

Security-Oriented Cloud Computing Platform for Critical Infrastructures

M. Mackay^a, T. Baker^b, A. Al-Yasiri^c

^a School of Computing and Mathematical Sciences, Liverpool John Moores University, UK

m.i.mackay@ljmu.ac.uk

^b School of Computing, Mathematics and Digital Technology, Manchester Metropolitan University, UK

t.baker@mmu.ac.uk

^c School of Computing, Science and Engineering, University of Salford, UK

a.al-yasiri@salford.ac.uk

Abstract - *The rise of virtualisation and cloud computing is one of the most significant features of computing in the last 10 years. However, despite its popularity, there are still a number of technical barriers that prevent it from becoming the truly ubiquitous service it has the potential to be. Central to this are the issues of data security and the lack of trust that users have in relying on cloud services to provide the foundation of their IT infrastructure. This is a highly complex issue, which covers multiple inter-related factors such as platform integrity, robust service guarantees, data and network security, and many others that have yet to be overcome in a meaningful way.*

This paper presents a concept for an innovative integrated platform to reinforce the integrity and security of cloud services and we apply this in the context of Critical Infrastructures to identify the core requirements, components and features of this infrastructure.

Keywords – Trust, Cloud Computing, Critical Infrastructures

1. Introduction

The rise of cloud computing is one of the most significant developments in modern computing and this growth in large distributed systems has made it possible to offer processing and storage resources as an on-demand service and on a highly scalable basis. Many companies are now offering extensive public cloud services including Amazon, Google, and Microsoft; while Enterprises (and increasingly those in the SME scope) are looking to deploy private cloud infrastructures based on this model. However, despite its massive popularity, there are still a number of technical barriers that may prevent Cloud Computing from becoming a truly ubiquitous service especially where the customer has strict or complex requirements over the security of the infrastructure. There are many aspects that contribute to this issue but the core areas are the lack of guarantees over the following:

- Data lock-in due to proprietary protocols

- Confidentiality concerns over shared resources
- Networking bottlenecks in and around core datacentres at peak times
- Loss of governance over mission critical data

Many of these issues stem from the fact that, while Cloud Computing had become very popular and even essential for day-to-day business, this area of computing is still relatively immature and so many aspects of the field, including industry-wide standardisation, are still in the process of being developed. The inescapable fact is that until a sufficient level of trust can be associated with moving services to the cloud, customers with high security requirements will shun these services in favour of a more controlled environment. Unfortunately, this covers a wide range of potential users, such as security firms, research and development groups, and ultimately the areas of government and military services but we focus in this paper on Critical Infrastructure providers as they represent a logical first step in this direction.

This paper really represents a first step towards convergence of the fields of Cloud Computing, Security, and Critical Infrastructures as the principles of the former are applied to the later (and vice versa) in an attempt to sufficiently protect assets deployed in the cloud and offer reasonable assurances over the performance and reliability of the service. This paper presents a concept for an innovative new platform to ensure the integrity of cloud services and identifies the core requirements, components, and features of such an infrastructure. We will first provide an analysis of the current approaches to security in clouds, including the emerging field of trusted cloud computing, which may make this a more compelling proposition for Critical Infrastructure providers.

Section 2 of this paper describes the state-of-the-art in the area of trusted platform computing and its

application to critical infrastructures. From there, we investigate the applications of this work to Cloud Computing and explore how such a system could be composed. Section 3 of this paper then presents an overview of the complex issue of data security in cloud platforms, including a discussion of data integrity and encryption issues and user access control and authentication mechanisms. In section 4, we discuss the issue of securing cloud networks and how to actively monitor and protect the internal and external infrastructure. Finally, in section 5 we draw these threads together to present our integrated approach to secured cloud services for critical infrastructures.

2. Trusted Computing for Critical Infrastructures

In this section we will describe the state-of-the-art in the area of trusted platform computing and its application to Critical Infrastructures. We will introduce trusted platforms and discuss the advantages and challenges of each approach, how it can be applied to Cloud Computing, and some of the major approaches that have already been put forward.

2.1 Trusted Computing Platforms

Trusted Computing (TC) is a specialised field of trusted systems whereby a device is made to behave in a consistent, predictable manner enforced through hardware and software techniques such as cryptography and automated authentication [1]. The main functionality of TC is to verify that only authorised code runs on a system and this is an important mechanism to help protect critical or sensitive systems. On the other hand, authentication is a major security addition as it allows clouds to authenticate themselves to known machines and provide a higher level of service to their users. However, TC is controversial because it is possible to restrict or even prevent access to protected resources through the stringent security measures imposed and therefore has significant potential for misuse if compromised [2]. Nevertheless, many major hardware and software vendors are now incorporating trusted computing elements in their offerings and some customers with high security requirements, such as the US Department of Defence, now require this functionality as part of any tender [3].

Trusted Computing encompasses six key technology concepts which are required for a fully trusted system according to the Trusted Computing Group specifications for a Trusted Platform Module (TPM): Endorsement keys, Secure input and output, Memory curtaining / protected execution, Sealed storage, Remote attestation and Trusted Third Party (TTP) support. This range of techniques aims to secure the system resources; provide guarantees over the integrity of running processes, and authenticates the

validity of communication endpoints for all transactions. Moreover, beyond its original prime application of Digital Rights Management, TC has also now found potential use in a range of uses such as distributed firewalls, Distributed Denial of Service Attack (DDoS) prevention, and mobile third party computing. Figure 1 represents a typical trusted platform module.

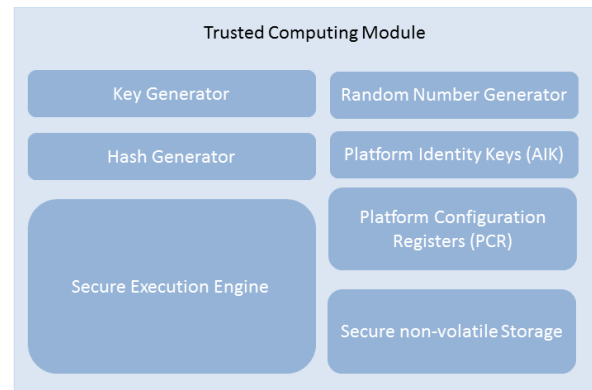


Figure 1 – Trusted Computing Platform

2.2 Critical Infrastructures

Critical Infrastructures (CIs) are a term widely used to distinguish services of a serious/sensitive nature such that they have the potential to cause massive disruption to dependent systems/services if they are compromised or destroyed [4]. The services most commonly associated with this term are:

- Energy distribution (electricity, gas and oil)
- Emergency services (police, fire, hospitals, etc)
- Security services (military, agencies)
- Transportation systems (railway, airports, harbours)
- Public utilities (water, waste, sewage, telecoms)

In the face of increasing threats from terrorism and indirect attacks, many governments and security agencies are looking at how to effectively protect and defend such services. In the UK, the Centre for the Protection of National Infrastructure (CPNI) was established in 2007 to provide this service and works closely with other agencies to enforce the highest level of protection [5]. Typical activities to provide CI protection involve; pre-emptive Analysis and Assessment, and Remediation, Indication and Warning as the event occurs, and Mitigation, Incident Response, and Reconstitution post-event as shown in figure 2.

In recent years, it has been found that an increasingly critical aspect of effectively defending CIs is to adequately protect their ICT infrastructure as an increasing number of attacks are partly or wholly conducted remotely through the Internet. One famous example was the recent cyber-attack on Iran's power generation systems by the highly modified Stuxnet virus [6]. As such, an increasing computing research

trend is in looking at how to effectively use secure mechanisms to protect CIs such as in the recently established PROTECT centre [7].

In this context, many network security mechanisms are being used or developed to enhance protection of CIs beyond the traditional deployment of firewalls and Intrusion Detection Systems (IDS). These include Intrusion Prevention Systems (IPS), network resilience mechanisms, user authentication, Demilitarised Zones (DMZs), system and protocol hardening, trusted systems, and many others [8]. More recently, the rise of Cloud Computing has introduced additional complexity to this picture as providers assess how customers with strong security requirements can be supported.

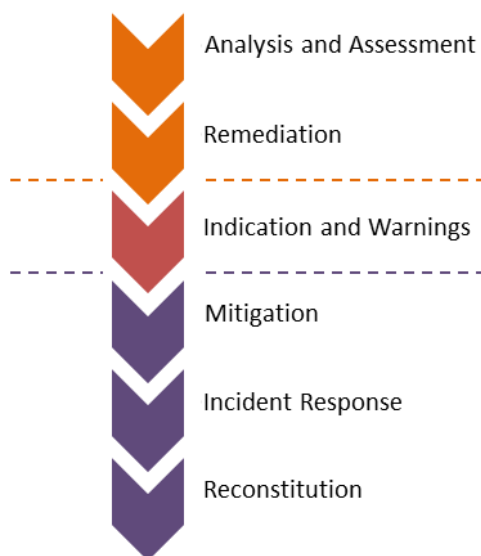


Figure 2 – Typical features of Critical Infrastructure protection

2.3 Trusted Cloud Computing

In the context of CI, the introduction of Trusted Computing elements into cloud computing could potentially provide several distinct advantages to harden the platform. This is still a very immature aspect of secure computing but there have already been a number of works aimed at designing a trusted cloud computing service [9].

One of the first, most natural, areas where trusted cloud computing can be applied is in protecting the underlying infrastructure including the datacentres and interconnection networks. Obviously, the effective deployment of encrypted data storage, memory curtaining, and protected execution environments, perhaps based on some specialised form of the TPM architecture, could make a significant contribution to securing cloud resources and isolating them in virtualised environments [10]. In addition, techniques such as watermarking could be used to protect shared modules and limit, or at least effectively detect, incursions. Finally, this will need to be combined with

secure end-to-end networking technologies and trust-based reputation systems to control access to cloud resources. This final element, reputation systems, are not typically part of trusted computing systems but can be combined with strong identity and access management (IAM) to establish trusted network zones and enforce role-based access control.

This approach has received increasing interest in recent years as a driver to enable the movement of CIs into the cloud. Moreover, this has the advantage of potentially being cheaper to implement, as it can be based on mass market elements, and is deployed and managed entirely by the cloud provider. Ultimately, this could lead to the development of a 'Trusted Layer as a Service' (TLaaS) model whereby the cloud provider offers trusted computing features to customers as an additional service layer and revenue opportunity.

3. Data Security in Cloud Platforms

Having looked at trusted platforms, we now put this in the broader context of security in Cloud Computing and discuss existing mechanisms to ensure data integrity, user access control, and overall cloud robustness. This section will highlight the current challenges in supporting Critical Infrastructures in Cloud Computing or other distributed systems.

3.1 Cloud Data Security

One of the major weaknesses traditionally associated with cloud computing is the lack of provision for data security and the perception that, by moving to the cloud, mission critical data can be exposed to attack. This is dependent, to a certain extent, on the form of cloud computing that is used and by extension the extent to which management is the responsibility of the user or the cloud provider but there are well-proven techniques that can be used to support this. As in traditional systems, the most obvious mechanism to employ is encryption whereby all data stored in the cloud is encrypted using a strong technique such as triple DES or AES to prevent unauthorised access. The trade-off here is the additional performance hit that is incurred during data access and so this is by no means universally employed by cloud providers [11]. A secondary issue is in attackers gaining access to the data through compromising shared resources, i.e. through the virtualised hardware, to snoop on active data flows within the virtual machine. The final issue to consider is the integrity of the data to ensure that resources stored in the cloud are not compromised and/or altered while in storage. Indeed, some cloud providers already implement complex data obfuscation mechanisms for this purpose but it is easy to see how trusted platforms could be employed to further address some or all of these issues [12].

3.2 Authentication, Authorisation, Accounting and User Control

Even if we can provide reasonable assurances over the security and integrity of data stored in the cloud, a second major issue to consider is the security of user access to these resources to prevent unauthorised use. In the absence of strong authentication mechanisms, attackers can gain control over resources in the cloud and potentially use this to compromise other services. Typically, access is provided on an all-or-nothing basis such that, once access is granted to the cloud, the user has the freedom to do a great deal of (un)intentional damage to start and stop virtual servers, or create other chaos inside the cloud environment [13]. This is obviously a more serious issue when it comes to public, hybrid, and community cloud services since private clouds will still reside behind existing corporate network security measures. However, in all cases it is a prime responsibility of the cloud provider to harden the service with appropriate Access Identity Mechanisms (AIM). Of course, virtual machines or services operated in the cloud should still be subject to the same hardening as with traditional systems. Again, a range of commonly employed techniques can be deployed here to enforce access control and authenticated user privileges such as Role-Based Access Control (RBAC) via mark-up languages such as SAML or XACML [14].

In the more extreme case, there is again a strong potential to introduce trust-based systems into cloud environments to ensure that only trusted individuals (and systems) can gain access to cloud resources.

3.3 Robustness

We next consider the issue of data robustness in the context of cloud computing and how distribution and replication strategies can affect the security of critical data. One of the traditional strengths of cloud computing is that data is not necessarily kept in a single, fixed location and that this data and any running service can migrate between cloud resources and be replicated to provide redundancy and resilience in the event of attack or failure. However, this issue is also not straightforward once one considers CIs as there will almost certainly be additional constraints imposed due to legal, managerial, or national security concerns of data storage.

Data replication, for example, is one such area where Critical Infrastructures introduce additional complexity. Of course, it is always a sound strategy to ensure that copies of critical data are securely replicated and stored in alternative physical locations to ensure that attacks/losses are mitigated but, in this case, special care must be taken to address legal and security concerns. For example, it may be necessary to stipulate exactly where data can (or cannot) be replicated to and how many copies are made [15]. Indeed, it may be the

case that the CI disregards automatic replication entirely in favour of their own local back-up system.

Moreover, the location of any CI service run within the cloud is a more sensitive issue in this case and restrictions may apply over what can be run from where. For example, CIs may restrict usable nodes in the cloud to 'friendly' countries, locations where certain legal systems apply, or where certain security assurances can be made that the node is trustable, certainly in the case of public cloud providers [16].

Finally, one lesser issue to consider with replication and migration is the network security aspect. This is a separate issue from securing the actual cloud network (see next section) as the networks over which the data passes may not be under the cloud provider or CI control. As such, special security mechanisms may be needed to secure 'live' data transfers to prevent interception, man-in-the-middle or other forms of network attack.

4. Securing the Cloud Network

The final aspect we consider is to look at network security in the context of cloud computing, both with regard to datacentre interconnection and virtual network integrity and security. Secure network connectivity is a vital aspect of supporting CIs moving into cloud computing as not only must the connection be made secure from attack but special care must be taken to prevent connection failures where reliable, consistent access to resources is mission critical [17]. In this context, we examine existing best-practice in the field of network management and how it can be applied to Cloud Computing, including encryption mechanisms to protect user data, and monitoring and IPS services, to secure the network infrastructure.

4.1 Network security approaches

As in the case of data security, traditional network security mechanisms have a strong role to play in hardening the cloud infrastructure against attack. This section will explore what mechanisms can be deployed to secure cloud computing and support CIs. Securing the network is clearly a critical aspect of any production cloud service and so any public or private provider will be expected to implement firewalls, monitoring, and other standard management mechanisms, to provide a sufficient level of security [18]. Moreover, more security-conscious providers may implement Intrusion Prevention Systems (IPS) that can determine a range of attack characteristics and resilient networking mechanisms to automatically react by triggering remedial measures [19].

The advantage of a cloud provider hosting a range of services is that, as long as the provider implements strong security measures, the customers all benefit from the underlying strength this provides and focus

only on securing their own services. The disadvantage here of course is that if these measures become compromised then all of the customer services are potentially also vulnerable. As such, CI customers may require additional levels of isolation from the rest of the infrastructure to minimise the chance of their networks getting compromised.

4.2 Network data encryption

Another major aspect of network security will be in securing CI data connections both into the cloud and within the cloud itself. As such, encryption protocols at Layer2 and/or Layer3 will be essential to providing an acceptable level of assurance over connection security. Transport Layer Security (TLS) connections are already a standard mechanism for providing a secure connection both into the cloud networks and, one assumes, between data centres. However, due to the potential implementation-specific vulnerabilities of TLS from certificate mismanagement, man-in-the-middle, or XML attacks, other approaches may be necessary [20]. The natural alternatives (or additions) here are IP Security (IPSec) which provides Layer3 security and, Application Layer security protocols such as Secure Shell (SSH) as shown in figure 3. IPSec is a particularly important mechanism as it offers both encryption and authentication, is a core aspect of the upcoming IPv6 protocol, and is used in Virtual Private Networking (VPN) [21].

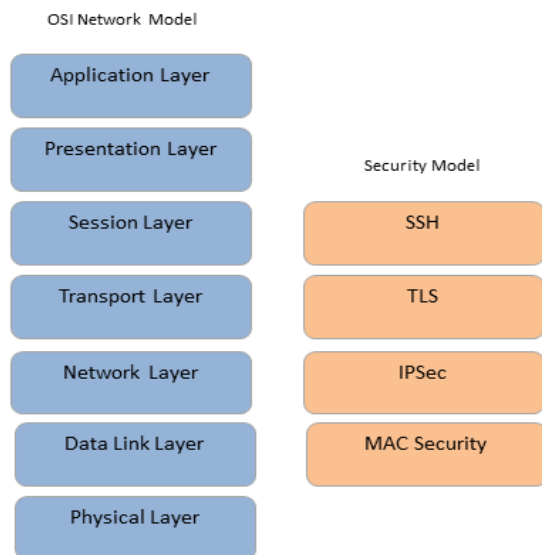


Figure 3 – Layers of secure network protocols

However, it is questionable whether any of these approaches alone will offer sufficient security for CIs to consider putting critical data in public/hybrid clouds and outside of their own private infrastructure unless further assurances can be offered over its integrity. As such, it may be necessary to consider incorporating multiple protocols to offer depth of defence or

provision dedicated PPP links to fully circumvent the public Internet.

4.3 Resilient Networking

With the cloud infrastructure and its communication reasonably secured, the other networking aspect of supporting CIs is to ensure that these secure connections are both consistent and reliable. In the event that CIs move services into the cloud, they will require a constant and dependable level of connectivity as any disruptions to the service can be both costly and seriously affect the CIs performance. There are two aspects to this, a) reinforcing the current best-effort IP routing mechanisms in the Internet with additional redundancy and b) preventing malicious denial of service attacks.

The first case is due to the fundamental properties of IP and the Internet as a best effort routing architecture which, while mostly reliable, offers no guarantee over the end-to-end connection. As such, connections may be dropped or even fail to establish due to technical faults, heavy load on intervening networks, routing errors, or any number of other issues. Clearly this is not sufficient where strict requirements exist over connectivity so a range of techniques can be employed to bolster the reliability of the network. For example, QoS mechanisms can be deployed to ensure that capacity is not exceeded thereby reducing the risk of connections being refused or dropped. Secondly, routing redundancy can be employed to ensure that there is always more than one route from the customer to the cloud thereby guaranteeing connectivity in the face of failures.

In the second case, connectivity to the cloud may be threatened by malicious activity, such as through Denial of Service (DoS) attacks as recently used by the Anonymous group in response to the arrest of Julian Assange in the wikileaks case [22]. DoS attacks attempt to overwhelm cloud provider infrastructure by making many independent requests to a specific point in the network (i.e. a specific web server or router). DoS attacks are often made more complex to diagnose and overcome by being distributed over many sources in the Internet via a botnet or some other mechanism. Research into effective countermeasures to detect and overcome DDoS attacks is still on going and, as evidenced in recent attacks, still not wholly effective.

5. An Integrated Approach to Secured Clouds

Based on our investigation of these converging technologies, we now present our approach to an integrated secure cloud platform that aims to embody all of the above principles. Our first step is to specify the threats and requirements of Critical Infrastructures and discuss how we will address this. From there we

move on to identify the major features and expected functionality of the platform.

5.1 Critical Infrastructure Threats

The threats that CIs face are similar to most corporate networks with the exception here that in most cases there are more strict requirements over the tolerances to faults and attacks. Some of the major threats in this case are therefore as follows:

- Hacking attacks
- DDoS attacks
- Insider attacks
- Equipment failures
- End-to-end issues
- Espionage
- Data loss or corruption

These threats may be innocent or malicious but the fundamental issue is that the CI provider is denied access to its data or services and that its proprietary or confidential data falls into the hands of another party. Many of these issues are common threats in the Internet but extra provisions need to be taken here if CI services are to be exposed in the cloud. With this in mind, we investigate the core requirements that must be met.

5.2 Critical Infrastructure requirements

The two primary concerns for CIs moving data into the cloud will be the security and integrity of the data and maximised service availability. While it is highly unlikely that CI providers will move mission critical services to public clouds, support systems and tertiary services can be more easily provisioned. In this context, the core requirements are listed below:

- **Real time support:** services must continue to provide a high degree of availability in the face of faults and intermittent connectivity. This requires the provider to detail policies and procedures for service backup and recovery. Critical infrastructure requires even more detailed description of the way the data backup is handled by the provider. For example, clients may need to know details of how removable media are managed and destroyed. This is specifically important when the media contains sensitive data.
- **Scalability:** the service must be able to cope with very large volumes of data sources being streamed at a variable rate. This requirement appears to be quite obvious at first. However, a closer look reveals that its implications span over other requirements. For example increasing the data volume may imply that various services and data could be duplicated on the provider's side. This requires guarantees that the data integrity is

maintained at all times. It also requires that appropriate levels of data and service isolation should be maintained.

- **Secure infrastructure:** the cloud platforms must be able to provide reasonable guarantees over the security of the data both when it is stored in the cloud and when it is in transit. This requires details on how data is protected in physical and virtual environments over the cloud. Concerns in this requirement should include
 - i. Definitions of the roles and duties of personnel who can access data and perform data management tasks.
 - ii. Descriptions of the procedures and processes used in data monitoring and logging.
 - iii. The levels of separation and isolation of applications, data and virtual network traffic, which are applied on the cloud.
 - iv. Details on how data is protected in virtual machines. For example details are needed to specify how many functions or services are hosted in one virtual server.
- **High assurance:** the cloud must also demonstrate reliability and very high levels of redundancy and resilience to minimise downtime. This should include details of other how utility programs and multi-tenanted applications are provided and isolated. A malicious activity carried out by one tenant may affect another tenant. This may cause the loss of valuable data or service interruption.
- **Minimal costs:** the cost of transitioning and maintaining the service on the cloud must be reduced in contrast to privately resources.
- **Dynamic provisioning:** the service must be able to react to on-demand performance requirements as the processing requirements increase (or decrease) over time including spikes and flash crowds. This requires descriptions of measures of performance and stability of service provision with respect to levels of user demands. Cloud platforms should also detail how resource usage is controlled and optimised using metering capability appropriate to the type of service provided.
- **Legal assurances:** the customer must be free to specify a fine degree of control over the service location and data replication strategy employed part of the Service Level Agreement (LSA). This requires specifying the level of data protection and privacy in compliance with relevant legislations and regulation including specifying the jurisdiction over the contract terms and data control.

While many of these requirements can be met intrinsically in cloud computing, there are several well-known weaknesses introduced in this approach that

are potentially critical, especially in the case of CIs. Perhaps the most critical of these is the relative lack of security and user authentication in typical cloud platforms and the limited control and monitoring of data replication and service location. Any usage of cloud services in this environment would therefore require a pre-existing SLA to be in place incorporating most of the above elements in addition to the usual conditions on overall performance.

5.3 Platform features and expected functionality

Finally, based on the work described in this paper, we will present the key features of our proposed platform which will support the migration of CIs to the cloud. The platform will be focussed on the provision of extra data integrity and security to minimise the risk of mission critical services being disrupted or taken down by equipment/network failures or attack. This will be focussed around 3 key services, Service Planning, End-to-End security, and Monitoring and Enforcement. These features will be presented as a 'toolbox' that will allow the platform to be adopted and deployed as components by a range of cloud providers as shown in figure 5.

Figure 5 – Key features of the proposed architecture


Service Planning

This aspect of the platform includes the planning and negotiation part of the service before the CI moves into the cloud. Primarily this will involve negotiating the SLA to provide assurances over the customer and cloud provider commitments. While most cloud providers will enforce some form of SLA with its customers, in the case of CIs this will involve extra stipulations and requirements. For example, as discussed earlier in this paper, CIs will have requirements over the location of data and services to trusted parties and enforce the legal and political authority it is served under.

End-to-End Security

The second feature of the platform will be the provision of end-to-end security in the form of data encryption and obscuring mechanisms and through the implementation of trusted computing aspects. This will be enforced, both while the data is stored on the server and communicated over the network, and will include a range of encryption techniques and authentication mechanisms to maximise the security of the platform. Moreover, the platform will include aspects of trusted computing such as memory curtaining, remote attestation and trusted third parties to provide assurances over the integrity of the service in a secure, virtualised TPM system.

Monitoring and Enforcement



The final feature of the platform is the monitoring and enforcement aspect, which will guard against attack and ensure that the stipulations of the SLA are met. The monitoring infrastructure will need to be distributed to provide assurances to both the customer and cloud provider that the SLA is being enforced and this can also be used to measure the effectiveness of the platform. This could for example include some form of accounting/auditing model to track the usage of applications and services in the cloud and model their usage. Moreover, an Intrusion Detection System (IDS) will be deployed to secure the cloud against attacks and this will be supplemented with resilient networking services to counter any recognised attack patterns.

6. Conclusions

There is a clear need for more demonstrably secure cloud platforms to drive the adoption of these services among Critical Infrastructure providers. This must include mechanisms to secure cloud services, the end-to-end networking interconnecting the users, and the cloud infrastructure itself. We have also shown that these issues are now starting to be tackled based on existing enterprise-strength technologies and that this approach can at least provide a reasonable level of assurance over the security of the platform.

Moreover, a more recent trend is in this area is to look towards more powerful cutting-edge mechanisms, such as trusted computing and resilient networking techniques, to increase the level of security offered and bring Cloud Computing towards a more demanding customer base such as Critical Infrastructure providers and government.

In this paper we have identified the key aspects of this development and highlighted the requirements that must be met before this can proceed. In this way, we can demonstrate how Cloud Computing and the end-to-end networking can reasonably be made secure enough to support Critical Infrastructure providers. Further, we have proposed an open architecture for CI

support in clouds and identified the key elements of a security 'toolbox' that cloud providers can implement and deploy to simplify this process.

This paper represents our first step in this area and extensive further work will be necessary to validate our approach against CI provider expectations, develop, integrate and evaluate the identified functionality, and demonstrate the strengths of the platform. Nevertheless, the authors are optimistic that the advantages of our approach will be self-evident and gain traction both within the academic and wider research community.

T. Baker (t.baker@mmu.ac.uk) (Corresponding Author) School of Computing, Mathematics, and Digital Technology, Manchester Metropolitan University, UK;
M. Mackay (m.i.mackay@ljamu.ac.uk) School of Computing and Mathematical Sciences, Liverpool John Moores University, UK; **A. Al-Yasiri** (a.alyasiri@salford.ac.uk) School of Computing, science and Engineering, University of Salford, UK.

References

- [1] The Trusted Computing Group, "Trusted Platform Module specification version 1.2", http://www.trustedcomputinggroup.org/resources/tpm_main_specification, March 2011.
- [2] R. Stallman, "Can you Trust your Computer", at <http://www.gnu.org/philosophy/can-you-trust.html>, last accessed 09/09/2011.
- [3] US Department of Defence, "Encryption of Sensitive Unclassified Data at Rest on Mobile Computing Devices and Removable Storage Media", <http://iase.disa.mil/policy-guidance/dod-dar-tpm-decree07-03-07.pdf>, Jul 2007.
- [4] J. Moteff, P. Parfomak, "Critical Infrastructure and Key Assets: Definition and Identification", Order Code RL32631, Congressional Research Service, Library of Congress, October 2004.
- [5] Centre for the Protection of National Infrastructure Homepage, <http://www.cpni.gov.uk/Templates/CPNI/pages/Default.aspx>, Last accessed 09/09/2011.
- [6] Stuxnet worm attacks Iran nuclear plant, <http://www.bbc.co.uk/news/world-middle-east-11414483>, September 2010, Last accessed 09/09/2011.
- [7] LJMU PROTECT Research Centre, <http://www.cms.livjm.ac.uk/pri/>, Last accessed 09/09/2011.
- [8] K. Zunnurhain, S. Vrbsky, "Security Attacks in Solutions in Clouds", 2nd IEEE Conference on Cloud Computing Technology and Science (CloudCom), Poster, November 2010.
- [9] Z. Shen, Q. Tong, "The Security of Cloud Computing System enabled by Trusted Computing Technology", 2nd International Conference on Signal Processing Systems (ICSPS), Vol2 pp11-15, July 2010.
- [10] F. Lombardi, R. Di Pietro, "Secure virtualization for cloud computing", Journal of Network and Computer Applications, Vol34, Iss4 pp1113-1122, July 2011.
- [11] R. Chow, P. Golle, M. Jakobsson, E. Shi, J. Staddon, R. Masuoka, J. Molina. "Controlling data in the cloud: outsourcing computation without outsourcing control", Proc of ACM workshop on Cloud computing security (CCSW '09), pp85-90, November 2009.
- [12] S. Ghemawat, H. Gobioff, S. Leung, "The Google file system", 19th ACM symposium on Operating systems principles (SOSP '03), Vol37, pp29-43 ACM, October 2003.
- [13] Cloud Security Alliance, "Guidance for Identity & Access Management v2.1", <https://cloudsecurityalliance.org/guidance/csaguide-dom12-v2.10.pdf>, April 2010.
- [14] D. A. Haidar, N. Cuppens-Boulahia, F. Cuppens, and H. Debar. "An extended RBAC profile of XACML", Proc of the 3rd ACM workshop on Secure Web Services (SWS '06). ACM, pp13-22, 2006.
- [15] D. J. Abadi, "Data management in the cloud: Limitations and opportunities", IEEE Data Eng. Bulletin, 32(1), March 2009.
- [16] M. P. Eisenhauer, "Privacy and Security Law Issues in Off-shore Outsourcing Transactions", Hunton & Williams - Atlanta Georgia, February 2005.
- [17] P. Schoo, et al, "Challenges for Cloud Networking Security", 2nd ICST Conference on Mobile Networks and Management (MONAMI 10), September 2010.
- [18] W. Jansen, T. Grance, "Guidelines on Security and Privacy in Public Cloud Computing", Draft Special Publication 800-144, National Institute of Standards and Technology (NIST) US Department of Commerce, January 2011.
- [19] K. Vieira, A. Schulter, C. B. Westphall, C. M. Westphall, "Intrusion Detection for Grid and Cloud Computing," IEEE IT Professional, vol.12, no.4, pp.38-43, July-Aug. 2010.
- [20] M. Rex, S. Santesson, "Transport Layer Security (TLS) Secure Renegotiation", IETF Internet Draft, Expires June 2010.
- [21] B. Kang, M. Balitanas, "Vulnerabilities of VPN using IPSec and Defensive Measures", International Journal of Advanced Science and Technology, Volume 8, July, 2009.
- [22] A. Pras, et al, "Attacks by 'Anonymous' WikiLeaks Proponents not Anonymous", CIIT Technical Report, December 2010, <http://www.ctit.utwente.nl/news/archive/2010/dec10/Attacks%20-anonymous.docx/> (last accessed 09/09/11).