University of Science and Technology

Faculty of Computer Science and Information Technology

Information Security Master Program

Course Title: Cloud Computing Security

Lecture (5): Privacy

Reference: Cloud Computing Security and Privacy, by Tim Mather, Subra and Lattif, -------Chapter (5)

Instructor: Prof. Noureldien Abdelrahman

5.1 What Is Privacy?

The concept of privacy varies widely among (and sometimes within) countries, cultures, and jurisdictions. It is shaped by public expectations and legal interpretations; as such, a concise definition is elusive if not impossible.

> ***Privacy rights or obligations are related to the collection, use, disclosure, storage, and destruction of personal data (or personally identifiable information (PII)).***

At the end of the day, privacy is about the accountability of organizations to data subjects, as well as the transparency to an organization's practice around personal information.

Likewise, there is no universal ***consensus about what constitutes personal data.*** For the purposes of this discussion, we will use the **definition adopted by the Organization for Economic Cooperation and Development (OECD):**

> ***Personal data is any information relating to an identified or identifiable individual (data subject).***

5.2 What Is the Data Life Cycle?

Personal information should be managed as part of the data used by the organization. Protection of personal information should consider the impact of the cloud on each of the following phases as detailed in Figure 7-1.

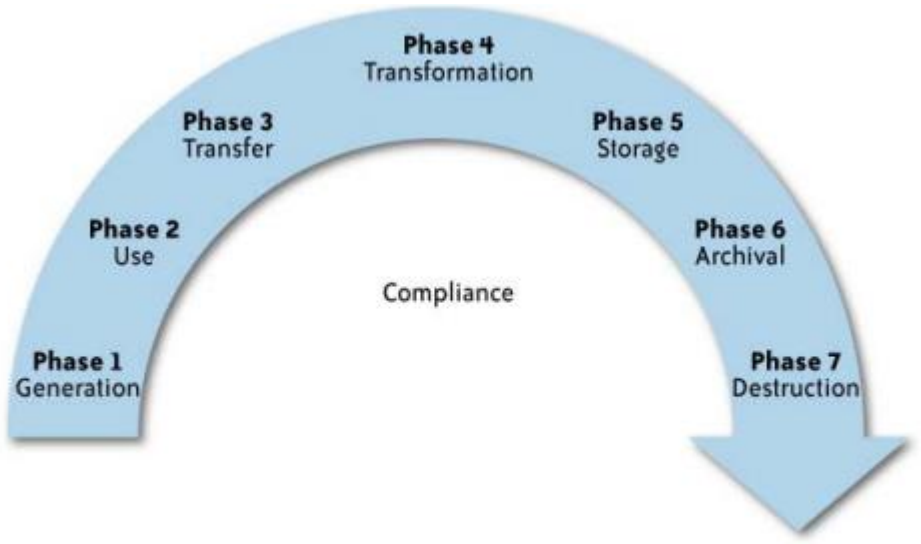

FIGURE 7-1. KPMG data life cycle

Each components within each of these phases need to be addressed to preserve PII privacy.

1. Generation of the information

• ***Ownership***: Who in the organization owns PII, and how is the ownership maintained if the organization uses cloud computing?

• ***Classification***: How and when is PII classified? Are there limitations on the use of cloud computing for specific data classes?

• ***Governance***: Is there a governance structure to ensure that PII is managed and protected through its life cycle, even when it is stored or processed in a cloud computing environment?

2. Use

• ***Internal versus external***: Is PII used only within the collecting organization, or is it used outside the organization (e.g., in a public cloud)?

• ***Third party***: Is the information shared with third parties (e.g., subcontractors or CSPs)?

• ***Appropriateness***: Is the use of the information consistent with the purpose for which it was collected? Is the use within the cloud appropriate based on the commitments the organization made to the data subjects?

• ***Discovery*:** Is the information managed in the cloud in a way that will enable the organization to comply with legal requirements in case **of legal proceedings**?

3. Transfer

• ***Public versus private networks***: When information is transferred to a cloud is the organization using public networks, and is it protected appropriately? (PII should always be protected to address the risk level and legal requirements.)

• ***Encryption requirements***: Is the PII encrypted? Some laws require that PII will be encrypted when transmitted via a public network (and this will be the case when the organization is using a public cloud).

• ***Access control***: Are there appropriate access controls over PII when it is in the cloud?

4. Transformation

• ***Derivation*:** Are the original protection and use limitations maintained when data is transformed or further processed in the cloud?

• ***Aggregation***: Is **data in the cloud aggregated so that it is no longer related** to an identifiable individual (**and hence is no longer considered PII**)?

• ***Integrity***: Is the integrity of PII maintained when it is in the cloud?

5. Storage

• ***Access control***: Are there appropriate controls over access to PII when stored in the cloud so that only individuals with a need to know will be able to access it?

• ***Structured versus unstructured***: How is the data stored to enable the organization to access and manage the data in the future?

• ***Integrity/availability/confidentiality***: How are data integrity, availability, and confidentiality maintained in the cloud?

• ***Encryption*:** Several laws and regulations require that certain types of PII should be stored only when encrypted. Is this requirement supported by the CSP?

6. Archival

• ***Legal and compliance***: PII may have specific requirements that dictate how long it should be stored and archived. Are these requirements supported by the CSP?

• ***Off-site considerations***: Does the CSP provide the ability for long-term off-site storage that supports archival requirements?

• ***Media concerns***: Is the information stored on media that will be accessible in the future? Is the information stored on portable media that may be more susceptible to loss? Who controls the media and what is the organization's ability to recover such media from the CSP if needed?

• ***Retention:*** For how long will the data be retained by the CSP? Is the retention period consistent with the organization's retention period?

7- Destruction

• ***Secure:*** Does the CSP destroy PII obtained by customers in a secure manner to avoid potential breach of the information?

• ***Complete***: Is the information completely destroyed? Does the destruction completely erase the data, or can it be recovered?

Hence, every organization should consider performing a Privacy Impact Assessment (PIA) before embarking on a cloud computing initiative that involves personal information.

5.3 What Are the Key Privacy Concerns in the Cloud?

Privacy advocates have raised many concerns about cloud computing. These concerns typically mix security and privacy. Here are some additional considerations to be aware of:

1. Access

Data subjects have a right to know what personal information is held and, in some cases, can make a request to stop processing it.

In the cloud, the main concern is the organization's ability to provide the individual with access to all personal information, and to comply with stated requests. If a data subject exercises this right to ask the organization to delete his data, will it be possible to ensure that all of his information has been deleted in the cloud?

2. Compliance

What are the privacy compliance requirements in the cloud? What are the applicable laws, regulations, standards, and contractual commitments that govern this information, and who is responsible for maintaining the compliance? How are existing privacy compliance requirements

impacted by the move to the cloud? Clouds can cross multiple jurisdictions; for example, data may be stored in multiple countries, or in multiple states within the United States. What is the relevant jurisdiction that governs an entity's data in the cloud and how is it determined?

3. Storage

Where is the data in the cloud stored? Was it transferred to another data center in another country? Is it commingled with information from other organizations that use the same CSP? Privacy laws in various countries place limitations on the ability of organizations to transfer some types of personal information to other countries. When the data is stored in the cloud, such a transfer may occur without the knowledge of the organization, resulting in a potential violation of the local law.

4. Retention

How long is personal information (that is transferred to the cloud) retained? Which retention policy governs the data? Does the organization own the data, or the CSP? Who enforces the retention policy in the cloud, and how are exceptions to this policy (such as litigation holds) managed?

5. Destruction

How does the cloud provider destroy PII at the end of the retention period? How do organizations ensure that their PII is destroyed by the CSP at the right point and is not available to other cloud users? How do they know that the CSP didn't retain additional copies? Cloud storage providers usually replicate the data across multiple systems and sites—increased availability is one of the benefits they provide. This benefit turns into a challenge when the organization tries to destroy the data—can you truly destroy information once it is in the cloud? Did the CSP really destroy the data, or just make it inaccessible to the organization? Is the CSP keeping the information longer than necessary so that it can mine the data for its own use?

6. Audit and monitoring

How can organizations monitor their CSP and provide assurance to relevant stakeholders that privacy requirements are met when their PII is in the cloud?

7. Privacy breaches

How do you know that a breach has occurred, how do you ensure that the CSP notifies you when a breach occurs, and who is responsible for managing the breach notification process (and costs associated with the process)?

If contracts include liability for breaches resulting from negligence of the CSP, how is the contract enforced and how is it determined who is at fault?

Many of these concerns are not specific to personal information, but to all types of information and a broader set of compliance requirements.

5.4 Impact of cloud computing on privacy principles

The following paragraphs describe analysis of the potential impact of cloud computing on the key OECD and other common privacy principles.

1. Collection Limitation Principle

This principle specifies that:

> *Collection of personal data should be limited to the minimum amount of data required for the purpose for which it is collected.*
>
> *Any such data should be obtained by lawful and fair means and, where appropriate, with the knowledge or consent of the data subject.*

In cloud there is no universally adopted privacy standard—instead, there are conflicting laws, regulations, and views on what privacy is and what it requires from organizations to protect it. Many organizations want to do what they perceive to be "the right thing"; however, their perception may be different from the law.

As a result, there may be different expectations regarding what privacy means between the organization and the CSP, and no agreed best practices.

It is essential that service-level agreements (SLAs) are initially defined before any information is provided or shared, because it is very hard to negotiate them later.

2. Use Limitation Principle

This principle specifies that:

> *Personal data should not be disclosed, made available or otherwise used for purposes other than those with the consent of the data subject, or by the authority of law.*

Cloud computing places a diverse collection of user and business information in a single location.

As data flows through the cloud, strong data governance is needed to ensure that the original purpose of collection and limitation on use is attached to the data. This is critical when organizations create a centralized database, because future applications can easily combine the data via expanded views that are utilized for new purposes never approved by data subjects.

The ability to combine data from multiple sources increases the risk of unexpected uses by governments. Governments in different countries could ask CSPs to report on particular types of behaviors or to monitor activities of particular types or categories of users. The possibility that a CSP could be obliged to inform a government or a third party about user activities might be troubling to the provider as well as to its users.

3. Security Principle

Security is one of the key requirements to enable privacy. This principle specifies that :

> *Personal data should be protected by reasonable security safeguards against such risks as loss or unauthorized access, destruction, use, modification, or disclosure of data.*

4. Retention and Destruction Principle

This principle specifies that:

> *Personal data should not be retained for longer than needed to perform the task for which it was collected, or as required by laws or regulations. Data should be destroyed in a secure way at the end of the retention period.*

How long data should be retained and when it should be destroyed is still a challenge for most companies.

Encryption can play a key role in the destruction process. Encrypted data can be destroyed even when organizations lose track of their data by destroying the encryption key—data can no longer be decrypted and hence is rendered inaccessible. This is especially beneficial when the data is kept by CSPs—encrypted data can be destroyed without the involvement of the CSPs.

The problem begins when there is a lack of clearly defined policies around data destruction in cloud computing. Virtual storage devices can be reallocated to new users without deleting the data, and then allocated to new users. Personal information stored in this device may now be available to the new user, potentially violating individual rights, laws, and regulations.

5. Transfer Principle

This principle specifies that:

> *Data should not be transferred to countries that don't provide the same level of privacy protection as the organization that collected the information*

In a cloud computing environment, infrastructure is shared between organizations; therefore, there are threats associated with the fact that the data is stored and processed remotely, and there is increased sharing of platforms between users, which increases the need to protect privacy of data stored in the cloud.

Another feature of cloud computing is that it is a dynamic environment; for example, service interactions can be created in a more dynamic way than in traditional e-commerce. Services can potentially be aggregated and changed dynamically by customers, and service providers can change the provisioning of services. In such scenarios, personal and sensitive data can move around within a single CSP infrastructure and across CSP organizational boundaries. The goal of integrated services provided by multiple CSPs is to enhance the possibility of data transfer to third parties. This transfer should be disclosed to the data subject prior to collection.

6. Accountability Principle

This principle states that:

> *An organization is responsible for personal information under its control and should designate an individual or individuals who are accountable for the organization's compliance with the remaining principles.*

Accountability within cloud computing can be achieved by attaching policies to data and mechanisms to ensure that these policies are adhered to by the parties that use, store, or share that data, irrespective of the jurisdiction in which the information is processed.