CHAPTER 1

INTRODUCTION

CLOUD computing has begun to emerge as a hotspot in both industry and academia. It represents a new business model and computing paradigm, which enables on demand provisioning of computational and storage resources. Economic benefits consist of the main drive for cloud computing due to the fact that cloud computing offers an effective way to reduce capital expenditure (CapEx) and operational expenditure (OpEx). The definition of cloud computing has been given in many literatures, but nothing has gained wide recognition.

We define cloud computing as: ”A large-scale distributed computing paradigm that is driven by economies of scale, in which a pool of abstracted, virtualized, dynamically scalable, managed computing power, storage, platforms, and services are delivered on demand to external customers over the Internet.”

1.1 Cloud architecture

The below figure depicts the general architecture of a cloud platform, which is also called cloud stack. Building upon hardware facilities, cloud services may be offered in various forms from the bottom layer to top layer. In the cloud stack, each layer represents one service model. Infrastructure-as-a-Service (IaaS) is offered in the bottom layer, where resources are aggregated and managed physically or virtually and services are delivered in forms of storage, network or computational capability. The middle layer delivers Platform-asa-Service (PaaS), in which services are provided as an environment for programming or software execution. Software as a Service (SaaS) locates in the top layer, in which a cloud provider further confines client flexibility by merely offering software applications as a service. Apart from the service provisioning, the cloud provider maintains a suite of management tools and facilities (e.g.,

service instance life-cycle management, metering and billing, dynamic configuration) in order to manage a large cloud system.

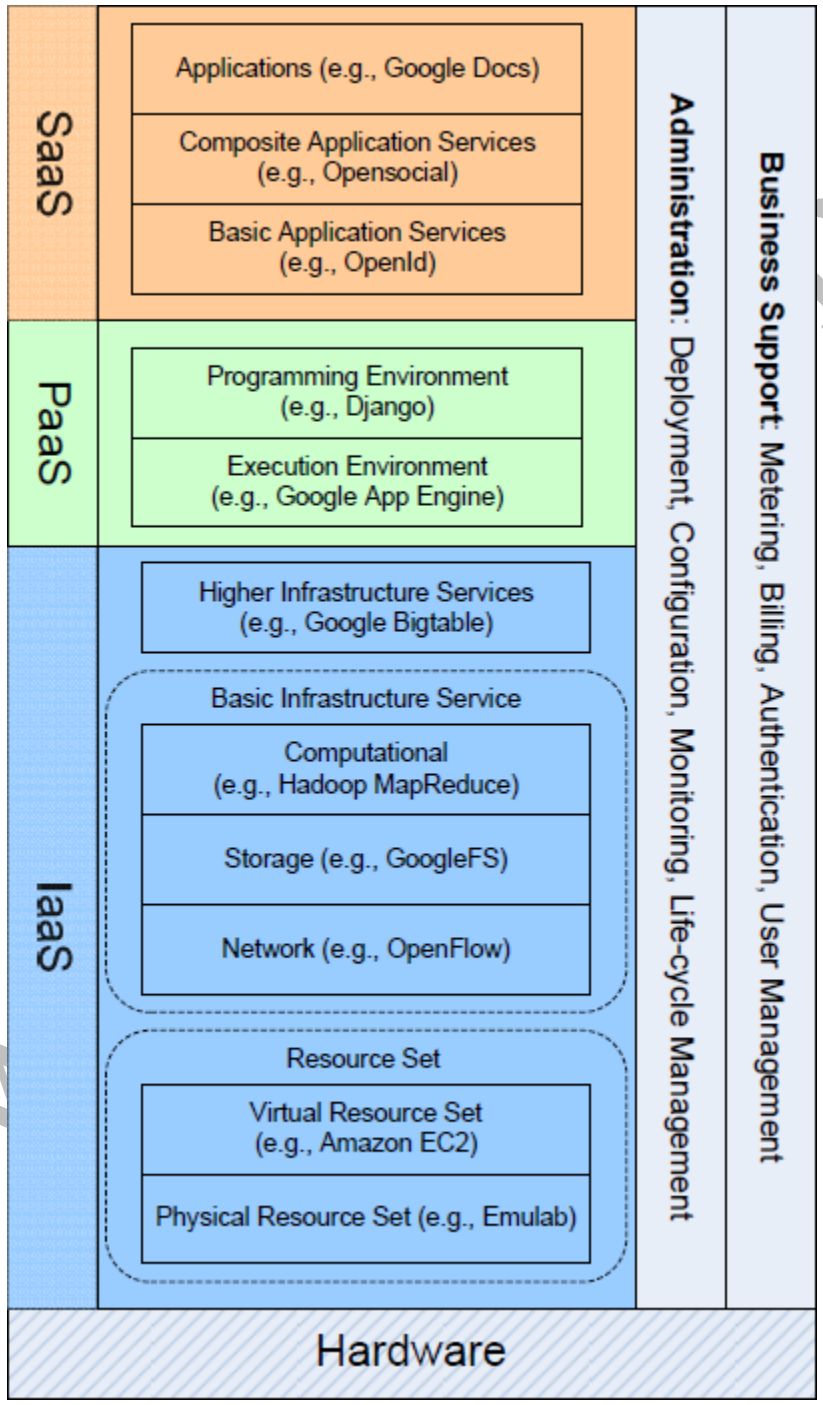

Fig. 1. Architecture of Cloud Computing

1.2 Cloud Characteristics and Security Challenges

The Cloud Security Alliance has summarized five essential characteristics [6] that illustrate the relation to, and differences from, traditional computing paradigm.

• **On-demand self-service** – A cloud customer may unilaterally obtain computing capabilities, like the usage of various servers and network storage, as on demand, without interacting with the cloud provider.

• **Broad network access** – Services are delivered across the Internet via a standard mechanism that allows customers to access the services through heterogeneous thin or thick client tools (e.g., PCs, mobile phones, and PDAs).

• **Resource pooling** – The cloud provider employs a multitenant model to serve multiple customers by pooling computing resources, which are different physical and virtual resources dynamically assigned or reassigned according to customer demand. Examples of resources include storage, processing, memory, network bandwidth, and virtual machines.

• **Rapid elasticity** – Capabilities may be rapidly and elastically provisioned in order to quickly scale out or rapidly released to quickly scale in. From customers' point of view, the available capabilities should appear to be unlimited and have the ability to be purchased in any quantity at any time.

• **Measured service** – The service purchased by customers can be quantified and measured. For both the provider and customers, resource usage will be monitored, controlled, metered, and reported. Cloud computing becomes a successful and popular business model due to its charming features. In addition to the benefits at hand, the former features also result in serious cloud-specific security issues. The people whose concern is the cloud security continue to hesitate to transfer their business to cloud. Security issues have been the dominate barrier of the development and widespread use of cloud computing.

There are three main challenges for building a secure and trustworthy cloud system:

• **Outsourcing** – Outsourcing brings down both capital expenditure (CapEx) and operational expenditure for cloud customers. However, outsourcing also means that customers physically lose control on their data and tasks. The loss of control problem has become one of the root causes of cloud insecurity. To address outsourcing security issues, first, the cloud provider shall be trustworthy by providing trust and secure computing and data storage; second,

outsourced data and computation shall be verifiable to customers in terms of confidentiality, integrity, and other security services. In addition, outsourcing will potentially incur privacy violations, due to the fact that sensitive/classified data is out of the owners' control.

• **Multi-tenancy** – Multi-tenancy means that the cloud platform is shared and utilized by multiple customers. Moreover, in a virtualized environment, data belonging to different customers may be placed on the same physical machine by certain resource allocation policy. Adversaries who may also be legitimate cloud customers may exploit the co-residence issue. A series of security issues such as data breach, computation breach, flooding attack, etc., are incurred. Although Multi-tenancy is a definite choice of cloud venders due to its economic efficiency, it provides new vulnerabilities to the cloud platform. Without changing the multi-tenancy paradigm, it is imperative to design new security mechanisms to deal with the potential risks.

• **Massive data and intense computation** – cloud computing is capable of handling mass data storage and intense computing tasks. Therefore, traditional security mechanisms may not suffice due to unbearable computation or communication overhead. For example, to verify the integrity of data that is remotely stored, it is impractical to hash the entire data set. To this end, new strategies and protocols are expected.

1.3 Supporting techniques

Cloud computing has leveraged a collection of existing techniques, such as Data Center Networking (DCN), Virtualization, distributed storage, MapReduce, web applications and services, etc.

Modern data center has been practically employed as an effective carrier of cloud environments. It provides massive computation and storage capability by composing thousands of machines with DCN techniques.

Virtualization technology has been widely used in cloud computing to provider dynamic resource allocation and service provisioning, especially in IaaS. With virtualization, multiple OSs can co-reside on the same physical machine without interfering each other.

MapReduce is a programming framework that supports distributed computing on mass data sets. This breaks large data sets down into small blocks that are distributed to cloud servers for parallel computing. MapReduce speeds up the batch processing on massive data, which

makes this become the preference of computation model for cloud venders. Apart from the benefits, the former techniques also present new threats that have the capability to jeopardize cloud security.

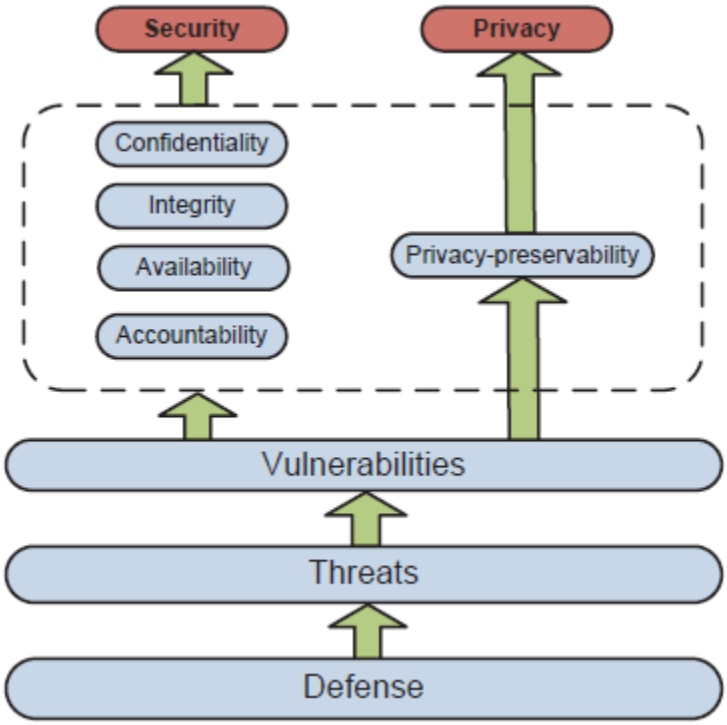

Fig. 2. Ecosystem of Cloud Security and Privacy

In this paper, we consider the cloud environment as a new computing platform to which the classic methodology of security research can be applied as well. Therefore, we determine to employ an attribute-driven methodology to conduct our review. We employ the ecosystem of cloud security and privacy in view of five security/privacy attributes (i.e., confidentiality, integrity, availability, accountability, and privacy-preservability), shown in Fig. 2,

1.4 Notation System

To better demonstrate the connection among vulnerability, threat, and defense mechanism. We employ the following notation system: let *Vi* denote a type of vulnerability, *Ti.j*

denote a type of threat that takes advantage of *Vi*, and *Di.j.k* denote a defense mechanism that deals with *Ti.j*. For instance, vulnerability *V*1 may be exploited by adversaries in order to create a threat model *T*1.1, which shall be patched by security solution *D*1.1.1.

1.5 Cloud Vulnerabilities

*1) V*1 – *VM co-residence:* In cloud computing, co-residence (or co-tenancy) means that multiple independent customers share the same physical infrastructure. Concretely, virtual machines belonging to different customers may be placed in the same physical machine. VM co-residence has raised certain security issues, such as Cross-VM attack and Malicious SysAdmin.

*2) V*2 – *Loss of Physical Control:* Cloud customers have their data and program outsourced to cloud servers. As a result, owners lose direct control on the data sets and programs. Loss of physical control means that customers are unable to resist certain attacks and accidents. For example, data or software may be altered, lost, or even deleted; in addition, it is difficult and impractical to ensure data/computation integrity and confidentiality with traditional methods.

*3) V*3 – *Bandwidth Under-provisioning:* A traditional DOS/DDOS attack does exist in cloud computing, and relative solutions have been given in prior researches. Specific to cloud computing, there is a new type of DOS attack that takes advantage of the current under-provisioned cloud-computing infrastructure. According to Cisco's design guide, a data center is usually designed to be under provisioned with a factor of 2.5:1 to 8:1, meaning that the actual network capacity is much less than the aggregate capacity of the hosts located in the same subnet.

*4) V*4 – *Cloud Pricing Model:* Cloud computing adheres to the pay-as-you-go pricing model that determines the cost of services in terms of metrics such as server hours, bandwidth, storage, etc. Since all cloud customers are financially responsible for the services they use, attackers always have incentives to harass the billing process by exploiting the pricing model. For example, Economic Denial of Sustainability (EDoS) attack manipulates the utility pricing model and causes unmanageable costs for cloud customers. The remainder of this paper is structured as follows: Sections II to VI discuss confidentiality, integrity, availability, accountability, and privacy in cloud computing, respectively; finally, the paper is concluded in Section VII.

CHAPTER 2

CLOUD CONFIDENTIALITY

When dealing with cloud environments, confidentiality implies that a customer's data and computation tasks are to be kept confidential from both the cloud provider and other customers. Confidentiality remains as one of the greatest concerns with regards to cloud computing. This is largely due to the fact that customers outsource their data and computation tasks on cloud servers, which are controlled and managed by potentially untrustworthy cloud providers.

2.1 Threats to Cloud Confidentiality

*1) T*1.1 – *Cross-VM attack via Side Channels:* Ristenpart et al demonstrates the existence of Cross-VM attacks in an Amazon EC2 platform. A Cross-VM attack exploits the nature of multi-tenancy, which enables that VMs belonging to different customers may co-reside on the same physical machine. Aviram regard timing side-channels as an insidious threat to cloud computing security due to the fact that a) the timing channels pervasively exist and are hard to control due to the nature of massive parallelism and shared infrastructure; b) malicious customers are able to steal information from other ones without leaving a trail or raising alarms. There are two main steps to practically initiate such an attack:

• **Step 1: placement**. An adversary needs to place a malicious VM on the physical server where the target client's VM is located. To achieve this, an adversary should first determine where the target VM instance is located; this can be done with network probing tools such as nmap, hping, wget, etc. An adversary should also be able to determine if there are two VM instances; 1) comparing Domain0's IP addresses to see if they match, and 2) measuring the small packet round-trip time can do this check. The correctness of co-resident checks can be verified by transmitting messages between instances via a covert channel. After all the prep work, a malicious VM instance must be created on the target physical machine by specifying a set of parameters (e.g., zone, host type); there are two basic strategies to launch such a VM: 1) brute-force strategy, which simply launches many instances and checks co-residence with the

target; 2) an adversary can exploit the tendency that EC2 launches new instances on the same small set of physical machines. The second strategy takes advantage of EC2's VM assigning algorithm by starting a malicious VM after a victim VM is launched so that they will likely be assigned to the same physical server; this approach surely has better success rate of placement.

• **Step 2: extraction**. After step 1, a malicious VM has co-resided with the victim VM. Since the malicious VM and the victim are sharing certain physical resources, such as data cache, network access, CPU branch predicators, CPU pipelines, etc., there are many ways an adversary can employ attacks: 1) measuring a cache usage that can estimate the current load of the server; 2) estimating a traffic rate that can obtain the visitor count or even the frequently requested pages; 3) a keystroke timing attack that can steal a victim's password by measuring time between keystrokes. As follow-up work, various covert channels are investigated and in-depth analysis is provided. Attackers can easily exploit L2 cache, due to its high bandwidth. Xu et al. have particularly explored the L2 cache covert channel with quantitative assessment . It has been demonstrated that even the channel bit rate is higher than the former work, the channel's ability to exfiltrate useful information is still limited, and it is only practical to leak small secrets such as private keys. Okamura et al. developed a new attack, which demonstrates that CPU load can also be used as a covert channel to encode information. Memory disclosure attack is another type of cross-VM attack. In a virtualized environment, memory deduplication is a technique to reduce the utilization of physical memory by sharing the memory pages with same contents. A memory disclosure attack is capable of detecting the existence of an application or a file on a co-residing VM by measuring the write access time that differs between deduplicated pages and regular ones.

*2) T*1.2 – *Malicious SysAdmin:* The Cross-VM attack discusses how others may violate confidentiality cloud customers that co-residing with the victim, although it is not the only threat. Privileged sysadmin of the cloud provider can perform attacks by accessing the memory of a customer's VMs. For instance, Xenaccess enables a sysadmin to directly access the VM memory at run time by running a user level process in Domain0.

2.2 Defense Strategies

Approaches to address cross-VM attack fall into six categories: a) placement prevention intends to reduce the success rate of placement; b) physical isolation enforcement c) new cache designs d) fuzzy time intends to weaken malicious VM's ability to receive the signal by eliminating fine-grained timers e) forced VM determinism ensures no timing or other non-deterministic information leaking to adversaries; f) cryptographic implementation of timing-resistant cache. Since c), d), e), and f) are not cloud-specific defense strategies, we do not include details in this section.

*1) D*1.1.1 – *Placement Prevention:* In order to reduce the risk caused by shared infrastructure, a few suggestions to defend the attack in each step are given in. For instance, cloud providers may obfuscate co-residence by having Dom0 not respond in traceroute, and/or by randomly assigning internal IP addresses to launched VMs. To reduce the success rate of placement, cloud providers might let the users decide where to put their VMs; however, this method does not prevent a brute-force strategy.

*2) D*1.1.2 – *Co-residency Detection:* The ultimate solution of cross-VM attack is to eliminate co-residency. Cloud customers (especially enterprises) may require physical isolation, which can even be written into the Service Level Agreements (SLAs). However, cloud vendor may be reluctant to abandon virtualization that is beneficial to cost saving and resource utilization. One of the left options is to share the infrastructure only with "friendly" VMs, which are owned by the same customer or other trustworthy customers. To ensure physical isolation, a customer should be enabled to verify its VMs' exclusive use of a physical machine. HomeAlone is a system that detects co-residency by employing a side-channel (in the L2 memory cache) as a detection tool. The idea is to silence the activity of "friendly" VMs in a selected portion of L2 cache for a certain amount of time, and then measure the cache usage to check if there is any unexpected activity, which indicates that the physical machine is co-resided by another customer.

*3) D*1.1.3 – *NoHype:* NoHype attempts to minimize the degree of shared infrastructure by removing the hypervisor while still retaining the key features of virtualization. The NoHype architecture provides a few features: i) the "one core per VM" feature prevents interference between VMs, eliminates side channels such as L1 cache, and retains multi-tenancy, since each chip has multiple cores; ii) memory partition restricts each VM's memory access on a assigned

range; iii) dedicated virtual I/O devices enables each VM to be granted direct access to a dedicated virtual I/O device. No- Hype has significantly reduced the hypervisor attack surface, and increased the level of VM isolation. However, NoHype requires to change hardware, making it less practical when consider applying it to current cloud infrastructures.

*4) D*1.2.1 – *Trusted Cloud Computing Platform: P*resent a trusted cloud-computing platform (TCCP), which offers a closed box execution environment for IaaS services. TCCP guarantees confidential execution of guest virtual machines. It also enables customers to attest to the IaaS provider and to determine if the service is secure before their VMs are launched into the cloud. The design goals of TCCP are: 1) to confine the VM execution inside the secure perimeter; 2) that a sysadmin with root privileges is unable to access the memory of a VM hosted in a physical node. TCCP leverages existing techniques to build trusted cloud computing platforms. This focuses on solving confidentiality problems for clients' data and for computation outsourced to the cloud. With TCCP, the sysadmin is unable to inspect or tamper with the content of running VMs.

5) Other opinions: retaining data control back to customer: Considering the customer's fear of losing the data control in cloud environments, Descher et al. [40] propose to retain data control for the cloud customers by simply storing encrypted VMs on the cloud servers. Encrypted VM images guarantee rigorous access control since only the authorized users known as key-holders are permitted access. Due to the encryption, the data cannot be mounted and modified within the cloud without an access key, assuring the confidentiality and integrity. This approach offers security guarantees before a VM is launched; however, there are ways to attack the VM during running time and to jeopardize the data and computation.

2.3 Summary and Open issues

Regarding confidentiality, cross-VM attack and malicious SysAdmin mainly threaten a cloud system; both threats take advantage of the vulnerability of virtualization and coresidence. Other tenants perform cross-VM attack, whereas the malicious SysAdmin is inside attack from cloud vender. Defending these threats is not a trivial task due to the following facts: 1) various side channels and other shared components can be exploited, and defending each of them is not an easy job; 2) There are a few open issues to be explored:

• Co-residency detection is considered as a promising technique since customers should be able to check whether the physical isolation is well enforced. HomeAlone has the ability to achieve accuracy of detection on L2 cache side channels. However, besides L2 cache, other side channels may be exploited as well. Therefore, in order to provide thorough detection of co-residence, a suite of detection methods targeting on various side channels should be developed.

• NoHype has opened another window to deal with cross- VM threat. However, current commodity hardware imposes limitations to implement NoHype. Additionally, live VM migration is not well supported by this new architecture. Therefore, before making a real step forward, researchers need to address the hardware changes to accommodate NoHype and to maintain more features for VM management.

CHAPTER 3

CLOUD AVAILABILITY

Availability is crucial since the core function of cloud computing is to provide on-demand service of different levels. If a certain service is no longer available or the quality of service cannot meet the Service Level Agreement (SLA), customers may lose faith in the cloud system. In this section, we have studied two kinds of threats that impair cloud availability.

3.1 Threats to Cloud Availability

*1) T*3.1 – *Flooding Attack via Bandwidth Starvation*: In a flooding attack, which can cause Deny of Service (DoS), a huge amount of nonsensical requests are sent to a particular service to hinder it from working properly. In cloud computing, there are two basic types of flooding attacks:

• **Direct DOS** – the attacking target is determined, and the availability of the targeting cloud service will be fully lost.

• **Indirect DOS** – the meaning is twofold: 1) all services hosted in the same physical machine with the target victim will be affected; 2) the attack is initiated without a specific target.

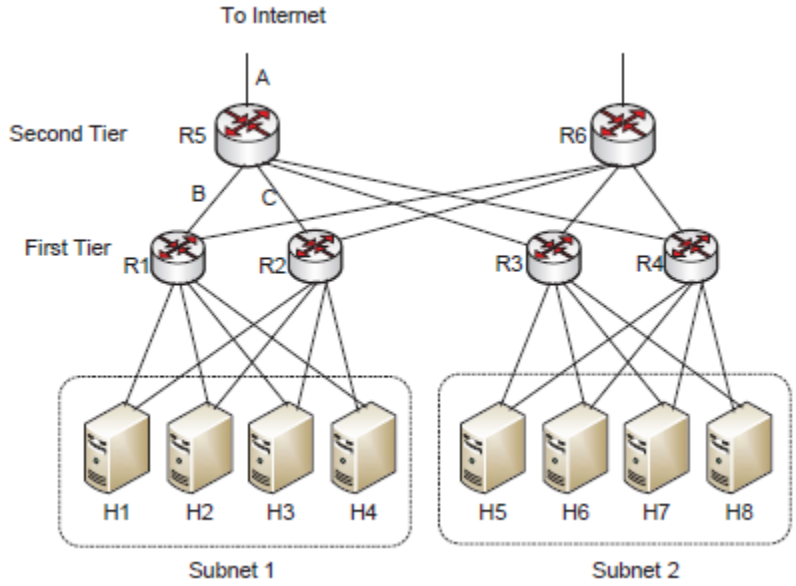

Fig. 3. Traditional Data Center Network Architecture

The authors in also point out that one of the consequences of a flooding attack is that if a certain cloud service is unavailable or the quality of service is degraded, the subscribers of all affected services may need to continue paying the bill. However, we argue that since cloud providers must have previously signed a Service Level Agreement (SLA) with their clients, a responsible party must be determined once the service level is degraded to some threshold since clients will be aware of that degradation. We will elaborate upon this problem (i.e., cloud accountability) in the next section. The nature of under-provisioning and public openness in a cloud system brings new vulnerability that can be exploited to carry out a new DOS attack to jeopardize the cloud service provision by saturating the limited network bandwidth. As shown in Fig. 3, links A, B, C are uplinks of router R5, R1, and R2, respectively. Suppose that link B is the active link and link C is the fail-over link (i.e., a link will be activated when the active link is down). Due to under-provisioning, the aggregate capacity of H1, H2, H3, and H4 (which form the subnet 1) is a few times larger than any capacity for links A, B, or C. In order to saturate link B, attackers (which may be a few hosts controlled by the adversary) in subnet 1 only need to generate enough traffic to target the hosts in another subnet (e.g., subnet 2). Once link B is saturated by the non-sense traffic, hosts in subnet1 are unable to deliver services to cloud users.
To initiate such a DOS attack (bandwidth starvation) effectively, there are a few steps:

1) **Topology identification** – Since only hosts in different subnets are connected by bottleneck links, an adversary needs to first identify the network topology. By exploiting the multiplexing nature of a router, the number of routers between two hosts can be determined; this helps selected hosts picture the topology.

2) **Gaining access to enough hosts** – The number of hosts to perform the attack is determined by the uplink's capacity, which can be estimated by some tools such as Pathload, Nettimer, or Bprobe.

3) **Carrying out the attack** – The author suggests employing UDP traffic because it will starve other TCP sessions.

*2) T*4.1 – *Fraudulent Resource Consumption (FRC) attack:* A representative Economic Denial of Sustainability (EDoS) attack is FRC, which is a subtle attack that may be carried out over a long period (usually lasts for weeks) in order to take effect. In cloud computing, the goal of a FRC attack is to deprive the victim (i.e., regular cloud customers) of their long-term

economic availability of hosting web contents that are publicly accessible. In other words, attackers, who act as legal cloud service clients, continuously send requests to website hosting in cloud servers to consume bandwidth, which bills to the cloud customer owning the website; seems to the web server, those traffic does not reach the level of service denial, and it is difficult to distinguish FRC traffic from other legitimate traffic. A FRC attack succeeds when it causes financial burden on the victim.

3.2 Defense strategy

*1) D*3.1.1 – *defending the new DOS attack:* This new type of DOS attack differs from the traditional DOS or DDOS attacks in that traditional DOS sends traffic to the targeting application/host directly while the new DOS attack does not; therefore, some techniques and counter-measures, for handling traditional DOSs are no longer applicable. A DOS avoidance strategy called service migration has been developed to deal with the new flooding attack. A monitoring agent located outside the cloud is set up to detect whether there may be bandwidth starvation by constantly probing the cloud applications. When bandwidth degradation is detected, the monitoring agent will perform application migration, which may stop the service temporarily, with it resuming later. The migration will move the current application to another subnet of which the attacker is unaware. Experiment results show that it only takes a few seconds to migrate a stateless web application from one subnet to another.

*2) D*4.1.1 – *FRC attack detection:* The key of FRC detection is to distinguish FRC traffic from normal activity traffic. Idziorek et al. propose to exploit the consistency and selfsimilarity of aggregate web activity . To achieve this goal, three detection metrics are used: i) Zipf ’s law are adopted to measure relative frequency and self-similarity of web page popularity; ii) Spearman’s footrule is used to find the proximity between two ranked lists, which determines the similarity score; iii) overlap between the reference list and the comparator list measures the similarity between the training data and the test data. Combining the three metrics yields a reliable way of FRC detection.

3.3 Summary and Open Issues

Service downgrade can be resulted by both internal and external threats. An internal threat comes from malicious cloud customers who take advantage of the bandwidth underprovisioning property of current DCN architecture to starve legitimate service traffic. On the other hand, external threat refers to the EDoS attack, which degrades the victim's longterm economic availability. Both DoS and EDos have appeared in other scenarios, however, the ways they employ to attack the cloud platform are novel and worthwhile to be investigated.

CHAPTER 4

CLOUD PRIVACY

Privacy is yet another critical concern with regards to cloud computing due to the fact that customers' data and business logic reside among distrusted cloud servers, which are owned and maintained by the cloud provider. Therefore, there are potential risks that the confidential data or personal information is disclosed to public or business competitors.

4.1 *Threats to Cloud Privacy*

In some sense, privacy-preservability is a stricter form of confidentiality, due to the notion that they both prevent information leakage. Therefore, if cloud confidentiality is ever violated, privacy-preservability will also be violated. Similar to other security services, the meaning of cloud privacy is twofold: data privacy and computation privacy.

4.2 Defense Strategies

Gentry proposed **Fully Homomorphic Encryption** (FHE, *D*2.5.1) to preserve privacy in cloud computing. FHE enables computation on encrypted data, which is stored in the distrusted servers of the cloud provider. Data may be processed without decryption. The cloud servers have little to no knowledge concerning the input data, the processing function, the result, and any intermediate result values. Therefore, the outsourced computation occurs 'under the covers' in a fully privacy-preserving way. FHE has become a powerful tool to enforce privacy preserving in cloud computing. However, all known FHE schemes are too inefficient for use in practice. While researchers are trying to reduce the complexity of FHE, it is worthwhile to consider alleviating the power of FHE to regain efficiency. Naehrig et al. has proposed somewhat homomorphic encryption ,which only supports a number of homomorphic operations, which may be much faster and more compact than FHE.

Pearson et propose **privacy manager** (*D*2.5.2) that relies on obfuscation techniques. The privacy manager can provide obfuscation and de-obfuscation service to reduce the amount of sensitive information stored in the cloud.

The main idea is to only store the encrypted form of clients' private data in the cloud end. The data process is directly performed on the encrypted data. One limitation is that cloud vendors may not be willing to implement additional services for privacy protection. Without provider's cooperation, this scheme will not work. Squicciarini explores a novel privacy issue that is caused by data indexing. In order to tackle data indexing and to prevent information leakage, the researchers present a three-tier data protection architecture to offer different levels of privacy to cloud customers. Itani presents a Privacy-as-a-Service so it may enable secure storage and computation of private data by leveraging the tamper-proof capabilities of cryptographic coprocessors. Which, in turn, protect customer data from unauthorized access. Sadeghi argue that pure cryptographic solutions based on fully homomorphic and verifiable encryption suffer high latency for offering practical secure outsourcing of computation to a distrusted cloud service provider.

They propose to combine a **trusted hardware token** (*D*2.5.3) with Secure Function Evaluation (SFE) in order to compute arbitrary functions on data when it is still in encrypted form. The computation leaks no information and is verifiable. The focus of this work is to minimize the computation latency to enable efficient, secure outsourcing in cloud computing. A hardware token is tamper-proof against physical attacks. If the token is under the assumption of being trusty, the clients' data processing may be performed in the token that is attached to a distrusted cloud server. The property of a token can guarantee that the data computation is confidential as well as being verifiable.

4.3 Open Issues

Regarding cloud privacy, there are some open issues to be studied in future researches:

• The authors think that accountability and privacy may conflict with each other. The enforcement of accountability will violate privacy in some degree, and extreme privacy protection (e.g., full anonymity to hide users' identity) will make accountability more challenging. An extreme example, a shared file, accessed by multiple users who, may hide their identities due to anonymity for the purpose of privacy protection. However, malicious users are tracked with difficultly because of the anonymous access. From the viewpoint of accountability, general approaches include information logging, replay, tracing , etc. These operations may not

be completed without revealing some private information (e.g., account name, IP address). We must seek a trade-off in which the requirement of one attribute can be met while simultaneously maintaining some degree of the other attribute.

• The assessment of attributes is another important issue since it provides a quantitative way to evaluate them. The goal is to determine how secure a cloud is or how much privacy can be offered. The meaning is twofold:

1) it will be helpful to compare different security approaches;

for example, to achieve 100% privacy, scheme A costs 100; scheme B can achieve 99% accountability with cost of 10. Apparently, scheme B is more practically efficient, although it sacrifices one percent of accountability. Without an assessment, it is difficult to compare two strategies quantitatively.

2) The quantitative clauses of the security/privacy requirements can be drafted into the Service Level Agreements (SLAs).

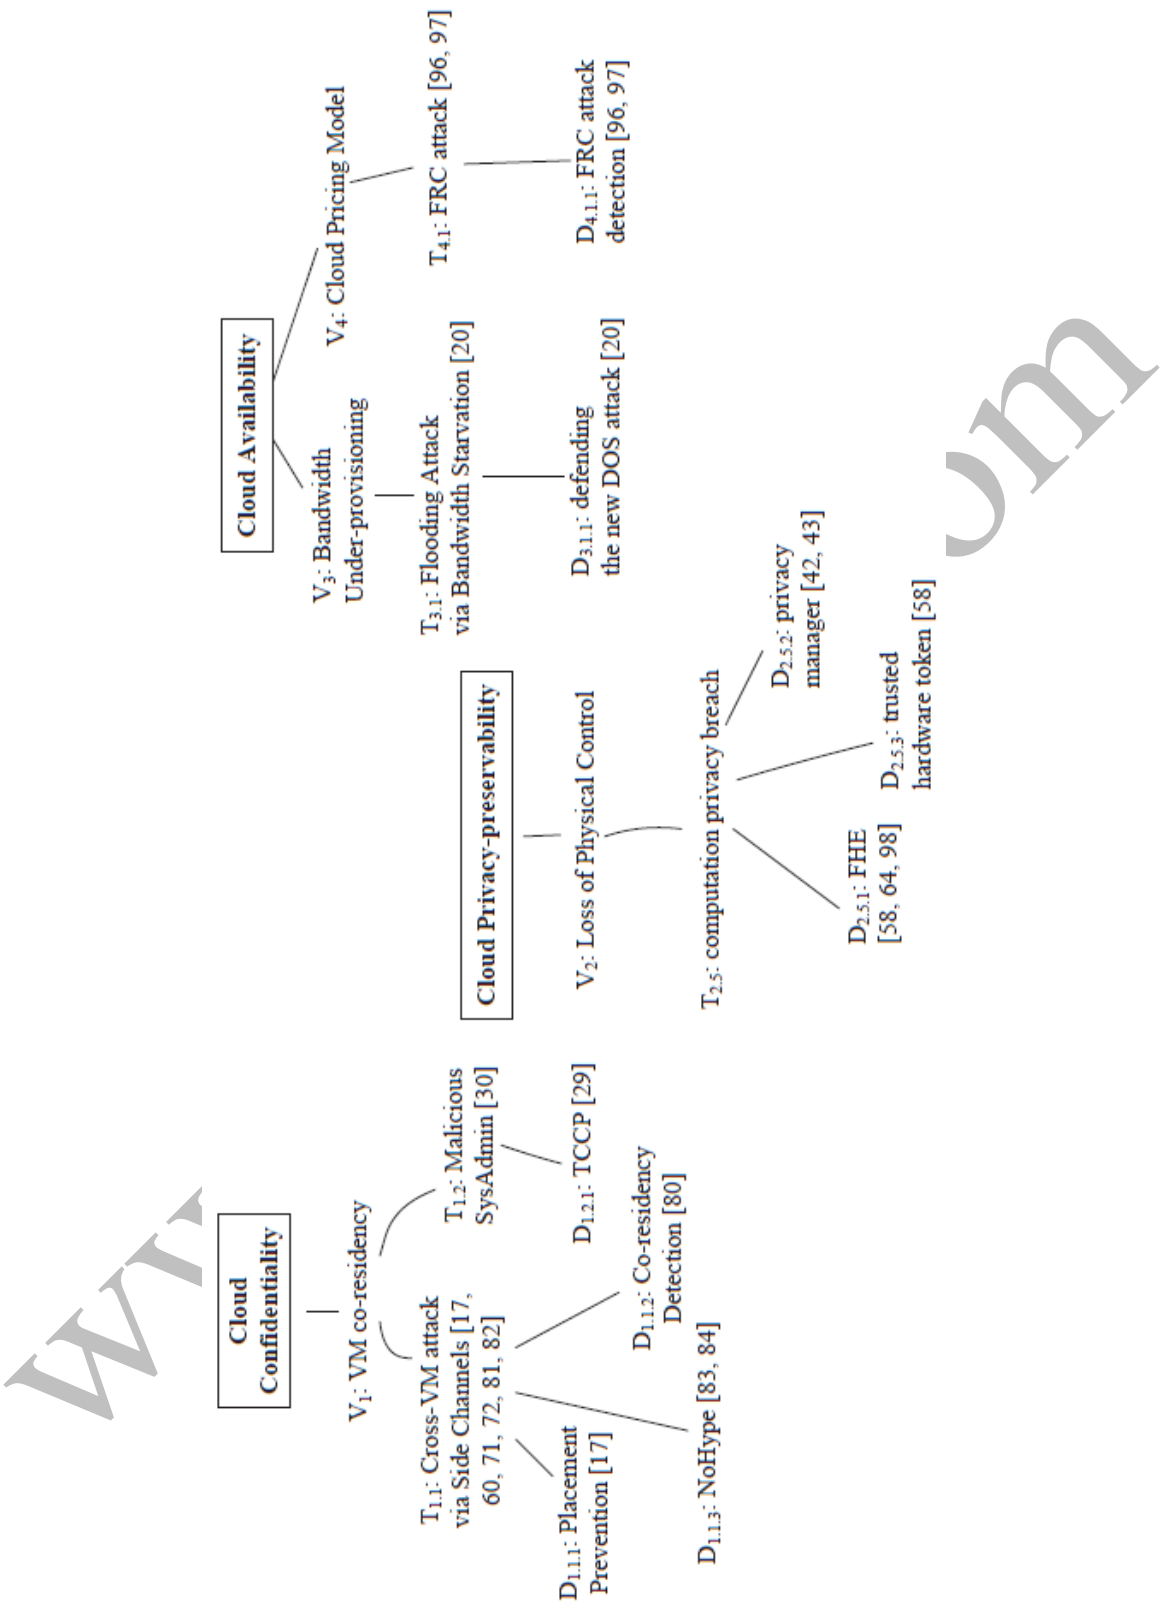

Fig. 4. A Summary of Research Advances in Cloud Security and Privacy

CONCLUSIONS

Throughout this paper, the authors have systematically studied the security and privacy issues in cloud computing based on an attribute-driven methodology, shown in Fig. 4 and have identified the most representative security/privacy attributes (e.g., confidentiality, integrity, availability, accountability, and privacy-preservability), as well as discussing the vulnerabilities, which may be exploited by adversaries in order to perform various attacks. Defense strategies and suggestions were discussed as well. Authors believe this review will help shape the future research directions in the areas of cloud security and privacy.

REFERENCES

[1] I. Foster, Y. Zhao, I. Raicu, and S. Lu, "Cloud computing and grid computing 360-degree compared," Grid Computing Environments Workshop, 2008. GCE'08, 2009, pp. 1-10.

[2] J. Geelan. "Twenty one experts define cloud computing," Virtualization, August 2008. Electronic Mag., article available at http:// virtualization.sys-con.com/node/612375.

[3] R. Buyya, C. S. Yeo, and S. Venugopal. "Market-oriented cloud computing: Vision, hype, and reality for delivering it services as computing utilities," CoRR, (abs/0808.3558), 2008.

[4] Luis M. Vaquero, Luis Rodero-Merino and Daniel Mor´an, "Locking the sky: a survey on IaaS cloud security," Computing, 2010, DOI:10.1007/s00607-010-0140-x.

[5] Google Docs experienced data breach during March 2009. http://blogs.wsj.com/digits/2009/03/08/1214/

[6] Cloud Security Alliance (CSA). "Security Guidance for Critical Areas of Focus in Cloud Computing V2.1," (Released December 17, 2009). http://www.cloudsecurityalliance.org/guidance/csaguide.v2.1.pdf

[7] Cloud Security Alliance (CSA). "Top Threats to Cloud Computing V 1.0," released March 2010.

[8] The security-as-a-service model. http://cloudsecurity.trendmicro.com /the-security-as-a-service-model/

[9] S.D. Di Vimercati, S. Foresti, S. Jajodia, S. Paraboschi, and P. Samarati, "A data outsourcing architecture combining cryptography and access control," Proc. 2007 ACM workshop on Computer security architecture, 2007, pp. 63-69.

[10] P. Mell and T. Grance. The NIST Definition of Cloud Computing (Draft). [Online] Available: www.nist.gov/itl/cloud/upload/cloud-defv15. pdf., Jan. 2011.

[11] G. Ateniese, R. Burns, R. Curtmola, J. Herring, L. Kissner, Z. Peterson, and D. Song, "Provable data possession at untrusted stores," In ACM CCS, pages 598-609, 2007.

[12] A. Juels and B. S. Kaliski, "PORs: Proofs of retrievability for large files," In ACM CCS, pages 584-597, 2007.

[13] Y. Dodis, S. Vadhan, and D. Wichs, "Proofs of retrievability via hardness amplication," In TCC, 2009.

[14] G. Ateniese, R. D. Pietro, L. V. Mancini, and G. Tsudik, "Scalable and efficient provable data possession," SecureComm, 2008.

[15] C. Erway, A. K¨upc¸ ¨u, C. Papamanthou, and R. Tamassia, "Dynamic provable data possession," Proc. 16th ACM conference on Computer and communications security, 2009, pp. 213-222.

[16] K.D. Bowers, A. Juels, and A. Oprea, "HAIL: A high-availability and integrity layer for cloud storage," Proc. 16th ACM conference on Computer and communications security, 2009, pp. 187-198.

[17] T. Ristenpart, E. Tromer, H. Shacham, and S. Savage, "Hey, you, get off of my cloud: exploring information leakage in third-party compute clouds," Proc. 16th ACM conference on Computer and communications security, 2009, pp. 199-212.

[18] V. Sekar and P. Maniatis, "Verifiable resource accounting for cloud computing services," in Proc. 3rd ACM workshop on Cloud computing security workshop, 2011, pp. 21-26.

[19] C. Hoff, "Cloud computing security: From DDoS (distributed denial of service) to EDoS (economic denial of sustainability)," [Online] Available: http://www.rationalsurvivability.com/blog/?p=66., 2008.

[20] H. Liu, "A New Form of DOS Attack in a Cloud and It's AvoidanceMechanism", Cloud Computing Security Workshop 2010.