring outside human interference.
- Broad network access: Resources should be available over the network, through both thin and thick clients
- Resource pooling: Resources are pooled to serve disparate customers on the same or different physical machines. Resources can be dynamically assigned according to customer demands. Resources can include computation, storage, and networking to name a few.
- Rapid elasticity: Users can acquire, release, and scale resources in an elastic manner, making the available resources appear unlimited from the clients' point of view.
- Measured service: The cloud management layer constantly monitors, controls, and reports resource use to both the provider and client, providing a metering capability.

Cloud providers usually build up large scale data centres and provide cloud users with computational resources in three delivery models, distinguished by their level of resource abstraction (Annapureddy, 2010). These are:

- software as a service (SaaS)
- platform as a service (PaaS)
- infrastructure as a service (IaaS).

Security is a major concern in cloud adoption. Critical security issues include data integrity, user confidentiality, and trust among providers, individual users, and user groups. Additionally, availability issues and real world impact would be the main concern for providers of critical infrastructure, depending upon the operations or services they are hosting (Hwang and Li, 2010). 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 SaaS, in which enterprises host and operate their applications over the internet so the customers can access it. One benefit of this model is customers do not need to buy any software licences or any additional equipment for hosting the application(s).

The virtual level includes PaaS and 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 with relative ease. IaaS provides a convenient option for organisations by migrating the IT infrastructure to the cloud provider. This means it is the responsibility of the cloud provider to tackle the issues of IT infrastructure management, such as configuring servers, routers, firewalls, to name a few. The physical level refers to the infrastructure upon which clouds are deployed. Table 1 conveys security requirements and threats for each service level.

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 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. There are two types of private cloud:
 - 1 On-premise: This is a cloud that has been integrated into an organisation's IT process. These clouds are better suited for organisations which desire greater configurability and control over their data infrastructure.
 - 2 External private cloud: This is a private cloud platform that is hosted by an external cloud provider, but with the guarantee of privacy.
- 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 (KPMG International, 2011).

Table 1 Security requirements and threats associated with each service level

<i>Level</i>	<i>Service level</i>	<i>Security requirements</i>	<i>Threats</i>
Application level	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 • Privacy breach • Session hijacking • Traffic flow analysis
Virtual level	PaaS	<ul style="list-style-type: none"> • Access control • Application security 	<ul style="list-style-type: none"> • Connection flooding • DDoS attack
	IaaS	<ul style="list-style-type: none"> • Cloud management control security • Communication security • Data security • Secure images • Virtual cloud protection 	<ul style="list-style-type: none"> • Defacement • Disrupting communications • Exposure in network • Impersonation • Programming flaws • Software modification • Software interruption • Session hijacking • Traffic flow analysis

Source: MacDermott et al., (2013)

The adoption of this innovative architecture may introduce a number of additional threats that vendors may not have considered. Based on the problem at hand, it is evident that sufficient security metrics need to be developed for protecting the sensitive data being stored in the cloud environment. The ability to clearly identify, authenticate, authorise, and monitor who or what is accessing the assets of an organisation is essential for protecting an information system from threats and vulnerabilities. The focus on cybercrime at a global level has led to ‘as-a-service’ models for illegal activity. The cybercrime market now affords potential criminals with a multitude of services, which means that deep technical expertise is not a prerequisite. Much like cloud computing, the service-based nature of cybercrime allows greater efficiency and flexibility when conducting business. Examples include the ability to rent services that offer financial return or that claim to be able to bring down entire sites or systems (Samani, 2013).

3 Cloud computing and critical infrastructure

Moving services to the cloud is a trend that has been present for many years now, with a constant increase in sophistication and complexity of such services. Today, even critical infrastructure operators are considering moving their services and data to the cloud; most

prominently are telecommunication operators, who aim to run their services as virtual network services. These services are usually composed from a set of components, each with individual resilience and scalability requirements. Hence, the problem of describing the blueprint of how to build a service from its components, including the components' requirements, and how to derive an actual deployment from such a blueprint needs to be resolved (Scholler et al., 2013).


Since ICT infrastructures have become an integral part of almost all organisations, cloud computing will have an important impact on them (OTE, 2012). Cloud services can offer efficient access to large ICT infrastructures that benefit from the economy of scale, which can be replicated and distributed globally. Cloud computing addresses at least two fundamental requirements of the UK energy industry community. First, accurate network simulations require highly variable quantities of computational resources depending on the contingent situation of energy delivery or on the type of energy delivered. Renewable energy output is typically much less predictable than the constant output offered by conventional generation sources, such as coal, oil, gas, or nuclear. For this reason, running simulations on the cloud allows for dynamic scaling of the required computational and data resources.

In some critical infrastructure, it is not just the data that is critical; it is the remote terminal units (RTUs) and programmable logic controllers (PLCs) running pumps or maintaining the turbines. Each infrastructure may vary its migration to the cloud environment, but what is clear is that when the data is in the cloud, its availability, and integrity are of utmost importance. Remote access into a SCADA system over the internet is commonplace for corporate staff to observe and analyse the information collected by the historians. In this instance, it is the data from the infrastructure that is critical. A further way in which infrastructure vendors could embrace the benefits of the cloud environment could involve storing the historian processes in an onsite private cloud.

Many operators do not have the infrastructure to support the growing need for accurate predictive and historical simulations imposed by the adoption of renewable energy sources and the ongoing development of smart grids. Cloud computing allows these operators to reduce or avoid over investment in hardware resources and their associated maintenance. Infrastructure vendors will inevitably take advantage of the benefits cloud computing has to offer.

3.1 Use case

Previous work of ours (MacDermott et al., 2013) has outlined a way in which critical infrastructure could utilise the cloud environment for improved performance and analysis of the automation processes. Most industrial plants employ networked historian servers storing process data and other possible business and process interfaces. PLCs in the control system generate a vast amount of data and logs. These communication logs are stored in the historian databases. This historical data is being logged 24/7, and in some cases can come from over 6,700 data points, so that it could be easily accessible by both operators and engineers (Fovino et al., 2010; Verba, 2008). These historian servers receive data from field control processors or front-end processors, which issue control commands to and poll data from devices in field networks. Figure 1 illustrates a high level arrangement of general critical infrastructure components.

Figure 1 Critical infrastructure components (see online version for colours)

The control network typically contains assets such as the human-machine interface (HMI) and other workstations, which run control system applications on conventional computer platforms. The field network devices directly monitor and control a physical process, such as refining, manufacturing, or electric power generation/transmission/distribution (Briesemeister et al., 2010).

One way for critical infrastructure to utilise the cloud environment would be for the historian database to send these historical processes to a private cloud. The use of a private cloud to audit the data from the system and process it more effectively would be valuable. This would overcome the challenges associated with processing vast data sets generated by the control systems. The cloud environment is suitable as it has massive storage and computational capabilities, is distributed and elastic, offering improved processing rates and efficiency compared to current methods.

This collection of data can be used to perform behavioural analysis and modelling of the flow of information. Looking for trends and subtle changes in the data would be beneficial in achieving state awareness. Behaviour modelling can take place without affecting the system in any way. In control system architectures, the major cyber-attack vector is the flow of network commands (Carcano et al., 2003). By processing the sensor data from the historian in a private cloud environment, the behaviour of the infrastructure could be analysed to discriminate between critical states due to cyber-attacks and critical states due to faults/physical attacks.

Another advantage of using the cloud in this context is to aggregate data from the IP enabled control devices, which have limited resources and cannot process data locally. Though there are benefits of this connection, one of the main concerns with utilising critical infrastructure services in the cloud environment is the threat of attack. Critical infrastructure providers will use cloud applications for their systems and related issues need to be investigated. Especially, with critical infrastructure in the cloud, legal

compliance will be a vital importance (Florian et al., 2013). Critical infrastructure data is normally highly sensitive and therefore subject to legal regulations for data security. Traditionally critical infrastructure secured data inside a closed environment with restricted external access. This makes it important to establish data access regulations to secure critical infrastructure data in the cloud.

3.2 Requirements for critical infrastructure

The cloud computing operational model is a major recent trend in the ICT industry, which has gained tremendous momentum. This trend will likely also reach the ICT services that support critical infrastructures, because of the potential cost savings and benefits of increased resilience due to elastic cloud behaviour. However, employing critical infrastructure services in the cloud introduces security and resilience requirements that existing offerings do not address well. For example, due to the opacity of cloud environments, the risks of deploying cloud-based critical infrastructure services are difficult to assess, especially at the technical level, but also from legal or business perspectives.

Though there are many benefits and advantages regarding the utilisation of cloud computing by critical infrastructure operators there are considerations that must be deliberated. Requirements in critical infrastructure regarding overall redundancy, data availability, authenticity, secure access, and low latency network connectivity are typically higher than in commercial applications. Critical infrastructure imposes much stronger requirements for security, reliability, and resilience on cloud computing environments.

- Control

A SLA defines how the consumer will use the services and how the provider will deliver them. Responsibilities of each party and remedies should be included. The SLA cannot be adapted or negotiated, which often deters organisations from pledging fully to the cloud. Additionally, there is no governing body, and no one enforcing this compliance. In this scenario, there may be data breaches that are often undisclosed to the affected party. Control is a common challenge as depending on which deployment model is chosen, control is not always in the hand of the owner. Private clouds allow organisations to shape how their data is stored and controlled, and what security measures are in place. Public clouds raise the questions: where is the data stored? Who has control? In this instance, public clouds do not seem the viable option.

- Data centric security approach

In general, a data-centric security approach must ensure that data protection mechanisms are deployed across all provided security solutions and that data owners have the full control over who has the right to access, use the data and what they are authorised to do with it. In addition, institutional security policies and access rules can be specified and mapped to the cloud environment. Requirement based security issues can be quite different for critical infrastructure applications and for common IT applications but need to be considered in combination for the given context.

- Protection

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, DDoS attacks, to name a few (Bhadauria, 2011). Tailored intrusion detection and prevention mechanisms are essential. Compared to other systems and services in the cloud environment, a critical infrastructure requires a much higher level of assurance. One of the risks in a multi-tenant environment is over provisioning of resources. Over provisioning resources results in resource contention and potential lack of availability, effectively creating a denial of service situation (Steiner, 2012). This could have an impact on users of the cloud service who depend on its continuity. An example of such a requirement is when critical services, such as emergency care, depend on these cloud services.

- Legal issues

Legal requirements include data protection and regulatory requirements. Issues also surround data being exchanged across multiple countries that have different laws and regulations concerning data traversal, protection requirements and privacy laws. Examples of such risks include, but not limited to, risks resulting from possible changes of jurisdiction and the liability or obligation of the vendor in case of loss of data and/or business interruption (Dahbur et al., 2011). There are also geographical requirements for healthcare data being stored, which could have legal ramifications if violated.

4 Methods for intrusion detection in the cloud environment

Based on the critical infrastructure requirements for cloud computing, we analysed the literature for solutions that may be applicable for protecting these services. There are many taxonomies detailing cloud vulnerabilities and attacks (Khajeh-Hosseini et al., 2010; Karnwal et al., 2012; Chonka and Abawajy, 2012; Aliyev et al., 2013; Dahbur et al., 2011; Grobauer et al., 2011). As such, it can be inferred that as the use of the cloud in organisations develops, so will the rate of DDoS 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 DDoS attack. Asymmetric application-level DDoS 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 (Cloud Security Alliance, 2013).

Unavailability of services due to cloud outages can cause monetary loss to cloud providers and operational loss to cloud users, which would be disastrous for critical infrastructure as another service could be dependent upon it. 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. The most common mechanism for protecting the cloud environment currently are intrusion detection systems (IDSs). These mechanisms require an extensive use of hardware, especially CPU and memory, and may cause unintentional resource exhaustion or a

bottleneck. IDSs can observe the traffic from each virtual machine (VM), generating alert logs and can manage cloud computing globally.

As IDSs generate substantial amounts of logs, administrators have to prioritise what to analyse first. It is often difficult to analyse logs as communications a vast number of systems and consumers generate large quantities of logs. Effective log and resource management is necessary, as an administrator may miss imperative alerts and events, thus endanger their system(s). In the cloud environment, where a vast amount of data is generated due to high network access rates, an IDS must be robust against noise data and false positives. Since cloud infrastructures have enormous network traffic, traditional IDSs are not efficient to handle such a large data flow.

The work of Hamad and Al-Hoby (2012) takes an innovative approach to tackling this security problem. They designed and implemented the 'Cloud Intrusion Detection Service' (CIDS). CIDS can be deployed by cloud providers to enable clients to subscribe to 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 intrusion prevention and detection system (IDS/IPS). The model outperforms current solutions used for service-based IDS but at the same time provides minimal overhead comparable to traditional IDS deployment for single network protection.

In the work of Dhage et al. (2011), it is conveyed that when there is only one IDS in the entire network, the load on it increases as the number of hosts increases. It is difficult to keep track of different kinds of attacks or intrusions, which are acting on each of the hosts 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 capable of working in an increasingly efficient manner. By proposing a model in which each instance of IDS has to monitor only a single user, an effort has been made to create a coordinated design, which will be able to gather appropriate information about the user, thus enabling it to classify intrusions in a better way.

The work of Lee et al. (2011) proposes a multi-level IDS and log management method. Their method is based on consumer behaviour for applying IDS effectively to the cloud system. A risk level is assigned to user behaviour based on analysis of their behaviour over time. By applying differentiated levels of security strength to users, based on the degree of anomaly, this 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 users suspected of being the highest risk to the system first.

Arshad et al. (2013) propose intrusion severity analysis for the cloud environment. Their work focuses on application-specific vulnerabilities. They establish a severity analysis for the overall architecture through the use of decision trees, which are a supervised learning approach. The main objective is to investigate the effectiveness of machine learning techniques for severity analysis for cloud environments. The fault model consists of arbitrary intrusions but excludes faults that occur at site level, i.e., if the hypervisor or domain is compromised.

In Mahmood and Agrawal (2012), PCANNA (principal component analysis neural network algorithm) is proposed to reduce the memory and CPU time required to detect an attack. Feature reduction is used to remove insignificant information from the high dimensional database of cloud traffic data. A back propagation algorithm is applied on

reduced cloud traffic data for classification. Dimensional reduction techniques compact and correlate similar alerts to detect intrusions that are more complex.

Our survey of existing methods has identified the current weaknesses with existing approaches:

- fail or overload with a high volume of traffic
- loss of accuracy
- inaccurate profile of usage
- require human intervention
- simply flags suspect behaviour
- ineffective log management
- cannot detect novel/unknown attacks.

While these weaknesses are challenging and need to be addressed, a common occurrence in the literature was the inefficiencies of algorithms for detecting attacks in the cloud environment. New algorithms optimised for detecting cloud attacks in an efficient manner are needed, and this is something we are currently exploring. Additionally, having a solution with the ability to adapt to varying computational and network loads in order to not be invasive is needed also. In the service-oriented architecture of the cloud, collaboration means data is coming from many different sources so existing IDS techniques will not be able to process data of this scale. 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.

4.1 Algorithms for detecting attacks in the cloud

We analysed algorithms for detecting attacks in the cloud environment, and looked in particular at adaptive threshold, random early drop (RED), robust random early drop (RRED), and CUSUM. Each of these are used by a variety of solutions in this area, but each have differing benefits and drawbacks. Adaptive threshold (Gore, 2012) detects anomalies based on violations of a threshold that is adaptively set based on recent traffic measurements. Seasonal variations and trends are taken care of by using an adaptive threshold whose value is set based on an estimate of the mean number of packets under consideration or the rate, either of which are computed from recent traffic measurements.

RED takes a different approach and monitors the average queue size and drops packets based on statistical probabilities. If the buffer is almost empty, all incoming packets are accepted. As the queue grows, the probability for dropping an incoming packet grows too. When the buffer is full, the probability has reached 1 and all incoming packets are dropped. In comparison, RRED was proposed to improve the TCP throughput against DDoS attacks, particularly Low-rate DDoS (LDDoS) attacks. Experiments have confirmed that the existing RED-like algorithms are notably vulnerable under LDDoS attacks due to the oscillating TCP queue size caused by the attacks. RRED algorithm can significantly improve the performance of TCP under Low-rate Denial-of-Service attacks.

Many algorithms, such as random sampling, do not take into account traffic dynamics. As a result, they cannot guarantee the sampling error falls within a prescribed error tolerance level. How to discover the evolving process of the network traffic and how to improve the accuracy of real-time detection is problematic. In regards to critical infrastructure services in the cloud environment, the confidentiality, integrity, and availability of the data are of utmost importance.

CUSUM (Cárdenas et al., 2011) involves the calculation of a cumulative sum (which is what makes it 'sequential'). Samples from a process X_n are assigned weights W_n and summed as follows:

$$S_0 = 0$$

$$S_{n+1} = \max(0, S_n + x_n - w_n)$$

When the value of S exceeds a certain threshold value, a change in value has been found. The formula only detects changes in the positive direction.

There are preventive measures in place to protect against such attacks, but they seem to be focusing on generic DDoS, where the characteristics mimic previous attacks of such nature. However, the rise in high volume and low rate DDoS is a problem. They could be spread out over a period of time, and have random high bursts, which can confuse the preventative measures. Algorithms for detecting DDoS attacks in the cloud environment often sample the packets and drop any they deem to be malicious. These can often be false positives.

Our research focuses on maintaining the availability of the data, as previously described, the service in question could be financial, organisational, or on demand. Protecting the cloud environment from DDoS attacks is on the increase. Cyber-crimes targeting organisations in the cloud environment are rising and the malicious nature is to prevent the availability of data. Having analysed these algorithms for their suitability, we have identified the following weaknesses in regards to detecting DDoS in the cloud environment:

- sample packets are often inaccurate
- vulnerability to unknown types of DDoS attacks
- does not always ensure the accuracy of estimation and tend to over sample at peak periods when efficiency and timeliness are crucial
- random sampling does not take into account traffic dynamics
- inefficient on low rate DDoS attacks
- prone to error on high-rate DDoS attacks.

There is an emerging need for the traffic processing capability of IDS, to match the high throughput of today's high-bandwidth networks. Recent research has shown that the vast majority of security solutions deployed today are inadequate for processing traffic at a sufficiently high rate to keep pace with the network's bandwidth. However, existing sampling algorithms are poorly suited for this task, especially because they are unable to adapt to the trends in network traffic. Satisfying such a criterion requires a sampling algorithm to be capable of controlling its sampling rate to provide sufficient accuracy at minimal overhead.

The adaptive threshold algorithm and CUSUM algorithm appear to be the most applicable for detecting attacks in the cloud environment. Our current work involves testing this hypothesis against real network data to determine the effectiveness of protecting critical infrastructure services in the cloud environment and to maintain availability of the data through predictive responsiveness. We seek to investigate how the parameters of the detection algorithm and the characteristics of the attack affect the performance. Currently, we have analysed these two predictive algorithms against values we have generated using numPY in Python. We have used the semantics of the algorithms and applied them syntactically. Matplotlib is used to plot the observed values on a graph to illustrate if they have exceeded the exponentially-weighted moving average (EWMA) threshold. Early detection of DoS attacks with increasing intensity would enable defensive actions to be taken earlier. Algorithms based on change point detection, such as CUSUM, can exhibit robust performance over a range of different types of attacks, without being more complex.

Next, we are to test the algorithms in a cloud environment and return a tuple of the start time, the end time, and returning a list of values per minute/second. In addition, comparing the performance of the detection algorithms in terms of three metrics: detection probability, false alarm ratio and detection delay.

4.2 Discussion

The scope of critical infrastructure transcends into industry, which is broadly defined and highly distributed. The term ‘critical infrastructure’ does not help security practitioners or policy makers, but the term ‘critical assets’ may. By calling a whole system critical can describe a class of objects or phenomenon in general terms, without specific details or attention to individual attributes. Not everything will be deemed critical, but the data can be. That which is critical, in terms of data, for one organisation may not be crucial for another, but when it is infrastructure services, or important historical processes it is. The criteria for data/or asset criticality can differ depending on infrastructure but we refer to it as essential for the functioning of operations. Data which is required to be viewed in real time can be deemed critical, or assets such as PLCs operating in real time could also be too. Currently there are no satisfactory solutions available to address the issues of intrusion detection in the cloud environments for critical infrastructure system protection. New algorithms optimised for detecting cloud attacks in an efficient manner are needed, and this is something we are currently exploring. We seek to investigate how the parameters of the detection algorithm and the characteristics of the attack affect the performance.

Our main aim is to derive new algorithms optimised for detecting cloud specific security threats to Cloud Computing infrastructures. This will occur through the analysis of various intrusion and anomaly detection algorithms and measurable cloud computing parameters, such as communication types, communication patterns, and access patterns, which can be used to detect and isolate abnormal behaviour and threats. Mechanisms can be in place to detect generic DDoS attacks; however, high-scale and small-scale attacks are still possible since these types of attacks differ from normal bandwidth hungry applications.

As cloud computing grows in popularity, new models are deployed to exploit its full capacity. One of these ideas is the deployment of cloud federations. A cloud federation is an association among different cloud service providers (CSPs) with the goal of sharing

resources and data. In order to cope with the resource capacity limits of a single cloud provider, the concept of federating multiple heterogeneous organisations is receiving increasing attention. Critical infrastructure vendors could adopt this approach. The effects of attacks can span from the loss of some data, to the potential isolation of parts of the federation (Macdermott et al., 2013). Protecting the federated cloud against cyber attacks is a key concern, since there are potentially significant economic consequences. Cloud federation and CSPs will benefit significantly if there is a comprehensive IDS that evolves based on their requirements. This is an area we aim to additionally consider.

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 are largely inefficient to be deployed in cloud computing environments due to their openness and specific essence. 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 infrastructures have enormous network traffic, traditional IDSs are not efficient to handle such a large flow of data. Due to the large data sets, classification techniques require a huge amount of memory and CPU usage. Failure is not an option in protecting critical infrastructure services in a federated cloud environment. We need a solution that can adapt itself to the ever-changing behaviour of the cloud environment. The cloud federation, and the CSPs, will benefit significantly if there is a comprehensive IDS that evolves based on their requirements. The security of applications and services provided in the cloud, against cyber attacks, is hard to achieve due to the complexity, heterogeneity, and dynamic nature of such systems.

5 Conclusions and future work

Cloud computing is being adopted in critical sectors such as transport, energy and finance. This makes cloud computing services critical in themselves. Adoption of cloud technologies allows critical infrastructure to benefit from dynamic resources allocation from managing unpredictable load peaks. Given the public awareness of critical infrastructure and its importance, the public wants to be assured that these systems are built to function in a secure manner. When cyber attacks and cyber disruptions happen, millions of users are potentially affected. A cyber disruption in this context means a temporary or permanent loss of service will mean users of the cloud service who rely on its continuity are affected. 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. Therefore recognising the signs of an attack quickly, and being able to limit the effect on operation is imperative. Cloud outages are unexpected events that occur within a cloud infrastructure and consequently affect availability of services placed in the cloud. Should this be critical infrastructure services, or important historical processes, this will not be good for the reputation of the infrastructure vendor, or for the cloud service provider.

Our future work involves developing methods and principles to tackle network availability attacks against the cloud environment. Network attacks, such as DDoS, aim

to use up or overwhelm the communication and computational resources and to result in delay or failure of communication. The detection of subtle changes in behaviour can be achieved through the CUSUM algorithm. Having identified current method weaknesses, we aim to expand and develop a tailored protective approach, which can help protect and detect intrusions in the cloud environment. The concept of cloud federations is an avenue we aim to explore further as this may be a suitable environment for developing resilient solutions for cloud protection.

References

- Aliyev, R., Seo, D. and Lee, H. (2013) 'DROP-FAST: defending against DDoS attacks using cloud technology', *The 2013 International Conference on Security and Management*.
- Annapureddy, K. (2010) 'Security challenges in hybrid cloud infrastructures', in *Aalto University, Seminar on Network Security*, p.T-110.5290.
- Arshad, J., Townend, P. and Xu, J. (2013) 'A novel intrusion severity analysis approach for clouds', *Future Generation Computer Systems*, Vol. 29, No. 1, pp.416-428.
- Bhadauria, R. (2011) *A Survey on Security Issues in Cloud Computing*, in arXiv preprint arXiv:1109.5388.
- Briesemeister, L. et al. (2010) 'Detection, correlation, and visualization of attacks against critical infrastructure systems', in *2010 Eighth Annual International Conference on Privacy Security and Trust (PST)*, IEEE, pp.15-22.
- Carcano, A. et al. (2003) 'A multidimensional critical state analysis for detecting intrusions in SCADA systems', *IEEE Transactions on Industrial Informatics*, Vol. 7, No. 2, pp.179-186.
- Cárdenas, A.A., Amin, S. and Lin, Z. (2011) 'Attacks against process control systems: risk assessment, detection, and response categories and subject descriptors', *6th ACM Symposium on Information, Computer and Communications Security*, pp.355-366.
- Chonka, A. and Abawajy, J. (2012) 'Detecting and mitigating HX-DoS attacks against cloud web services', *2012 15th International Conference on Network-Based Information Systems (Xml)*, pp.429-434.
- Cloud Security Alliance (2013) *The Notorious Nine Cloud Computing Top Threats in 2013*, 2013, Cloud Security Alliance.
- Dahbur, K., Mohammad, B. and Tarakji, A.B. (2011) 'A survey of risks, threats and vulnerabilities in cloud computing', *Proceedings of the 2011 International Conference on Intelligent Semantic Web-Services and Applications - ISWSA '11*, pp.1-6.
- Dekker, D.M.A.C. (2013) *Critical Cloud Computing – A CIIP Perspective on Cloud Computing Services*, Greece.
- Department of Homeland Security (2009) *Recommended Practice: Improving Industrial Control Systems Cybersecurity with Defense-In-Depth Strategies*, USA.
- Dhage, S.N. et al. (2011) 'Intrusion detection system in cloud computing environment', *Proceedings of the International Conference & Workshop on Emerging Trends in Technology - ICWET '11, (ACM)*, pp.235-239.
- Florian, M., Paudel, S. and Tauber, M. (2013) 'Trustworthy evidence gathering mechanism for multilayer cloud compliance', in *8th International Conference for Internet Technology and Secured Transactions (ICITST)*, pp.529-530.
- Fovino, I.N. et al. (2010) 'Modbus/DNP3 state-based intrusion detection system', in *2010 24th IEEE International Conference on Advanced Information Networking and Applications*. IEEE, pp. 729-736.
- Gore, S. (2012) 'Anomaly detection algorithm using multi-agents', *International Journal of Scientific & Technology Research*, Vol. 1, No. 3, pp.54-57.

- Grobauer, B., Walloschek, T. and Stocket, E. (2011) 'Understanding cloud computing vulnerabilities', *IEEE Computer and Reliability Societies*, April, Vol. 9, No. 2, pp.49–57.
- Hamad, H. and Al-Hoby, M. (2012) 'Managing intrusion detection as a service in cloud networks', *International Journal of Computer Applications*, Vol. 41, No. 1, pp.35–40.
- Hurst, W. (2011) 'Towards a framework for operational support in critical infrastructures', in *12th Annual Postgraduate Symposium on Convergence of Telecommunications, Networking and Broadcasting (PGNet 2011)*, Liverpool, UK.
- Hurst, W. et al. (2013) 'Protecting Critical Infrastructures through behavioural observation', *International Journal of Critical Infrastructures*, Vol. 10, No. 2, pp.174–192.
- Hwang, K. and Li, D. (2010) 'Trusted cloud computing with secure resources and data coloring', *IEEE Internet Computing*, Vol. 14, No. 5, pp.14–22.
- Karnwal, T., Sivakumar, T. and Aghila, G. (2012) 'A comber approach to protect cloud computing against XML DDoS and HTTP DDoS attack', *2012 IEEE Students' Conference on Electrical, Electronics and Computer Science*, pp.1–5.
- Khajeh-Hosseini, A., Sommerville, I. and Sriram, I. (2010) *Research Challenges for Enterprise Cloud Computing*, arXiv preprint arXiv: 1001.3257.
- KPMG International (2011) *The Cloud Changing the Business Ecosystem*.
- Lee, J., Park, M. and Eom, J. (2011) 'Multi-level intrusion detection system and log management in cloud computing', *2011 13th International Conference on Advanced Communication Technology (ICACT)*, No. 1, pp.552–555.
- MacDermott, Á. et al. (2013) 'Protecting critical infrastructure services in the cloud environment', in ACPI, *Proceedings of the 12th European Conference on Information Warfare and Security*, Jyväskylä, Finland: ACL, UK, pp. 336–343.
- Macdermott, Á. et al. (2014) 'Security as a service for a cloud federation', in *The 15th Post Graduate Symposium on the Convergence of Telecommunications, Networking and Broadcasting (PGNet2014)*, pp.77–82.
- Mahmood, Z. and Agrawal, C. (2012) 'Intrusion detection in cloud computing environment using neural network', *International Journal of Research in Computer Engineering and Electronics*, Vol. 1, No. 1, pp.1–4.
- Mell, P. and Grance, T. (2011) *The NIST Definition of Cloud Computing Recommendations of the National Institute of Standards and Technology*, Gaithersburg, MD 20899-8930.
- OTE (2012) *Discussion on the Challenges for the Development of a Context for Secure Cloud computing for Critical Infrastructure IT*, Greece.
- Samani, R. (2013) *Cybercrime Exposed White Paper*, McAfee, UK, [online] <http://www.mcafee.com/uk/resources/white-papers/wp-cybercrime-exposed.pdf>
- Scholler, M., Stiemerling, M. and Ripke, A. (2013) 'Resilient deployment of virtual network functions', in *5th International Congress on Ultra Modern Telecommunications and Control Systems and Workshops (ICUMT)*, Almaty, pp.208–214.
- Steiner, T. (2012) *An Introduction To Securing a Cloud Environment*, SANS Institute, <http://www.sans.org/reading-room/whitepapers/cloud/introduction-securing-cloud-environment-34052>.
- Ten, C.W., Manimaran, G. and Liu, C.C. (2010) 'Cybersecurity for critical infrastructures: attack and defense modeling', *IEEE Transactions on Systems, Man and Cybernetics, Part A: Systems and Humans*, Vol. 40, No. 4, pp.853–865.
- Verba, J. (2008) 'Idaho national laboratory supervisory control and data acquisition intrusion detection system (SCADA IDS)', *International Conference on Technologies for Homeland Security (HST)*, Vol. 208, pp.469–473.

Comment [t1]: Author: Please cite the reference in the text or delete from the list if not required.

Comment [t2]: Author: Please provide the access details (date when the site was accessed/visited).

Comment [t3]: Author: Please provide the access details (date when the site was accessed/visited).