ORIENTAL JOURNAL OF
COMPUTER SCIENCE & TECHNOLOGY
An International Open Free Access, Peer Reviewed Research Journal
Published By: **Oriental Scientific Publishing Co., India.**
www.computerscijournal.org

ISSN: 0974-6471
December 2015,
Vol. 8, No. (3):
Pgs. 222-227

High Level Security Trust Zones for Private Clouds

R. VINOTH, B. GURUPRASATH and B. N. KARTHIK

A.V.C College of Engineering,Mannampandal,Mayiladuthurai, 609 305, India.

(Received: September 15, 2015; Accepted: December 10, 2015)

ABSTRACT

The vulnerability of Cloud Computing Systems (CCSs) to Advanced Persistent Threats (APTs) is significant. So a cloud architecture reference model that incorporates a wide range of security controls and best practices, and a cloud security assessment model – Cloud-Trust – that estimates high level security metrics to quantify the degree of confidentiality and integrity offered by a CCS or cloud service provider (CSP) is used. Cloud-Trust is used to assess the security level of four multi-tenant IaaS cloud architectures equipped with alternative cloud security controls and to show the probability of CCS penetration (high value data compromise) is high if a minimal set of security controls are implemented. CCS penetration probability drops substantially if a cloud defense in depth security architecture is adopted that protects virtual machine (VM) images at rest, strengthens CSP and cloud tenant system administrator access controls, and which employs other network security controls to minimize cloud network surveillance and discovery of live VMs.

Key words: Cloud Computing, VM, CSP, IaaS, CCSs, APTs.

INTRODUCTION

Cloud computing enables a new business model that supports on-demand, pay-for-use, and economies-of-scale IT services over the Internet. The Internet cloud works as a service factory built around virtualized data centers.1 Cloud platforms are dynamically built through virtualization with provisioned hardware, software, networks, and datasets. The idea is to migrate desktop computing to a service-oriented platform using virtual server clusters at data centers.

However, a lack of trust between cloud users and providers has hindered the universal acceptance of clouds as outsourced computing services. To promote multi tenancy, we must design the cloud ecosystem to be secure, trustworthy, and dependable.2 In reality, trust is a social problem, not a purely technical issue. However, technology can enhance trust, justice, reputation, credibility, and assurance in Internet applications. To increase the adoption of Web and cloud services, *cloud service providers* (CSPs) must first establish trust and security to alleviate the worries of a large number of users.

A healthy cloud ecosystem should be free from abuses, violence, cheating, hacking, viruses, rumors, pornography, spam, and privacy and copyright violations. Both public and private clouds demand “trusted zones” for data, *virtual machines* (VMs), and user identity, as VMware and EMC3 originally introduced.

Virtualization, the basis for most CCSs, enables CSPs to start, stop, move, and restart computing workloads on demand. VMs run on computing hardware that may be shared by cloud tenants. This enables flexibility and elasticity, but introduces security concerns. The security status of a CCS depends on many factors, including security applications running on the system, the hypervisor (HV) and associated protection measures, the design patterns used to isolate the control plane from cloud tenants, the level of protection provided by the CSP to cloud tenant user data and VM images, as well as other factors.

Cloud-Trust can assess the relative level of security offered by alternative CSPs or cloud architectures. Cloud tenants can use it to make decisions on which CSP security options or cloud security features to implement. Cloud-Trust is based on CCS unique attack paths that cover the essential elements of an IaaS cloud architecture. It is based on a Bayesian network model of the CCS, the class of APT attack paths spanning the CCS attack space, and the APT attack steps required to implement each attack path. It provides two key high-level security metrics to summarize CCS security status quantitatively:

- Probability an APT can access high value data
- Probability the APT is detected by cloud tenant or

CCS security monitoring systems

Trust Zones in Cloud Services

Trust zone (TZ) as a combination of network segmentation and identity and access management (IAM) controls. These define physical, logical, or virtual boundaries around network resources. Cloud TZs can be implemented using physical devices, virtually using virtual firewall and switching applications, or using both physical and virtual appliances.IAM systems use usernames, passwords, and access control lists (ACLs), and may use Active Directory Domain Controllers[1], Federated Trusts[2], and multifactor authentication mechanisms using time limited codes or X.509 certificates. IAM servers can also use hardware information to make access decisions. For example, devices without a pre-validated MAC address can be prevented from joining a network. Routers using ACLs and IP address white listing can prevent an unauthorized device from accessing network resources. These are examples of hardware based TZ enforcement.

The security of TZ implementations depend on correctly configuring domain controllers, firewalls, routers, and switches that are used in segmenting and restricting access to portions of the cloud network and on "locking down" secure communications between users and domain controllers to prevent SOAP interface or signature wrapping attacks[3]. Misconfiguration of IAM servers,

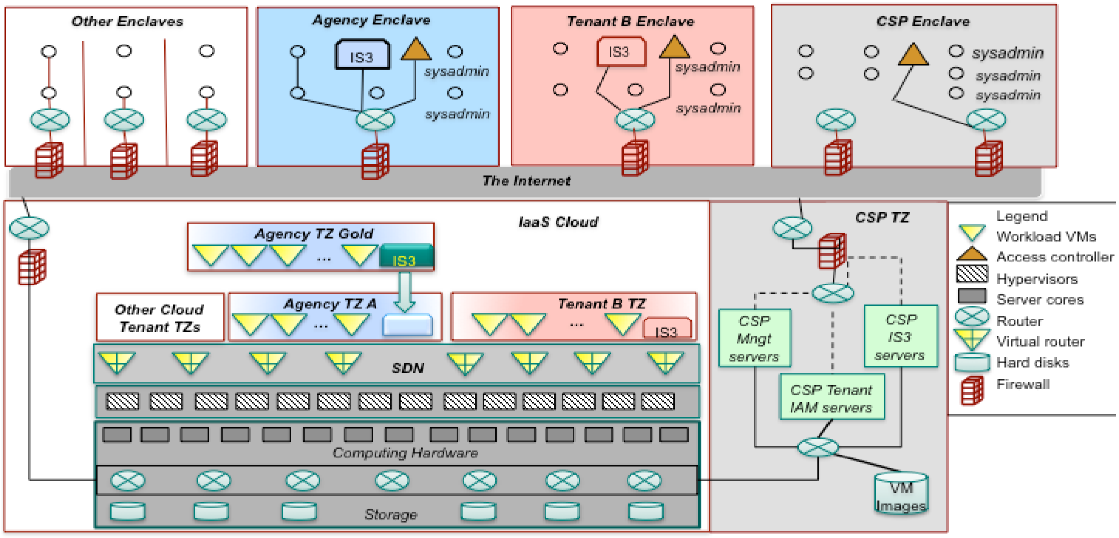

Fig.1: CCS Reference Model

domain controllers and other network devices can introduce vulnerabilities in the cloud network and let attackers enter restricted TZs. Careful configuration management is a key factor that must be taken into account in assessing cloud security status.

CCS Reference Model and Architecture

CCS depends on both CSP and tenants. The CSP TZ is segregated from tenant TZs and contains cloud management servers, SDN controller servers, CSP tenant IAM servers, and CSP Information System Security System (IS3) servers. CSP sys-admins communicate with CSP management systems through a separate firewall and Internet port to isolate CSP communications traffic. It is a best practice to isolate CSP management and monitoring systems from cloud tenant VMs, as illustrated in Fig. 1 [4].This cloud reference model is based on this best practice and design tenets developed by the Defense Information Systems Agency (DISA) for securing enterprise networks[5].

The cloud system that detects and prevents the actions of malware and bad actors is called as the Information System Security System (IS3). IS3 systems can generate lots of data and have high false alarm rates. A cloud IS3 includes IDSs, host based security systems, fire -walls, IAM servers, reverse proxy web servers, syslog servers, and SIEM servers

The SIEM aggregates event data produced by security devices, network infrastructures, systems and applications. Event data is combined with contextual information about users, assets, threats and vulnerabilities. The data is normalized, so events, data and contextual information from disparate sources can be correlated and analyzed for specific purposes, such as network security event monitoring, user activity monitoring and compliance reporting. Fig. 1 shows the location of IS3 servers used by the CSP, the Agency, and other tenants. We assume tenants provide their own IS3s to monitor and manage their TZs.

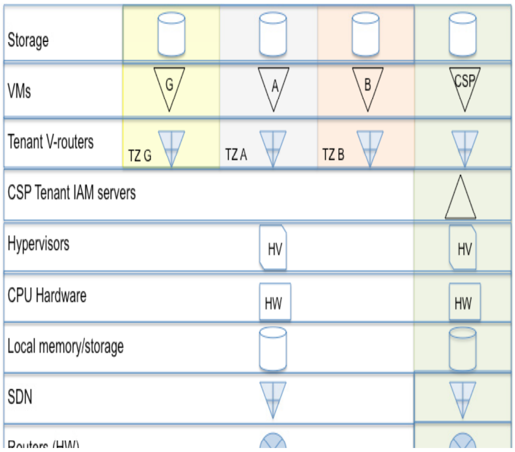

Fig. 2: CCS Node Classes

System protection and risk reduction involve numerous actions not performed directly on the CCS. These include physical protection measures, vetting employees, security awareness training, maintaining a vulnerability management data base, and participating in national vulnerability organizations and fora (e.g., SANS). We do not include employee training or vetting activities in Cloud-Trust, but note they are important for securing CCSs and CSPs.

CCS Node Classes

The abstracted view of an IaaS CCS is shown in Fig.3. It is the starting point for Cloud-Trust, and is based on the types of nodes in a CCS. These are labeled node classes, because many individual nodes of each type or class will be present in the CCS. To simplify the analysis assume all nodes in each node class are identical in terms of their security properties .Therefore,it is not essential to distinguish between individual elements in each node class, and it is possible to define a Bayesian network model in which the nodes of the network are CCS node classes, and not individual system components of the CCS. This Bayesian network model forms the basis of Cloud-Trust.

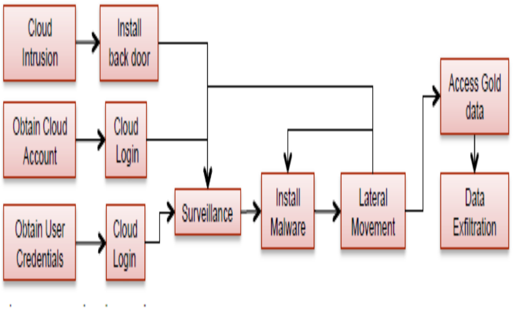

Fig. 3: Attack Stages

The columns in Fig. 2 indicate the TZs node classes belong to. The types of nodes classes are indicated in the first column. Node classes reflect the segregation of CSP and tenant network paths. The CCS architecture shown in Fig. 2 also has the feature that VM traffic within a TZ can be confined in that zone and segregated if all intra-TZ message traffic is routed by the V routers. This functionality is consistent with SDN or virtual networking capabilities provided by leading HV vendors and CSPs.

The attacker's objective is assumed to be the data store in TZ Gold in the upper left hand corner . The APT will have to traverse the network of node class objects from bottom to top to gain such access if the attack starts from outside the cloud. Using such node class diagrams, a cyber attack against an IaaS cloud can be represented by a directed graph of edges and nodes

The types node classes included in the node class diagram depend on the specifics of the cloud architecture examined. To find the set of edges that represent technically feasible cyber attacks investigate specific CCS vulnerabilities identified in the literature. These are used to develop a set of attack paths that span the set of all feasible paths through the CCS infrastructure to the APT target.

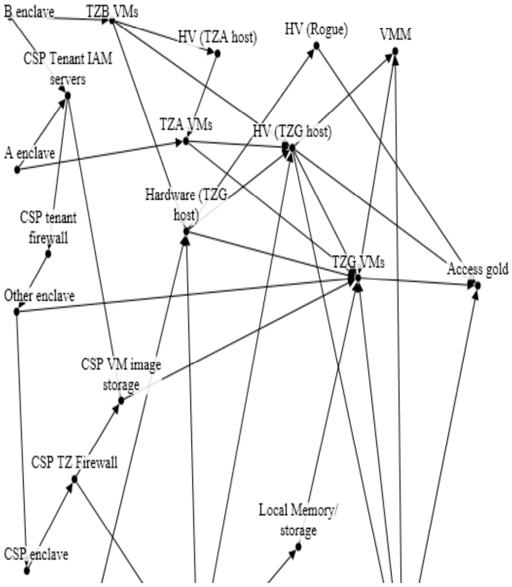

Fig. 4: IaaS CCS Infiltration Bayesian Sub-Network

CCS Attack Paths

CCS attacks can be divided into outsider or insider attacks. Outsiders can gain access to the cloud using three attack paths.

The first exploits weaknesses in cloud access control mechanisms. Such weaknesses may exist in firewalls or IAM servers used by the CSP or cloud tenants.

The second starts by stealing valid credentials of a cloud user at some location outside the cloud .

The third outsider attack path starts with the attacker using valid credentials and prior legitimate access to the cloud.

Insider attack paths start inside the cloud when the attacker already exploits credentials for at least one cloud TZ, for example the CSP TZ. Agency VMs operating in the same TZ run on the same physical machines and HVs.

The attack paths are defined in two variants. The first called a "Stuxnet" variant where the APT requires little or no command and control (C2) by the external human attacker. In this case the APT has the surveillance information it needs to conduct all stages of the attack, or capabilities needed to independently do surveillance. The second attack variant is one where the APT has much less capability and information about the CCS environment. In this case assume it must communicate with an external control authority and be updated with new capabilities during the attack.

Nested Virtualization

A nested virtualization attack[6] uses an additional unauthorized HV to access sensitive data and credentials. The additional HV could be inserted either between the normal HV and the physical hardware, or between a guest OS and the normal HV. In the former case, the additional HV will provide an attack surface that spans all of the VMs on the original HV. In the later case, the additional HV could be confined to a specific guest OS.

The target for the attack is a VM running in TZ G or is a VM image with stored TZ G credentials that is at rest. Finding the VM image at rest, or finding the physical machine that the target VM is or will be spun up would be accomplished by surveillance of Agency VM operations. Either target is likely to begin with the attacker gaining access to the CSP management enclave in order to perform sufficient surveillance. An insider working for the CSP can do the surveillance.

Attacking the VM image and inserting the unauthorized HV provides the advantage that the operation can be performed before the image is loaded into the CSP infrastructure

Targeting the image at rest, the attacker would 'wrap' it with an additional HV . Targeting the physical machine would require that the attacker either be able to reboot the machine and cause it to load the attacker's HV first, and then load the CSP's HV, or implement a 'blue pill' rerouting of a live HV without rebooting[6].

Once an attacker has successfully nested a HV at either layer, one of the main advantages, in addition to gaining access to memory and other sensitive resources, is that the rest of the stack would function 'normally'. The guest VMs continue to run on virtualized infrastructure, and the original HV thinks it is running on CSP hardware.

Once the attacker has succeeded in injecting a HV that it controls, it has gained a stealthy point of access to sensitive VM data and credentials. However, unless the attack is completely autonomous, it may require additional surveillance and C2 activities. The HV may therefore have to beacon to another node to complete the attack.

Nested virtualization attacks exploit the fact that both the intended hosts and guests might not have mechanisms available to verify the other parties. The guests are supposed to run on a virtualized platform and may not be able to detect that they are not running directly on a CSP sanctioned HV. Similarly, both the CSP HV and the CSP hardware provide interfaces that do not discriminate between consumers of their resources. Absent specific restrictions, an additional attacker controlled HV could be a consumer that is as accepted as a guest OS, or CSP controlled HV.

Bayesian Network Model

It is possible to apply Bayesian network statistics to the attack paths described above. Attack paths have been used to understand the vulnerability status of information systems[7]. They have also been used to develop probabilistic measures of enterprise network security[8,9]. Extend this approach to CCSs by constructing an acyclic directed graph using the attack paths defined above[10]. We apply these attack paths to the CCS node classes defined in Fig. 2. The resulting directed graph is shown in Fig.4

Cloud-Trust relies on conditional probabilities that represent the probability that a vulnerability in an individual CCS component can be exploited by an APT, if other CCS components have already been compromised

These conditional probabilities correspond to the directed edges shown in Fig. 4. This approach enables us to factor in the contributions that specific CCS security features can have in reducing the vulnerabilities of nodes in the CCS and which then can contribute to a reduction in the overall security profile of an IaaS cloud. This model of CCS architectures includes the security features and controls the CSP provides, what the CSP permits the customer or cloud tenant to provide, and what the cloud tenant actually provides.

The complete security model consists of two Bayesian sub-networks: an infiltration sub-network and an exfiltration sub-network. Only the Cloud-Trust infiltration sub-network is shown in Fig. 5. The infiltration subnetwork characterizes the probability that an APT will be able to access the gold data, while the exfiltration network characterizes the likelihood that the APT can exfiltrate the accessed gold data.

CONCLUSION

Thus it is demonstrated how Cloud-Trust can be used to assess the security status of IaaS CCSs and IaaS CSP service offerings, and how it is used to compute probabilities of APT infiltration (high value data access) and probabilities of APT detection. These quantify two key security metrics: IaaS CCS confidentiality and integrity. Cloud- Trust also produces quantitative assessments of the value and contribution of specific CCS security controls and can be used to conduct sensitivity analyses of the incremental value of adding specific security controls to an IaaS CCS, when there is uncertainty regarding the value of a specific security control

ACKNOWLEDGMENTS

This work was supported by the Institute of information and infrastructure protection (I3P), the Department of Homeland Security (DHS) National Cyber Security Division, and the RAND Corporation, and was performed under DHS contract #2006-CS-001-000001. The authors thank Inette Furey of DHS and Martin Wybourne of Darmouth College for sponsoring this research and Richard George (Johns Hopkins University Applied Physics Laboratory), Dr Shari Pfleeger (Dartmouth University), and Kartik Gopalan (Binghamton University (SUNY)) for useful conversations.

REFERENCES

1. M.Walla, "Kerberos Explained," May, 2000.[Online].Available:http://technet.microsoft. com/en-us/library/bb742516-aspx. [Accessed:12-Jan-2014].
2. Microsoft,"Federation trusts," Aug22, 2005. [Online]. Available: http: //technet. microsoft. com/en-us/ library/ cc738707(v=ws.10). aspx.[Accessed:12-Jan-2014].
3. J.Somorovsky,M.Heiderich,M.Jensen,N. Grusehka, J.Schwenk , and L.Lo Iacono, " All your Clouds are belong to us: Security analysis of Cloud management interfaces," in *Proceeding of the 3rd ACM workshop on Cloud computing security workshop*,2011,pp.3-14
4. V.J.Wintler, *Securing the Cloud:Cloud computing security techniques and tactics*. Elsevier.2011
5. *Network Infrastructure Technology Overview*,Version 8,Release 3,Defense Information Systems Agency,August 27,2010.
6. J.Rutkowska, and A.Tereshkin, "Bluepilling the Xen Hypervisor," presented at the 2008 Blackhat conference, Las Vegas,NV,Aug.2008.
7. O.Sheyner et.al., "Automated Generation And Analysis Of Attack Graphs." *Secur. Priv.2002 Proc. IEEE Symp.On.* pp. 273,284 (2002).
8. A.Singhal and X.Ou, *Security Risk Analysis of Enterprise Networks Using Probabilistic Attack Graphs*, National Institute of Science and Technology Interagency Report 7788, Gaitheersburg, MD,August 2011.
9. M.Frigault, et.al., "Measuring network security using dynamic Bayesian network," in *proceedings of the 4th ACM workshop on Quality of protection,* pp. 23–30 (2008).
10. J.Ben-Gal, " Bayesian networks," *Encycl. Stat.Qual.Relaib* 2007.