Privacy and Compliance

Chapter 12

Privacy and Compliance

Organizations are subject to legal requirements that govern how data is collected, stored, processed, and shared.

Unit Goals

- Provide an overview of common legal and compliance risks.
  - Identifying legal risks
  - Identifying privacy risks
  - Managing identity in the cloud

Identifying Legal Risks

- An organization cannot rely upon cloud service provider to ensure compliance with laws and regulations.
- Provider may have some responsibility as a data controller or custodian (depending upon the provider's role in processing data)
- Legal liability lies with the organization or individual owning the data.

Legal Considerations

Data location and jurisdiction

- Data in the cloud may be stored or processed in multiple data centers located anywhere in the world.
- Figure below illustrates a hypothetical situation in which the customer is located in North America, the cloud provider in Europe, and the provider's data centers in various locations around the world.

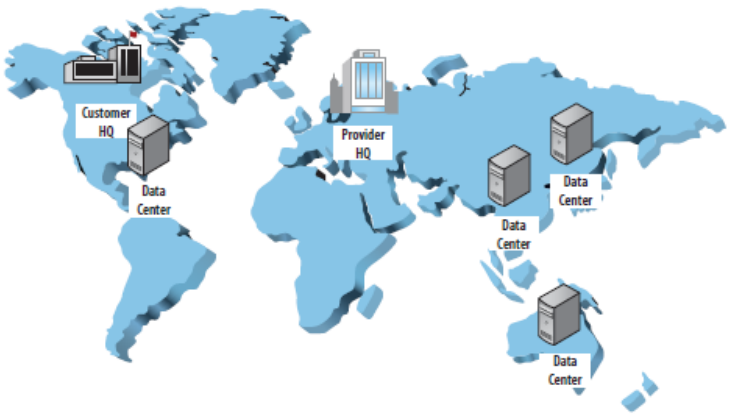

FIGURE 12.1 A map showing potential differences in geographic locations among customers, providers, and data centers

Legal Concerns

Having data in multiple locations can raise legal concerns.

- Data in the cloud could potentially be subject to the laws of:
  - The location of the physical servers
  - The location of the service provider's headquarters
  - The location of the data owner
  - The locations the data passes through between the provider's servers
- Risk can be mitigated by contractually obligating the service provider to keep data within appropriate geographic locations.

Data Isolation

- Data isolation may be required by regulation.
- In traditional computing, data can be isolated
  - Physically (e.g., a separate server) or
  - Logically (e.g., a separate virtual server, file store, or database).

In the cloud, multitenancy is common

- More difficult to ensure data isolation compliance.
- In a multitenant cloud environment, isolation is logical.

Data Isolation

Can occur at:

- Hypervisor level, by isolating virtual machines, or
- Database level, which may involve:
  - Isolation at the row level in a shared database is accomplished by uniquely identifying each tenant and associating each row with the owner's unique key.
  - Isolation at schema level is implemented by using a shared database with separate tables for each tenant.
- The greatest level of isolation occurs by providing tenants with individual databases.

Data Destruction

- What happens to data in the cloud after the contract between an organization and cloud service provider expires or is terminated.
- Organization must have assurances (via contract or terms of service) that its data will be deleted.

Bankruptcy

- If a cloud service provider files for bankruptcy, there is a risk that data may be exposed.
- Data may even be considered a corporate asset and sold, depending upon the provider's terms of service.

Terms of Service

- Cloud service providers do not always individually contract with customers.
- They may have published terms of services and privacy policies that apply to all customers.
- Even if the terms appear to be compatible with the data owner's needs, there is always the risk that the provider will change the terms of service, possibly without sufficient notification.
- May result in civil or even criminal liability for the data owner.
- Additionally, certain categories of information have specific legal requirements that may impact the use of cloud services.

Health Insurance Portability and Accountability Act (HIPAA) Example

Health information

- HIPAA primarily affects health care providers and health plans (covered entities);
- However, compliance is also required of business associates that have access to electronic protected health information (EPHI).
  - Would include cloud service providers.
- The covered entity and the cloud service provider are required to enter into a business associate agreement that defines the compliance obligations of each.

Privileged information

- Certain professionals, such as doctors and lawyers, have legal obligations to keep client information confidential,
- Laws vary by state and country.

Personally identifiable information (PII)

- Information can be used to uniquely identify an individual.
  - Types of data categorized as PII depend upon the jurisdiction
  - As do the security requirements and limits on use.

PII examples include:

- contact information
- financial information
- online account usernames
- government-issued identity documents (e.g., SSN, passport)
- biometric data.

Records Management

- Both public and private sector organizations can be subject to records retention requirements.
- Following conditions related to records management and retention can lead to risk:
  - Original metadata being associated with archived records
  - Provider-based record retention periods shorter than those required by the organization
  - Destruction of records after retention expires

Risks Related To Records Production

Lawful access and compelled disclosure

- Government agencies seeking information may choose to compel disclosure of information from service providers instead of directly from the data owners.
  - Detrimental to data owners because neither the provider nor the government is required to notify them.
  - Some countries have gag orders that prevent service providers from providing any notification at all.
- Table 12.1, next slide, provides examples of laws regulating governmental access.

TABLE 12.1 Examples of laws regulating governmental access

Law	Jurisdiction
Anti-Terrorism Act of 2001	Canada
Directive 2006/24/EC	European Union
USA PATRIOT Act	United States
Electronic Communications Privacy Act	United States
Convention on Cybercrime	International
Mutual legal assistance treaties	Various

Various Private litigation

- Information may also be compelled from service providers in private litigation.
- Carries similar risk but adds the additional risk of impacting compliance.
- While a data owner may be able to successfully resist or quash a subpoena, the service provider may have no obligation to try.

Electronic discovery (e-discovery)

- Electronically stored information (ESI) is subject to production in the discovery phase of litigation.
- In addition to data files and records, metadata is considered to be ESI and subject to discovery.

Organizations should consider the likelihood of litigation when selecting a cloud service provider to ensure that the provider's technical and business processes would not negatively impact the organization.

- For example, a cloud provider's archival process may not maintain the original metadata.

Risk Mitigation

- Mitigations for these risks include due diligence on the part of the organization and a service-level agreement (SLA) that takes the organization's compliance requirements into account.
- Additionally, organizations concerned with records retention should consider using records or document management software in the cloud instead of a more traditional file system.

Software Licensing

- Traditional software licensing models are not always compatible with cloud computing.

The three traditional software licensing models:

Per user

- Each user granted a license.
- May prove costly.

Per device

- Each device (or each processor in a device) granted a license.
- Some software vendors that use per-device license models may consider virtual servers to be the equivalent of physical servers
  - Does not solve the licensing problems that arise due to dynamic scaling.

Enterprise Licensing

- Organization granted a license, regardless of number of users or devices.
- Often, most cost effective
  - Particularly for large organizations
  - But these licenses may not translate to the cloud.
- For example, even if an organization has an enterprise license for a database product, a cloud service provider may still charge per instance for hosting the database.

Using the Right Cloud Services Provider

- Some software licensing risks can be eliminated by using the appropriate vendor.
- Software vendors that are also cloud service providers may have simplified licensing for customers that use their software in their cloud.
- Additionally, some service providers, such as Amazon, have partnered with major software vendors such as Adobe, Citrix, Microsoft, SUSE, and Oracle to provide a clear licensing structure to their customers.

Using the Right Cloud Services Provider

- Consequences of being out of compliance depend on software license agreements and applicable laws and regulations.
  - An application with built-in antipiracy protections may simply stop working if the maximum number of licenses reached.
  - Generally, underlicensing software is considered the same as piracy.
- Many major software vendors are members of the Business Software Alliance (BSA), an antipiracy watchdog group.
- Should a BSA software audit uncover licensing violations, an organization's liability could be up to US$150,000 per title plus additional fines.

Software Licensing and the Sarbanes-Oxley (SOX)

- Requires that publicly traded US companies have adequate internal controls that ensure reliable financial reporting.
- Because violations in software licensing can lead to large fines that can negatively impact financial statements, companies subject to SOX may have additional risk when moving to the cloud due to software license complexity.
- Where possible, investigate using cloud-friendly software licensing that supports:

Concurrency

- Licensing based on the number of users allowed to use the software at once can be more cost effective for organizations in which many users need access to an application but are unlikely to need to use it at the same time.

Licensing Considerations

Mobility

- In the cloud, applications and operating systems move between virtual environments, such as from host to host, data center to data center, and even from cloud to cloud.
- Flexibility
- Subscription or pay-as-you-go license models may be attractive for organizations using public cloud services that are not heavily invested in traditional software licenses, particularly for IaaS services.

Auto-scaling

- Number of servers may increase or decrease dynamically to provide sufficient quality of service and may overrun per-device or per-processor licenses.

Audit

- Most regulations impacting data security and privacy require periodic auditing for compliance and security.
  - Organizations may also choose to schedule their own internal audits.
- Should ensure that SLAs support these audit requirements
  - Must be diligent in monitoring and enforcing SLAs to avoid falling out of compliance.

Organizations subject to audit should consider the following:

- How will accounts for auditors be provisioned?
- Will appropriate audit logs be available?
  - How long can they be retained and how are they secured?
- What are the cloud service provider's policies on vulnerability management and security monitoring?
- Has the cloud service provider undergone an independent audit?

Identifying Privacy Risks

- Privacy requirements vary between geographic locations, not only from country to country but even from state to state.
- Data may become subject to the laws of countries in which the service provider or data center is located.
- Table 12.2 shows examples of privacy legislation in several countries.
- Because of the global nature of commerce, some efforts have been made to facilitate the transfer of data between countries.

TABLE 12.2 Examples of privacy legislation

Law	Jurisdiction	Applicability
Personal Information Protection and Electronic Documents Act (PIPEDA)	Canada	Protection of PII in commercial activities
EU Data Protection Directive	European Union	Processing of PII and movement of data between member states
UK Data Protection Act	United Kingdom	Protection of PII
Children's Online Privacy Protection Act (COPPA)	United States	Collection and use of children's personal information
Family Educational Rights and Privacy Act (FERPA)	United States	Student educational records
Gramm-Leach-Bliley Financial Services Modernization Act (GLBA)	United States	Consumer data related to financial products and services
Privacy Act of 1974	United States	Collection and use of PII by federal agencies
Video Privacy Protection Act	United States	Use of PII associated with video rentals
Swiss Federal Act on Data Protection	Switzerland	Processing of PII by private individuals and federal authorities

Safe Harbor

- The US-EU Safe Harbor Framework was created to allow the transfer of personal data between resources in the United States and the European Union, which have different restrictions on privacy.
- It allows individual US organizations to comply with the EU Directive on Data Protection.
- A similar framework exists between the United States and Switzerland to facilitate compliance with the Swiss Federal Act on Data Protection.
- According to the US Department of Commerce, in order to participate, US organizations must comply with the seven Safe Harbor principles:
- Notice Individuals must be notified about the information collected, used, and disclosed

Opt Out

- Choice Individuals must be provided with the opportunity to opt out of having their personal information disclosed to a third party.

Transfer to third parties

- Organizations must ensure that the third party subscribes to the Safe Harbor principles or is compliant with the EU Directive on Data Protection.

Access

- With some exceptions, individuals must be allowed to access and manage their personal information.

Security

- Reasonable protections must be implemented to protect personal information.
- Data integrity
- Data must be reliable and accurate.

Enforcement

- Certification must be maintained annually to remain in the program, and there must be mechanisms in place both to effectively handle complaints and violations and to verify compliance.

Managing Identity in the Cloud

- The purpose of identity management is to manage the life cycle of users and other entities that need trusted access to organizational resources.
- Identity management also goes hand in hand with privacy, and identity records for users generally contain PII that may be subject to privacy regulations.
- Prior to discussing the characteristics of identity management systems, it is necessary to understand the three main elements of identity and access control:

Authentication

- Authentication is the process of verifying an entity's identity by validating one or more factors:
  - something you know
  - something you have
  - something you are.

Authentication

- A user ID–password combination (something you know) is currently the most widely used form of authentication.
- Other forms include security tokens or smart cards (something you have) and biometrics (something you are).

Authorization

- Authorization is the process of determining whether an entity is allowed to access a resource and with what level of permissions based on access control lists.

Role-Based Access Control

- Using role-based access control (RBAC) is an effective way of managing access for a large number of users.
- Instead of being assigned permissions directly, users are assigned to role-based groups and permissions are managed at the group level.

Accounting

- Accounting is the process of tracking resource usage for operational, security, and compliance purposes.
- Operationally, accounting can be used for capacity monitoring and billing.
- Monitoring of access logs and the ability to generate audit trails are often required by security policy and regulation.
- An organization should consider its security, privacy, and compliance needs when evaluating an identity management system.
- One of the primary characteristics a system should support is the ability to assign users to roles to support separation of duties, association of users with business roles, and role-based access controls.
- Additionally, organizations should consider requirements such as self-service functions (e.g., password reset, user data update) and access to user data.

Managing identity in the cloud

Identity provisioning

- Identity provisioning is the process of creating and deactivating user accounts (deactivation may also be called deprovisioning).
- In IaaS and SaaS deployments, service providers may have proprietary provisioning processes that may add complexity to business processes, particularly if each offering an organization uses has different methods of provisioning.

Federated Identity Management

- In discussing federation, we refer to service providers and identity providers.
- A service provider is an application or service, and an identity provider is an authentication authority.
- An organization may be its own identity provider (e.g., via the organization's directory services) or it may use an external source (e.g., OpenID, Google, icrosoft Windows Live ID).
- Federation allows users in different security domains to share services without having identities in each domain.
- Identity providers provide information (i.e., identity attributes) to service providers, taking the burden of authentication off of individual service providers and placing it with a trusted identity provider.

Single Sign-On

- Implementing SSO allows an organization's users to authenticate once and access multiple applications, as shown in Figure 12.2.
- This improves efficiency by streamlining the authentication process, reduces IT overhead by reducing account administration duties, and improves security by requiring the user to remember only one password. (Increasing the number of passwords a user must remember increases the likelihood that the user will write them down.)
- SSO can be configured using Kerberos in both Windows and Unix/Linux environments, using smart cards, and through standards such as OpenID, Security Assertion Markup Language (SAML), and Web Services Federation Language (WS-Federation).

Web Application

- F I G U R E 1 2 . 2 Simple SSO diagram showing a user authenticating to a SSO server and accessing both email and web applications

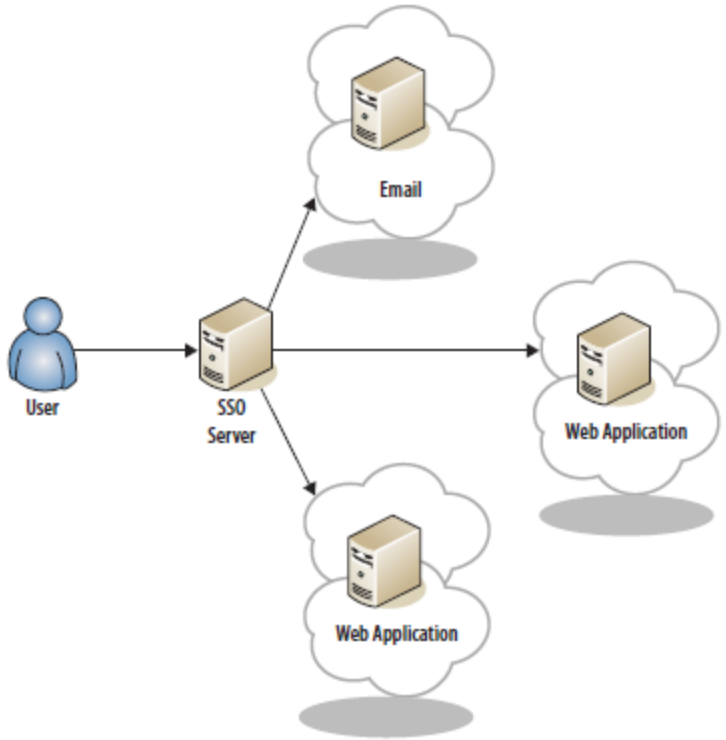

FIGURE 12.2 Simple SSO diagram showing a user authenticating to a SSO server and accessing both email and web applications

Credential management

- Many security standards and data protection laws have requirements for credential management, particularly user accounts and passwords.
- An organization must ensure that its compliance needs are met for requirements such as secure transmission of passwords, strong password policies, password storage, and self-service password reset.
- Complexity may be reduced through the use of federated identity management and single sign-on (SSO).
- When choosing a cloud service provider, organizations should consider their existing environment and standards supported by potential vendors.

Questions???