UNIVERSITETET I AGDER
FACULTY OF ENGINEERING AND SCIENCE

Reflection Report on Summer School on Cloud Security, Metochi, August 22-29, 2015

Mohamed Abomhara

Department of Information and Communication Technology,
University of Agder, Norway;
Mohamed.abomhara@uia.no

Abstract

In this report, a summary of lectures given at the COINS summer school is presented and the way they can be useful to my PhD research is discussed. Five topics related to cloud security were presented at this summer school session, namely software-defined networks (SDN) security, virtual machine introspection, end-to-end defense against kernel rootkits in the cloud environment, cloud application security, privacy in the cloud environment. These topics will be addressed in different sections separately.

1 Software-Defined Networks (SDN) security

In this lecture, Sandra Scott-Hayward explained the fundamentals of SDN security. The presentation was divided into a morning and evening session. The presenter gave the definition of SDN and pointed out what

COINS Summer School, 1–12.

the OpenFlow protocol is, the implementation and challenges of SDN, attacks and vulnerabilities in SDN as well as solutions to SDN security issues.

I am not very familiar with this research field. However, the idea of SDN and the OpenFlow protocol is quite interesting. Below, some common security issues that we attempt to mitigate and that relate to confidentiality, integrity, availability and non-repudiation are addressed.

1.1 Introduction: What is SDN?

Software-Defined Networks (SDN) are a new technology designed to make networks more agile and flexible. Today's networks are often more static, slow to change and dedicated to single services. With SDN, it is possible to create a network that handles many different services dynamically, allowing providers and customers to consolidate multiple services onto one common infrastructure.

1.1.1 SDN Architecture

SDN is comprised of three different layers as shown in figure 1. The bottom layer (infrastructure) contains the network equipment, such as router switches and wireless access points. The second layer (control) is responsible for configuring the infrastructure layer, and it does so by receiving service requests from the third layer (application). The application layer is where the cloud applications, management applications and business applications place their demand for the network onto the control layer.

SDN is designed to be open, so it would be possible to have multiple different vendors' equipment in the infrastructure layer, multiple vendors' control components in the control layer, and multiple vendors' applications in the application layer.

1.1.2 SDN Evolution

The concept of SDN is not new and has undergone many developments, including Ipsilon GSMP in 1996, the separation of control introduced by the IETFForCE group in 2000, Secure Architecture for the Networked Enterprise (SANE) in 2006 and the OpenFlow protocol in 2008. This presentation mainly focuses in the OpenFlow protocol.

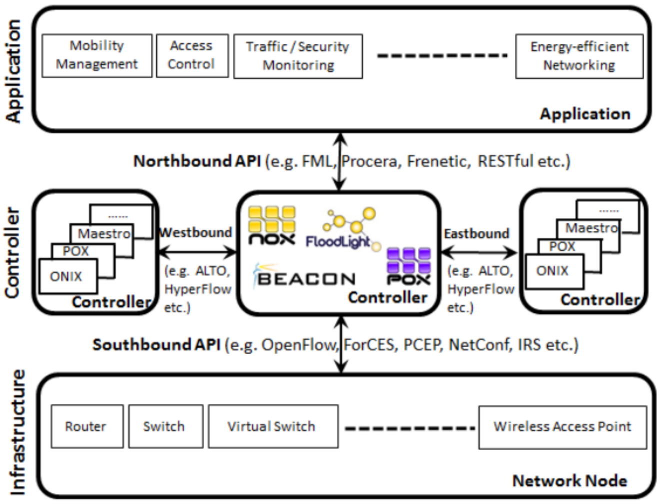

Figure 1 IoT model: key concepts and interactions

1.2 What is OpenFlow?

OpenFlow is one of the new standards for SDN. It is a protocol used to control the forwarding behavior of switches from multiple vendors in a software-defined network. It provides a way to control the behavior of switches throughout the network dynamically and programmatically.

The main idea is to introduce a model with layers (Figure 2) of programmable hardware, operating systems and applications to the network industry. Each network device will have an open API, making the network more programmable. Once there is an API, another layer could be added above, called the control layer (software layer). The control layer communicates downwards to the API layer that in turn conveys all instructions in real time down to the network. Subsequently, an application can be built at the top in the application layer.

1.3 SDN implementation and challenges

In establishing a more secure and ready SDN, there are many challenges to overcome to achieve robust security, efficiency and compliance. The challenges include:

- Performance vs flexibility

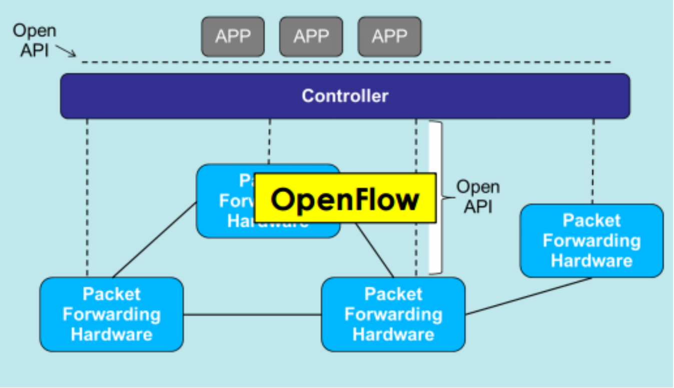
Open API
APP
APP
APP
Controller
OpenFlow
Open API
Packet Forwarding Hardware
Packet Forwarding Hardware
Packet Forwarding Hardware
Packet Forwarding Hardware

Figure 2 IoT model: key concepts and interactions

- Scalability
- Security (authentication, authorization and attack prevention)
- Interoperability

More about these challenges is available in the cited work [1]

1.4 Attacks and vulnerability

The speaker mentioned different attacks including confidentiality, integrity, availability, authentication and non-repudiation. Specific focus is directed toward:

1. Unauthorized access
2. Data leakage
3. Data modification
4. Denial-of-Service (DoS)
5. Configuration issues.

1.5 Solutions to SDN security issues

Sandra Scott-Hayward presented a number of proposed solutions for security problems in SDN. She explained how threat modeling serves as a foundation for the analysis and specification of security requirements, and how threat modeling involves understanding system complexity and identifying all possible threats to the system. The identified threats should be further analyzed based on their criticality and likelihood, and decisions are made as to whether to mitigate the threats or accept the associated risks.

1.6 Demo

During the lecture, Sandra presented a demo on:

- How packets move through an OpenFlow network
- Packets in flight using Mininet network emulation and Wireshark.
- SDN control plane attacks.
- Operation checkpoint and appregister.
- Controller security rating
- Packet walkthrough of an example HTTP request and reply in an OpenFlow enabled network.

1.7 Questions and Answers

Question: What is the Trusted Domain Controller in SANE?

Answer: A trusted domain control is a domain that is trusted by the system to authenticate users or services. For example, if an application is authenticated by a DC, then this authentication is accepted by all domains that trust the authenticating domain.

Question: As shown in the figure, the DC is located within a trusted boundary. Does it mean that it is owned by a trusted party?

Answer: Yes, it is assumed that the DC is owned by a trusted organization.

Offline question (during break)

Question: What is SDN a solution for?

Answer: SDN is a solution for:

- High demand on resources.
- Unpredictable traffic patterns.
- Rapid network reconfiguration.
- Incorporating business rules.

2 Everything You Need to Know about VMI

In this presentation, Zhiqiang Lin briefly introduced virtual machine introspection. The main aim of this lecture was to provide students with the necessary knowledge of VMI. It started with the basic concept and continued with the principles behind VMI and the enabled applications. The presentation consisted of five sessions. In the first session, Zhiqiang presented a roadmap of the lecture and introduced virtual machine introspection: What is VM? Why do we need VM? (reasons for virtualization), definitions of hypervisor, security monitoring system and the advantages and disadvantages of VM inspection. The different types of hypervisors were also discussed: bare metal, hosted, native and emulation hypervisors.

In the second session, Zhiqiang talked about the semantic gap challenge, what can be observed from hypervisors (for example, what can be observed from native hypervisors and emulation-based hypervisors), and what is needed from hypervisors. The session was concluded by pointing out what can be observed and what hypervisors are useful for, as well as the fact that the gap must be bridged in order to provide effective monitoring services.

In the third session, the speaker explained approaches to bridge the semantic gap. We also explored what other researchers have developed in the past. Many researches have proposed various approaches to address the problems with the semantic gap. Seven different approaches were discussed and compared in terms of speed, performance, flexibility and practicality, etc. The Zoom-in Binary Code Analysis Approach was recommended by the speaker for the cloud develops.

In the fourth session, Zhiqiang discussed possible VMI applications. The talk was summarized in one sentence: anything can be done with VMI. Zhiqiang gave examples of applications that were classified into security and non-security applications.

1. Security applications: a hypervisor offers the capability to implement any type of detection, prevention and recovery system desired as well as forensic analysis.
2. Non-security applications: e.g. virtual machine management, high performance computing and autonomous computing.

In the final session (hands-on lab), the speaker showed a demo of how this tool is used to detect and prevent activities. He demonstrated how to use the kernel debugging tool (Red Hat crash utility) to inspect kernel states and also how to use the volatility tool to perform memory introspection. The lecture was concluded by presenting future directions on hypervisor developments.

This topic is not relevant to my PhD research, but it is new knowledge about VMI together with how virtual machine introspection-based architecture can be used for intrusion detection in cloud security.

2.1 Questions and Answers

Question: Could you repeat the definition of VMI?

Answer: It is a softlayer program used to check what is going on in the hardware and software layers. With VMI, you could do multiplexing, isolation and migration.

3 End-to-End Defense against Kernel Rootkits in a Cloud Environment

In this section, the presenter (Sachin Shetty) mainly introduced some basic knowledge about creating rootkits and how a malice code can be created. First, Sachin briefly introduced kernel rootkits, what rootkits are, how to hijack a system call using rootkits, why we should consider rootkits, how rootkits can provide an attacker with camouflaged access as well as how to defend against this kind of attack.

Kernel rootkits and cloud computing were also explained along with how kernel rootkits achieve various goals, especially concealing certain malicious processes from security monitoring, antivirus software, intrusion detection and VMI (virtual machine introspection). The goals by modifying certain kernel data structures and codes. Moreover, Sachin talked about how to defend against kernel rootkits in cloud Infrastructure as a Service (IaaS), Platform as a Service (PaaS), and Software as a Service (SaaS).

In this part of the presentation, the threat model and assumptions for the kernel-level rootkits detection system were reviewed. The threat model was developed based on an assumption where an attacker aims to change one of the system call table functions using a rootkit. The attacker must figure out where the system call table is stored. An attacker can install rootkits in OS kernel managing VMs by exploiting zero-day vulnerabilities in the kernel, applying software in VMs and gaining control over multiple VMs to steal confidential data, modify memory, etc.

The speaker also presented the design, implementation and evaluation of the RootkitDet system, which is an end-to-end defense against kernel-level rootkits. It is also efficient and practical in the cloud environment. More about RootkitDet is available in the cited work [2].

4 Cloud application security

Daniel Hedin talked about cloud computing and its key enabling technology. The talk was initiated with the question: What is cloud computing? He went on to introduce aspects of the cloud (SaaS, IaaS, the internet, utility storage, etc.), leading to the question of what cloud apps are. The most frequent key words appearing in descriptions of the cloud and cloud apps are availability, collaboration and simplicity.

There are many cloud apps, among which cloud computing and cloud storage are two important ones. The presenter provided examples of cloud apps, including Vivino, Dropbox and Asana. He additionally explained the differences between cloud apps and web apps. The focus of this lecture was on web apps.

4.1 What are web apps?

Daniel defined web apps as apps delivered to a client over http as a web 'page' (html, css and JavaScript). The provider runs the app backend that provides core functionality like login, sessions, storage, etc. The app communicates with the provider and other resources in the cloud via AJAX for instance (to send data, receive data, update the 'page' dynamically) and the provider may use cloud resources to realize the app backend (for storage, authentication or other services to increase the product appeal).

4.2 Security in the cloud

Daniel mentioned there are numerous security challenges in cloud applications, which are related to data confidentiality, data integrity, data availability and non-repudiation. The emphasis in this lecture was on confidentiality. Daniel spoke about how to ensure that user information imparted to the applications is safe. He elaborated on user data confidentiality, what happens when a user enters sensitive data into the system (e.g. user login data) as well as how to guarantee that sensitive data are only sent to target addresses and not stolen or compromised.

Moreover, information flow and the two types of information flow were addressed: explicit flow and implicit flow.

4.3 Demo

During the lecture, Daniel showed a demo on:

- Content injection.
- 3rd party code injection.
- Cross Site Scripting (XSS).
- Accidental data leaks
- There was also an exercise on information flow control for the audience.

The talk concluded with Daniel's personal views. He claimed that access control alone is not enough to solve the problem of information flow. He suggested that access policies should be defined in terms of what information is allowed to flow where. Moreover, information flows during computing should be analyzed and flows that violate policies must be disallowed.

This topic is a popular research area and relates to my PhD research. Because I am working on access control for cloud-based health information, Daniel's suggestions about defining the right policies of access control (what information is allowed to flow where) and analyzing information flow will be highly considered in my work.

5 Privacy in the cloud environment

In this seminar, Fatema Kaniz presented privacy in the cloud environment. Privacy was first defined, followed by a talk on global laws of privacy and data security, the differences between privacy and security, privacy breach payoff, and the latest privacy challenges in the cloud environment. This topic is related to my PhD research and I have talked with Fatema about access control policies and means of verifying and testing policies.

5.1 Global data privacy regulations

Fatema gave examples of privacy laws, including the Privacy Act of 1974, the EU Data Protection, the EU Data Privacy Directive, HIPPA and C6 Canada, among others. She stated that by January 2015, the number of countries with data privacy laws increased to 10%, signifying that data privacy is becoming a very intense and important topic worldwide. In addition, she talked about privacy policy principles. Fatema conveyed that according to Australian Law Professor Graham Greenleaf, there are ten common elements to four international privacy instruments: the OECD Guidelines, the Council of Europe Convention, the EU Data Protection Directive, and the APEC Privacy Framework. The elements include data collection, data quality, purpose of collection (at time of collection), notice of purpose and rights at the time of collection, use limitations (including disclosure), security through reasonable safeguards, openness to personal data practices, access (individual rights of access), correction (individual rights of correction) and accountability (data controllers accountable for implementation). A comparison of global data privacy regulations was presented via PowerPoint.

5.2 What is personal data?

Personal data is defined as any data that can be used to identify an individual (either from the data or from the data in conjunction with other information within). Privacy can be described as the state of individuals establishing for themselves when, to whom, and to what degree discretely recognized data concerning them is communicated or used by other individuals or organizations. Not all individual information should be treated the same. On one hand, non-sensitive information can be presented publicly with no concern, such as public information, names, and public phone books. However, in most situations, sensitive data must not be disclosed or read by unauthorized entities. Any information, if lost, compromised, or disclosed without authorization and that can be used to cause harm to an individual or organization is considered sensitive. For instance patient and private business data such as trade secrets must be hidden from unauthorized entities. Privacy in information systems encompasses information access, collection and storage.

5.3 Privacy vs Security

Fatema demonstrated the difference between privacy and security based on Borkings work. She showed various privacy criteria (e.g. reporting of processing, data quality conservation and so) and compared them with information security goals (confidentiality, availability and integrity). For more about the comparison please refer to slide number 19 [3].

5.4 Privacy breach payoff

Privacy breaches may impact individuals and/or organizations in different ways. A number of negative consequences can occur due to the lack of, or poor privacy protection. Examples include harm to the person whose data are used or disclosed inappropriately, damage to an organizations reputation, financial loss of individuals and/or organizations, loss of business due to negative publicity and many others.

5.5 Does the cloud introduce new privacy challenges?

The speaker presented some of the many privacy challenges with the cloud environment as follows:

- Data can be stored in multiple locations and shared among multiple providers, while multiple providers can have different terms of service and policies.
- Information disclosure by insiders due to the lack of organizational control over employees.
- Service providers do not typically have control over the physical location of data.
- Conflicting laws from different jurisdictions.

Moreover, Fatema shared the possible methods of minimizing risk to cloud consumers by highlighting that consumers should be aware of what data they are sharing (avoid oversharing personal information), and with whom they are sharing personal data. Consumers should also be mindful of where their data is going and how it will be processed and shared.

The mechanisms for privacy in the cloud were presented as well, including access control (authentication and authorization), legal compliance and location control. Three types of access control models were briefly explained: identity-based access control, role-based access control (RBAC) and attribute-based access control (ABAC). The advantages and disadvantages of each model type were presented along with a talk on the privacy-preserving authentication technique and how it facilitates anonymity, unlinkability and minimum disclosure. In addition, Fatema spoke about the policy-based authorization system and adding privacy protection to this system. She offered examples of policy systems such as the Privacy-Preserving Advanced Authorization System (P-PAAS) infrastructure, which actually entails her PhD research.

Fatema concluded with proposing a future research direction. She hopes that in future there will be more work on the trustworthy monitoring mechanism for data access and sharing, compliance verification techniques, identifying privacy metrics, trustworthy ways of verifying privacy metrics and technical means of satisfying the right of data portability.

5.6 Questions and Answers

Question: slide number 19 (Privacy VS Security), I think the information security goals are missing non-repudiation? Because, for privacy criterion (reporting of processing), if you want to report for ever process then auditing and logging is required?

Answer: I am not quite sure, because this work is not mine. I just adapted to show the difference between privacy and security. but more can be found in the original paper [4].

References

[1] S. Sezer, S. Scott-Hayward, P.-K. Chouhan, B. Fraser, D. Lake, J. Finnegan, N. Viljoen, M. Miller, and N. Rao, "Are we ready for sdn? implementation challenges for software-defined networks," *Communications Magazine, IEEE*, vol. 51, no. 7, pp. 36–43, 2013.

[2] L. Zhang, S. Shetty, P. Liu, and J. Jing, "Rootkitdet: Practical end-to-end defense against kernel rootkits in a cloud environment," in *Computer Security-ESORICS 2014*. Springer, 2014, pp. 475–493.

[3] F. Kaniz, "Privacy in the cloud environment," *COINS summer school 2015*, 2015. [Online]. Available: https://coinsrs.no/wp-content/uploads/2015/08/COINS-lecture-Kaniz-2015.pdf

[4] J. J. Borking, "Why adopting privacy enhancing technologies (pets) takes so much time," in *Computers, Privacy and Data Protection: an Element of Choice*. Springer, 2011, pp. 309–341.