Edith Cowan University

Research Online

2014

Cloud Security meets Telemedicine

Michael N. Johnstone
Edith Cowan University

Cloud Security meets Telemedicine

Michael N Johnstone

Edith Cowan University Security Research Institute, Perth, Australia

Abstract

Medical systems are potentially one domain where security is seen as an impediment to patient care and not as an essential part of a system. This is an issue for safety-critical systems where reliability and trust are essential for successful operation. Cloud computing services offer a seamless means to allow medical data to be transferred from patient to medical specialist, whilst maintaining security requirements. This paper uses a case study to investigate the use of cloud computing in a mobile application to assist with diagnostics for patients with Parkinson Disease. It was found that the developers of the app ignored security requirements and standards, preferring to focus on functionality.

Keywords: Information Systems Security; Cloud Computing; Telemedicine; Applications Development

1 Introduction

Confidentiality, Integrity and Availability are considered to be the core building blocks of information security. It would be expected, therefore, that software systems would be developed as a matter of course with these building blocks in mind as software is ubiquitous (cf. the Internet of Şthings Ť). This is especially important in safety-critical domains such as medical systems. Unfortunately most software, according to [1], is insecure. Johnstone [2] notes that this is due to the tension between functional requirements (which are visible to customers-those who commission and fund systems) and security requirements (which often are not visible). A somewhat darker view is held by [3], who suggest that security requirements are often omitted from requirements specifications altogether.

Medical systems appear especially problematic as, given that their primary focus is patient care, security is either assumed or ignored [4]. Several well-reported cases, such as Stanford HospitalŠs loss of 20,000 ER patient records [5] and an Australian pathology laboratoryŠs loss of patient data ([6], highlight the embarrassment and loss of trust that occurs when medical data is leaked (a breach of confidentiality). Clearly the nature of the data and its intended use must determine which of the core tenets of information security would be applicable. For a real-time heart-rate monitor in an operating theatre, both Integrity and Availability would be critical, Confidentiality less so. In an on-line web-based patient record input system, Confidentiality and Integrity would be dominant, with Availability being perhaps not as important.

According to a recent IBIS report [7] health and allied systems are Şpoised to become AustraliaŠs biggest industry division and employer well before 2050...In this division, superfast broadband will be vital in driving healthcare costs down by faster diagnostics, preventive health systems [and] partial self-diagnostic services...Ť

Software engineering as a discipline is still maturing, so it is not unreasonable that software development that focusses on security concerns is still in its infancy. There is certainly evidence of an evolution from object-orientation in the 1980s, component-based software engineering in the 1990s, service-oriented architectures in the 2000s to cloud computing now. Given that cloud computing in its most basic form provides a façade for data storage and retrieval, it can provide seamless access to data which could make data management simpler and thus potentially improve information technology security management, especially as both patients and medical practitioners make increased use of wireless

transmission of data and Internet-based applications.

This paper describes the issues involved with medical systems and the concomitant standards that apply to the development and use of such systems, explains the theory behind cloud computing and how this may benefit medical systems, uses a case study to illustrate the effectiveness of cloud-based data storage and retrieval for medical data, specifically an iPad app which provides diagnostic tests for patients and neurologists, and finally considers some security weaknesses of cloud computing relevant to medical systems.

2 Methods

The research question being examined is ŞTo what extent are security concerns addressed in the development of cloud-based medical applications?Ť. Galliers [8] provides a useful taxonomy which enables the correct selection of a research method (or approach, to use Galliers term). According to Table 1, the object of interest is methodology as the focus is on how system developers prioritise and implement security requirements in the domain of medical systems.

On the spectrum of modes, those based on interpretivism were rejected on the basis that the researcher was not directly involved in the development (except to provide guidance when requested). This immediately discounted modes such as action research, which require the researcher to be actively immersed in the research. Therefore several observational research modes present as candidates, viz: field experiment, case study and survey. Given that field experiments require control of a limited set of variables and that surveys can only report what is said, not what was actually done, case study was deemed to be the most appropriate research approach for this work.

The researcher was able to gather project artefacts from the system development (which included project management documents, product documentation and versions of the product itself). The researcher had full access to the development team and the client and thus was able to make extensive field notes of meetings.

3 Results

Before describing the results of the study, it is worthwhile outlining relevant medical standards and presenting an overview of cloud computing which will contextualise the results, given that the system being studied is a medical application running on a tablet device and transmitting data via cloud services.

3.1 Security Standards pertinent to Medical Systems

Mizukura et al. [9] proposed a home health care network based on the IEEE 11073 standard. ISO/IEEE 11073 is actually a family of health informatics standards, for example, 11073-10407 specifies the behaviour of blood pressure monitors. Mizukura et al. field-tested a health monitoring application that was designed to capture heath data from elderly patients in their own environment and transmit such data across a network to relevant medical practitioners.

Significant progress has been made on issues to do with precisely how to transfer medical data. For example, ISO/IEEE 11073-20601 [10, p1] defines an abstract model of personal health data as well as the appropriate transport independent transfer grammar required to set up logical connections between systems. Such standards are being implemented by manufacturers of telemedical equipment [11].

ISO 27799 [12] recognises the problem and states ŞThe need for effective IT security management in healthcare is made all the more urgent by the increasing use of wireless and Internet technologies in healthcare delivery. If not implemented properly, these complex technologies will increase the risks to the confidentiality, integrity and availability of health information.Ť. ISO 27799 provides guidance about what sort of health data needs to be protected, but like many standards, is descriptive, rather than prescriptive. For example, it declares that personal health information (such as that collected by the iPad app described later in this section) needs to be protected, but does not specify the precise means of protection that would meet the standard.

HL7 V2 is an OSI level 7 (hence the name) ANSI standard protocol for communication between health service providers in Australia. It offers security checks, participant identification, availability checks, negotiating exchange mechanism negotiation and provides a standard data structure. Whilst HL7 messages are text-based (and thus perhaps lend themselves to encoding in XML and transmission via SOAP), the HL7 protocol does support the transfer of picture data via Base64 encoding of the binary data stream. Whilst there is a clearly defined structure or ontology for HL7 messages, there is no innate formatting or defined dependency between the data in the observation (OBX) segments of a results message.

HL7 supports the encoding of patient identifiers, provider identifiers and observations (medical test results), whilst at the lowest level, recognises message start and end identifiers as well as individual unique message identifiers (to ensure that collisions do not oc-

Object	Modes for traditional empirical approaches (observations)							Modes for newer approaches (interpretations)			
	Theorem proof	Laboratory experiment	Field experiment	Case study	Survey	Forecasting	Simulation	Game/role playing	Subjective/ argumentative	Description/ interpretive	Action research
Society	No	No	Possibly	Possibly	Yes	Yes	Possibly	Yes	Yes	Yes	Possibly
Organizational group	No	Possibly (small groups)	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Individual	No	Yes	Yes	Possibly	Possibly	Possibly		Yes	Yes	Yes	Possibly
Technology	Yes	Yes	Yes	No	Possibly	Yes	Yes	Yes	Yes	Possibly	No
Methodology	No	No	Yes	Yes	Yes	No	Yes	Yes	Yes	Yes	Yes
Theory Building	No	No	No	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Theory Testing	Yes	Yes	Yes	Possibly	Possibly	No	Possibly		No	Possibly	Possibly
Theory Extension	Possibly	Possibly	Possibly	Possibly	Possibly	No	No		No	Possibly	Possibly

Table 1: A Taxonomy of IS Research Approaches (adapted from [8]).

cur) and acknowledgments (although it is more correct to say that an HL7-compliant application recognises the message identifiers and processes them accordingly).

Having examined some relevant standards and protocols for the transmission and storage of medical data, it is now appropriate to discuss how those data could be stored using cloud-based services.

3.2 An Overview of Cloud Computing

As mentioned previously, cloud computing represents an evolution in the provision of software services, rather than a revolution. However, as with any ŞnewŤ concept, there is sometimes confusion as to what it actually represents and what benefits might accrue from the use of such technology. In this section cloud computing is defined, the architecture of a cloud-based system explained and various models of cloud computing are discussed.

Conventionally, cloud computing appears to be focussed on large-scale storage of information across multiple servers. NIST [13, p2] provide a succinct definition of cloud computing that encompasses more than just distributed storage: ŞĚa model for enabling ubiquitous, convenient, on-demand network access to a shared pool of configurable computing resources (e.g., networks, servers, storage, applications, and services) that can be rapidly provisioned and released with minimal management effort or service provider interactionŤ.

Badger et al. [14] claim that cloud computing has essential characteristics that differentiate it from earlier models of distributed computing: on-demand self-service, broad network access, resource pooling, rapid elasticity and measured service.

Badger et al. also suggest a range of service and deployment models that provide coverage of the cloud computing landscape. The service models include not only the familiar software as a service (SaaS), but also platform as a service (PaaS) and Infrastructure as a service (IaaS). PaaS encompasses software platforms (such as .NET), database engines and operating systems. IaaS provides CPUs, virtualisation (if required) and block storage. An example architecture is shown in figure 1. Clearly, one of the main advantages of a multi-layered architecture is the ability to fine-tune resource pooling in the middle layers to effect a change in performance without the service user being aware of the change (apart from the observed performance boost). Whilst resource pooling is usually a benefit as it provides redundancy, it will be shown in a later section that there are security implications with the complex architectures that are
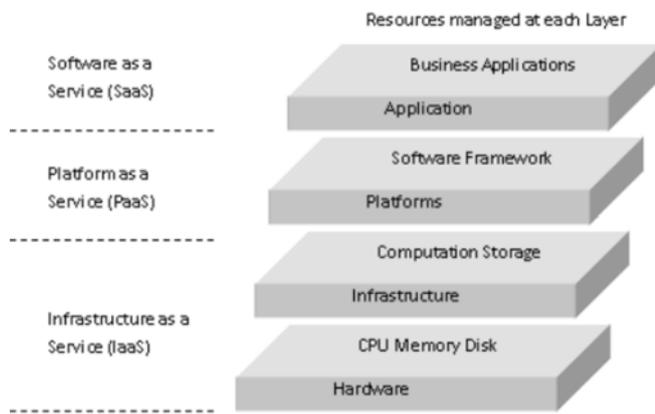
Figure 1: Cloud Computing Architecture (adapted from [15]).

Figure 2: Sample Test from the Parkinson Disease Tester App.

used to deliver cloud services.

The deployment models proposed by Badger et al. are private, community, public and hybrid clouds. Clearly, a private cloud exists for the use of one consumer (business) exclusively. The cloud may be used by many business units within the same enterprise but the service provision may, in fact, be outsourced to a third party (which is likely and therefore the infrastructure is also likely to be remote from the consumer). A community cloud is similar except that the consumer in this case is a group of interested parties that are not from the same enterprise. The service may be managed by one of the parties in the community or by a third party. A hybrid cloud, as the name implies, can use a combination of any of the three aforementioned deployment models. The models remain distinctive but are linked by standards or proprietary systems that permit data and/or application portability.

Having defined cloud computing and discussed various cloud service and deployment models in general, what follows is a case study which uses cloud-based services to share medical data between interested parties.

3.3 Case Study: An iPad App to assist with the management of Parkinson Disease

The brief for this system was to provide a proof-of-concept iPad application (app) that allowed patients with Parkinson Disease to perform several tests which provide diagnostic information and allow the test results to be shared with a neurologist, hence facilitating management of the disease. The major benefit is that patients need not travel to see their neurologist to perform the tests. Given that a large proportion of patients are in the 70-79 age group [16], the ability for the system to link to a neurologist and transfer data seamlessly was a prime requirement. By conforming to AppleŠs Human Interface Guidelines, it was expected that this require-
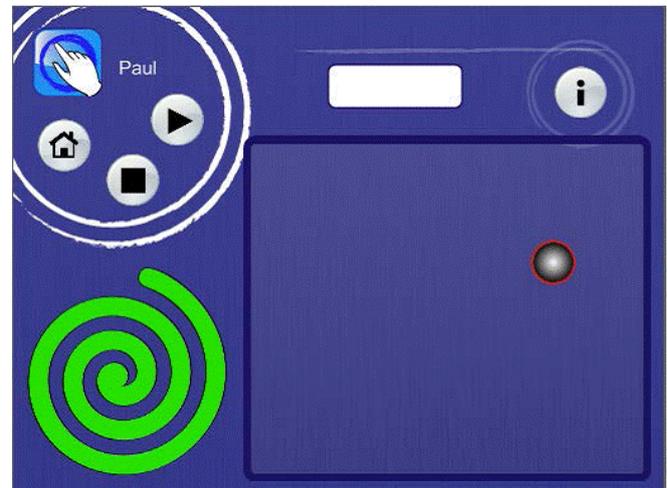
ment could be met, provided that the cloud services could be implemented for an iPad. Figure 2 and Figure 3 show some sample screen shots of the app.

The basic requirements were to provide an app that allows a patient to perform two diagnostic tests, allow one or more of those tests to be saved and stored locally, to provide summary statistics and relevant graphical feedback to a patient so that s/he may track his/her progress and (critically for this discussion on information security) allow the sharing of patient data between one or more parties (usually the patientŠs neurologist) in a seamless and transparent way.

Cloud computing was deemed to be a potential solution to the last requirement. The sharing requirement had to allow a patient to send a share message to another iPad user, for the second user to respond in the affirmative (or negative) to allow the transfer of patient data and for the first user to have the ability to rescind the original sharing request and thus break the connection. Several cloud providers were investigated including iCloud, Google, Nuvolabase and Moai. iCloud was an obvious first choice as the target device was an iPad, but this service is meant to be a personal cloud service used across many devices. It is not designed for sharing files with multiple users. The other cloud services were evaluated, the result being that the Moai cloud, despite being targeted at the gaming community, met all of the functional requirements and thus was selected as the cloud service for this application.

The app works by storing local data using SQLite (chosen because it does not require a database server), the data is stored in the cloud as a JSON (JavaScript Object Notation) file and accessed using RESTful services (common to most Web applications). JSON is a text-based standard (see RFC 4627) for defining and
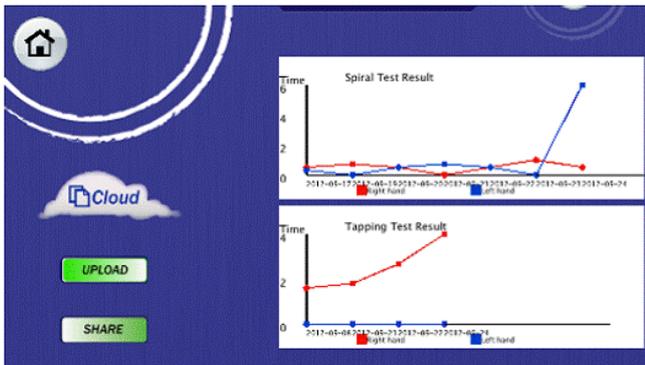
Figure 3: Sample Output from the Parkinson Disease Tester App.

sending structured data between a Web application and a server. JSON provides the usual data types (number, string, Boolean, array and object). A JSON structure for the iPad app test data object is:

```
"test": [
        {
"type": "spiral",
"number": "20",
    "date": ''20121009",
    "value-array", [4,6,5,3,3,7]
        },
        {
"type": "countdown",
"number": "6"
    "date": ''20121009",
    "value": ''13"
        }
    ]
```

In this structure there are two types of test and whilst they share some common data elements such as a test number and a test date, the actual results are different. JSON is able to characterise these different structures easily and thus is a good choice for representing the ontology of a test.

4 Discussion

This section evaluates cloud computing in terms of the tenets of information security, addresses issues specific to the case study and concludes by considering some practical (security) and legal barriers to the acceptance and use of cloud computing services for medical data.

Considering cloud computing in its simplest (conventional) incarnation that of distributed file storage and retrieval, there are several aspects of confidentiality, integrity and availability that are worth discussing. If patient data is stored unencrypted on a cloud-based

file system this would appear, at first glance, to be a breach of confidentiality as the file is stored in plain text (plain text in this context does not necessarily refer to ASCII text, but to any non-encrypted form of data, for example human-readable XML records, Microsoft Word documents or the data referred to in the case study). Confidentiality is maintained by two means. First, the user does not know which physical location stores the data and second, the data may be split into several parts across several locations. The semblance of a single file is maintained by the cloud façade as part of SaaS (recall figure 1). Integrity appears problematic by virtue of the benefits which assure confidentiality, that is, the separation of the file into multiple parts across multiple locations. Provided that the PaaS layer is intact, the marshalling of the file from its parts into a whole is transparent to an end-user of the cloud. Availability is, of course, handled by the IaaS layer.

This describes the scenario where all of the components of cloud computing work seamlessly to provide the services expected of them. From a security perspective, it is worth examining how standard attacks on confidentiality, integrity and availability might affect the provision of cloud services. A standard attack on availability is denial-of-service (DoS). Figure 4 indicates an alarming trend. DoS attacks are increasing, not in complexity, but in their size. This means that a DoS attack on a cloud service provider will almost certainly result in a loss of availability. The wider problem is that the outcome of a DoS attack may affect integrity if a file is partially constructed. There may be the opportunity for the data to be modified or for data to be leaked (a breach of confidentiality) because of a failure in the other service provision layers.

Turning now to the specifics of the case study, a well-trodden mitigation pathway for problems of confidentiality is encryption. Certainly the data being transferred from an iPad to the Moai cloud could be encrypted before transmission from a patientŠs iPad and storage and decrypted on retrieval on a neurologistŠs iPad. Whether a public key infrastructure or private keys are chosen is perhaps not an issue as long as the key length prohibits the data being compromised during its effective lifetime, notwithstanding the key transmission safety issues inherent in the sharing of private keys. Integrity issues are often dealt with by the use of cyclic redundancy checks or hashing. Both techniques are feasible with the data being transferred from the iPad to the cloud. It requires that the PaaS layer be capable of forming and sending a re-transmission request if data were found to be corrupt. Availability issues appear the most insoluble in this scenario because of the ease by which DoS attacks can be mounted. It is possible that IPv6, with its significantly

larger address space (as compared to IPv4) may provide a successful mitigation strategy.

Rather than using JSON as the messaging format, HL7 could provide a better alternative. One possible mapping of an observation or OBX record (diagnostic test result) equivalent to the aforementioned JSON representation is:

```
OBX-2 (Value type) NM
        // a number
OBX-3 (Observation ID) 6^Countdown^LN
        // the type of test
OBX-5 (numeric) 13
        // the actual value
OBX-6 (units) s^Seconds^ISO+
        // units of the value
OBX-14 (date/time of observation)
        20121009+1000
```

On its own, the use of HL7 over JSON does not appear to provide significant benefits. In terms of message transfer between a patient and a neurologist the overheads for HL7 are greater but there are some security advantages as mentioned in a previous section. Remembering that the app is designed to share data in a one-to-many relationship, the real benefit to using HL7 is realised when several health care providers wish to share data about the same patient. In this scenario, using a common protocol designed for health data makes the translation and interpretation of the data relatively straightforward.

The benefit of the cloud in terms of hiding the physical location/structure of the file becomes problematic when confidentiality is breached at the lower levels of the cloud architecture (Figure 1). This is largely because, especially in a hybrid deployment, multiple stakeholders across multiple domains may share physical data space. In contrast to a more conventional model of data storage where a stakeholder has access to contiguous space, in the shared (cloud) space, parts of files may be juxtaposed with data from other stakeholders. This leads to questions about effective access controls and authentication mechanisms at the higher levels of the cloud architecture. The implementation of such controls and mechanisms is non-trivial.

Problems surrounding confidentiality, integrity and availability can be solved in a physical sense, but there remains some interesting legal issues that need to be solved before medical data can be securely stored on a cloud-based service. These issues relate to the location of cloud data (or, in fact, the transmission of any such data across national boundaries) and the extent to which the liability of a cloud provider is limited. Given that
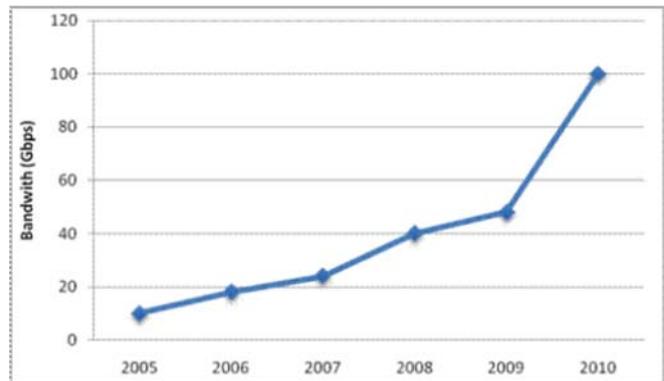

Figure 4: Largest Single Denial of Service Attack [17, p5].

a user of a cloud service does not know the physical location of the data, the law regarding who can access the data and in what circumstances becomes murky. For example, if data are transferred via network links inside the USA, the Patriot Act can be invoked which allows US Federal authorities to capture the data. Interestingly, Amazon will not guarantee that its cloud service will not transfer data via the USA.

Also, cloud providers such as Google and Amazon have specific no-liability clauses in their contracts. As Calloway [18] points out, whilst limited liability clauses are not new, their use in cloud service agreements is problematic. The value placed on medical data suggests that such contracts be examined carefully and the extent to which the cloud provider has reduced liability fully articulated before such services are employed for safety-critical data. This problem is reasonably well-known as Snooks [19, p4], with respect to the Australian Commonwealth Privacy Act notes that ŞAgencies engaging cloud service providers need to take appropriate contractual measures to ensure personal information is protected, regardless of whether or not the provider (and any subcontractors) are based in Australia or overseas.Ť. Whilst the advice given by Snooks is well-founded, e.g. contractually prohibiting a provider from transferring data outside Australia, precisely how this obligation would be or could be enforced is unclear.

Security is recognised as an issue for cloud computing. Lo et al. [20] for example, propose a cooperative form of an intrusion detection system as a means of detecting attacks (particularly DoS attacks) on cloud-based systems. Simply, they place an intrusion detection system in each cloud and allow the systems to transmit data about attacks between them, thus allowing the other cloud(s) to have a priori knowledge of a potential attack. Neisse et al. [21], in discussing trust in cloud infrastructure, point out that encryption is not a perfect answer as, given that malicious or negligent infrastructure providers have full control of the

cloud infrastructure, they are able to give themselves unrestricted access to the data. Neisse et al. provide a system, called BonaFides, which remotely monitors and attests to the integrity of crucial system files, i.e. it provides assurance of claims, although it it does not guarantee data integrity. The former is a problem with systems such as DropBox. DropBox does encrypt data as they are uploaded, which implies a level of security; however, the system is searchable, which means that the infrastructure provider must have the keys to decrypt the data. This problem could, of course, be solved by encrypting the data prior to upload.

As pointed out preciously, cloud services are not just about data storage (disk space) but can also include CPU and bandwidth services. The nature of cloud computing means that conventional security procedures and mechanisms may not be appropriate or effective. As an example consider the physical decommissioning and destruction of hard disks containing sensitive information, as mandated by several standards. ISO 27799, (2008, p30) states that ŞIn addition to following the guidance given by ISO/IEC 27002, organizations [sic] processing health information applications shall securely overwrite or else destroy all media containing health information application software or personal health information when the media are no longer required for use.Ť. This proves somewhat difficult to action when the disks are not owned by the client, but by an infrastructure provider and impossible when the disks on which the data are stored are not co-located, especially when space on the same disks may be allocated to another client, who would, no doubt, be concerned about the potential destruction of their data.

In summary, the difficulties in guaranteeing cloud service security were discussed and a case study which highlighted some of the problems likely to be encountered in the hosting of medical data was examined. Legal issues to do with distributed hosting of data or services were also discussed.

5   Conclusion

This study explored the problems of using a nascent technology, cloud computing, to store medical data. The complex nature of cloud services was revealed and a case study that described the implementation of an iPad app that transferred medical data was articulated and discussed.

Specifically, this study used a case study to show how a medical data could be generated, stored and shared using cloud computing. It was argued that cloud services provided benefits in that the cloud façade hid the complexity of data transfer and storage from the end-user, as compared to conventional database or file-based storage techniques. It was shown that the cloud could inadvertently be responsible for security breaches under certain circumstances. In terms of the research question posed earlier, the case study provides evidence that functionality takes precedence over security requirements.

A limitation of this work is that it used only a single case study with a specific cloud platform, therefore it would be unwise to conclude that all cloud platforms or deployment models suffer from identical security problems. Further work would involve extending this idea to see how well other cloud providers dealt with the security issues outlined in this paper. Other avenues of research to be explored are DoS-resistant network protocols (to address the problem of availability) and homomorphic encryption of databases (to address issues of confidentiality and integrity). When these issues are solved, cloud computing will be an cost-effective and efficient vehicle for the transport, storage and retrieval of medical data.

Acknowledgements

References

1. Shostack A, Stewart A. The New School of Information Security. Upper Saddle River, NJ: Addison Wesley; 2008.

2. Johnstone MN. Security Requirements Engineering-The Reluctant Oxymoron. Proceedings of the 7th Australian Information Security Management Conference, Edith Cowan University, Perth Western Australia, 1st-3rd December 2009.

3. Wysopal C, Nelson L, Dai Zovi,D, Dustin E. The Art of Software Security Testing. Upper Saddle River, NJ: Addison Wesley; 2007.

4. Williams T. When trust defies common security sense. Health Informatics Journal. 2008; 14(3), 211-221. Sage Publications London.

5. Moisse K. Stanford Hospital Patient Records Leaked Online. 2011. Retrieved from http://abcnews.go.com/blogs/health/2011/09/09/stanford-hospital-patient-records-leaked-online/

6. Caldwell A, Earley D. Patients' medical records leaked online by pathology lab Sullivan Nicolaides. 2009. Retrieved from http://www.news.com.au/technology/patients-medical-records-leaked/story-e6frfro0-1225699562788

7. IBISWorld. A Snapshot of AustraliaŠs Digital Future to 2050. 2012. IBIS Publishing.

8. Galliers RD. Choosing Appropriate Information Systems Research Approaches: A Revised Taxonomy. Proceedings of The Information Systems Research Arena of the 90s: Challenges, Perceptions and Alternative Approaches, Copenhagen, Denmark. 1990.

9. Mizukura I, Tamura T, Kimura Y, Yu W. New Application of IEEE 11073 to Home Health Care. The Open Medical Informatics Journal. 2009;3: 44-53.

10. ISO/IEEE 11073-20601. Health informatics Ů Personal health device communication Ů Part 20601: Application profile Ů Optimized exchange protocol. New York, NY: Institute of Electrical and Electronics Engineers, Inc. 2010.

11. Biotronik. BIOTRONIK Home Monitoring EHR DataSync Documentation of the BIOTRONIK IEEE 11073-10103 XML Structure: Technical information for software developers and system architects. Berlin, Germany: BIOTRONIK SE & Co. KG. 2011.

12. ISO 27799. Health informatics Ů Information security management in health using ISO/IEC 27002. Geneva, Switzerland: International Organisation for Standardisation. 2008.

13. Mell P, Grance T. The NIST Definition of Cloud Computing.. NIST Special Publication 800-145. 2011.

14. Badger L, Grance T, Patt-Corner R, Voas J. Cloud Computing Synopsis and Recommendations. NIST Special Publication 800-146. 2012.

15. Zhang Q, Cheng L, Boutaba R. Cloud computing: state-of-the-art and research challenges. Journal of Internet Services and Applications. 2010; 1(1): 7-18.

16. Brown I. Does caffeine protect against ParkinsonŠs disease? A preliminary study, Nutrition & Food Science. 2002;32(6): 227-30.

17. Arbor Networks. Worldwide Infrastructure Security Report. 2010; VI.

18. Calloway TJ. Cloud Computing, Clickwrap Agreements, and Limitation on Liability Clauses: A Perfect Storm? Duke Law & Technology Review 2012; 11:163-174.

19. Snooks A. Negotiating the cloud Ű legal issues in cloud computing agreements. Australian Government Information Management Office. 2012.

20. Lo C-C, Huang C-C, Ku J. A Cooperative Intrusion Detection System Framework for Cloud Computing Networks. In proceeding of: 39th International Conference on Parallel Processing, ICPP Workshops 2010, San Diego, California, USA, 13-16 September 2010; 280-284.

21. Neisse R., Holling D, Pretschner A. Implementing Trust in Cloud Infrastructures. Proceedings of the 2011 11th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing (CCGRID '11). 2011; 524-533.

Correspondence

Mike Johnstone
Edith Cowan University Security Research Institute, Perth, Australia
m.johnstone@ecu.edu.au