UNIT- V

SECURITY IN THE CLOUD

Topics Covered:

- **Security Overview**
- **Cloud Security Challenges and Risks**
- **Software-as-a-Service Security**
- **Security Governance**
- **Risk Management**
- **Security Monitoring**
- **Security Architecture Design**
- **Data Security**
- **Application Security**
- **Virtual Machine Security**
- **Identity Management and Access Control**
- **Autonomic Security**

⇒ Cloud Security Overview

→ Security is a principal concern when ***entrusting*** an organization's critical information to geographically dispersed cloud platforms not under the direct control of that organization.

- **Information Life Cycle Management —** "Understand cloud provider policies and processes for *data retention* and *destruction* and how they compare with internal organizational policy. Be aware that data retention assurance may be easier for the cloud provider to demonstrate, but data destruction may be very difficult. Perform regular backup and recovery tests to assure that logical segregation and controls are effective."

- **Application Security —** "IaaS, PaaS and SaaS create differing trust boundaries for the software development lifecycle, which must be accounted for during the development, testing and production deployment of applications."

- **Storage —** "Understand cloud provider storage retirement processes. Data destruction is extremely difficult in a multi-tenant environment and the cloud provider should be utilizing strong storage encryption that renders data unreadable when storage is recycled, disposed of, or accessed by any means outside of authorized applications."

→ With cloud computing providing **SaaS**, secure software is a critical issue. From the cloud consumer's point of view, using SaaS in the cloud reduces the need for secure software development by the customer. The requirement for secure software development is transferred to the cloud provider.

Cloud Information Security Objectives

The Data and Analysis Center for Software (DACS) requires that software must exhibit the following three properties to be considered secure:

- **Dependability** — Software that executes predictably and operates correctly under a variety of conditions, including when under attack or running on a malicious host.

- **Trustworthiness** — Software that contains a minimum number of vulnerabilities or no vulnerabilities or weaknesses that could sabotage the software's dependability. It must also be resistant to malicious logic.

- **Survivability (Resilience)** — Software that is resistant to or tolerant of attacks and has the ability to recover as quickly as possible with as little harm as possible.

Seven complementary principles that support information assurance are **confidentiality, integrity, availability, authentication, authorization, auditing, and accountability.**

Confidentiality, Integrity, and Availability

Confidentiality, integrity, and availability are sometimes known as the *CIA triad* of information system security, and are important pillars of cloud software assurance.

Confidentiality

Confidentiality refers to the prevention of intentional or unintentional unauthorized **disclosure of information.** Confidentiality in cloud systems is related to the areas of intellectual property rights, covert channels, traffic analysis, encryption, and inference:

- **Intellectual property rights** — Intellectual property (IP) includes inventions, designs, and artistic, musical, and literary works. Rights to intellectual property are covered by copyright laws, which protect creations of the mind, and patents, which are granted for new inventions.
- **Covert channels** — A *covert channel* is an unauthorized and unintended communication path that enables the exchange of information. Covert channels can be accomplished through timing of messages or inappropriate use of storage mechanisms.
- **Traffic analysis** — *Traffic analysis* is a form of confidentiality breach that can be accomplished by analyzing the volume, rate, source, and destination of message traffic, even if it is encrypted. Increased message activity and high bursts of traffic can indicate a major event is occurring. Countermeasures to traffic analysis include maintaining a near-constant rate of message traffic and disguising the source and destination locations of the traffic.
- **Encryption** — *Encryption* involves scrambling messages so that they cannot be read by an unauthorized entity, even if they are intercepted. The amount of effort (*work factor*) required to decrypt the message is a function of the strength of the encryption key and the robustness and quality of the encryption algorithm.
- **Inference** — *Inference* is usually associated with database security. Inference is the ability of an entity to use and correlate information protected at one level of security to uncover information that is protected at a higher security level.

Integrity

The concept of cloud information *integrity* requires that the following three principles are met:

- Modifications are not made to data by unauthorized personnel or processes.
- **Unauthorized modifications** are not made to **data** by authorized personnel or processes.
- The data is internally and externally consistent — in other words, the internal information is consistent both among all sub-entities and with the real-world, external situation.

Availability

Availability ensures the **reliable and timely access to cloud data** or cloud computing resources by the appropriate personnel. Availability guarantees that the systems are functioning properly when needed. In

addition, this concept guarantees that the security services of the cloud system are in working order. A denial-of-service attack is an example of a threat against availability.

The reverse of confidentiality, integrity, and availability is disclosure, alteration, and destruction (DAD).

Cloud Security Services

Additional factors that directly affect cloud software assurance include authentication, authorization, auditing, and accountability.

Authentication

Authentication is the testing or reconciliation of evidence of a **user's identity**. It establishes the user's identity and ensures that users are who they claim to be. For example, a user presents an identity (user ID) to a computer login screen and then has to provide a password. The computer system authenticates the user by verifying that the password corresponds to the individual presenting the ID.

Authorization

Authorization refers to **rights** and **privileges** granted to an individual or process that enable access to computer resources and information assets. Once a user's identity and authentication are established, authorization levels determine the extent of system rights a user can hold.

Auditing

To maintain operational assurance, organizations use two basic methods: system audits and monitoring. These methods can be employed by the cloud customer, the cloud provider, or both, depending on asset architecture and deployment.

- A ***system audit*** is a one-time or periodic event to evaluate security.
- ***Monitoring*** refers to an ongoing activity that examines either the system or the users, such as intrusion detection.

IT auditors typically audit the following functions:

- System and transaction controls
- Systems development standards
- Backup controls
- Data library procedures
- Data center security
- Contingency plans

Audit logs should record the following:

- The transaction's date and time
- Who processed the transaction
- At which terminal the transaction was processed
- Various security events relating to the transaction

In addition, an auditor should examine the audit logs for the following:

- Amendments to production jobs
- Production job reruns
- Computer operator practices
- All commands directly initiated by the user
- All identification and authentication attempts
- Files and resources accessed

Accountability

Accountability is the ability to determine the actions and behaviors of a single individual within a cloud system and to identify that particular individual. Accountability is related to the concept of *non repudiation*, where in an individual cannot successfully deny the performance of an action.

Relevant Cloud Security Design Principles

11 security design principles:

Least privilege

The principle of *least privilege* maintains that an individual, process, or other type of entity should be given the minimum privileges and resources for the minimum period of time required to complete a task.

Separation of duties

Separation of duties requires that completion of a specified sensitive activity or access to sensitive objects is dependent on the satisfaction of a plurality of conditions.

Defense in depth

Defense in depth is the application of multiple layers of protection wherein a subsequent layer will provide protection if a previous layer is breached.

Fail safe

Fail safe means that if a cloud system fails it should fail to a state in which the security of the system and its data are not compromised.

Economy of mechanism

Economy of mechanism promotes simple and comprehensible design and implementation of protection mechanisms, so that unintended access paths do not exist or can be readily identifyfied and eliminated.

Complete mediation

In complete meditation, every request by a subject to access an object in a computer system must undergo a valid and effective authorization procedure.

Open design

There has always been an ongoing discussion about the merits and strengths of security designs that are kept secret versus designs that are open to scrutiny and evaluation by the community at large.

Least common mechanism

This principle states that a minimum number of protection mechanisms should be common to multiple users, as shared access paths can be sources of unauthorized information exchange.

Psychological acceptability

Psychological acceptability refers to the ease of use and intuitiveness of the user interface that controls and interacts with the cloud access control mechanisms.

Weakest link

As in the old saying “A chain is only as strong as its weakest link,” the security of a cloud system is only as good as its weakest component.

Leveraging existing components

To increase cloud system security by leveraging existing components is to partition the system into defended subunits.
Then, if a security mechanism is penetrated for one sub-unit, it will not affect the other sub-units, and damage to the computing resources will be minimized.

Secure Cloud Software Requirements

→ *Software requirements engineering* is the process of determining customer software expectations and needs, and it is conducted before the software design phase.

→ Department of Defense Data and Analysis Center for Software (DACS) state that all software shares the following three security needs:

- ✓ It must be **dependable** under anticipated operating conditions, and remain dependable under hostile operating conditions.
- ✓ It must be **trustworthy** in its own behavior, and in its inability to be compromised by an attacker through exploitation of vulnerabilities or insertion of malicious code.
- ✓ It must be **resilient** enough to recover quickly to full operational capability with a minimum of damage to itself, the resources and data it handles, and the external components with which it interacts.

NIST Security Principles National Institute of Standards and Technology's

Principle 1 — Establish a sound security policy as the "foundation" for design.
Principle 2 — Treat security as an integral part of the overall system design.
Principle 3 — clearly delineates the physical and logical security boundaries governed by associated security policies.
Principle 6 — Assume that external systems are insecure.
Principle 7 — Identify potential trade-offs between reducing risk and increased costs and decreases in other aspects of operational effectiveness.
Principle 8 — Implement layered security; ensure there is no single point of vulnerability.
Principle 9 — Isolate public access systems from mission-critical resources (e.g., data, processes, etc.).
Principle 10 — Use boundary mechanisms to separate computing systems and network infrastructures.
Principle 11 — Minimize the system elements to be trusted.
Principle 12 — Implement least privilege.
Principle 13 — Authenticate users and processes to ensure appropriate access control decisions both within and across domains.
Principle 14 — Use unique identities to ensure accountability.

⇒ Cloud Security Challenges and Risks

Cloud computing presents an organization with its own set of security issues.

❖ **Confidentiality**

Confidentiality is the prevention of the intentional or unintentional unauthorized disclosure of contents. Loss of confidentiality can occur in many ways. For example, loss of confidentiality can occur through the intentional release of private company information or through a misapplication of network rights.

Some of the elements of telecommunications used to ensure confidentiality are as follows:

- Network security protocols
- Network authentication services
- Data encryption services

❖ **Integrity**

Integrity is the guarantee that the message sent is the message received and that the message is not intentionally or unintentionally altered. Loss of integrity can occur through an intentional attack to change information (for example, a website defacement) or, more commonly, unintentionally (data is accidentally

altered by an operator). Integrity also contains the concept of non repudiation of a message source, which we will describe later.

Some of the elements used to ensure integrity include the following:

- Firewall services
- Communications security management
- Intrusion detection services

❖ **Availability**

This concept refers to the elements that create reliability and stability in networks and systems. It ensures that connectivity is accessible when needed, allowing authorized users to access the network or systems.

Some of the elements that are used to ensure availability are as follows:

- Fault tolerance for data availability, such as backups and redundant disk systems
- Acceptable logins and operating process performance
- Reliable and interoperable security processes and network security mechanisms

There are also several other important concepts and terms that apply equally well to traditional IT computing and cloud computing:

Identification — The means by which users claim their identities to a system. Most commonly used for access control, identification is necessary for authentication and authorization.

Authentication — The testing or reconciliation of evidence of a user's identity. It establishes the user's identity and ensures that the users are who they say they are.

Accountability — A system's capability to determine the actions and behaviors of a single individual within a system and to identify that particular individual. Audit trails and logs support accountability.

Authorization — The rights and permissions granted to an individual or process that enable access to a computer resource. Once a user's identity and authentication are established, authorization levels determine the extent of a user's system rights.

Privacy — The level of confidentiality and privacy protection given to a user in a system. This is often an important component of security controls. Privacy not only guarantees the fundamental tenet of confidentiality of company data, but also guarantees the data's level of privacy, which is being used by the operator. We examine privacy risks in more detail in the following section.

Privacy and Compliance Risks

Privacy laws attempt to provide protection to an individual from unauthorized disclosure of the individual's personally identifiable information (PII).

- Names
- Postal address information, other than town or city, state, and zip code
- Telephone numbers
- Fax numbers
- Electronic mail addresses
- Social security numbers
- Medical record numbers
- Health plan benefi ciary numbers
- Account numbers
- Certifi cate/license numbers

- Vehicle identifi ers and serial numbers, including license plate numbers
- Device identifi ers and serial numbers
- Web Universal Resource Locators (URLs)
- Internet Protocol (IP) address numbers
- Biometric identifi ers, including fi ngerprints and voiceprints
- Full face photographic images and any comparable images

Threats to Infrastructure, Data, and Access Control

This includes the following:

- Communications and network security as it relates to voice, data, multimedia, and facsimile transmissions in terms of local area, wide area, and remote access networks
- Internet/intranet/extranet in terms of firewalls, routers, gateways, and various protocols

- A threat is simply any event that, if realized, can cause damage to a system and create a loss of confidentiality, availability, or integrity.
- A *vulnerability* is a weakness in a system that can be exploited by a threat. Reducing the vulnerable aspects of a system can reduce the risk and impact of threats on the system.

Common threats to both cloud and traditional infrastructure include the following:

- **Eavesdropping** — Data scavenging, traffic or trend analysis, social engineering, economic or political espionage, sniffing, dumpster diving, keystroke monitoring, and shoulder surfing are all types of eavesdropping to gain information or to create a foundation for a later attack. Eavesdropping is a primary cause of the failure of confidentiality.
- **Fraud** — Examples of fraud include collusion, falsified transactions, data manipulation, and other altering of data integrity for gain.
- **Theft** — Examples of theft include the theft of information or trade secrets for profit or unauthorized disclosure, and physical theft of hardware or software.
- **Sabotage** — Sabotage includes denial-of-service (DoS) attacks, production delays, and data integrity sabotage.
- **External attack** — Examples of external attacks include malicious cracking, scanning, and probing to gain infrastructure information, demon dialing to locate an unsecured modem line, and the insertion of a malicious code or virus.

The most common types of attacks are,

- ***Inappropriate System Use***
- ***Eavesdropping***
- ***Network Intrusion***
- ***Denial-of-Service (DoS) Attacks***
- ***Session Hijacking Attacks***
- ***Fragmentation Attacks***

It is important for the information security professional to understand and identify other types of attacks

- Back-Door
- Spoofing
- Man-in-the-Middle
- TCP Hijacking
- Trojan Horses and Malware

Cloud Computing Security Challenges

The introduction of cloud services presents many challenges to an organization. When an organization migrates to consuming cloud services, and especially public cloud services, much of the computing system infrastructure will now be under the control of a third-party Cloud Services Provider (CSP).

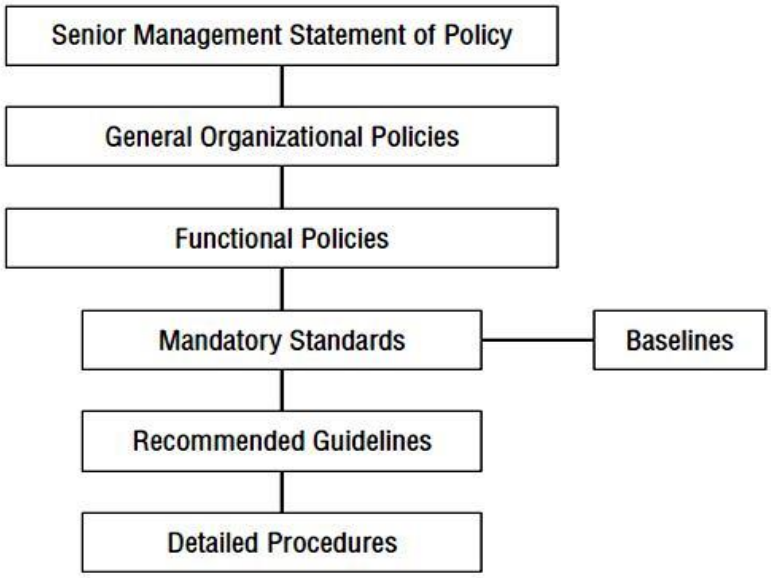

Figure 5-1: Security policy hierarchy

The main tasks of a CSIRT (computer security incident response team) are as follows:

- Analysis of an event notifi cation
- Response to an incident if the analysis warrants it
- Escalation path procedures
- Resolution, post-incident follow-up, and reporting to the appropriate parties

Five immutable laws of virtualization security must be understood and used to drive security decisions:

Law 1: All existing OS-level attacks work in the exact same way.
Law 2: The hypervisor attack surface is additive to a system's risk profile.
Law 3: Separating functionality and/or content into VMs will reduce risk.
Law 4: Aggregating functions and resources onto a physical platform will increase risk.
Law 5: A system containing a "trusted" VM on an "un trusted" host has a higher risk level than a system containing a "trusted" host with an "un trusted" VM.

VM Security Recommendations

The following security implementation techniques are required for most computer systems, and are still best practices for virtualized systems. These areas include physical security, patching, and remote management techniques.

Hardening the Host Operating System

Some of these techniques include the following:

- Use strong passwords, such as lengthy, hard to guess passwords with letters, numbers, and symbol combinations, and change them often.
- Disable unneeded services or programs, especially networked services.

- Require full authentication for access control.
- The host should be individually firewalled.
- Patch and update the host regularly, after testing on a nonproduction unit.

Limiting Physical Access to the Ho

When attackers can access a host they can do the following:

- Use OS-Specific keystrokes to kill processes, monitor resource usage, or shut down the machine, commonly without needing a valid login account and password
- Reboot the machine, booting to external media with a known root password
- Steal files using external media (floppy, CD/DVD-RW, USB/flash drives, etc.)
- Capture traffic coming into or out of the network interfaces
- Remove one or more disks, mounting them in a machine with a known administrator or root password, potentially providing access to the entire contents of the host and guest VMs
- Simply remove the entire machine

Using Encrypted Communications

Encryption technologies, such as Secure HTTP (HTTPS), encrypted Virtual Private Networks (VPNs), Transport Layer Security (TLS), Secure Shell (SSH), and so on should be used to provide secure communications links between the host domain and the guest domain, or from hosts to management systems.

- Disabling Background Tasks
- Updating and Patching
- Enabling Perimeter Defense on the VM
- Implementing File Integrity Checks
- Maintaining Backups

VM-Specific Security Techniques

A fundamental requirement for a successful virtualization security process is recognizing the dynamic nature of virtual machines.

Hardening the Virtual Machine

This hardening can include many steps, such as the following:

- Putting limits on virtual machine resource consumption
- Configuring the virtual network interface and storage appropriately
- Disabling or removing unnecessary devices and services
- Ensuring that components that might be shared across virtual network devices are adequately isolated and secured
- Keeping granular and detailed audit logging trails for the virtualized infrastructure.

Harden the Hypervisor
Root Secure the Monitor
Implement Only One Primary Function per VM
Firewall Any Additional VM Ports
Harden the Host Domain

These include but are not limited to:

- Remove unnecessary accounts and groups.
- Disable unnecessary services.
- Remove unnecessary binaries, libraries, and fi les.
- Firewall network access to the host.
- Install monitoring or Host Intrusion Detection Systems.
- Ensure that the Host Domain is not accessible from the Guest Domains.
- Ensure that monitoring or remote console interfaces for the Host Domain are not accessible via the Guest Domains.
- Ensure that the Guest Domains cannot directly affect any network storage or other resources that the Host Domain relies on for boot, configuration, or authentication.

Use Unique NICs (network interfaces) **for Sensitive VMs Disconnect Unused Devices**

Securing VM Remote Access

Standard practices for remote administration include the following:

- Strong authentication practices should be employed:
- Two-factor authentication
- Strong passwords ○

One-time passwords

- Private/public PKI key pairs
- Use encrypted communications only, such as a SSH or VPNs.
- MAC address or IP address filtering should be employed.
- Telnet access to the unit should be denied, as it does not encrypt the communications channel.

⇒ Cloud Computing Security Architecture Design

✓ The security posture of a cloud system is based on its security architecture. While there is no standard definition for security architecture, the Open Security Alliance (OSA) defines *security architecture* as "the design artifacts that describe how the security controls are positioned, and how they relate to the overall IT Architecture. These controls serve the purpose to maintain the system's quality attributes, among them confidentiality, integrity, availability, accountability and assurance".

✓ A *security architecture* as "a cohesive security design, which addresses the requirements (e.g., authentication, authorization, etc.) and in particular the risks of a particular environment/scenario, and specifies what security controls are to be applied where.

Architectural Considerations

General Issues: A variety of topics influence and directly affect the cloud security architecture. They include such factors as compliance, security management, administrative issues, controls, and security awareness.

1) **Compliance -** The cloud vendor should provide transparency to the client by supplying information about storage used, processing characteristics, and other relevant account information. The accessibility of a client's data by the provider's system engineers and certain other employees.

2) **Security Management -** Security architecture involves effective security management to realize the benefits of cloud computation. Proper cloud security management and administration should identify management issues in critical areas such as access control, vulnerability analysis, change control, incident response, fault tolerance, and disaster recovery and business continuity planning.
3) **Controls -** The objective of cloud security controls is to reduce vulnerabilities to a tolerable level and minimize the effects of an attack. To achieve this, an organization must determine what impact an attack might have, and the likelihood of loss.

- **Deterrent controls** —reduce the likelihood of a deliberate attack.
- **Preventative controls** — Protect vulnerabilities and make an attack unsuccessful or reduce its impact. Preventative controls inhibit attempts to violate security policy.
- **Corrective controls** — reduce the effect of an attack.
- **Detective controls** — Discover attacks and trigger preventative or corrective controls. Detective controls warn of violations or attempted violations of security policy and include such controls as intrusion detection systems, organizational policies, video cameras, and motion detectors.

4) **Complementary Actions -**Additional activities involved in cloud security management include the following:

- Management and monitoring of service levels and service-level agreements.
- Acquisition of adequate data to identify and analyze problem situations through instrumentation and dashboards.
- Reduction of the loss of critical information caused by lack of controls.
- Proper management of data on an organization's distributed computing resources. Data centralized on the cloud reduces the potential for data loss in organizations with large numbers of laptop computers and other personal computing devices.
- Monitoring of centrally stored cloud information, as opposed to having to examine data distributed throughout an organization on a variety of computing and storage devices.
- Provisioning for rapid recovery from problem situations.

DATA SECURITY- Information Classification

The following classification terms are typical of those used in the private sector and are applicable to cloud data**:**

Public data — Information that is similar to unclassified information; all of a company's information that does not fi t into any of the next categories can be considered public. While it's unauthorized disclosure may be against policy, it is not expected to impact seriously or adversely the organization, its employees, and/or its customers.

Sensitive data — Information that requires a higher level of classification than normal data. This information is protected from a loss of confidentiality as well as from a loss of integrity due to an unauthorized alteration. This classification applies to information that requires special precautions to ensure its integrity by protecting it from unauthorized modification or deletion. It is information that requires a higher-than-normal assurance of accuracy and completeness.

Private data — This classification applies to personal information that is intended for use within the organization. Its unauthorized disclosure could seriously and adversely impact the organization and/or its employees. For example, salary levels and medical information are considered private.

Confidential data — This classification applies to the most sensitive business information that is intended strictly for use within the organization. Its unauthorized disclosure could seriously and adversely impact the organization, its stockholders, its business partners, and/or its customers.

❖ Classification Criteria

Several criteria may be used to determine the classification of an information object:

Value — Value is the number one commonly used criteria for classifying data in the private sector. If the information is valuable to an organization or its competitors, then it needs to be classified.

Age — The classification of information might be lowered if the information's value decreases over time. In the U.S. Department of Defense, some classified documents are automatically declassified after a predetermined time period has passed.

Useful life — If the information has been made obsolete due to new information, substantial changes in the company, or other reasons, the information can often be declassified.

Personal association — If information is personally associated with specific individuals or is addressed by a privacy law, it might need to be classified. For example, investigative information that reveals informant names might need to remain classified.

❖ Information Classification Procedures

There are several steps in establishing a classification system. These are the steps in priority order:

- ✓ Identify the appropriate administrator and data custodian. The data custodian is responsible for protecting the information, running backups, and performing data restoration.
- ✓ Specify the criteria for classifying and labeling the information.
- ✓ Classify the data by its owner, who is subject to review by a supervisor.
- ✓ Specify and document any exceptions to the classification policy.
- ✓ Specify the controls that will be applied to each classification level.
- ✓ Specify the termination procedures for declassifying the information or for transferring custody of the information to another entity.
- ✓ Create an enterprise awareness program about the classification controls.

Employee Termination

It is important to understand the impact of employee terminations on the integrity of information stored in a cloud environment.

- ✘ The removal of access privileges, computer accounts, authentication tokens.
- ✘ The briefing on the continuing responsibilities of the terminated employee for confidentiality and privacy.
- ✘ The return of company computing property, such as laptops.
- ✘ Continued availability of data. In both the manual and the electronic worlds, this may involve documenting procedures or filing schemes, such as how documents are stored on the hard disk and how they are backed up.
- ✘ Employees should be instructed whether or not to "clean up" their PC before leaving.
- ✘ If cryptography is used to protect data, the availability of cryptographic keys to management personnel must be ensured.

Security Awareness, Training, and Education

Security awareness is often overlooked as element affecting cloud security architecture because most of a security practitioner's time is spent on controls, intrusion detection, risk assessment, and proactively or reactively administering security.

- Improving awareness of the need to protect system resources
- Developing skills and knowledge so computer users can perform their jobs more securely
- Building in-depth knowledge, as needed, to design, implement, or operate security programs for organizations and systems.

In general, a computer security awareness and training program should encompass the following seven steps:

1. Identify program scope, goals, and objectives.
2. Identify training staff.
3. Identify target audiences.
4. Motivate management and employees.
5. Administer the program.
6. Maintain the program.
7. Evaluate the program.

Training and Education

Training is different from awareness in that it provides security information in a more formalized manner, such as classes, workshops, or individualized instruction. The following types of training are related to cloud security:

- Security-related job training for operators and specific users
- Awareness training for specific departments or personnel groups with security-sensitive positions
- Technical security training for IT support personnel and system administrators
- Advanced training for security practitioners and information systems auditors
- Security training for senior managers, functional managers, and business unit managers

Trusted Cloud Computing

Trusted cloud computing can be viewed as a computer security architecture that is designed to protect cloud systems from malicious intrusions and attacks, and ensure that computing resources will act in a specific, predictable manner as intended. A trusted cloud computing system will protect data in use by hypervisors and applications, protect against unauthorized access to information, provide for strong authentication, apply encryption to protect sensitive data that resides on stolen or lost devices, and support compliance through hardware and software mechanisms.

Secure Execution Environments and Communications

In a cloud environment, applications are run on different servers in a distributed mode. These applications interact with the outside world and other applications and may contain sensitive information whose inappropriate access would be harmful to a client.

- **Secure Execution Environment**

Configuring computing platforms for secure execution is a complex task; and in many instances it is not performed properly because of the large number of parameters that are involved. Another major concern in secure execution of code is the widespread use of "unsafe" programming languages such as C and C++ instead of more secure languages such as object-oriented Java and structured, object-oriented C#.

- **Secure Communications :** As opposed to having managed, secure communications among the computing resources internal to an organization, movement of applications to the cloud requires a reevaluation of communications security. These communications apply to both data in motion and data at rest.

- **Confidentiality** — Ensures that only those who are supposed to access data can retrieve it. Loss of confidentiality can occur through the intentional release of private company information or through a misapplication of network rights. Some of the elements of telecommunications used to ensure confidentiality are as follows:

  Network security protocols
  Network authentication services
  Data encryption services

- **Integrity** — Ensures that data has not been changed due to an accident or malice. Integrity is the guarantee that the message sent is the message received and that the message is not intentionally or unintentionally altered. Integrity also contains the concept of non repudiation of a message source. Some of the constituents of integrity are as follows:

  Firewall services
  Communications Security Management
  Intrusion detection services

- **Availability** — Ensures that data is accessible when and where it is needed, and that connectivity is accessible when needed, allowing authorized users to access the network or systems. Also included in that assurance is the guarantee that security services for the security practitioner are usable when they are needed. Some of the elements that are used to ensure availability are as follows:

  Fault tolerance for data availability, such as backups and redundant disk
  systems Acceptable logins and operating process performances
  Reliable and interoperable security processes and network security mechanisms

APIs(Application Programming Interfaces)

Common vulnerabilities such as weak antivirus software, unattended computing platforms, poor passwords, weak authentication mechanisms, and inadequate intrusion detection that can impact communications must be more stringently analyzed, and proper APIs must be used.

For example, in using IaaS, a cloud client typically communicates with cloud server instances through Representational State Transfer (REST) client/server model or Simple Object Access Protocol (SOAP) APIs.

Virtual Private Networks

Another important method to secure cloud communications is through a virtual private network (VPN). A VPN is created by building a secure communications link between two nodes by emulating the properties of a point-to-point private link. A VPN can be used to facilitate secure remote access into the cloud, securely connect two networks together, or create a secure data tunnel within a network.

Remote Access VPNs

A VPN can be configured to provide remote access to corporate resources over the public Internet to maintain confidentiality and integrity.

Network-to-Network VPNs

A VPN is commonly used to connect two networks, perhaps the main corporate LAN and a remote branch office LAN, through the Internet. This connection can use either dedicated lines to the Internet or dial-up connections to the Internet.

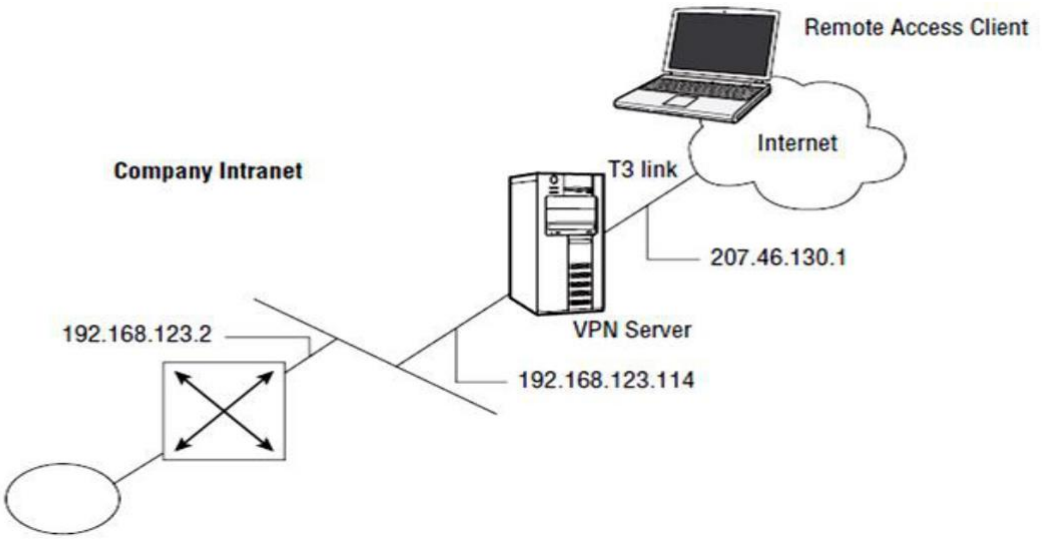

Figure 6-2: Remote access VPN configuration

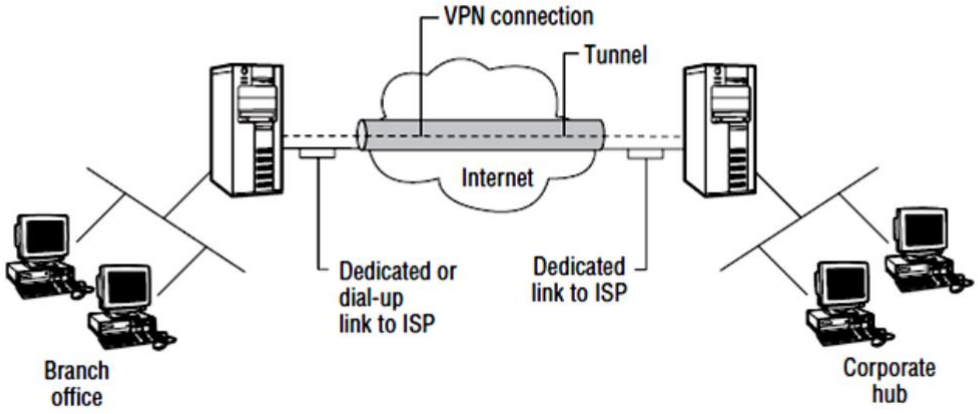

Figure 6-3: A network-to-network VPN configuration

Public Key Infrastructure and Encryption Key Management

To secure communications, data that is being exchanged with a cloud should be encrypted, calls to remote servers should be examined for imbedded malware, and digital certificates should be employed and managed.

- Digital certificates
- Certificate authority (CA)
- Registration authorities
- Policies and procedures
- Certificate revocation
- Non repudiation support
- Time stamping
- Lightweight Directory Access Protocol (LDAP)
- Security-enabled applications

Identity Management and Access Control

Identity management and access control are fundamental functions required for secure cloud computing. The simplest form of identity management is logging on to a computer system with a user ID and password.

True identity management, such as is required for cloud computing, requires more robust authentication, authorization, and access control. It should determine what resources are authorized to be accessed by a user or process by using technology such as biometrics or smart cards, and determine when a resource has been accessed by unauthorized entities.

Identification and authentication are the keystones of most access control systems. Identification is the act of a user professing an identity to a system, usually in the form of a username or user logon ID to the system.

Authentication is verification that the user's claimed identity is valid, and it is usually implemented through a user password at logon. Authentication is based on the following three factor types:

- **Type 1** — Something you know, such as a personal identification number (PIN) or password
- **Type 2** — Something you have, such as an ATM card or smart card
- **Type 3** — Something you are (physically), such as a fingerprint or retina scan

Passwords

Because passwords can be compromised, they must be protected. In the ideal case, a password should be used only once. This "one-time password," or OTP, provides maximum security because a new password is required for each new logon. A password that is the same for each logon is called a *static password*. A password that changes with each logon is termed a *dynamic password*.

Tokens

Tokens, in the form of small, hand-held devices, are used to provide passwords.

Memory Cards

Memory cards provide nonvolatile storage of information, but they do not have any processing capability. A memory card stores encrypted passwords and other related identifying information. A telephone calling card and an ATM card are examples of memory cards.

Smart Cards

Smart cards provide even more capability than memory cards by incorporating additional processing power on the cards. These credit-card-size devices comprise microprocessor and memory and are used to store digital signatures, private keys, passwords, and other personal information.

Biometrics

An alternative to using passwords for authentication in logical or technical access control is *biometrics*. Biometrics is based on the Type 3 authentication mechanism — something you are. Biometrics is defined as an automated means of identifying or authenticating the identity of a living person based on physiological or behavioral characteristics.

The following are typical biometric characteristics that are used to uniquely authenticate an individual's identity:

- **Fingerprints** Fingerprint characteristics are captured and stored. Typical CERs are 4–5%.
- **Retina scans** The eye is placed approximately two inches from a camera and an invisible light source scans the retina for blood vessel patterns. CERs are approximately 1.4%.

- **Iris scans** A video camera remotely captures iris patterns and characteristics. CER values are around 0.5%.
- **Hand geometry** Cameras capture three-dimensional hand characteristics. CERs are approximately 2%.
- **Voice** — Sensors capture voice characteristics, including throat vibrations and air pressure, when the subject speaks a phrase. CERs are in the range of 10%.
- **Handwritten signature dynamics** The signing characteristics of an individual making a signature are captured and recorded. Typical characteristics including writing pressure and pen direction. CERs are not published at this time.

Other types of biometric characteristics include facial and palm scans.

Implementing Identity Management

Realizing effective identity management requires a high-level corporate commitment and dedication of sufficient resources to accomplish the task. Typical undertakings in putting identity management in place include the following:

- Establishing a database of identities and credentials
- Managing users' access rights
- Enforcing security policy
- Developing the capability to create and modify accounts
- Setting up monitoring of resource accesses
- Installing a procedure for removing access rights
- Providing training in proper procedures

Access Control

Access control is intrinsically tied to identity management and is necessary to preserve the confidentiality, integrity, and availability of cloud data.

Three things that must be considered for the planning and implementation of access control mechanisms are threats to the system, the system's vulnerability to these threats, and the risk that the threats might materialize. These concepts are defined as follows:

- **Threat** — An event or activity that has the potential to cause harm to the information systems or networks
- **Vulnerability** — A weakness or lack of a safeguard that can be exploited by a threat, causing harm to the information systems or networks
- **Risk** — The potential for harm or loss to an information system or network; the probability that a threat will materialize

Controls

Controls are implemented to mitigate risk and reduce the potential for loss. Two important control concepts are *separation of duties* and the principle of *least privilege*

Control measures can be administrative, logical (also called technical), and physical in their implementation.

- Administrative controls include policies and procedures, security awareness training, background checks, work habit checks, a review of vacation history, and increased supervision.
- Logical or technical controls involve the restriction of access to systems and the protection of information. Examples of these types of controls are encryption, smart cards, access control lists, and transmission protocols.

- Physical controls incorporate guards and building security in general, such as the locking of doors, the securing of server rooms or laptops, the protection of cables, the separation of duties, and the backing up of files.

Models for Controlling Access

Controlling access by a subject (an active entity such as an individual or process) to an object (a passive entity such as a file) involves setting up access rules. These rules can be classified into three categories or models.

Mandatory Access Control

The authorization of a subject's access to an object depends upon labels, which indicate the subject's *clearance*, and the *classification or sensitivity* of the object. *Rule-based access control* is a type of mandatory access control because rules determine this access (such as the correspondence of clearance labels to classification labels), rather than the identity of the subjects and objects alone.

Discretionary Access Control

With discretionary access control, the subject has authority, within certain limitations, to specify what objects are accessible. For example, access control lists (ACLs) can be used. An *access control triple* consists of the user, program, and file, with the corresponding access privileges noted for each user. This type of access control is used in local, dynamic situations in which the subjects must have the discretion to specify what resources certain users are permitted to access.

Nondiscretionary Access Control

A central authority determines which subjects can have access to certain objects based on the organizational security policy. The access controls might be based on the individual's role in the organization (role-based) or the subject's responsibilities and duties (task-based). Access control can also be characterized as *context-dependent* or *content-dependent*. Context-dependent access control is a function of factors such as location, time of day, and previous access history.

⇒ Autonomic Security

- ✓ Autonomic computing refers to a self-managing computing model in which computer systems reconfigure themselves in response to changing conditions and are self-healing.
- ✓ The promise of autonomic computing will take a number of years to fully materialize, but it offers capabilities that can improve the security of information systems and cloud computing in particular.
- ✓ The ability of autonomic systems to collect and interpret data and recommend or implement solutions can go a long way toward enhancing security and providing for recovery from harmful events.

Autonomic Systems

Autonomic systems are based on the human autonomic nervous system, which is self-managing, monitors changes that affect the body, and maintains internal balances.

Examples of events that would have to be handled autonomously include the following:

- Malicious attacks
- Hardware or software faults
- Excessive CPU utilization
- Power failures

- Organizational policies
- Inadvertent operator errors
- Interaction with other systems
- Software updates

IBM introduced the concept of autonomic computing and its eight defining characteristics5 as follows:

1) **Self-awareness** — An autonomic application/system "knows itself" and is aware of its state and its behaviors.
2) **Self-configuring** — An autonomic application/system should be able configure and reconfigure itself under varying and unpredictable conditions.
3) **Self-optimizing** — An autonomic application/system should be able to detect sub-optimal behaviors and optimize itself to improve its execution.
4) **Self-healing** — An autonomic application/system should be able to detect and recover from potential problems and continue to function smoothly.
5) **Self-protecting** — An autonomic application/system should be capable of detecting and protecting its resources from both internal and external attack and maintaining overall system security and integrity.
6) **Context-aware** — An autonomic application/system should be aware of its execution environment and be able to react to changes in the environment.
7) **Open** — An autonomic application/system must function in a heterogeneous world and should be portable across multiple hardware and software architectures. Consequently, it must be built on standard and open protocols and interfaces.
8) **Anticipatory** — An autonomic application/system should be able to anticipate, to the extent possible, its needs and behaviors and those of its context, and be able to manage itself proactively.

Autonomic Protection

Autonomic self-protection involves detecting a harmful situation and taking actions that will mitigate the situation. These systems will also be designed to predict problems from analysis of sensory inputs and initiate corrective measures.

An autonomous system security response is based on network knowledge, capabilities of connected resources, information aggregation, the complexity of the situation, and the impact on affected applications.

Autonomous protection systems should, therefore, adhere to the following guidelines:

- Minimize overhead requirements.
- Be consistent with security policies.
- Optimize security-related parameters.
- Minimize impact on performance.
- Minimize potential for introducing new vulnerabilities.
- Conduct regression analysis and return to previous software versions if problems are introduced by changes.
- Ensure that reconfiguration processes are secure.

Autonomic Self-Healing: Autonomic self healing systems can provide the capability to detect and repair software problems and identify hardware faults without manual intervention. The objective of the autonomous self-healing process is to keep the elements operating according to their design specifications.