



CLOUD SECURITY AND ANALYTICS

• Feon Jaison • Dr. K Suneetha

Cloud Security and Analytics

Cloud Security and Analytics

Feon Jaison

Dr. K Suneetha



BOOKS ARCADE

KRISHNA NAGAR, DELHI

Cloud Security and Analytics

Feon Jaison
Dr. K Suneetha

© RESERVED

This book contains information obtained from highly regarded resources. Copyright for individual articles remains with the authors as indicated. A wide variety of references are listed. Reasonable efforts have been made to publish reliable data and information, but the author and the publisher cannot assume responsibility for the validity of all materials or for the consequences of their use.

No part of this book may be reprinted, reproduced, transmitted, or utilized in any form by any electronic, mechanical, or other means, now known or hereinafter invented, including photocopying, microfilming and recording, or any information storage or retrieval system, without permission from the publishers.

For permission to photocopy or use material electronically from this work please access booksarcade.co.in

BOOKS ARCADE

Regd. Office:

F-10/24, East Krishna Nagar, Near Vijay Chowk, Delhi-110051

Ph. No: +91-11-79669196, +91-9899073222

E-mail: info@booksarcade.co.in, booksarcade.pub@gmail.com

Website: www.booksarcade.co.in

Year of Publication 2023

International Standard Book Number-13: 978-81-19199-26-6



CONTENTS

Chapter 1. Fundamental Information Security	1
— <i>Feon Jaison</i>	
Chapter 2. Security in Cloud	11
— <i>Bhuvana J</i>	
Chapter 3. Multi-Layer Security Strategy	16
— <i>Dr. Lalit Kumar</i>	
Chapter 4. Next Generation Model in Cloud Security Zero Trust	25
— <i>Neetha SS</i>	
Chapter 5. Introduction to FedRAMP	37
— <i>Dr. K Suneetha</i>	
Chapter 6. Cloud Migration.....	42
— <i>Dr. Srikanth V</i>	
Chapter 7. Developing a DevSecOps Mentality.....	50
— <i>Dr. Mir Aadil</i>	
Chapter 8. Centralizing Common Cloud Services	53
— <i>Dr. Manju Bargavi Sankar Krishnamoorthy</i>	
Chapter 9. Federated Identity Management	79
— <i>Dr. Suchithra R</i>	
Chapter 10. Infrastructure Security	86
— <i>Dr. Murugan R</i>	
Chapter 11. Security of Virtual Servers	93
— <i>Dr. Ananta Charan Ojha</i>	
Chapter 12. Data Security and Storage.....	101
— <i>Dr. A Rengarajan</i>	
Chapter 13. Identity and Access Management	109
— <i>Dr. Pawan Kumar</i>	
Chapter 14. IAM Standards, Protocols, and Specifications for Consumers	122
— <i>Dr. Taskeen Zaidi</i>	
Chapter 15. Security Management in the Cloud.....	135
— <i>Karthikeyan Palniswamy</i>	
Chapter 16. Access Control.....	145
— <i>Raghavendra R.</i>	

CHAPTER 1

FUNDAMENTAL INFORMATION SECURITY

Feon Jaison

Assistant Professor, Department of Computer Science and Information Technology,
Jain (Deemed to be University) Bangalore, Karnataka, India
Email Id-feon.jaison@jainuniversity.ac.in

A collection of risks and assaults typical of public cloud systems before introducing words and ideas that address fundamental information security inside clouds. The security measures used to thwart these attacks are established by the cloud security mechanisms. Foundational Phrases and Ideas: Information security is a sophisticated collection of methods, tools, laws, and behaviors that work together to safeguard the confidentiality, integrity, and accessibility of computer systems and data. IT security procedures are designed to ward against dangers and disruptions brought on by both criminal intent and inadvertent human mistake. The sections that follow identify and explain cloud computing-related core security terminology.

Confidentiality: Something being made available exclusively to authorized people is a sign of confidentiality (Figure 1.1). Confidentiality in cloud settings mostly refers to limiting access to data while it is being stored and transferred.

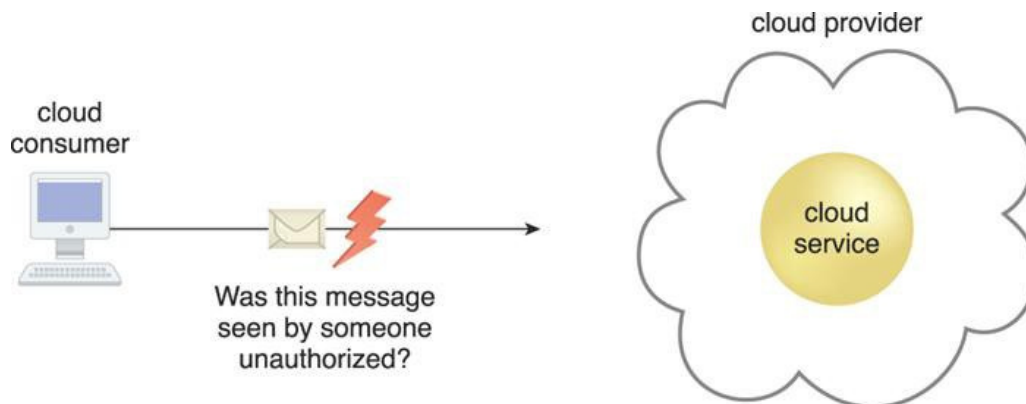


Figure 1.1: The communication sent from the cloud user to the cloud service is only regarded as secret if it is not accessed by or read by an unauthorized person.

Integrity: The quality of integrity is not having been changed by an unauthorized person (Figure 1.2). Whether a cloud consumer can be certain that the data it sends to a cloud service matches the data that cloud service receives is a critical problem that concerns data integrity in the cloud. Integrity may extend to how cloud services and cloud-based IT resources handle data storage, processing, and retrieval.

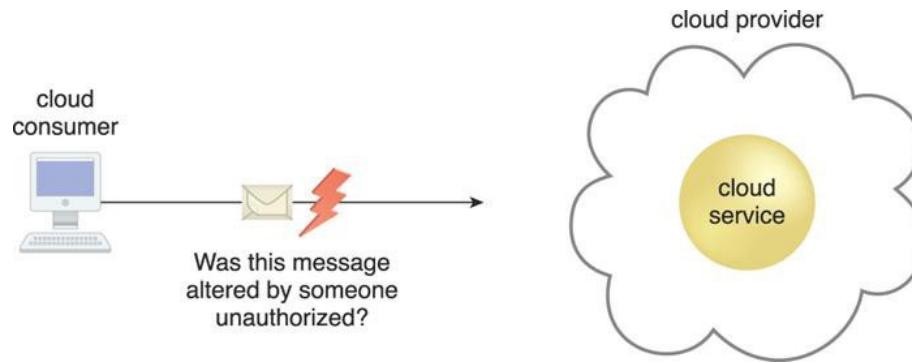


Figure 1.2: The message sent from the cloud consumer to the cloud service is deemed to have integrity if it hasn't been changed.

Authenticity: Having been offered from a trusted source, something is said to be authentic. Non-repudiation, which refers to a party's incapacity to dispute the authenticity of an encounter, is included in this idea. Non-repudiable interaction authentication demonstrates that these interactions are specifically connected to an authorised source. For instance, after receiving a non-repudiable file, a user may not be able to access it without simultaneously creating a record of that access.

Availability: The quality of being useable and available throughout a certain time frame is availability. In normal cloud systems, both the cloud provider and the cloud carrier may be accountable for the cloud services' accessibility. The cloud consumer also shares in the availability of a cloud-based solution that extends to users of cloud services.

Threat: A security threat is a prospective security breach that might test defences in an effort to violate privacy or inflict damage. Threats that are launched manually or automatically are intended to take advantage of known vulnerabilities. An assault happens when a threat is carried out.

Vulnerability: A weakness that can be taken advantage of due to inadequate security measures protecting it or current security controls being defeated by an assault is known as a vulnerability. Several factors may contribute to IT resource vulnerabilities, such as configuration problems, security policy holes, human mistake, hardware or firmware issues, software bugs, and insufficient security architecture.

Risk: Risk is the potential for loss or injury as a result of engaging in an activity. A risk's threat level and the number of potential or known vulnerabilities are often used to quantify it. The likelihood that a threat would attempt to exploit a resource's vulnerabilities and the expected loss in the event that a resource is compromised are two metrics that may be used to assess a resource's risk.

Security measures: Security controls are preventative procedures taken in response to security risks in order to lessen or eliminate risk. The security policy, which provides a collection of guidelines and practises outlining how to execute a system, service, or security strategy for the greatest protection of sensitive and crucial IT resources, often contains information on how to deploy security countermeasures.

Defending techniques: Typically, countermeasures are explained in terms of security mechanisms, which are parts of a protective architecture that safeguards IT resources, data, and services.

Security Guidelines: A security policy lays forth a list of security guidelines. Security policies often go into further detail about the application and enforcement of these rules and laws. Security rules, for instance, might dictate where and how security controls and processes are used.

Agents of Threat: An entity that presents a danger because it has the ability to launch an attack is known as a threat agent. Threats to cloud security might come from people or software, both inside and internationally. The sections that follow discuss the corresponding threat agents. Figure 1.3 outlines the function a threat agent plays in regard to weaknesses, threats, and risks as well as the protections put in place by security policies and security procedures.

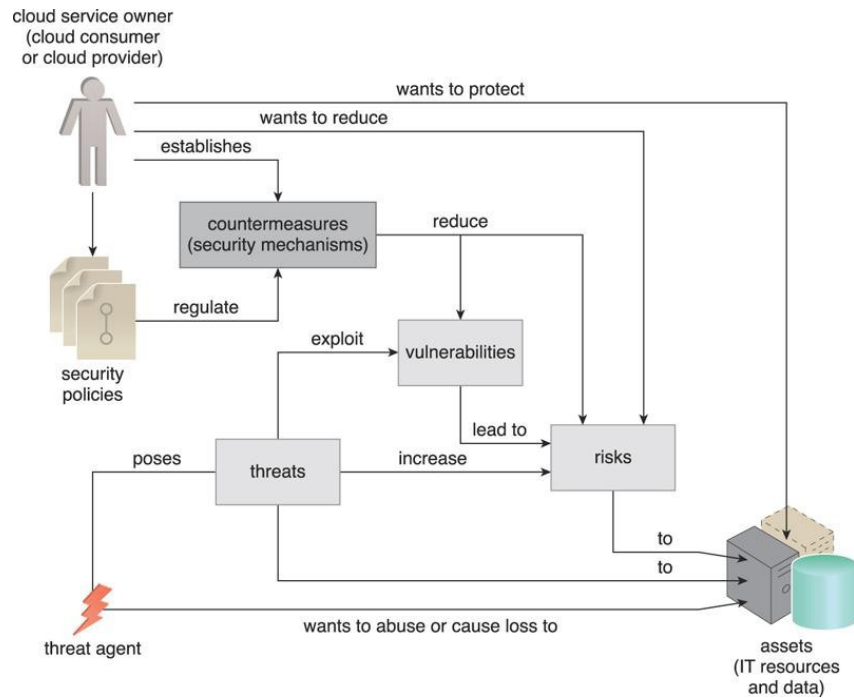


Figure 1.3: How threats, vulnerabilities, and hazards brought on by threat agents are addressed by security policies and security procedures.

Unknown Assailant: A non-trusted cloud service user without authorization to use the cloud is an anonymous attacker. It often takes the form of an external software application that conducts network-level assaults across open networks. Lack of knowledge about security procedures and countermeasures may prevent anonymous attackers from developing successful assaults. Hence, in order to conduct crimes while maintaining their anonymity or making it difficult to be caught, anonymous attackers sometimes turn to circumventing user accounts or stealing user credentials.

Untrustworthy Service Agent: The network communication that moves via a cloud may be intercepted and forwarded by a malicious service agent. It often manifests as a service agent with corrupted or malicious logic (or as a malware posing as a service agent). It could also take the form of an outside application that can remotely intercept and perhaps tamper with communication content.

Dependable attacker: trusted attacker targets cloud providers and the cloud tenants with whom they share IT resources and seeks to use valid credentials to access resources in the same cloud environment as the cloud consumer. Unlike non-trusted anonymous attackers,

trusted attackers often initiate their assaults from inside a cloud's trust boundaries by misusing valid credentials or by stealing private and sensitive data. The use of cloud-based IT resources by trusted attackers, also referred to as malicious tenants, is possible for a variety of exploits, such as the hacking of lax authentication procedures, the cracking of encryption, the spamming of email accounts, or the launch of widespread attacks like denial-of-service campaigns.

Negative Insider: Human threat actors operating on behalf of or in connection with the cloud provider are known as malicious insiders. Usually, they are third parties that have access to the cloud provider's facilities or are current or former workers. Since the hostile insider could have administrator access rights to cloud consumer IT resources, this sort of attack agent has a very high potential for harm.

A workstation and a lightning bolt are symbols used to denote a generic kind of human-driven assault. This general symbol merely indicates that an attack was launched from a workstation; it makes no mention of a particular threat agent. A non-trusted threat agent known as an anonymous attacker often launches assaults from beyond the cloud's perimeter. A malicious service agent eavesdrops on network traffic in an effort to alter or misuse the data. A trusted attacker may get access to cloud-based IT resources by posing as an authorised cloud service user with valid credentials. A person attempting to misuse access rights to cloud premises is considered a malicious insider.

Cloud Security Threats: A number of typical threats and weaknesses in cloud-based setups and explains the functions of the threat agents described above. Security measures that are used to combat these dangers.

Traffic listening: When a malicious service agent passively intercepts data being transported to or within a cloud (often from the cloud customer to the cloud provider) for unauthorised information gathering purposes, this is known as traffic eavesdropping (Figure 1.4). The objective of this assault is to immediately jeopardise the data's confidentiality and, maybe, the relationship between the cloud customer and cloud provider's secrecy. The passive nature of the assault makes it easier for it to go unnoticed for a long time.




Figure 1.4: shows how a malicious service agent that is externally positioned eavesdrops on traffic by intercepting a message that a cloud service customer sends to the cloud service. Before the communication is dispatched through its original route to the cloud service, the service agent creates an illegal duplicate of it.

Unsavory Intermediary: The hazard posed by hostile intermediaries occurs when communications are intercepted and changed by a malicious service agent, possibly jeopardising the confidentiality and/or integrity of the message. Before sending the message to its intended recipient, it could additionally include dangerous info. A typical instance of the malicious intermediate attack is shown in Figure 1.5.




Figure 1.5: A communication sent by a cloud service user to a cloud service (not displayed) that is housed on a virtual server is intercepted and altered by the malicious service agent. The mail contains malicious data, which compromises the virtual server.

The malicious intermediate attack may also be carried out via a malicious cloud service consumer software, albeit this is less prevalent. **Disruption of Service: Denial of service (DoS)** attacks aim to overwhelm IT resources to the point where they become inoperable. One of the following methods is often used to initiate this kind of attack: Imitation messages or repetitive communication requests artificially raise the demand on cloud services. Numerous cloud service requests are issued, each of which is intended to demand excessive memory and processing resources; the network is flooded with traffic in order to impair its responsiveness and cripple its performance. Successful DoS attacks result in server failure or deterioration, as seen in Figure 1.6.




Figure 1.6: Cloud Service Consumer A uses Virtual Server A to run a cloud service, which receives various messages from Cloud Service Consumer A. As a result, Virtual Servers A and B experience disruptions because the capacity of the underlying real server is overloaded. As a consequence, authorised users of cloud services, such Cloud Service Consumer B, are prevented from interacting with any cloud services that are hosted on Virtual Servers A and B.

Lack of Authorization: The inadequate authorization attack happens when access is mistakenly or excessively provided to an attacker, giving the adversary access to IT resources that are typically secured. This often happens as a consequence of the attacker acquiring direct access to IT resources that were set up with the expectation that only reputable consumer apps would use them (Figure 1.7).




Figure 1.7: A database that was constructed under the presumption that it would only be accessed via a Web service with a published service contract is now accessible to Cloud Service Consumer A. (as per Cloud Service Consumer B).

Weak passwords or shared accounts used to safeguard IT resources may lead to a variant of this attack known as weak authentication. Depending on the scope of IT resources and the scope of access the attacker acquires to those IT resources inside cloud systems, these sorts of assaults may have major effects (Figure 1.8).

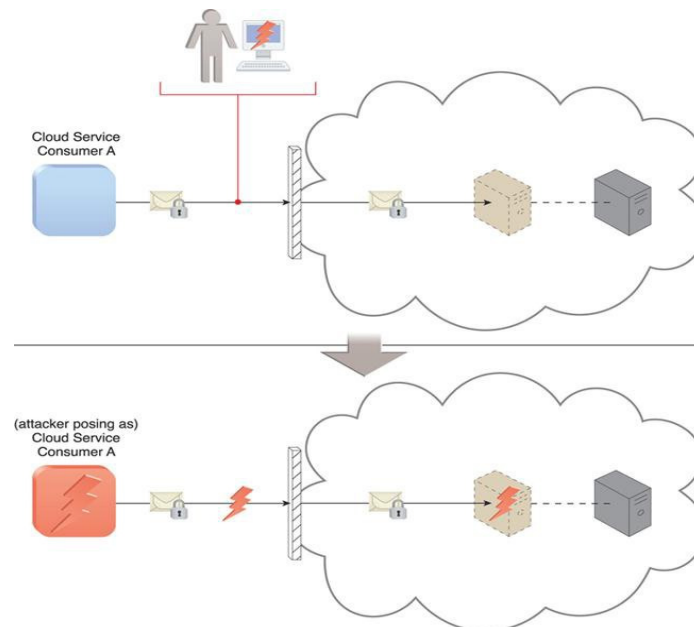


Figure 1.8: shows how an attacker deciphered Cloud Service Consumer A's lax password. In order to access the cloud-based virtual server, a malicious cloud service consumer (owned by the attacker) is created to pretend to be Cloud Service Customer A.

Attack by Virtualization: Using virtualization, several cloud users may access IT resources that have the same underlying hardware but are conceptually separate from one another. There is an inherent danger that cloud customers might misuse this access to attack the underlying physical IT resources since cloud providers provide cloud users administrative access to virtualized IT resources (such as virtual servers). A virtualization attack takes use of holes in the virtualization platform to compromise its availability, confidentiality, and/or integrity. Figure 1.9, in which a trusted attacker successfully gains access to a virtual server to compromise its underlying physical server, serves as an illustration of this issue. Such an assault might have serious consequences with public clouds, because a single physical IT resource may be supplying virtualized IT services to several cloud users.




Figure 1.9: An authorized user of a cloud service conducts a virtualization attack by exploiting its administrative access to a virtual server to take advantage of the underlying hardware

Doubling Up on Trust Boundary: If several cloud service users share the same physical IT resources inside a cloud, their trust boundaries will overlap. Bad cloud service users may target shared IT resources in an effort to compromise other users of those resources or other cloud users that share the same trust boundary. As a result, the attack may have an effect on some or all of the other cloud service users, and/or the attacker may employ virtual IT resources to target other users who also happen to be inside the same trust boundary. Two cloud service users are shown in Figure 1.10 sharing virtual servers hosted by the same physical server, which causes their respective trust boundaries to overlap.




Figure 1.10. An authorized cloud service consumer carries out a virtualization attack by abusing its administrative access to a virtual server to exploit the underlying hardware.

As a result of the cloud's confidence, Cloud Service Consumer A is granted access to a virtual server, which it then assaults with the goal of also assaulting the virtual server utilised by Cloud Service Customer B. Malicious service agents that intercept network traffic often conduct traffic eavesdropping and malicious intermediary attacks. When a specific IT resource is overwhelmed with requests in an effort to disable or make it inaccessible, the result is a denial-of-service assault. A virtualization attack takes use of flaws in virtualized systems to get unauthorized access to underlying physical hardware. This attack takes place when access is provided to an attacker incorrectly or too widely, or when weak passwords are used. Attackers may use cloud-based IT resources shared by several cloud users if there are overlaps in the trust boundaries. This section offers a broad checklist of concerns and recommendations related to cloud security. The factors mentioned below are not in any particular sequence.

Implementations with flaws: Beyond runtime errors and failures, poor cloud service deployment design, implementation, or configuration may have unfavorable effects. Attackers may compromise the integrity, confidentiality, and/or availability of cloud provider IT resources and cloud consumer IT resources hosted by the cloud provider if the software and/or hardware of the cloud provider has built-in security flaws or operational weaknesses. A cloud service that is poorly built and causes a server to go down is shown in Figure 1.11. While a trustworthy cloud service user mistakenly discloses the issue in this instance, an attacker might have easily found and used it.

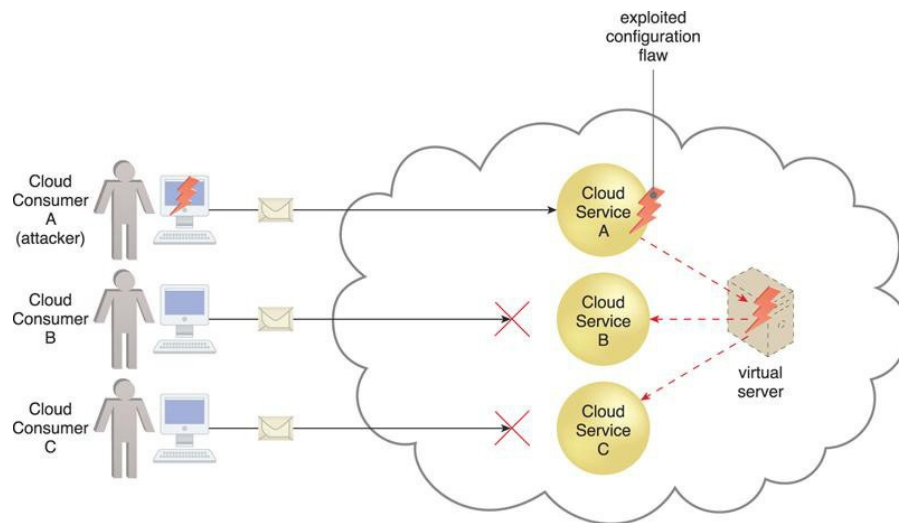


Figure 1.11: Cloud Service, A configuration error in Cloud Service A is triggered by a message from Consumer A, which ultimately leads to the virtual server hosting Cloud Services B and C to crash.

Disparity in Security Policies: It may be necessary for a cloud customer to recognise that the public cloud provider's information security policy may not be the same as their own, or even closely related. To make sure that any data or other IT assets being moved to a public cloud are effectively safeguarded, this incompatibility has to be evaluated. The cloud user could not be given enough administrative control or sway over security rules that apply to the IT resources they have leased from the cloud provider, even when leasing raw infrastructure-based IT services. This is mostly due to the fact that the cloud provider continues to legally own and be accountable for those IT resources. Any efforts to standardize the protection of cloud user assets would be made more difficult by the presence of extra third parties in

certain public clouds, such as security brokers and certificate authorities, who may implement their own unique set of security rules and procedures.

Contract: Customers of cloud services should carefully review the contracts and SLAs offered by cloud providers to make sure that the security policies and other relevant assurances are adequate in terms of asset protection. The amount of obligation borne by the cloud provider and/or the degree of indemnification the cloud provider may demand must be stated in unambiguous terms. The risk to cloud consumers decreases with the amount of obligation that the cloud provider assumes.

What assets belong to cloud providers and consumers is another facet of contractual duties. A technological architecture made up of artefacts owned by both the cloud consumer and the cloud provider will be created by a cloud consumer that installs its own solution on infrastructure provided by the cloud provider. How is responsibility assigned in the case of a security breach (or other kind of runtime failure)? Additionally, how can the mismatch be resolved if the cloud user is free to apply their own security standards to their solution while the cloud provider requires that the security policies for their supporting infrastructure be different (and maybe incompatible)?

Sometimes switching to a new cloud provider with more agreeable contract conditions is the best course of action. **Management of Risk:** Cloud users are recommended to do a formal risk assessment as part of a risk management plan when evaluating the possible implications and difficulties related to cloud adoption. Risk management is a cycle-based procedure used to improve tactical and strategic security. It consists of a number of coordinated actions for monitoring and managing risks. Generally speaking, risk assessment, risk treatment, and risk control are the three key tasks (Figure 1.12).

Risk evaluation: The cloud environment is examined in the risk assessment step to find any possible flaws or vulnerabilities that attackers may exploit. You may ask the cloud provider to disclose statistics and other details regarding previous attacks (both successful and failed) made using their cloud. According to the likelihood of occurrence and the degree of effect in connection to how the cloud consumer intends to use cloud-based IT resources, the detected risks are quantified and categorised.

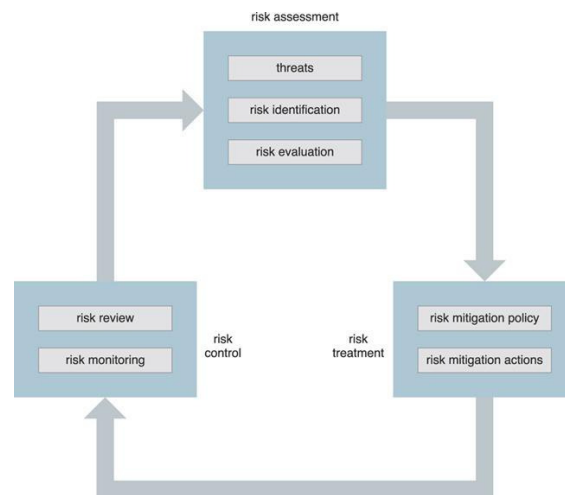


Figure 1.12: The ongoing risk management process, which may be started from any of the three phases

Risk Treatment: During the risk treatment stage, mitigation strategies and plans are created with the goal of effectively addressing the risks that were identified during the risk assessment. Some risks can be reduced, some can be removed, some can be outsourced, and some can even be included in the insurance and/or operational loss budgets. It's possible that the cloud service provider will consent to take responsibility as part of its legal requirements.

Risk Control: The risk control stage is connected to risk monitoring, a three-step procedure that entails gathering information about relevant events, analysing those events to evaluate the efficacy of prior assessments and interventions, and identifying any need for policy change. This step may be completed or shared by the cloud provider, depending on the kind of monitoring necessary. The risk assessment step may be used to identify and record the threat agents and cloud security risks. The appropriate risk treatment may include documentation and references to the cloud security procedures.

A list of the main points: Users of the cloud should be aware that by implementing subpar cloud-based solutions, they may be posing security threats. Forming evaluation criteria for selecting a cloud provider vendor requires a thorough grasp of how a cloud provider establishes and applies proprietary, and potentially incompatible, cloud security regulations. The legal contracts that cloud customers and cloud providers sign must expressly specify and agree upon who is responsible for future security breaches. It's crucial for cloud users to examine the risks that have been found after learning about any possible security-related problems that may be unique to a particular cloud environment.

ATN analysts discover a set of dangers based on an evaluation of its internal applications. The myTrendek application, which was adapted from OTC, a business that ATN just bought, carries one such risk. This programme has a function that examines Internet and phone activity and allows for a multi-user mode with various access levels. Hence, various rights may be granted to administrators, managers, auditors, and common users. Both internal and external users, including business partners and contractors, are included in the application's user base.

There are many security issues with the myTrendek programme when used by internal staff: Complex passwords are not required or enforced for authentication. European rules (ETelReg) mandate that some kinds of data acquired by the application be erased after six months; communication with the application is not encrypted;

The lack of secrecy offered by the application and the poor authentication threat force ATN to rethink their intention to deploy this application to the cloud through a PaaS environment. An additional risk analysis finds that local laws may clash with ETelReg if the application is moved to a PaaS environment hosted by a cloud that is located outside of Europe. Since that the cloud provider doesn't care about ETelReg compliance, ATN may be subject to financial fines as a consequence. ATN makes the decision not to go forward with its cloud migration strategy in light of the findings of the risk assessment.

CHAPTER 2

SECURITY IN CLOUD

Bhuvana J

Associate Professor, Department of Computer Science and Information Technology,
Jain (Deemed to be University) Bangalore, Karnataka, India
Email Id-j.bhuvana@jainuniversity.ac.in

Information Technology Security, or IT Security, is the process of putting into place safeguards and systems to securely protect and preserve information (department and personal data, conversational information, still images, moving images, and multimedia presentations, including those that haven't even been thought of) using the various technologies created to create, store, use, and exchange such information. The management of people, process, and technology with rigorous rules that protect data and applications running in the cloud is cloud security. In order to fully safeguard the data, a government agency must first analyse how it processes and stores data before developing a tailored strategy. Using the finest cloud security procedures is essential for every contemporary department since departments seldom have the budget to take a significant harm to their image. Similar to how security has developed for all new technology and breakthroughs, cloud security has changed. A cloud incident response strategy must be in place in the unfortunate event that a government agency suffers from such a breach in order to lessen the effects of suspicious behavior and limit damage. Every tragic occurrence is terrible enough, but how the department responds after one will often decide the department's future. The cost of a cyber-breach is often determined by the department's response strategy. By identifying security components that will work smoothly with their cloud demand and streamlining security procedures, government departments may now include security as a part of the migration to IaaS services.

Cloud service providers and experts in cloud security have both benefited greatly from the use of cloud computing by government departments. Empanelment of Cloud Service Providers (CSPs) elaborates on the security requirements that the prospective CSPs must comply with while seeking for empanelment. Appointment to provide their services to government agencies. Government Departments must use the CSPs' services in accordance with their needs. Strong/robust tenant isolation provided by a virtual private cloud enables the logical separation of infrastructure (server, storage, and network) from other cloud service provider services.

The Government Community Cloud enables the physical separation of the Cloud Service Provider's Public and Virtual Private Cloud offerings from the infrastructure (server, storage, and network). A public cloud service makes accessible storage, hypervisor separation, role-based permissions, and software controls to ensure cloud security. Other cloud deployment models, such as GCC or VPC, may be taken into consideration if the Departments need a higher degree of workload and data separation or isolation amongst cloud consumers. Because of the inherent advantages of the cloud, which allow government agencies to concentrate primarily on their applications, cloud security has always been a key consideration when analysing the cloud. While the empanelment outlines the security standards that the empaneled CSPs must meet, Government Departments would also need to implement specific procedures to roll-out

their applications/services in a secure manner. This paper highlights several cloud security best practises that Departments may use as they move towards cloud enablement.

If the Cloud Service Provider (CSP) does not properly manage the responsibility of addressing IT and Cyber security parameters / controls at each layer, as it should be placed in the Cloud environment, Departments rely on CSP security and control to maintain the secure environment and mitigate potential risk. Thus, Departments must guarantee that the Security Service Level Agreements (SLAs) essential for CSP to adhere to the requisite security services are in place.

Requirement for Cloud Security: While cloud computing services are a fantastic alternative for government agencies, the technology does carry certain concerns. Since the Indian government started using the cloud, several departments have been gradually transferring to the authorised cloud service providers. Due to the concentration of important data in one place, CSPs are a common target for malicious activities.

Government departments must work with MSPs and CSPs to safeguard their crucial data and make sure the essential security measures are in place, either directly or through their SIs. In addition to MeitY-imposed rules and compliances, a security fabric has to be combined at the level of the data center and cloud. Insider attacks are one issue that many CSPs are becoming worried about. These are several security issues that have been addressed, including certain OWASP Cloud Security threats:

Data Breach: While cloud computing services are relatively new and important, data breaches of various kinds have been a problem for years. The issue of "Is the cloud secure given that critical departmental data is being housed online rather than on premises" is one of the most frequent ones that government departments encounter. Cloud would provide improved security measures and required certifications to the User Departments. According to the MeitY empanelment of Cloud Service Providers (CSP), all CSPs apply security measures in accordance with ISO 27001, 27017, etc. Nevertheless, owing to the Government Department users' non-enforcement of security standards, it may result in data breaches.

Poor management of cloud accounts: Account assaults and account hijackings have become a whole new set of problems as a result of the development and use of the cloud in many enterprises. Attackers may now remotely access sensitive or essential data that is stored on the platform or cloud using the department's cloud login credentials. Moreover, attackers can display false information and change data using credentials that have been hijacked. Thus, it is necessary to establish suitable cloud account management procedures. A Managed Service Provider (MSP) may sometimes have access to a government department's cloud account; hence, the proper controls should be put in place for this circumstance as well.

Internal Threat: While an insider danger may seem implausible, it does occur when it comes to government departments. Users of government departments have the ability to abuse or get access to sensitive information such as citizen and financial data by using their permitted access to the department's cloud-based services. So, in order to prevent security difficulties, it is crucial for government departments to establish a safe plan for their cloud installation and access and to make sure that the appropriate access control mechanism is in place.

Adherence to Regulations: A country's or region's perception of data security may differ from that country's or region's perception. So, when selecting a cloud provider, data ownership and governance become crucial considerations. According to MeitY's appointment, every appointed cloud service provider will provide cloud services from Indian data centre

facilities and guarantee data residence in the nation. The Government Department is the owner of the data.

Unsafe APIs: Operators may modify their cloud platform via application programming interfaces (API). While APIs enable customers to modify cloud service features to suit their requirements, they also have an impact on encryption, authentication, and access/control provision. Better services are made possible by the expansion of APIs, but security threats also rise. APIs provide developers with the tools they need to create programmes and combine applications. The communication that occurs between apps is where an API's vulnerability rests. Also, they provide a chance for exploitable security issues.

Denial-of-Service assaults: Denial-of-service attacks don't merely try to breach the security perimeter; instead, they try to get a foothold and extract important information over time, unlike other types of assaults. Instead, they try to prevent the legal users of the Department from accessing the services and systems. In certain circumstances, DoS is also used as a cover for criminal activity and a targeted assault to bring down security devices like WAF (Web Application Firewalls).

A lack of due diligence: While the vulnerabilities mentioned above are purely technical, this specific security hole appears when a government department lacks a defined plan for its resources and cloud-related rules. The Government Departments must keep a close eye on internal controls for cloud services. There are several service parameters that must be set up so they don't cause problems with operations, reputation, or compliance. When a Government Department ignores certain cloud setups at the user level, it might represent a serious security risk.

Joint Responsibility: The Cloud Service Provider and the Cloud Consumer both share responsibilities for cloud security. The customer must take the appropriate steps to secure their data as part of this partnership between them and the supplier. Major international Cloud Service Providers do have established methods to protect their side, but users are responsible for fine-grained restrictions. The basic line is that customers and providers have mutual obligations, and if users don't fulfil their obligations, their data may be compromised. Please see Section 5 Cloud Security as a Shared Responsibility Model for more information.

Loss of Data: Data on cloud platforms may disappear due to a natural catastrophe, data erasure, or criminal service provider activity. Without a recovery strategy, losing important data or information may be disastrous for enterprises. Additionally, the Open Web Application Security Project (OWASP) has outlined a number of cloud security issues, including User Identity Federation, Business Continuity and Resilience, Service and Data Integration, Multi-tenancy and Physical Security, and Infrastructure Security, which have been addressed in the specifications established as part of MeitY's empanelment of Cloud Service Providers. These specifications hold CSPs responsible for their obligations regarding cloud security.

Cloud security vs on-premises data centre security: The key that powers departmental operations nowadays is data. Such information aids in monitoring performance, gaining valuable insights, and enhancing security. Data is also a key component in defining and outlining different IT security rules, whether they are implemented on-premises or in the cloud. Although some departments choose to gather and handle their data internally, others choose to migrate to the cloud because of the services' availability and scalability. Data administration has been simpler thanks to cloud technologies, which have also improved data security. Departments are able to create and maintain comprehensive and efficient cloud security frameworks that can monitor and address emerging risks because the cloud

guarantees on-demand infrastructure access. It's crucial to make a distinction between standard IT security and cloud security. Both techniques have benefits and drawbacks, and understanding both can help the Department make better operational decisions.

Security for On-Site Data Centers: The gadgets must be installed, purchased, and maintained on-site for an IT framework to be functional. The conventional IT infrastructure makes sure that data is gathered, stored, and processed for a variety of purposes. Also, the departments are able to establish data security measures thanks to conventional IT infrastructure. This indicates that government departments should be allowed to choose the appropriate security components for the architecture, methods for managing network controls, and responses to incoming threats. Moreover, the department would be responsible of maintaining a disaster recovery plan as well as identifying and countering potential threats. Traditional IT methods allow the department more control over how each device is used on a regular basis. The data is stored on departmental property, and monitoring, controlling, and daily data management are all available. But, an on-site solution would need educating current Department resources on cutting-edge security systems.

Yet, the capital expense needed to build and maintain the security components is the major issue with conventional IT systems. Capital expenses for the departments would also increase due to asset replacement for security components that have reached the end of their useful lives. The department is charged with managing and overseeing security-related compliances and certifications, which are costly and call for internal resources inside the agency. Hence, weaknesses in security-related procedures might result in vulnerabilities. Larger internal staff are also needed to administer the hardware, conduct security incident responses, and monitor traditional IT systems. While there may be more control over data processing as a result, there are significant financial repercussions.

Security in the Cloud: Cloud computing, in contrast to conventional IT systems, refers to on-demand access to infrastructure and services. Here, the MSP configures and manages a portfolio of cloud platforms and services, which the CSP is responsible for making accessible to the client or user. Depending on their skill level, users may be able to setup and administer the cloud platform without the need for an MSP. The government departments can access the software, hardware, and other essential infrastructure they need to carry out their everyday operations thanks to cloud computing. Also, the cloud guarantees simpler system security and data administration. The Government may simply outsource the data security requirements to a well-known and respected Managed Service Provider rather than managing every component of data security control on-site.

Infrastructure that is located on-site may be more vulnerable to simple mistakes and oversights that increase the risk of cyberattacks. Also, the majority of cloud developers have deeper knowledge of sophisticated security and data governance approaches. In order to achieve real-time risk mitigation, the Departments will be able to prepare the necessary measures. Concerns about security are a significant factor in the hesitation to transfer more data onto the cloud. This is a comparison of on-premises and cloud configurations with security at the forefront:

Although the government department's on-premise setup gives it total control over the setup, with cloud computing, many features like infrastructure provisioning and maintenance, compliance and certifications, and technology refresh are handled by the cloud service provider, freeing the government department to concentrate on the delivery of applications. Previous While cloud providers continue to spend in effectively hardening their security and

privacy profiles and standards, limiting factors like worries about data security and privacy are waning.

Considerations for Cloud Security Design: To safeguard systems, applications, and platforms and to enhance overall security architecture, the security design principles are the fundamental building blocks for the adoption and implementation of cloud security. The following are the main design tenets that should be taken into account while using cloud technology:

Security at all levels: Make sure that numerous security measures are applied to all layers of their architecture, including the physical, network, data, and application layers. This will provide complete protection of the applications and data that departments host on the cloud platform.

Protect data while it is at rest and in transit: Identify the data, classify it according to its sensitivity and criticality, and establish its levels. By using the security mechanisms that are already in place, such as access control, tokenization, encryption, etc., this may be avoided.

Monitoring and Auditing: Ensure that monitoring, auditing, and alerting are set up to record system changes in real time for the department. Moreover, log integration and metric collecting may investigate, act, and reply automatically.

Access management and controls: Ensure that the concept of selected privileges is implemented, and enforce the division of responsibilities with the proper access and permission. Any unwanted access and information loss/theft may be stopped via centralised identity and access control.

Security event readiness: The department/CSP must set up the system for any unexpected security incident. To find the security flaws and problems, regular vulnerability and security testing must be performed. To record the reaction of the Cloud systems at various tiers, several drills may be carried out.

Automate recommended practices for security: Automating software, hardware, and application-based security systems using AI, ML, and bots will increase the capacity to protect environments and apply the necessary controls to thwart attacks and improve cloud security.

Cloud Vendor Lock-in: Departments must ensure that no vendor lock-in occurs when a cloud services provider hosts an application or data because there are no common standards among cloud providers for data exports and migration, making it challenging to move data from one cloud provider to another or to an on-premise data center.

CHAPTER 3

MULTI-LAYER SECURITY STRATEGY

Dr. Lalit Kumar

Assistant Professor, Department of Computer Science and Information Technology,
Jain (Deemed to be University) Bangalore, Karnataka, India
Email Id-k.lalit@jainuniversity.ac.in

It is crucial for a government department moving to the cloud to adhere to security procedures at different levels in order to achieve a safe cloud adoption. The goal is to draw attention to a few best practices that may be included into the architecture of the government department's cloud requirement in order to boost security and trust in using the cloud.

Multi-Layer Security Strategy: To embrace or understand when selecting the cloud for service delivery, as well as best practices to follow at different stages of a cloud deployment architecture. A Multi-Layer Security Strategy: Physical facilities via the Department's installation and design of IT infrastructure components are all included in the lowest security levels of IT security. These fundamental elements, including as networking, computation, and storage security, constitute the foundation upon which cloud computing is constructed (Figure 3.1).

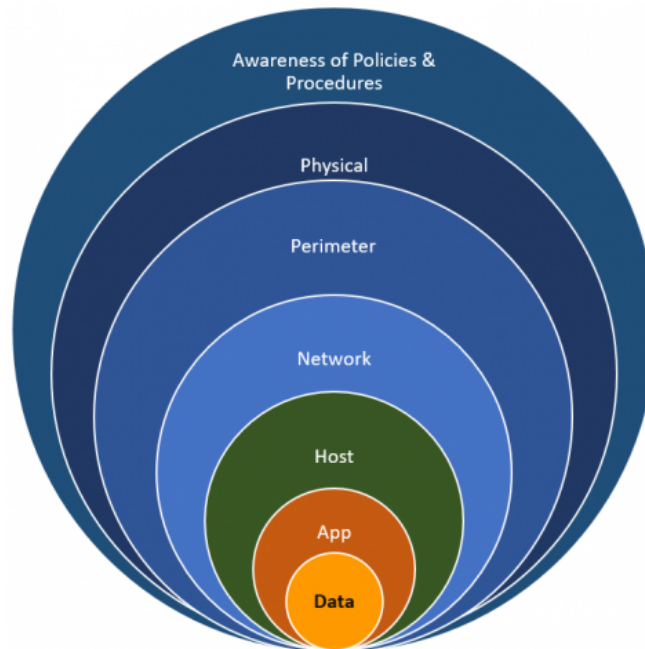


Figure 3.1: Multi-Layer Security Strategy

Understanding the CSP's infrastructure security requirements is crucial. Under the shared security model, it is the provider's job to manage the security of the private cloud platform and its underlying physical, abstraction, and orchestration layers. Although we outline the best practices for cloud security in this part across all the levels listed above, we'll also touch on privacy as a crucial factor to take into account when it comes to cloud-based security.

Data: The techniques used for cloud data security don't really call for any novel approaches. Data security in a cloud is quite comparable to that in a conventional data center. The following data security techniques may all be used in cloud computing: identity and authentication, encryption, access control, secure deletion, data masking, and integrity checking.

For the cloud to succeed, control over the data must be maintained. Nowadays, thanks to virtualization and the cloud, data may be logically controlled by the government department but physically exist on infrastructure that is owned and operated by another organization. To guarantee that businesses can maintain data security, new methodologies and strategies are necessary because of this change in control. To get past any reservations about cloud data security, government departments may develop specific data protection measures. Departments must be aware of the types of data they manage, so they can classify and upload pertinent data to the cloud while taking precautions against unauthorized data access, deletion, and backup vulnerabilities, data leaks, a management interface that has been compromised, malware assaults, etc.

Departments should create a data use policy that explains the different forms of data access, who has access to the data, and what circumstances may constitute proper data usage. For policy infractions and the ensuing impact-based penalties, there should be protections. One of the most important and fundamental aspects of cloud data security is access control, and it is the responsibility of government departments to ensure that administrative and technical controls are in place to limit who has access to what data. The installed cloud infrastructure's physical security must be ensured by the CSP. This section includes some of the procedures related to cloud data security.

Data in transit and at rest should be protected and secured using encryption. CSP may implement a variety of encryption types, including Full Disk Encryption (FDE), Format Preserving Encryption (FPE), Application layer Encryption, File Encryption, Database Encryption, etc.) To safeguard the contents of data in transit, select encryption of sensitive data before transferring to the cloud and/or utilize encrypted connections (HTTPS, SSL, TLS, FTPS, etc.).

Departments may easily encrypt sensitive data before storing it to secure data at rest. Take into account using encryption settings controlled by the service provider. Use customer-managed keys wherever feasible since they provide greater control. Take into account adding a second layer of data protection by using data categorization (Restricted, Confidential, Sensitive and Unclassified). While duplicating data from one location to another, ensure its integrity. Establish a data use policy (data access control management, repercussion of policy violations, correct usage of data etc.) To maintain data availability, frequently backup (Full, Incremental, Differential) data and undertake periodic recovery operations to verify accuracy. Verify that data-level monitoring is in place and that logs adhere to all applicable departmental compliance standards, if any.

Application: Application hosting takes place on separate virtual machines. Since cloud environments are housed on shared resources, applications and sensitive data are more susceptible there. Thus, further security precautions or controls are needed to protect the client environments. Cloud service providers make sure that users and departments only have access to the data that is shared on the Cloud model and to which they are allowed.

The use of micro-service architectures further improves security. Developers may instead create more, smaller virtual machines, each devoted to a distinct function or service, since

minimizing the usage of real servers by customers is not a necessity. As a result, the attack surface of each individual VM is reduced, and granular security measures are supported.

DevOps is a cutting-edge approach to application development that focuses on automating every step of the creation and deployment of applications. Version control management, change management, and increased security operations are all made possible by DevOps in numerous ways. Not simply development and operations teams are involved in DevOps. If IT security plays an integrated role across the whole life cycle of the application development, the agility and responsiveness of a DevOps strategy may be fully used. Each Department may inquire. In the past, security was confined to a single team during the latter phase of development. Since development cycles used a waterfall technique, it wasn't as much a concern. Continuous Integrations and Continuous Deployments are made possible by DevOps, which guarantees quick and frequent development cycles (often lasting weeks or days). Nevertheless, obsolete and archaic security methods may negatively impact even the most successful DevOps operations.

Web application security is the practice of securing online services and applications that may be accessed via a browser. It defends against several security risks that take use of flaws in both core and non-core programs. Sensitive data, content management systems, and other management and administrative applications are common targets of assaults. To prevent access into department applications, important technologies are being used. By monitoring and filtering HTTP traffic between a web application and the Internet, a web application firewall aids in the protection of web applications. It generally guards against attacks like cross-site scripting, SQL injection, file inclusion, cross-site forgery, etc. that target online applications. It is put into practise with a clear set of guidelines or regulations for safety.

CSPs provide cloud APIs to software developers so they may create user interfaces for interacting with the offered cloud services. The Cloud API security layer may be improved by implementing security tools and appliances on the Cloud provider's interface and turning on authentication and access control. This is because the addition of another layer necessitates protection against vulnerabilities and assaults. DevSecOps entails integrating the infrastructure and application security teams from the very beginning of the application development lifecycle. Moreover, it entails automating security gates to prevent the DevOps process pipeline from becoming sluggish. The intended goals may be achieved with the use of the appropriate tools, such as selecting an integrated development environment (IDE) with security features. Effective DevOps security, however, calls much more than simply new technologies; it builds on DevOps's cultural shifts to incorporate and integrate the work of security teams as soon as possible. DevSecOps is assuring a smooth software development life cycle, which is significantly changing the IT sector (SDLC). DevSecOps asks for security integration throughout all phases of the software development process chain, addressing security issues from the very beginning of every step, bucking the conventional pattern of having security as a segregated process. Specifically for security automation and setup of cloud assets, the DevSecOps approach to cloud security necessitates extensive preparation that necessitates cultural change in an IT context. The following variables control how well DevSecOps is implemented in a cloud environment:

Code analysis: Comprehensive code reviews are necessary for ongoing programme enhancements. Automated testing - Automated testing guarantees time and effort savings. By the effective execution of repeating test cases, automated testing, a key component of the DevSecOps methodology, speeds up and simplifies the testing process.

Change management: Encouraging teamwork so that each team is informed of the activities of every other team. Early disclosure of security-related actions to developers may aid in the prompt addressing of potential vulnerabilities.

Compliance Monitoring: Compliance continues to be very important to the development of a business. Code production and source code changes are made easier by regulations. This facilitates auditing in real time.

Threat investigation: Threat investigation is crucial to determining an organization's level of security preparedness. It's crucial for businesses to have a tight and Continuous monitoring of potential threats, regular security checks, and code reviews are used to solve ongoing security concerns.

Employee Training: By providing teams with the necessary domain expertise, organising hands-on training sessions and certification programmes strengthens the organisation.

Many government departments would find it difficult to decide where to place security for multi-tier applications: Multi-tier systems have complicated design approaches, but as the quantity and complexity of security mechanisms increases, performance may suffer and the behaviour of the programme may become unpredictable. Thus, more thorough security evaluation is required when creating such sophisticated multi-tier systems in order to ensure security. Priority should be placed on ensuring security provision and permitted access at the application level, after which the database may trust the application to authenticate and provide end users access to database data. Database access should only be permitted via the application. When creating multi-tier applications, make sure auditing and logging are done at the application level. Crucial things to keep in mind Incorporate security into the application's original design process. The potential to include cloud-based security early is presented by the development of cloud native apps.

Integration of the deployment process with security testing. To keep control over all private and public encryption keys, there is encryption key management. The new architectural alternatives and cloud-based services should be understood by departments. Instead of only trying to enforce current standards strictly on model, the departments should adapt their standards and security policies to support them.

When deploying an application to the cloud, a web-facing application should be placed in a DMZ (De-militarized) zone, and a database server should be placed in a protected zone. Online applications and internet portals should use web application firewalls. Information sharing and application integration over protected API routes. Security measures for APIs and interfaces. Record and keep track of API requests. Automate security measures by using software-defined security. When event-driven security is available, such as antivirus software, use it to automate the discovery and correction of security flaws.

Host/Compute: A workload is a processing unit that may be a container, a virtual machine, or another abstraction. Workloads use up memory and execute on a CPU. Workloads comprise a variety of processing activities, from conventional apps running in a virtual machine under a common OS to workloads with high GPU demands. For the majority of tasks, it is advised that virtual machines may be handled as if they were real machines. Yet, it's crucial to keep in mind that virtual computers are just as susceptible to problems like data loss or corruption, hardware malfunctions, viruses, and hackers as real machines are. Here are descriptions of several cloud computing offering types:

The Virtual Machine Manager (hypervisor) is in charge of separating the operating system from the underlying hardware. In order to ensure isolation and facilitate high-performance activities, hypervisors might rely on the capabilities of the underlying hardware. Due to continual hardware and software improvements to strengthen isolation, memory assaults on virtual machines are becoming more and more challenging. Since hardware isolation for VMs and safe execution environments continue to evolve, virtual machines on today's hypervisors are typically an effective security control.

Code execution environments known as containers share the resources of an operating system while currently executing inside that OS. A container is a contained space where different processes may operate while accessing the kernel and other features of the underlying OS, as opposed to a virtual machine (VM), which is an accurate abstraction of an operating system. Many containers may operate on a single virtual machine or on hardware without an operating system.

The term "serverless" describes a scenario in which cloud users only use accessible functionality and don't control any of the underlying physical or virtual machines. In retrospect, though, they continue to make use of tools like virtual machines, containers, or specific hardware platforms. Several workloads, often from many tenants, will always be operating on any given processor and memory. The same physical computing node is shared by many tenants, and various segregation options are available on various hardware stacks.

This is a description of several cloud computing security best practices, provide High Availability for key workloads at all deployment levels, including the compute, firewall, and load balancers. Departments should make sure that security scans are activated as soon as the new application server is installed and that the servers are added to continuous monitoring. During the construction of the VM image, include security tests and policies. Disable post-application setup remote access. For all Virtual Machines (VM), Containers, and VM images, implement the proper role-based access restrictions and strong authentication. Use cloud platform pre-certified VM images instead, when pre-certification would need ongoing work. Patching VM images is preferred over patching active instances. Make that all operating system, database, and other patches are current. To protect VM security in the event of a breach, use Bitlocker, LUKS, etc. to ensure VM level encryption. Make sure the virtual machine has its operating system hardened. Take routine VM snapshots and store them in a safe location. Use file integrity monitoring wherever possible to verify authenticated modifications and spot unauthorised file changes. Keep audit logs and other logs separate from workloads. Install antivirus software on the virtual machine and make sure regular patches are applied. Do routine vulnerability assessments and penetration tests (VAPT) on the departments cloud infrastructure. Plan for Disaster Recovery for Business-Critical Apps (DR).

Network: Virtual networks come in many forms, from simple VLANs to fully functional Software-Defined Networks (SDNs). The network layer must also provide security for the data while it is in transit. A cloud service provider must comprehend the department's data transmission and reception network traffic strategy. Department to confirm CSP has put adequate security safeguards in place for separating and communicating between internal and external networks. CSP will make sure that networks with varying degrees of sensitivity are segmented properly. Nowadays, SDN is used by the majority of cloud computing platforms to virtualize the networks. Several common networking restrictions are eliminated by SDN, which isolates the network management strategy from the underlying physical infrastructure. For instance, a department may overlay many virtual networks over the same physical hardware with adequate isolation and segregation of all traffic.

In order to facilitate agility and orchestration, SDNs are also established via software settings and API calls. Although virtual networks operate on physical networks, they vary from physical networks in that abstraction enables more substantial changes to networking behaviour that have an influence on security procedures and technological advancements. Data transfers to the cloud need to be secure. Government agencies must make sure that data is protected while it is transferred to the cloud. Understanding the CSP's data migration processes is necessary since doing so typically results in more cost-effective and secure data transmission than using "manual" techniques like Secure File Transfer Protocol (SFTP). For instance, setting up a personal SFTP server on a virtual machine inside the same provider is probably less safe and dependable than delivering data to an object storage service provided by a CSP through an API.

Many alternatives for in-transit encryption exist, depending on the capabilities of the cloud platform. Client-side encryption, or encrypting data before transmitting it to the cloud, is one method. Network encryption (TLS, SFTP, etc.) is another. The third option is proxy-based encryption, where there is an encryption proxy in a trusted area between the CSP and the cloud consumer, and the proxy manages the encryption before data transfer to the CSP. The majority of CSP APIs use Transport Layer Security (TLS) by default as this is an essential security capability.

Before merging the data, it is a good idea to segregate and scan it. The faster rate of change in the cloud necessitates the offloading and external collection of logs. Government Departments would be protected from losses by a procedure like gathering logs in an auto-scale group before the cloud controller terminates such unnecessary instances. It is necessary to protect CSP cloud computing resources from the denial of service attack, which is often an external danger to public cloud services. A distributed denial-of-service attack (DDoS) is a malicious effort to obstruct a server, service, or network's regular traffic by saturating the target or its surrounding infrastructure with a torrent of Internet data. Thus, it is crucial that the department use CSP's Anti-DDoS services.

The following are some procedures that the government departments may take into account while handling network security in the cloud: things to keep in mind To access Cloud infrastructure and services, use a virtual private network (SSL or site to site). Refrain from turning off any personal firewalls on any Department computers that are networked. If appropriate, use IP Whitelisting to enable connections from certain IPs and block all others. Pre-certifying more firewall ports, load balancers, and VLANs. Compared to conventional data centres, separate virtual networks and cloud accounts lower security threats. If feasible, communication between workloads in the same virtual subnet must be restricted using a firewall policy. It is important to reduce reliance on virtual appliances that limit flexibility or result in performance bottlenecks.

Establish regulations and internal security measures to prevent unapproved traffic monitoring, contracts with third parties, and alterations to consumer networks. Departments need to make sure that all new network segments are added to continuous monitoring after deployment and registered for security checks. It is necessary to establish an automated attack response and gather more details about the breach. Among of the procedures covered by such an automated response include IP blocking, connection termination, and signature analysis.

Use a SIEM or regularly check network traffic logs to acquire real-time security warnings produced by network and application devices. To improve network isolation, leverage SDN capabilities for numerous virtual networks and various cloud accounts/segments. Identity and Access: Identity and access management (IAM) is the process of establishing and controlling

the conditions under which each unique network user is granted (or denied) access rights. These users might either be inside to the Department or external (like citizens) (e.g. employees). One digital identity per person is IAM systems' primary goal. It is crucial to preserve, modify, and track that digital identity throughout each user's "access lifetime" after it has been created. The following would constitute identity management and access control security:

Making use of multifactor authentication (MFA). MFA may be used, allowing conditional access policies to be set up and LDAP or AD authentication can be used for authentication. Find any possible weaknesses that might compromise the Department's identity. Automatic reactions may be set up to recognise suspicious behaviour linked to the identities of the Department. It is crucial to look into suspicious situations and take the necessary steps to put things right. MFA provides the potential of lowering account takeovers since using cloud services that depend only on a single factor (password) has significant dangers. MFA may be given through the following methods:

Hard tokens are tangible objects that need to be inserted into a reader or create one-time passwords for human entrance. Highest degree of security is ensured by them. As contrast to hard tokens, soft tokens are software programmes that operate on end devices. Soft tokens are another feasible alternative, however they could be compromised if the user's device is hacked, and therefore this risk has to be taken into account in any threat modelling. Out-of-band Passwords are often provided as text messages to a user's phone, where they are inputted as one-time passwords produced by a token. Message interception must be taken into account in any threat model, particularly with SMS.

Using Access Control techniques: For any Department employing cloud services, access control of cloud resources is essential. For instance, a recognised access management approach like Role-Based Access Control (RBAC) aids in controlling end users' access to cloud resources, what they may do with the resources that have been assigned to them, and the regions to which they have access. Assigning tasks to distinct roles in the cloud helps eliminate misunderstanding, which often results in human and automated mistakes, increasing the risk of security breaches. The security team for the Department must be able to assess possible dangers to cloud resources. The security team must thus enable the needed rights in order to have the necessary insight into Cloud resources. To provide rights to pertinent people, user groups, and apps that fall within a certain scope, the Department or its security team may use a variety of techniques. A subscription, a collection of resources, or a single resource may be included in a role's scope.

The Department is in charge of granting security teams the required authorizations so they may carry out their operational duties. Review of the built-in roles for responsibilities assigned is part of this process. If the built-in roles are unable to adequately address the Department's unique requirements, custom roles may also be established. If the Departments fail to implement data access restrictions, their end users may be given greater rights than they need. Constant observation of suspicious activity. Identity monitoring systems are responsible for quickly identifying questionable activities and then setting out notifications to urge further action.; It's best practise for departments to have a way to spot brute force attacks, sign-in attempts from various locations, strange IP addresses, and sign-ins from infected systems to their cloud deployments. Departments must institutionalise a thorough strategy that outlines the procedures for maintaining user identities and authorizations for Cloud services.

Crucial things to keep in mind: Department and CSP to prohibit the usage of root and generic accounts for Cloud Management and operations by adopting Identity-as-a-Service (IDaaS) for single sign-on implementation. Departments should adopt a role-per-group access paradigm for systems. Installation of multi-factor authentication and distinct cloud accounts; Department to guarantee resource-level access restrictions to access Web, application, and cloud resources. Departments to regulate user information and access levels on a regular basis. Physical and Perimeter: Controlling network traffic entering and leaving a data centre network is a key component of perimeter protection. The use of several, interconnected defences is one of the best practises. The perimeter protection provided by a firewall, which filters out potentially harmful or unknown traffic that may constitute a threat based on a set of rules about the types of traffic and permitted source/destination addresses on the network, serves as the first layer of defence beyond a router, which connects the internal and external networks. In addition, data centre providers use intrusion detection or intrusion prevention systems (IDS/IPS), which scan traffic after it has gone through the firewall for suspicious activity.

Unauthorized physical access to hardware may result in a disruption of the cloud service. By employing availability measures, CSP should safeguard its data centre infrastructure and take resilience into account. If a cloud service provider has not put in place sufficient secure or remote working environments from internal and external sources, the danger escalates.

Government Departments must request confirmation from the chosen CSPs that the required security measures are in place. In order to provide confidence, CSPs must conduct pertinent audits and compile evaluation reports. They may also prove compliance with the security requirements listed Standards applicable for Security. The CSP will be responsible for ensuring perimeter security and the physical security of the Data Center in line with the standards established for MeitY's empanelment of Cloud Service Providers. In order to prevent illegal or coerced admission into the data center's premises, it is the responsibility of the CSPs to guarantee that suitable safeguards are in place there, including security guards, protected fences, security scanners, biometric access, CCTV monitoring, Access Logs, etc.

A Crucial Factor to Have in Mind: The appointed cloud service provider is in charge of maintaining the integrity and security of the physical data centre as well as the deployment of the cloud's IT infrastructure (computing, networking, storage, and security). Some security measures relating to the physical security of the data centre that the CSPs must maintain include:

Physical infrastructure must be stored in safe locations; CSP must make sure that perimeter and interior security measures are in place to guard against unwanted access. For sites holding critical infrastructure, this involves installing physical entrance restrictions that guarantee access to only authorised employees. Protection from environmental risks - CSPs must offer protection from hazards including floods, earthquakes, lightning, fire, natural catastrophes, civil unrest, and other dangers that might impair a data center's ability to function. Data Center IT infrastructure security controls - CSP must make sure that the required safeguards are in place to stop asset loss, theft, compromise, and damage. Safety against equipment failure - CSP must make sure that the required safeguards are in place to carry out preventative maintenance on all data centre equipment in order to prevent service interruptions brought on by detected equipment breakdowns. Data centre asset removal/theft procedure - CSP must make sure that the necessary safeguards are in place to prevent the removal or theft of sensitive assets.

Secure disposal or reuse of data center equipment: CSP must make sure that all required safeguards are in place for the proper disposal of any data centre equipment, particularly any devices like storage media that may hold crucial data. Security measures for DC employees - Appropriate measures must be put in place for every employee working at a CSP's premises, including any temporary employees. Ensuring Backup, Redundancy, and Continuity Plans - CSP must set up the necessary mechanisms to perform routine backups of stored data, equipment redundancy, and continuity plans for dealing with scenarios that can result in equipment failure. On the other hand, in order to guarantee fully secure access to and use of the Cloud, Government Departments would need to assure physical security of their endpoints from both a physical and logical standpoint.

Cloud Security Evaluation: Government Departments may concentrate on doing a security assessment for their Cloud initiatives while institutionalising norms surrounding Cloud security, as described in the section above. Throughout each stage of the security assessment, Departments must pose a number of important questions to their Cloud Service Providers and to themselves.

CHAPTER 4

NEXT GENERATION MODEL IN CLOUD SECURITY ZERO TRUST

Neetha SS

Department of Computer Science and Information Technology,
Jain (Deemed to be University) Bangalore, Karnataka, India
Email Id-neetha.s.s@jainuniversity.ac.in

Zero Trust is the Next-Generation Cloud Security Model: An analyst with Forrester Research Inc. originally used the phrase "zero trust" in 2010 when they initially presented the concept's model. A few years later, Google introduced zero trust security to its network, which sparked an increase in interest in adoption among the tech industry. The next-generation IT security architecture known as "zero trust security" mandates rigorous identity verification for every device and person seeking to access resources on a private network, regardless of whether they are within or outside the network perimeter. Zero trust is a holistic and all-encompassing approach to network security that includes a variety of technologies and concepts rather than a particular solution.

The castle-and-moat model of IT network security is the traditional approach. Castle-and-moat security makes it impossible for outsiders to enter the network, but everyone within enjoys implicit trust. The drawbacks of this strategy are that once an attacker has access to the network, they may freely control anything that is present there. The fact that organisations and departments do not keep all of their data in a single location makes this restriction and vulnerability in castle-and-moat security systems worse. The diffusion of information among cloud service providers in the modern day makes it even more challenging to have a single security control for a full network and connections, no matter what setting they are in.

Zero trust security states that nobody, within or outside the network, is trusted by default, and verification is required for anybody attempting to access network resources. To avoid data breaches, an extra security measure has been implemented. Zero Trust Model Principles: Zero Trust assumes that all users are unreliable. The main idea behind a zero-trust network is that no computers or users should be automatically trusted since it is assumed that there are possible attackers both within and outside the network. The Zero Trust Model's guiding principles, according to Forrester Research, are as follows:

By segmenting and activating Layer 7 policy, you may make sure that only valid traffic or application communication is permitted. Use a least-privileged approach to access and carefully enforce access control. This indicates that users should only be granted access when necessary and on a need-to-know basis. By doing this, users' exposure to delicate network components is reduced. Examine and record each cloud traffic event. Alternatively, it would be thought that an attacker might easily get access to a department's network. Microsegmentation is an idea used in zero trust networks. The process of microsegmenting involves dividing security perimeters into smaller areas while maintaining independent access for various network segments. For instance, a network that uses microsegmentation and has files stored in a single data centre may have hundreds of distinct and secure zones. Without separate and specific authorisation, a programme or person having access to one of these zones cannot access the others.

MFA, or multi-factor authentication, is another essential component of the zero-trust architecture. Simply said, MFA requires more than one piece of evidence for user authentication; access is not granted by inputting a password alone. Users of cloud services

must input a code delivered to another device, such as a cell phone, in addition to their password, allowing two-factor authentication. Zero trust requires stringent restrictions on device access in addition to user access control.

Zero trust systems must keep track of the variety of devices attempting to connect to their network and verify that each one is permitted. As a result, the network's attack surface is significantly reduced. For enterprises of any size or kind, implementing zero trust security via cloud-based architecture is more economical and adaptable. IT departments may benefit from greater security without losing usability by doing away with the maintenance-related costs of on-premise hardware.

Crucial things to keep in mind: Start with passive application discovery and keep an eye on the network. Provide cooperation with key stakeholders who are familiar with typical traffic patterns and intersystem interactions as well as the finding of the existing linkages. After relationship appropriation and application behaviour have been validated in accordance, enforcement rules should be put into place. Based on how data moves across the network and how users and apps access sensitive data, create a zero-trust architecture. This will help in figuring out how to partition the network. It could also help security teams locate where access restrictions should be placed between the boundaries of various network parts. Employ monitoring and real-time auditing to keep track of all privileged sessions and metadata. Audit each system individually to get a complete view of goals and results.

Standards that Apply to Security: There is a thorough list of NIST Controls referenced from the paper "Indian Governmental Cloud Selecting Framework" that was released and is available on the MeitY Portal to help with planning on information security management for government departments. The department may adhere to the globally accepted information security standards listed below when implementing a cloud platform.

ISO 27001: This standard offers an Information Security Management System best practise (ISMS). This management system standard is intended to handle sensitive data for a business as well as its predetermined rules and procedures. The company is more susceptible to data breaches and cyberattacks when there are no ISMS in place. As a result, this system is an essential part of an organization.

ISO 27002: The ISO 27002 Standard provides recommendations for corporate Information Security Management System (ISMS) procedures, including the choice, implementation, and maintenance of administration of controls while taking the organisational environment's information security risk into account.

ISO 27017: The ISO 27017 standard is meant to help cloud-based enterprises adopt controls. This standard applies to businesses who keep data in the cloud as well as businesses that provide cloud services to other businesses that may have sensitive data.

ISO 27018: Specifically created to safeguard Personally Identifiable Information (PII) stored and/or processed in the cloud, ISO 27018 is for cloud computing organisations. This standard's main audience is cloud providers, not cloud customers. Customers may feel more confident when dealing with businesses that handle sensitive information thanks to this standard.

PCI DSS: The Cloud Service Provider (CSP) and its users / customers are both accountable for cloud security. The PCI DSS will apply to that environment if card payment data is kept, transferred, or processed there. Often, this will include validating both the infrastructure of the CSP and the client's use of the environment. A client is still responsible for making sure

that their cardholder data is safe in accordance with relevant PCI DSS rules, despite the provider and client sharing management of security measures. A collection of logical, procedural, and physical security criteria known as the Payment Card Industry Data Security Standard (PCI DSS) are necessary for businesses that handle credit and debit card transactions as well as payment apps. This standard must be followed by all entities that store, handle, or transfer cardholder information.

Sector-specific norms:

SOC 2/SSAE 16: Standards for Attestation Engagements Statements No. 16. Organizations that operate data centres are required by SSAE 16 to provide written reports outlining the controls at businesses that offer consumers services. The SOC 2 report focuses on service provider controls related to security, processing integrity, availability, confidentiality, and privacy of a system. Customers may view their data at any time and are guaranteed to have it kept private and safe while it is being stored and transported.

Multi-cloud/Hybrid Cloud Environment Cloud Security: It is difficult to secure a hybrid IT infrastructure that utilises many Clouds. Government agencies planning to utilise a combination of on-premise and cloud technologies would think the hybrid approach is more secure than using just internal systems. Hence, increased security would play a key role in increasing their use of hybrid or multi-cloud services. Many CSPs did not, until recently, have the required safeguards for compliance and security that the government departments would anticipate. Now, this situation has drastically altered.

Using several, cloud and storage services inside a same heterogeneous architecture is referred to as multicloud. The spread of cloud assets, software, apps, etc. across several cloud-hosting environments is sometimes referred to as a heterogeneous environment. This multicloud environment aims to remove reliance on any one CSP by adopting a standard multicloud design that employs two or more public clouds, or numerous private clouds. The chosen CSP must safeguard the Cloud infrastructure, but the Government Department is in charge of protecting any data that it uploads to the Cloud. Hence, in order to ensure that the chosen CSPs/MSPs satisfy the relevant regulatory and security criteria, the Government Department must ultimately do the appropriate due diligence while making its selection.

In a multi-cloud context, protecting the network perimeter is less important than protecting data wherever it resides, whether it is at rest or in transit. Regarding the department's obligation to protect its data, in a multi-cloud environment, the focus switches from safeguarding the network perimeter to securing the data itself, whether it is at rest or in transit. The goal in a multi-cloud scenario is to thoroughly comprehend data flows and safeguard them according to their level of sensitivity. The following are some precautions to keep in mind when contemplating a multi-cloud deployment or environment:

The primary security policy of the Department. Department internal security teams or MSPs would need to concentrate the security control in order to spot risks across a hybrid multiple Cloud platform and efficiently integrate security methods to meet demands of each of the Cloud platforms access to data. To improve the department's security capabilities, information about all security procedures and technologies adopted must be communicated among the designated points of contact who are in charge of each Cloud platform. A multi-cloud architecture's safe integration is made possible by having a unified mechanism for security enforcement that ensures a consistent approach to Cloud platforms. Scaling Cloud security may be facilitated by using third-party services for automation.

Develop a security strategy for a hybrid multi-cloud environment. Although the Departments must make sure that their applications are current, it is equally crucial to make sure that their security features are regularly updated to match the changing security needs of their IT environment. Cyberattacks are always looking for security gaps to exploit and using creative methods to compromise systems nowadays. Security professionals must regularly assess the multi-security clouds via real-time reporting in order to monitor risks to a multi-cloud architecture.

The applications secure communications. While the communications inside and between apps in a multi-cloud architecture are safe, many Departments could overlook protecting the communications that are intended to regulate how the applications behave. This is referred to as the control plane, and an effective multi-cloud security strategy should include the need to encrypt communications that fall under the control plane's purview. The security teams for the department must make sure that the communications controlling virtual machines and containers are encrypted. These interactions are often left unencrypted and open, making it possible for malevolent parties to exploit these weaknesses.

Check to see whether departmental employees are adhering to security policy. Unauthorized data and services might be accessed by certain end users, which would constitute one of the greatest security breaches. Unrelated individuals who are given access to confidential or sensitive data run the risk of exposing it to security lapses and even cyberattacks. In these situations, Departments must ensure that any acquired software is patched and protected before distributing it to their staff. Workers must also get training on how to follow the strict security procedures intended to stop security breaches.

Cloud Access Security Brokers (CASBs) may also be used to maintain security controls across the Department's hybrid or multicloud installations. CASBs (or Cloud Security Gateways) may monitor DNS requests, integrate with an existing network gateway or monitoring tool, or monitor network traffic to find internal uses of cloud services. When learning about the services that people are using, the majority of these technologies provide user activity monitoring on authorized services, mostly via API connections (where accessible) or inline interception (man in the middle monitoring). Several provide controls to efficiently govern the usage of sensitive data in Cloud services (SaaS, PaaS, and IaaS), provide security alerts, including DLP.

Security policies must be maintained in accordance with the policy definitions regardless of where the application is deployed. A centralised policy management system that covers several Clouds and data centre facilities may be used to accomplish this. The capacity to consistently enforce, monitor, and manage the policies is a need for this system. To enable centralised security monitoring, incident management, and event analysis, infrastructure and application logs from the many CSP environments that interact to offer services should be collected in one place. This will provide for a unified picture of all emerging risks across the data assets of the companies. This system will be built on machine learning and artificial intelligence, and include capabilities for user and entity behaviour analysis (UEBA). As a result, the time and effort needed to find anomalies would be decreased, which would shorten the time it takes to resolve the problem and lower the cost of a data breach.

A centralised identity access management system (IDAM solution) should be used by all end users. Several user identities may be accommodated and managed by this method. The ability to support several user types, a single sign-on process, and the incorporation of the provision for multi-factor authentication. Data from all the various environments would be made accessible even after a catastrophe, when this information is most required, thanks to a

centralised data backup system. A centralised analytics and monitoring system must be used to manage security infrastructure across diverse settings. This tool may also be used to orchestrate IT infrastructure, including security tools and technologies. A central HSM or a key management system may be used to handle the encryption keys and certificates.

Security Operation Centre (SOC) will be empowered by automation technologies thanks to automated processes and playbooks. Having an automated procedure, for instance, to approve modification requests. Hence, by following a few guidelines, the Departments may adopt a multi-cloud deployment in a safe manner and take use of the numerous services provided by the CSPs that MeitY has appointed.

Cloud Security Governance: Departments must make sure their cloud computing initiatives are well integrated into their whole information security programme in addition to considering whether to migrate applications onto the cloud. This entails comprehending the process as it relates to the ICT Resilience Lifecycle, which takes into account both the general governance process and the prevention aspects, such as risk management and information security, as well as the reaction elements, such as incident management and continuity planning.

Governance: Departments must take into account their entire IT security governance processes when they decide to deploy apps in cloud environments. Establishing a broad understanding of how the cloud fits into the essential information security procedures, as well as the procedures' aims and objectives and a plan for implementing them, should be the first step in this process.

Risk management: Cloud security needs must be in line with an internal data categorization system and a department-wide awareness of risk levels.

Incident management: When applications are deployed in the cloud, each Cloud Service Provider for the Department must be connected into the Department's centralised issue response processes.

Business continuity plans must account for assets that have been migrated to the cloud and must be updated and tested often to reflect new cloud architecture and provider models. Security professionals need to define the scope and limitations of all security functions that might apply to cloud environments as part of this effort. They also need to come up with a strategy for enhancing and tracking the performance of all cloud stakeholders, such as service providers, users, and technical staff. Lastly, they should provide senior management the controls necessary to oversee the whole cloud computing programme as well as the instruments required to obtain insight into cloud security, such as a security-level dashboard (Figure 4.1).

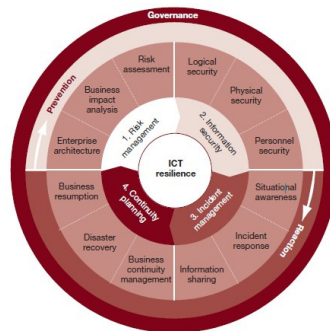


Figure 4.1: ICT Resilience Lifecycle.

The shared responsibility paradigm for cloud security. The issue of "Who is responsible for security in the cloud, the Cloud Consumer or the Cloud Service Provider?" is one of the most important when considering the cloud.

In order to secure sufficient security measures, cloud computing adheres to the shared responsibility paradigm. The Cloud Service Provider will always be in charge of protecting physical infrastructure and the virtualization platform itself. While comprehending the underlying threats, the Cloud customer, in this instance the Government Department, is in charge of formulating and institutionalizing appropriate security procedures. Choosing whether to employ dedicated hosting instead of a shared host, for example, or when to correctly configure the virtual network and firewalls.

Together with the best practices for cloud security, the following additional indicators are ones that Government Departments should timely monitor with regard to their cloud deployments: Install the most recent fixes for the operating system and installed apps. Verify that the MSP/CSP solution complies with all applicable legal, organisational, and privacy standards. MeitY has onboarded CSPs that fulfil a certain set of technical and regulatory standards via its empanelment RFP. Please see <https://meity.gov.in/content/gi-cloud-meghraj> for the list of conditions that CSPs must satisfy in order to be appointed by MeitY. The Department may assess any additional criteria based on the needs of the project.

To operate the cloud service, use specific PCs equipped with MFA, strong password policies, access-controlled rights, and encrypted communication channels. Refrain from giving the MSP/CSP access to sensitive systems that are not under their control or account credentials. Use measures to safeguard information when it is being transferred between the Department's end and the Cloud Service Provider. Think about fully encrypting sensitive Department data while it is at rest and keeping custody of the encryption keys. Even if periodic scanning and audits are a service supplied under contract with the MSP/CSP, take into account frequent scanning and monitoring for non-standard or suspicious code, files, or folders on hosts, and assure regular audits.

Conduct recurring cloud auditing activities (every six months) to make sure the CSP is adhering to project requirements from the Department. Use anti-malware and other security solutions on the infrastructure and/or assets of the Department. Think of devices that can both detect and treat infections. All executables downloaded to Departments' infrastructure should be examined before execution, and anti-malware applications and other security technologies should be maintained and updated. Make sure MSPs and CSPs regularly analyse network and system logs for any unusual behaviour or traffic that could point to a possible breach.

Establish a backup and recovery strategy for any important data. Make sure MSP/CSP additionally uses data backup and recovery strategies. To lessen the effects of data or system loss and to hasten the recovery process, perform and test frequent backups. This data should be retained on a different device, and backups should be kept offline since network storage may also be impacted. Contractually reserve the right to get a copy of a compromised virtual server for forensic examination by the Department. To maintain security in the cloud environment, both the CSP and the cloud user have certain obligations; however, in some instances, these obligations overlap. Many security issues that cloud users are now experiencing stem from a lack of awareness of these shared responsibility areas.

Moreover, it's probably reasonable to conclude that consumers, not providers, bear the most of the liability for cloud security in more domains. Government agencies will be far better able to maintain a high degree of cloud security if they are aware of the fundamental boundaries of duty. In essence, the CSP has to guarantee that the infrastructure created inside

their platform is innately reliable and safe. On the other hand, the cloud consumer is in charge of some customisable Cloud features including network setup, account access, application administration, compute configuration, and data encryption. To further understand the duties between CSP and the government departments while setting cloud security for their infrastructure, you may refer to the shared security model, which is shown below Figure 4.2 and Figure 4.3).




Figure 4.2: Shared Responsibility Model in Cloud

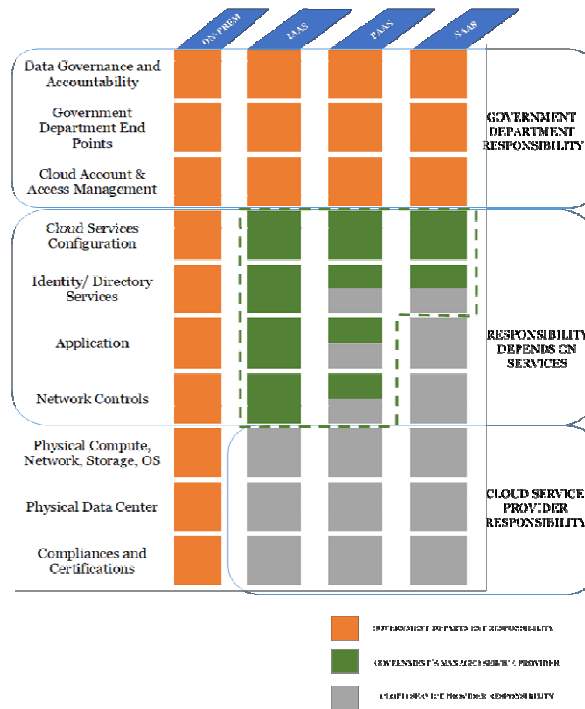


Figure 4.3: Shared Responsibility Model in Cloud (including MSP).

The term "cloud security" refers to the measures taken to protect the data, infrastructure, and applications that are essential to the usage of cloud computing. This includes organisational policies, technology, and controls. Cloud-based apps and their data are becoming more dispersed in order to increase agility and lower costs for government departments. Private clouds, public clouds (hybrid or dedicated), and software as service (SaaS) applications are all examples of this trend.

The worry about data exposure has grown, and cloud security is now a top issue. So, the issue is to strike a compromise between the Department's demand for agility and simultaneously strengthening the security of apps and their data as they move across different Cloud platforms. To avoid assaults that try to exfiltrate data—through a lateral attack or from an external location it is crucial to get the required awareness across all places where apps and data are stored.

The application team, network team, security team, compliance team, or the IT infrastructure team may be in charge of various facets of cloud security. Nonetheless, the CSP and the Government Department also share responsibility for cloud security. As the Cloud is hosted locally in the Departments' own data centres in an on-premises configuration, they are alone in charge of its security. The infrastructure, physical network, hypervisor, operating systems, virtual network, service configuration, firewalls, identity and access management, etc. are all covered by this. In this case, the Departments are in charge of both the data's ownership and security.

The CSP is the owner of the physical network, hypervisor, and infrastructure for an IaaS service in any Cloud deployment type. The workload, its applications, virtual network, access to the cloud environment, and the deployed data, on the other hand, are owned by the Departments.

In a SaaS configuration, CSPs are primarily in charge of the platform's security, which includes physical security, infrastructure security, and application security. Data belonging to customers cannot be owned by CSPs, and therefore cannot be held liable for how apps are utilised. In order to avoid and reduce the danger of malicious data exfiltration, unintentional disclosure, or virus infiltration, the Department must first focus on security. In this instance, the CSP is solely in charge of application security, while the Departments are in charge of any required environmental setups and cloud-based data.

The need of CSPs to maintain and adhere to the most recent regulatory compliances and certifications, along with having the certification updated on a regular basis, is one of the important elements that should be taken into account when comparing a cloud environment to one that is on-premise. In contrast to cloud computing, where the CSP is in charge of maintaining compliance with the standards required by government departments, such governance procedures may be a cost centre for the departments in an on-premise configuration. The responsibility for security of data, applications, and infrastructure is more the duty of the chosen CSP than the Department itself when government departments go from an on-premise configuration to an IaaS, PaaS, or SaaS service. Yet, the Department will always be responsible for ensuring the security of its own data, regardless of the platform utilised. The Government Departments must be certain that their CSPs have implemented the necessary security procedures before they can securely allow their apps.

A Department must also have the required tools in place in order to assess, manage, and protect the risks efficiently and enable data security in order to make up for what is not covered by the CSP. For a SaaS offering, these tools might be able to provide visibility into activities within the SaaS application, detailed analytics on the service usage to prevent risk

to data and violations of compliance requirements, policy controls to call for enforcement and even quarantine in the event of a violation, real-time threat intelligence on known and also to detect unknown threats to prevent new malware insertion points, and policy controls to require enforcement and even quarantine in case a violation occurs (Figure 4.4).

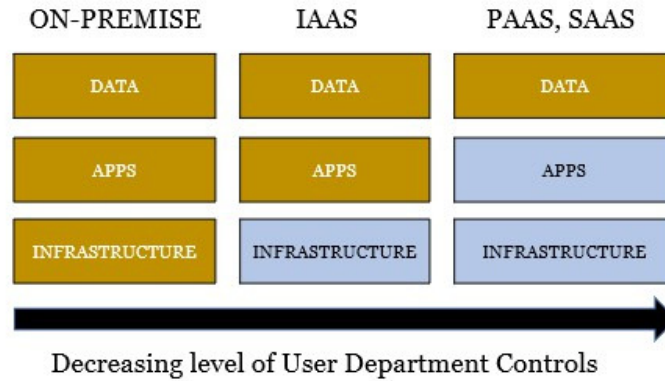


Figure 4.4: User Department Control

The host Operating System (OS), the virtualization layer, as well as the physical security of its data centre facilities, are all layers that the CSP monitors and controls in order to provide a safe cloud. The Department or its MSP is required to setup and administer the security controls for the guest OS, additional applications (including updates and security patches), and the firewall as well in order to ensure security inside a specific Cloud environment. Thus, it is essential that CSPs and government agencies choose the cloud for their deployments after fully understanding their roles and obligations with regard to security. Workspace-as-a-Service (WaaS), among others, and Communications-as-a-Service (UCaaS). These extra services cross over with the three fundamental service models and muddle the distinction between SaaS, PaaS, and IaaS, further complicating maintenance and security obligations. The most common cloud service types, however, are SaaS, PaaS, and IaaS. Each differs in how they are used and secured. The shared security model, shown in Figure 4.5, is a typical way to depict this. These models specify who is in charge of technology, security, data, etc.

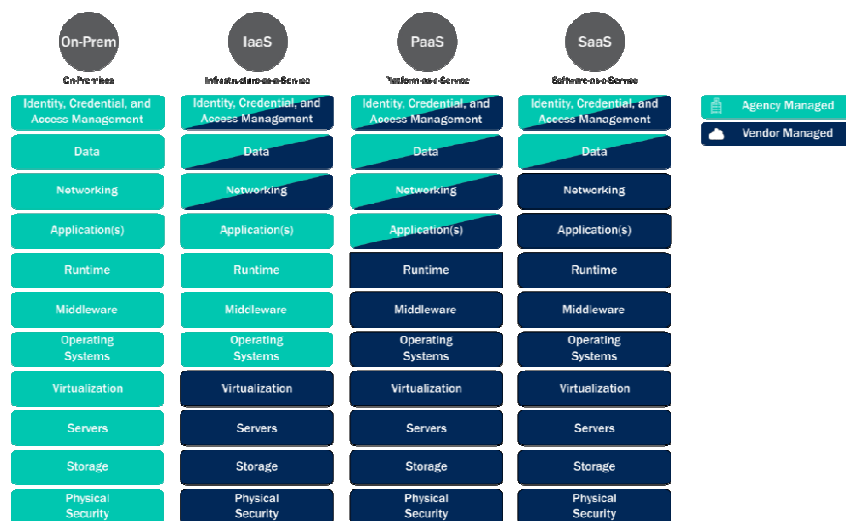


Figure 4.5: Responsibilities for Different Service Models

The shared security model (Figure 4.5) demonstrates that the service provider bears a significant portion of the burden of protecting a SaaS product. This implies that the organisation using the service is putting greater faith in the service provider, nevertheless. In contrast, IaaS places a greater burden on the agency, while sharing certain obligations with the cloud service provider (CSP) and other duties. From one vendor to the next, CSPs may define this shared security relationship differently. Agencies are required to recognise and comprehend the separation of duties between themselves and their CSP. Agencies should carefully create service level agreements (SLA) with each of their CSPs to specify expectations and obligations. When CSPs improve their service offerings, agencies may discover that they need to modify their security posture in order to keep up with them. Agencies need to make sure they are aware of the security posture of the CSP(s) they have chosen both at the beginning and over time.

A sub-agency may utilise services supplied by a parent agency, for example. Agencies may also use services offered by other agencies. These services may vary from IaaS environments that the sub-agency is given access to by the parent agency to SaaS apps like email. Between the parent agency and the sub-agency, coordination of duties and responsibilities must be recognised in several circumstances, including but not limited to Configuration management, identity, access, and credential management (ICAM), log monitoring, and analysis.

Alternatives for Cloud Services: SaaS, PaaS, and IaaS are the three main types of cloud service, as was before noted. While building effective designs, agencies should take into account the security implications of each kind of cloud service since each one has distinctive qualities. Agencies should be aware that CSPs that provide IaaS services also often provide PaaS and SaaS services, while CSPs that provide PaaS services also provide SaaS services. A single CSP being used by an organisation to provide several cloud service models is thus not unusual. A company may have certain CSP services operating on-premises, in satellite or distant offices, in data centres, and/or in the cloud. In addition, some CSPs provide the option to host their services on-premises utilising pre-packed hardware and virtualization. The subsections below go into further information about each cloud service.

Software-as-a-Service: SaaS products are often specialised and focus on a particular business requirement, such as email, document management, or human resources operations. SaaS offers may be programmes or application programming interfaces (APIs) that can be combined with other services, but they are normally supplied over the web. Little shared obligations exist between the service provider and the hardware and software, but both agencies and service providers must protect any application or API access to these environments. Others will not have choices for authentication integration and will have their own identity world. Certain SaaS providers will be able to interface with current identity access providers. Certain SaaS solutions may be included in the service portfolios of IaaS and PaaS providers.

Platform-as-a-Service: Vendors in PaaS provide platforms for developing solutions, including web servers and databases. Although certain PaaS functionalities are often supplied with IaaS, they may also be purchased alone. Agencies may concentrate on developing services for mission requirements rather than purchasing, installing, and maintaining server infrastructure or the application or database server, which is one benefit of PaaS over IaaS. As a result, an agency may concentrate on managing platform resources as well as creating and implementing services and solutions rather than maintaining the underlying infrastructure.

Infrastructure as a Service:IaaS environments will provide a wide range of capabilities and services that may be utilised to create and coordinate solutions. To use these resources while creating solutions, agencies should be aware of and take into account cloud native characteristics. These characteristics include flexibility, scalability, and the virtualization of resources like operating systems, networks, containers, etc.

Forms of Deployment:There are four alternative approaches to install the services mentioned above in the cloud. The various cloud deployment types and associated NIST definitions are as follows:

Private:A company with several clients is given exclusive access to the cloud infrastructure (e.g., an agency with multiple business units). The organisation, a permitted third party, or a combination of them may own, manage, and operate it. The infrastructure may be located on-site at the company or off-site with the cloud service provider.

Community: The cloud infrastructure is made available to a certain group of customers who have similar worries (e.g., mission, security requirements, policy, and compliance considerations). One or more organisations, a permitted third party, or a combination of these bodies may own, manage, and operate it. Either on-site or off-site infrastructure is possible.

Public: The broader public is permitted to utilise the cloud infrastructure. One or more organisations, a permitted third party, or a combination of these bodies may own, manage, and operate it. Off-site is where the infrastructure is located.

Cloud infrastructure that combines two or more of the aforementioned deployment methods is known as a hybrid (i.e., Private, Community, or Public). In order to retain the interoperability of data and applications, several deployment models are linked through a standardised or proprietary technology provided by the provider.

Many people believe that government cloud services fall within the category of community cloud models when referring to them. Government cloud installations could come with certain benefits over public cloud services, such as US residents working at the CSP data centre, but there might also be some drawbacks. New security tools and features are often first made available to the public model by CSPs. These same security tools and functionalities could not be made available to government cloud installations for weeks, months, or even years. Moreover, certain tools provided by CSPs in a public cloud deployment may never have certain functionality enabled in the related government deployment. Also, only U.S. locations are allowed for federal cloud installations. Certain organisations may need to reach a worldwide audience, and a public cloud deployment is the most effective way to do this.

Multi-cloud environments are likely to be used by multi-cloud agencies. Agencies using multiple clouds must optimise their setups while keeping situational awareness and appropriate security procedures in each CSP they use. Agencies have the option of protecting each of these services separately or by maintaining a comprehensive picture of their security posture across all the services they use. To administer security policies in a centralised manner, agencies are advised to adopt solutions that provide a comprehensive picture of their application and infrastructure across all CSPs. For security investigation across different CSPs, agencies may also choose to employ technologies provided by third-party suppliers and CSPs. Based on their unique requirements, agencies should decide which of these technologies will enhance their security posture the most. Agencies should assess the advantages and drawbacks of both independent solutions made for multi-cloud systems and

security technologies provided by CSPs. Agencies should employ security solutions that are compatible with several CSPs wherever feasible.

Agencies should consider how to maintain situational awareness and sound security procedures while monitoring each cloud service they utilise. Finding uniformity in the security data across the many cloud services a company utilises is crucial. Due to differences in field names and the quantity of fields that each service provider makes accessible, data normalisation of logs by type will aid in achieving parity. Agencies must decide whether, and which logs will be backhauled, they will consolidate logs to a single place for analysis. Certain logs, such as authentication logs when employing an integrated identity access provider across several CSPs, will have a centralised location. For log management, agencies are required to be aware of and adhere to OMB Memorandum (M)-21-31.

Agencies must decide how they will implement authentication and access control for each service when intending to deploy cloud services. If they have more than one CSP, they must decide which one will host the identity provider before making any decisions about where their identity provider will be located (e.g., on-premises, in a CSP). If at all feasible, agencies should utilise phishing-resistant multi-factor authentication (MFA)^{14, 15}, and they should think carefully about whether to deploy convenience features like single sign-on.

Agencies should be aware of the possibility of vendor lock-in while operating in a multi-cloud environment. When a tenant is dependent on the resources and services provided by a CSP, vendor lock-in takes place. In certain circumstances, opting to design systems that impose vendor lock-in might provide several benefits. In contrast, in some circumstances, agencies may need to design solutions with little vendor lock-in so that they may be quickly deployed across several services with little modifications to configurations and deployment settings.

CHAPTER 5

INTRODUCTION TO FEDRAMP

Dr. K Suneetha

Professor, Department of Computer Science and Information Technology,
Jain (Deemed to be University) Bangalore, Karnataka, India
Email Id-k.suneetha@jainuniversity.ac.in

The OMB Memorandum "Security Authorization of Information Systems in Cloud Computing Environments," sometimes known as the FedRAMP Memo16, created FedRAMP in 2011. FedRAMP offers a risk-based, cost-effective strategy for the Federal Government's adoption and usage of cloud services. With a focus on security and the protection of government information, FedRAMP gives agencies the ability to adopt cutting-edge cloud technology. By offering a standardised method for security and risk assessments for cloud technologies and federal agencies, FedRAMP is a government-wide initiative that encourages the use of safe cloud services across the Federal Government. According to the FedRAMP Memo, FedRAMP is appropriate for:

1. Executive departments and agencies that purchase both for-profit and nonprofit cloud services from information systems that support their assets and operations, including systems that are supplied or managed by other departments or agencies, contractors, or other sources.
2. All NIST-defined cloud service models, such as Infrastructure as a Service, Platform as a Service, and Software as a Service; all cloud deployment types, such as Public Clouds, Community Clouds, Private Clouds, and Hybrid Clouds.

According to the FedRAMP Memo, each Executive department or agency must also:

1. For any Executive department or agency usage of cloud services, employ FedRAMP for conducting risk assessments, obtaining security authorizations, and granting Authorization to Operate (ATO).
2. When initiating, evaluating, granting, and revoking security authorizations for cloud services, use the FedRAMP Program Management Office (PMO) procedure and the FedRAMP security authorization standards that have been authorised by the Joint Authorization Board (JAB) as a baseline.
3. Verify that contracts that are relevant to CSPs adequately require them to adhere to FedRAMP security authorisation requirements.
4. Create and put into place, in line with DHS recommendations, a capacity for incident response and mitigation for security and privacy issues involving cloud services.
5. Verify that contract conditions for contractor reviews and inspections are applicable to CSPs and that procurement requirements address sustaining FedRAMP security authorisation requirements.
6. Demand that CSPs route their traffic in accordance with DHS recommendations so that the service complies with the standards of the Trusted Internet Connections (TIC) programme.
7. On or before April 30 of each year, agencies must submit to the Federal Chief Information Officer (CIO) (1) a written certification from the Executive department or agency CIO and Chief Financial Officer (CFO), and (2) a list of all cloud services that

- they have determined do not meet FedRAMP security authorization requirements, along with an explanation of why and suggested solutions.
8. Benefits: Reduces discrepancies, financial inefficiencies, and duplication of work.
 9. Creates an alliance between the public and commercial sectors to foster creativity and the development of safer information technologies.
 10. Allows agencies to use security authorizations on a government-wide scale, establishing clear standards and procedures for security authorizations that enable the Federal Government to speed up the adoption of cloud computing.
 11. Goals: Increase the adoption of secure cloud computing by government organisations.
 12. Strengthen the framework for securing and approving cloud technology by the government.
 13. Create and maintain enduring relationships with FedRAMP participants.
 14. Provide direction to vendors and agencies on how to use safe cloud solutions.

In order to better fulfil the goals of our programme, FedRAMP is always looking for methods to modernise and automate. The Open Security Control Assessment Language (OSCAL), a collection of formats expressed in XML, JSON, and YAML, was created by FedRAMP in collaboration with NIST and business. Control catalogues, control baselines, system security plans, assessment plans, and outcomes are all represented in these formats in a way that is machine-readable. To speed up the creation and evaluation of authorization packages, OSCAL is being applied to FedRAMP baselines and security package components. FedRAMP has made available open source tooling, such as the OSCAL Generator and Conversion tools¹⁸, to help customers get started with OSCAL. FedRAMP will continue to give ongoing process improvement a high priority in order to build on the foundation laid in Fiscal Year 2021 and benefit all stakeholders. Benefits will have the following effects on important stakeholder groups:

1. Agencies will have a better understanding of risk management, which will allow them to make more-informed decisions when approving cloud service products. As a consequence, their companies will be able to implement new services more quickly.
2. To reduce the work and time required for authorizations, CSPs and Third Party Assessment Organizations (3PAOs) will have automated processes for self-testing, developing, submitting, and remediating security packages. Also, CSPs will have automated channels for continuous monitoring, which will speed up the response to cybersecurity issues.
3. FedRAMP will get better packages at the beginning of an authorization lifecycle, which will lead to fewer obstacles throughout the review process. Package evaluations will be simplified, less burdensome on stakeholders, and lead to quicker decision-making using automated forms.
4. Stakeholders in FedRAMP: Their Positions and Duties
5. FedRAMP involves four stakeholder groups: CSPs, 3PAOs, federal agencies, and the JAB.

Providers of cloud services

The Federal Government is one of the biggest consumers of cloud computing, and CSPs provide agencies with cutting-edge technologies that enable them to achieve their pressing mission demands while saving time and money. CSPs that provide cloud services to the federal government should apply for a FedRAMP authorization, commit to learning about the programme, and utilise the FedRAMP templates to stay in compliance with the program's standards for shared responsibility. All Federal Civilian Executive Branch (FCEB) agencies accept FedRAMP as the standardised security architecture for all cloud goods and services.

CSPs must undertake ongoing monitoring of each allowed service and only go through the FedRAMP Authorization procedure once for each CSO. To increase efficiency across the government, all agencies examine the same deliverables for continual monitoring. The FedRAMP PMO assists CSPs in navigating the FedRAMP procedure and comprehending the requirements by offering training, advice, and advisory assistance. To ensure alignment and compliance with the shared responsibility standards set by FedRAMP, CSPs that provide CSOs for federal use should be dedicated to understanding FedRAMP and using FedRAMP templates.

Organizations engaged in independent evaluation

By evaluating the security of a CSO, Third Party Assessment Organizations (3PAOs) play a crucial part in the authorisation process. They provide initial and recurrent examinations of cloud systems as impartial third parties in accordance with federal security standards. The 3PAO evaluations serve as the foundation for the Federal Government's informed, risk-based authorisation choices for the use of cloud-based goods and services. A Readiness Assessment Report (RAR), which is necessary for the JAB Authorization procedure, is produced by 3PAOs during FedRAMP evaluations. While a RAR is not required for agency authorizations, it is strongly advised. The 3PAOs create a Security Assessment Plan (SAP) and Security Assessment Report for both JAB and agency authorizations (SAR). An Authorizing Official (AO) within the government must be consulted in order to approve the SAP and SAR.

Federal Organizations: FedRAMP enables federal agencies to utilise cloud services to update their technology and further their missions in a safe manner. To achieve this, organisations assess the security of cloud services using FedRAMP's defined baselines. To examine the security posture and approve the CSO for any cloud services they desire to employ, agencies collaborate with CSPs. Agencies and CSOs are urged to complete FedRAMP training and create system-level security artefacts using FedRAMP templates in order to build a uniform approach to the deployment of cloud computing across the federal government. After being "Authorized" in the FedRAMP Marketplace, agencies may evaluate and reuse CSO security packages by providing their own permission to use.

The thing. The "do once, use many" philosophy of FedRAMP allows agencies to increase the range of secure cloud services that are accessible to the Federal Government. Board of Joint Authorization: The JAB serves as FedRAMP's main governing and decision-making body. The Chief Information Officers from the Department of Defense (DoD), Department of Homeland Security (DHS), and General Services Administration (GSA) make up the Joint Advisory Board (JAB).

Creating and maintaining a current list of the FedRAMP security authorization criteria. Endorsing 3PAO certification standards. Examining cloud service authorization requests depending on the priority queue. Issuing provisional authorizations for cloud services, which Executive departments and agencies may use as a starting point for approving security authorizations and the corresponding ATO for usage. Ensuring that temporary authorizations are routinely examined and updated. Executive departments and agencies should be informed of any changes to provisional authorizations, including their withdrawal. Creating and disseminating requirements for the authorization package review priority queue. More information about the goals and duties of the board may be found in the JAB Charter.

FedRAMP Security Considerations: The purpose of FedRAMP is to provide federal agencies and cloud technology a standardised method for security and risk assessment. CSPs and

agencies need to be mindful of continuous security concerns and standards even after permission.

Constant Watching: After getting authorisation, a system's security posture is inevitably going to alter. This might be as a result of new vulnerabilities being found, modifications to the cloud service provider's hardware or software, or both. For the purpose of making risk-based choices, ongoing assessment and authorization provide government agencies adopting cloud services a way to track changes to the security posture of a system. The monitoring of the parts of the environment that CSPs do not monitor that are often covered by different authorizations is still the responsibility of the agencies employing cloud environments (See Section 3.1 for how the layers of the cloud service models work with various roles and responsibilities).

If a CSP receives a FedRAMP Authorization, the FedRAMP technique it should use is described in the FedRAMP Continuous Monitoring Strategy Guide (via agency authorization or JAB provisional authorization).²⁰ To allow informed risk-based decision-making, the CSP must continually monitor the cloud service offering to identify changes in the security posture of the system. The manual gives the CSP instructions on how to use FedRAMP to continually monitor their systems. The FedRAMP Handbook for Multi-Agency Continuous Monitoring is one of the extra continuous monitoring guideline publications offered by FedRAMP. In order to share the burden of ongoing oversight, lessen the leveraging agencies' reliance on the original authorising agency, and work together with the CSP and other member agencies to ensure that the cloud service continues to meet their needs, FedRAMP strongly encourages agencies to make use of this guide. In order to establish a uniform strategy to manage the security posture of CSOs in the continuous monitoring phase, agencies should also think about employing the FedRAMP Continuous Monitoring Performance Management Handbook. With the CSP's approval, agencies may integrate security artefacts from vendors into agency governance, risk, and compliance (GRC) capabilities to boost efficiency through automation and tooling and guarantee that the cloud service security posture is visible to agency risk management framework (RMF) stakeholders and approving officials.

Incident management: An "incident" is defined as "an occurrence that (A) actually or imminently jeopardises, without lawful authority, the integrity, confidentiality, or availability of information or an information system; or (B) constitutes a violation or imminent threat of violation of law, security policies, security procedures, or acceptable use policies" by the Federal Information Security Modernization Act of 2014 (FISMA),²³ at 44 U.S.C. 3552(b)(2). In this paper, the phrases "security event" and "information security incident" are used synonymously.

A CSP starts the continuous monitoring phase after obtaining a FedRAMP Agency ATO or Provisional-ATO (P-ATO) for its service offering. In order to guarantee that all incident management is transparent and that all stakeholders are informed of the current status and remediation activities, clear and prompt incident communication to relevant stakeholders is a vital component of continuous monitoring. The processes FedRAMP stakeholders should follow for reporting information about information security events, including how to react to published Emergency Directives, are outlined in the FedRAMP Incident Communications Procedures²⁴ document. Every suspected or verified event that compromises the confidentiality, integrity, or availability of a cloud service or the data/metadata that it stores, processes, or transmits must be reported to FedRAMP by CSPs. Agencies and other impacted customers may take action to preserve crucial data, maintain a regular level of efficiency, and make sure a thorough resolution is attained quickly by reporting actual and suspected events.

Authorization Limits: All components of an information system must be approved for operation by an authorising authority and exclude separately authorised systems, to which the information system is linked," is how NIST defines the security authorization boundary. For CSPs to create the "authorization border" connected to their CSO and support their FedRAMP Authorization package, FedRAMP offers guidelines.

Authentication Limit: A diagrammatic representation of a CSO's internal services, systems, and other devices, as well as linkages to other services and systems, is provided by an authorization border. An authorization boundary diagram accounts for all technologies, external and internal services, leveraging systems, and government data that CSP is in charge of. The authorization border is a crucial element connected to the OMB circular A-130, Managing Information as a Strategic Resource, and NIST special publication (SP) 800-37, Guidance for Applying the Risk Management Framework (RMF) to Federal Information Systems.

The Authorization Boundary Guidelines document²⁶ is currently being updated by FedRAMP to reflect advances in cloud computing technology and federal information security law that are pertinent to FedRAMP. The main adjustments will be: (1) Scoping and establishing the Authorization Boundary in the cloud; (2) identifying data kinds, including federal data and federal metadata in the cloud; and (3) using interconnectivity, external, and corporate services.

But only for the high baseline, FedRAMP does provide U.S. /U.S. Territories or geographic areas under U.S. jurisdiction requirements for the data centres. Agencies should be aware that there are no implied or stated safeguards for federal agencies that guarantee their data will only remain inside the US or that their resources will only be created in areas that operate within the US for FedRAMP low and moderate baselines. Via SLAs or memorandums of understanding, agencies must set these parameters and expectations with their CSPs and address any issues with areas outside of the United States, U.S. Territories, or geographical areas under U.S. authority (MOUs).

CHAPTER 6

CLOUD MIGRATION

Dr. Srikanth V

Associate Professor

Department of Computer Science and Information Technology,
Jain (Deemed to be University) Bangalore, Karnataka, India
Email Id-vsrikanth743@gmail.com

The compute plane and discusses factors that agencies should take into account while creating, implementing, and maintaining cloud-based digital services. Agencies should do the following to guarantee a smooth and safe transition to cloud services:

1. Create cloud-optimized software: To develop a safe and effective cloud environment, decide which services and capabilities to use right away.
2. Develop a plan for moving to the cloud: Create a strategy for your agency to move data and services from an on-premises environment to the cloud.
3. Use a "DevSecOps" strategy (development, security, and operations): Using coding and integrating support staff, develop dependable automated digital services.
4. Consolidate standard cloud services: Determine which CSPs will be applied across the organisation and consolidate administration and purchasing.
5. Invest in People: Organizations need to develop specific capabilities for cloud migrations.

Software Development for the Cloud: Agencies may combine services to serve their goal by using the flexibility of the cloud. As early as feasible in the Software Development Life Cycle, agencies should endeavour to integrate security measures into their cloud-based digital services (SDLC). Agencies will be able to create architectures that are scalable, repeatable, dependable, and adhere to the zero trust ethos if they support DevSecOps with automated security testing. To create digital services, agency teams must collaborate throughout this process. IT teams may support DevSecOps in conjunction with centralised SaaS to allow security testing of products prior to deployment. Digital services that use the cloud might include IaaS, PaaS, and SaaS. According to Section 3's discussion, various levels of the system architecture are handled by different parties in these service models as well as the on-premises model. Agencies must verify the services and tasks that their suppliers are doing and are not performing

Switch to the Cloud for Software: Software may be more dependable, scalable, and predictable when it is moved from an on-premises data centre to the cloud by organisations. Without having to invest in a new data centre, cloud services enable organisations to swiftly grow capacity when necessary and have disaster recovery accessible in other locations. Before trying to migrate bigger services, agencies might first move smaller, internal projects and tools to the cloud to build knowledge and confidence with new environments. By moving to the cloud, outdated digital services may be redesigned to support modernisation or ambitious advancement.

There are many well-known advantages of the cloud, but one that agencies should take into account is the ease with which zero-trust architectures and more secure apps may be built there. Identity, Devices, Networks, Apps, and Data are the five zero trust pillars that can be

addressed by CSPs, who can also provide the visibility required to start developing cross-pillar interaction. Agencies may often speed up an ATO, simplifying the migration process, by checking for the proper FedRAMP clearance level for cloud services. DevSecOps teams or other administrators may be in charge of correctly setting these services, creating useful ICAM roles, and encrypting sensitive data using a Key Management System (KMS). Further advice on Cloud Security Posture Management is provided.

In order to manage their cloud installations safely, agencies should take into account the security benefits of leveraging APIs or data services. The same data may be accessed via services from CSPs and other suppliers without requiring organisations to develop, validate, and maintain complicated software. APIs offered by CSPs and others are often supported by a large team of developers and other specialists who specialise in these systems. The purpose of an agency may suffer as a result of the time and money required to establish a comparable team inside the organisation.

Migration strategy for the cloud: Cloud migration is the process of moving business operations and missions to the cloud. For a lot of agencies, this means switching from antiquated infrastructure that may no longer be able to satisfy their needs to more modern infrastructure that enables a more flexible and reasonably priced solution for an agency's application. Cloud technologies, by their very nature, require a shift in thinking from on-premises solutions. Certain cloud services, such as infrastructure as code (IaC) concepts, could function differently than they would under on-premises conditions. Depending on the elasticity of service demand, some possibilities include dynamic resource supply and decommissioning or temporal-based maintenance to replace risky infrastructure.

The size of the application ecosystem, the age of the legacy systems and applications, the user base, and the quantity of data being transferred all affect how much preparation is necessary for a cloud migration. Agencies should consider the amount and age of data in their application ecosystem since shifting to the cloud may become more challenging as data accumulates over time each time a company wants to relocate their application. Before making a choice, they should weigh the benefits, risks, and challenges of using cloud-based solutions.

Potential Problems with Cloud Migration: The transition from on-premises to the cloud includes several particular issues related to people, money, and data that are not present in other large-scale software initiatives. The typical obstacles that agencies have while converting to the cloud are listed below.

Funding: Before cost savings may be achieved, it may take some time for the application infrastructure and data to exist across many settings. There are additional expenses related to data transport. Moving data out of a CSP may be more expensive but putting data into one is often cheap or even free, depending on the CSP, the architecture, and the technique.

Onboarding: During onboarding, the team should spend more time learning about the new technologies that will help their application migrate successfully.

Infrastructure Support a team that has never moved a database or application to the cloud may need assistance setting up servers, network support, their application, and servers in the cloud.

Staffing: As a project develops, a specialised staff that is focused on assisting the migration effort may be required.

Policy Support: While current application/project ATOs are often pushed to their limit by cloud migration, they may need to be upgraded or replaced by new ATOs.

Change management: In addition to the technological changes needed to switch to a cloud architecture, process adjustments will also be necessary. Recognizing this and making room to rework the procedures helps lessen some of the stress associated with change.

Typical Cloud Migration Obstacles: Agencies should take into account technological issues associated with data transfer in addition to ordinary difficulties. It takes longer to move, verify, and support large volumes of data. If there are additional needs that result in little to no application downtime or if the underlying data changes often, migration challenges multiply.

Technological Difficulties in Moving to the Cloud:

Data Integrity: The migration process must guarantee both the security of the data while it is being transferred and its integrity after it has arrived at its ultimate storage location.

Reducing Downtime: Although many agency apps are accessible during regular business hours, weekends may be used for downtime. Some applications may have more exacting criteria for downtime. Preparation and, in many circumstances, an incremental deployment of the application in the cloud are necessary to minimise downtime while replacing a system.

Network Support: An agency should be aware of the latency and throughput characteristics of the network when a significant volume of data is moving across the agency's network infrastructure in support of a data migration. Decisions on the best way to transfer the data to the cloud vendor's environment may be guided by these measures.

Also, users on home networks and developers who must transfer data and programmes around may have problems with bandwidth.

Advantages of Moving to the Cloud: As many business and mission operations are cloud-centric in nature, cloud services offer agencies a variety of operational and financial benefits. In SP 800-14529, NIST identifies the five fundamental aspects of cloud computing as on-demand self-service, wide network access, resource pooling, quick flexibility, and measurable service. The ability to provide hardware in accordance with tenant demands marks a significant departure from conventional hardware management and procurement practises. Virtual machines (VMs) may be chosen by tenants as an alternative to hardware reservations. Moreover, renters have the option to use platforms provided by the CSP rather of creating any servers at all (virtual or bare metal). Agencies would still be responsible for the security of their systems, but part of the mundane job of health monitoring and patch management may be transferred to the CSP. Instead than being contained in a single place like a server room or data centre on-site, provisioned resources may instead be spread over a number of geographical locations and availability zones within regions. Agencies should think about their own resources and requirements while exploring various cloud services to decide if cloud services are suitable for implementation. While not exhaustive, some major advantages of cloud migration.

Advantages of Moving to the Cloud:

More Support: Agencies have access to a larger selection of cloud suppliers and support.

Design Flexibility: Managed services like document storage, database storage with replication, and application interfaces for automation are offered by cloud services.

Scalable performance refers to the capacity to increase the number of computers in a pool of resources used by an application. Cloud services allow a wide range of horizontal scalability. Distributed systems must be scalable.

Availability: Cloud services can handle failures of the application's underlying infrastructure such that running code may be transferred with little disruption.

Cost-effectiveness: CSP services enable organisations to focus financial resources on activities that are most important to their missions.

Agencies utilising off-premises cloud data and infrastructure are better equipped to manage and recover from negative occurrences at agency headquarters. Disaster recovery and business continuity (e.g., natural disasters).

Cybersecurity: CSPs often provide alternatives for various security-related features so that individual customers don't have to develop their own support for it. But, it is essential that organisations understand their alternatives, put them into practise, and customise the ones that are best for them. One approach is to transfer data across an agency's network. Other options for data migration into the cloud might include transferring data on discs and sending them by land or air to the CSP.

Plans for Cloud Migration: some of the key cloud migration tactics that have gained popularity among business partners. While transferring an application, agencies can need to use many tactics. As not all applications are created to operate in a cloud environment, agencies must take their unique requirements into account when they move. For instance, a CSP may not be able to give the speed that an application needs since it depends on the low latency offered by a local network.

Cloud Migration Methods:

Details of the Cloud Migration Strategy:

Rehost: This method duplicates the application architecture using a "lift and shift" methodology, moving the initial configuration to cloud servers.

Refactor and restructure:

With the justification that it will be able to use cloud native services from a code and architectural standpoint, this strategy restructures the programme into use cases.

Revision/New Platform:

A portion of a programme will be upgraded and migrated to use cloud native services. Using cloud native managed databases is a common choice owing to its ease of maintenance.

Rebuild: To rebuild an application, the old version must be thrown away and the new version must be created using cloud infrastructure. For this to work, a cloud native solution must be built or the application must be placed inside it.

Replace:

By transferring the use cases to a SaaS environment with a third-party provider, this strategy renders the legacy application unnecessary. Agencies should carefully examine which approach is best for them before choosing between the Rehost strategy, Refactor, and Revise strategies. Due to the deprecation of old systems, it is sometimes required to migrate an application to the cloud; however, doing a Refactor at the same time is not practical. The best

course of action in such situation may be to pursue the Refactor once the Rehosting is finished. The Refactor should still be taken into account since cloud native services from an IaaS or PaaS may do so much to simplify systems, boost performance, and save costs associated with hosting.

Agencies may need to take into consideration the intricacies of moving various kinds of services to and across cloud environments when making the switch to the cloud. For instance, a company can decide to relocate development procedures. DevSecOps may be utilised in this situation to fulfil the particular scalability and flexibility requirements of on-demand infrastructure as well as to sustain newly integrated cloud-native systems over time. For instance, a company could opt to use containerization to make it easier for users of each service to orchestrate their computing needs.

Scenarios for Cloud Migration

Giving general advice on how to accomplish the conversion is difficult since every cloud migration is as unique as the original application. Following the steps listed below, however, may improve your odds of success.

Plan: Choose the strategy, CSP, and service type, as well as the application's path plan.

Design: Construct the application's architecture, paying particular attention to the system's distributed nature. Try out the CSP's cloud-native features.

Pilot: Produce a Minimal Viable Product (MVP) to show that the cloud-based application will function.

Migrate: Prepare the cloud version for production, including transferring any necessary data.

Maintain: Continue to enhance the cloud application, whether from the standpoint of a product feature or from the viewpoint of performance.

The usual migration scenarios for agencies are described in the following subsections. These scenarios exclude the security functionality that is standard to the environment since their attention is on the ways that application design changes when shifting to a cloud environment.

Scenario 1: Cloud-Based PDF Storage (IaaS): InPhase1 of this cloud migration, the agency wants to begin storing new uploaded files in the cloud but has not transferred all the older files. In this scenario, the agency will need an additional layer to manage the identification of stored files' locations.

The agency should research how to properly redirect newly uploaded files to the cloud environment and should redirect users via a reverseproxy to the proper file location, since files may now be split between on-premises and cloud. Finally, the agency will also need to care fully test all assumptions in a development environment to prepare for the migration. Figure 6.1 presents an overview of the architecture for Phase1.

A company is moving an internal application with 10,000 users that stores and uploads millions of Portable Document Format (PDF) files, totaling 1 Petabyte of data (1,000 Terabytes). The programme employs a local datacenter with several server racks for data storage.

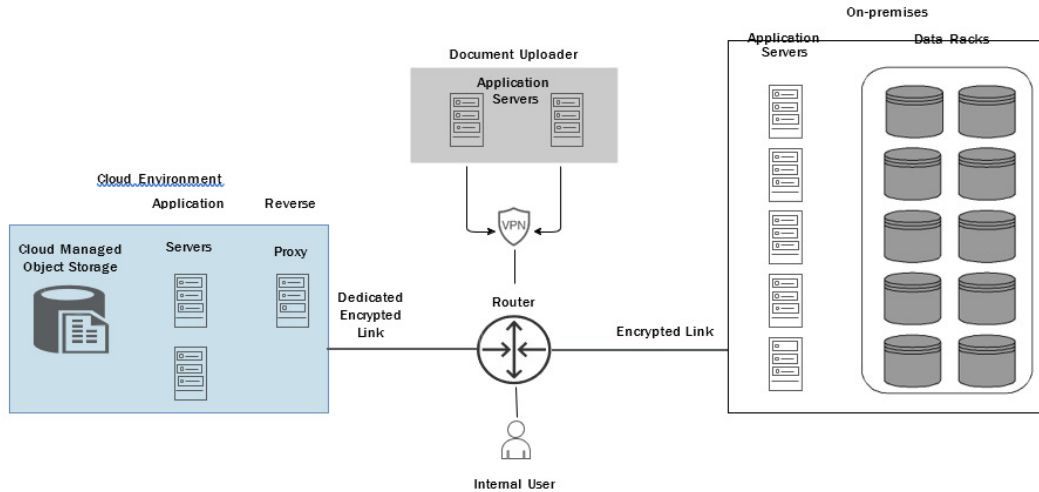


Figure 6.1: Scenario 1 Notional Phase 1 Architecture

The agency wishes to transfer the older data to cloud storage during Phase 2 of this cloud migration. They will need to schedule the 1 petabyte of data transmission over the network in consultation with the network team. The distributed data will be gathered by application servers in the on-premises environment, create a set of integrity checksums for later verification and send the traffic to the cloud environment across encrypted channels. If feasible, the agency may think about using hard drives or other forms of storage to move all data to the CSP. It's possible that this method is more effective than sending all the data across the network. Figure 6.2 displays these modifications.

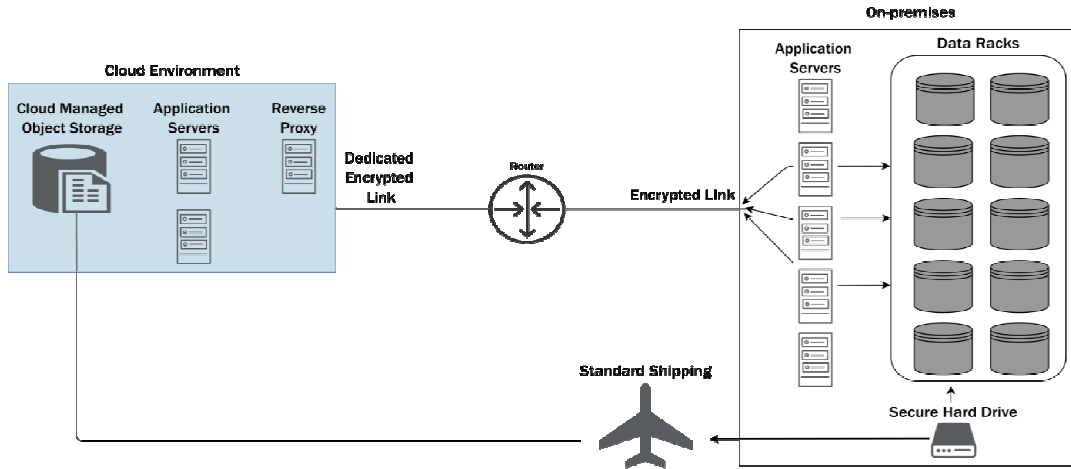


Figure 6.2: shows the fictitious Phase 2 architecture in Scenario 1 with out-of-band data transfer.

The data is checked when it enters cloud storage to make sure it is accurate. The organization has to make sure users and file uploaders can utilize the cloud environment without any problems once the data have been transferred. The on-site data centre may now be shut down or used for anything else.

Website switches to a PaaS Service in Scenario 2: With a fresh design and a contemporary content management system, an agency intends to move an on-premises, outdated website

infrastructure. The organization has been hosting thousands of pages on a locally managed, antiquated content management system for the last 20 years (CMS).

In this scenario, the legacy infrastructure is noticeably dated and many of the webpages require redesign. The agency decides to use a PaaS to build the next enhancement of their CMS. Figure 6.3 shows the architecture of some of the webpages during the migration and redesign.

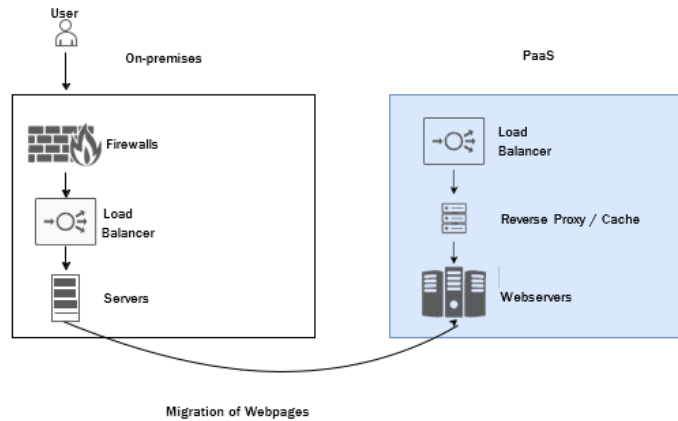


Figure 6.3: Scenario 2 – Notional Migration of a Website to PaaS

After the relocation and redesign, the firm realized that a content delivery network would be most appropriate for the website's public and seldom changing content (CDN). By using a CDN, the organization will be able to cache the majority of the material closer to the user, resulting in quicker upload times. The agency will test the CDN, move files there iteratively, and set it up to handle user traffic. Agencies need to evaluate the data that will be stored in a CDN service. Agencies may benefit from the extra security capabilities that many CDNs provide, including web application firewalls (WAFs) and mitigations for Distributed Denial-of-Service (DDoS) attacks. As the majority of the data will be public, caching outside the approved boundaries is permissible. Certain data will have CUI requirements, thus it should either not be cached or the agency should utilize a CDN provider that has been approved. Figure 6.4 illustrates one instance of website migration to PaaS.

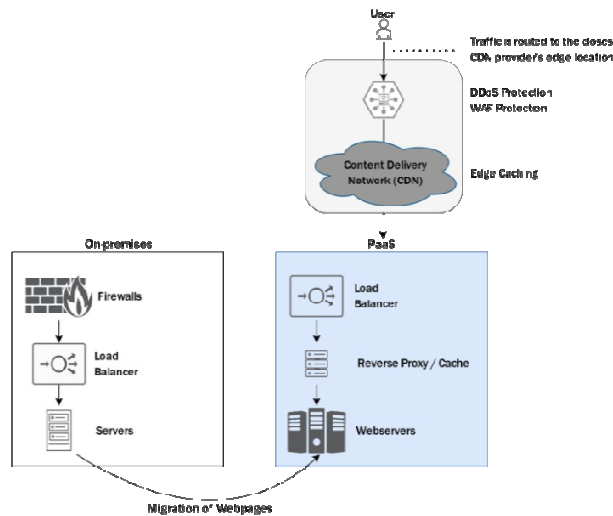


Figure 6.4: Scenario 2 – Notional Website with CDN.

The agency website will operate in a PaaS environment with a CDN entry point, as depicted in Figure 6.5, and the on-premises infrastructure will be discontinued.

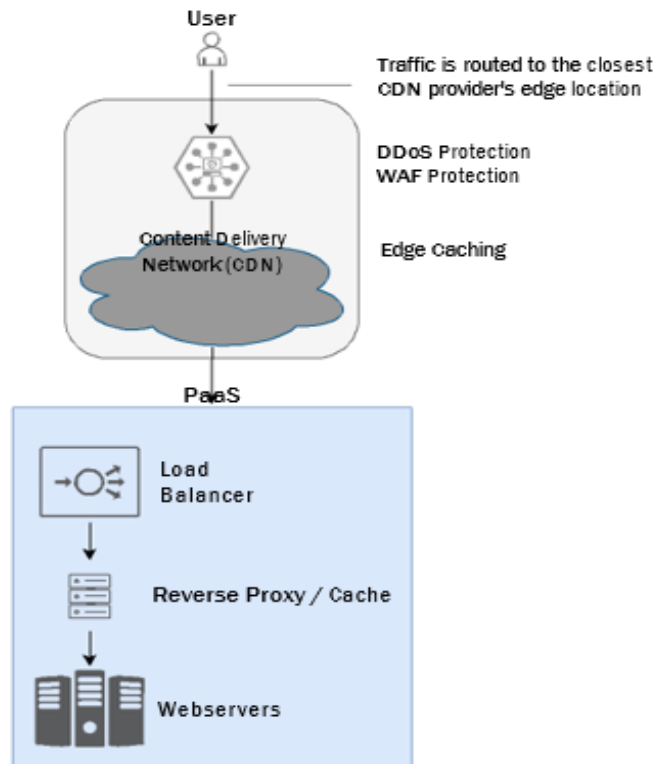


Figure 6.5: Scenario 2 – Notional Final Architecture of the New Website

Scenario 3: Monitoring Services for Public Facing Applications: In order to make sure that it is consistently providing services for its consumers, an organisation must check the uptime of its websites that are accessible to the general public. The organisation has to look at performance monitoring tools that can manage the globally spread systems since it has several websites that are hosted in various places. The organisation chooses synthetic monitoring, which entails automating conceivable user behaviours in order to see how the system reacts and gather data on uptime based on such requests. The organisation investigates the technical aspects and financial tradeoffs of establishing their own monitoring infrastructure in a PaaS or IaaS system as opposed to a SaaS solution created to produce the synthetic traffic and gather the related data. The group decides to use a SaaS solution.

CHAPTER 7

DEVELOPING A DEVSECOPS MENTALITY

Dr. Mir Aadil

Assistant Professor, Department of Computer Science and Information Technology,
Jain (Deemed to be University) Bangalore, Karnataka, India
Email Id-aadilbgsbu@gmail.com

Creating a DevSecOps Mindset: A software development methodology called DevSecOps, which stands for Development, Security, and Operations, closely combines developing code with testing, securing, and deploying that code. Figure 9 depicts the conventional DevSecOps cycle. It may enable teams to work together across the conventional roles of developers, security engineers, operation engineers, and quality assurance specialists by removing organisational barriers between them. This is accomplished by assembling cross-functional teams with members in various positions taking full responsibility for the successful creation, introduction, and upkeep of their service. Agencies should create, protect, and deploy cloud applications using DevSecOps as their main strategy. To harness automation and create scalable, dependable, and predictable digital services, DevSecOps often applies Infrastructure as Code (IaC), security testing, and the concept of least privilege.

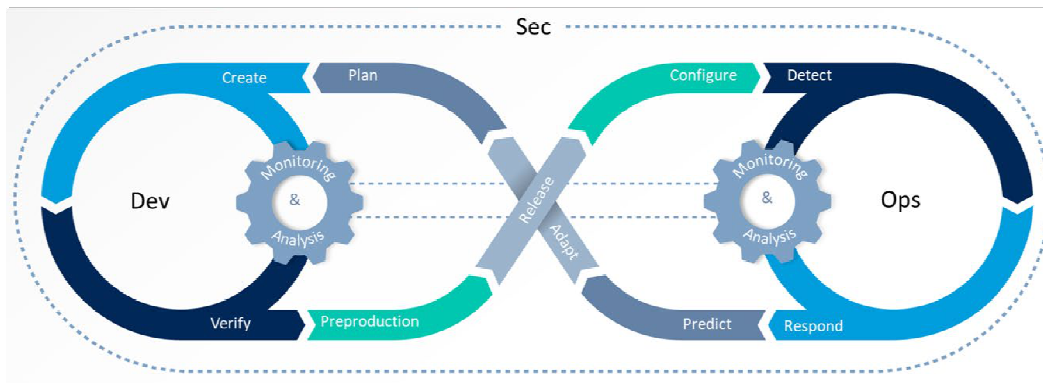


Figure 7.1: DevSecOps Loop

Constant Delivery and Continuous Integration: Code integration, building, and testing are routine tasks that are automated using CI to speed up the process and eliminate human error. As the project progresses, this tooling is extended earlier in the product lifecycle. Development teams may pick from a wide range of tools, some of which are SaaS packages that have received FedRAMP approval, to generate, test, and save source code. These could be included in the services offered by IaaS and PaaS companies. Software for source code management may also impose rules for code review and check-in that further cut down on human mistake and strengthen the system's non-repudiation.

Delivering code that has been automated created, tested, and integrated at regular intervals is known as continuous delivery (CD). To decide when the code is prepared for production, it expands on the CI pipelines. These procedures are collectively known as CI/CD. Agencies may generate more dependable software when it comes time for deployment by establishing the pass and fail criteria for testing, being prepared for deployment early in the SDLC, and

then employing automated procedures to verify that the requirements are satisfied. This method not only aids in the early detection of deployment difficulties, but it also boosts stakeholder support by encouraging an agile workflow that permits smaller and more frequent course adjustments since stakeholders can see the partly functional product.

Code for Infrastructure: Development teams may create their infrastructure as machine-readable specification files that perform automatic and documented provisioning, runtime adjustments, and decommissioning of their digital services in addition to building their apps in code. In order to assess changes to the resources required by IaaS or PaaS before checking in the code, this process is known as Infrastructure as Code (IaC). Moreover, it enables the mass manufacture of cloud infrastructure, allowing environments to expand automatically and fixes to be deployed swiftly. Since separate patches are required for every servers, when the infrastructure is manually updated, it is vulnerable to deviating from its initial setup.

IaC may provide a wide range of advantages: Automating compliance checks using IaaS or PaaS features, such as enforcing encryption at rest for storage containers; Automating deployment of ICAM policies and granular access controls; Facilitating security testing, patch deployments, and updates; Increasing zero trust maturity through enforcing encryption on networks and storage; Removing the need for a User Interface (UI) at each device, further reducing opportunities for human error. IaC, like other software, has the ability to degrade an environment and maybe even bring new unforeseen vulnerabilities to one that was previously safe. Agencies should execute security code audits for production deployments and/or monitor IaC code for configuration errors to lower the risk of exposure.

Testing for Security, Automated: Application security testing is a further element that may be included into the DevSecOps pipeline. It is essential to include security at an earlier stage of the software development lifecycle by using this testing as part of the DevSecOps pipeline. Application security testing makes use of a mix of static code analysis to find common coding errors such possible SQL Injection vulnerabilities and dynamic testing to examine how the code functions.

With testing, organisations are able to address possible security problems while they are still simpler to address than after they have been deployed into production. Testing contributes to improve zero trust maturity across the CI/CD process.

Automated security testing is just one line of protection against application vulnerabilities during development. In order to make sure that applications are used and to enhance the likelihood that vulnerabilities are discovered before they can be exploited, a combination of manual expert analysis, third-party security testing, public vulnerability disclosure programmes, and bug bounty programmes is used. For more on bug bounties and vulnerability disclosure schemes, see Binding Operational Directive 20-0131. A possible design for a CI/CD system with two locations for security testing. Developers would check their work into the proper repository for both the infrastructure and the application. When the application has been built by the build system, testing will start. Any unsuccessful tests would be noted in the monitoring system, and the developer would be informed of the findings, either through an alert or a status page. The application may be deployed into a development environment for further testing after all build-related problems have been fixed. The programme may be moved to production after all problems are fixed and is then ready for usage.

Infrastructure testing may also be automated. Security scanners will be able to discover possible problems early, preferably before they are accessible on the internet, with the help of definitions of the structure in IaC.

Principle of Least Privilege Organizations should make sure that each member of the DevSecOps team has just the rights necessary to carry out his or her duties. The scope and duration of access for each individual to carry out the responsibilities of their jobs and positions are appropriately sized according to the least privilege concept. This reduces the chance of abuse by a malevolent actor (internal or external) by reducing their ability to transfer or raise privileges. Based on the actions carried out within a certain interval, some CSPs may enable various rights inside the infrastructure. Someone who manages operations (also known as On Call) may be given access to other jobs that give them control over production. After that shift is completed, their responsibilities may subsequently be eliminated. A "break glass" technique is an option to provide temporary access to mend something that is shattered.

Other security best practises, such as establishing more granular access rights throughout the team and enforcing frequent revocations of unnecessary access, may also help to reduce the danger of roles that are constantly increasing. Rules for blocking access when a team member departs are also essential.

The development of attribute-based access control (ABAC) is connected to least privilege. By imposing checks surrounding the user's identity, the qualities of the resource being accessed, and the environment, ABAC goes beyond role-based constraints. After then, roles become a characteristic of the user's identity. The question "Can this user access the data?" is another crucial factor to look for. Information regarding the device the user is using, such as whether it is an agency device and if its patches are current, is also often checked as an environment-based property. Mixing many qualities might increase the likelihood that the user is who they claim to be and is authorised to carry out the desired activity. By integrating more than one pillar in access choices, ABAC is a crucial part of an established zero trust architecture.

Separation of responsibilities has historically been used to prevent insider threats and uncover innocent errors by requiring more than one person to carry out crucial functions. As an example, the group responsible for development and coding is distinct from the group responsible for production deployment. This strategy conflicts with DevSecOps since these duties are now distributed throughout a team.

A two-person integrity check method using code reviews is referred to as a replacement procedure. This implies that before a change is committed and merged into the main repository, it must first be reviewed and approved by another authorised team member. This is helpful to find problems before deployment in both the application code and the IaC. Code reviews may be made mandatory in many source repositories. Moreover, routine, everyday activity should not be performed using the repository administrator accounts that permit this option to be deactivated.

CHAPTER 8

CENTRALIZING COMMON CLOUD SERVICES

Dr. Manju Bargavi Sankar Krishnamoorthy
Professor, Department of Computer Science and Information Technology,
Jain (Deemed to be University) Bangalore, Karnataka, India
Email Id-b.manju@jainuniversity.ac.in

Consolidating Standard Cloud Services: Their organization may assist developers as they move, build, and deploy cloud-based apps by overseeing and maintaining shared services. Agencies enable developers to devote more time to the objective and less time to overhead or maintenance duties by offering shared services. Here, these services are divided into four categories: Agency PaaS, development tools and services, services geared at the general public, and security services. By minimising administrative cost, sharing certain services at the agency level may assist teams in implementing cloud native practises more quickly. This gives them more time to consider additional areas of overhead that can be reduced. A company may go from using entire virtual machines (VMs) for web servers to utilising containers for servers, and then from using containers to "Functions-as-a-Service".

Specialized roles are employed differently when agencies transition from conventional on-premises servers to IaaS, which represents another progression. Teams won't need a dedicated administrator for the server that houses the database, but they will still want the specific skills to know how to produce "An access control method where subject requests to perform operations on objects are granted or denied based on assigned attributes of the subject, assigned attributes of the object, environment conditions, and a set of policies that are specified in terms of those attributes and conditions," is how NIST defines ABAC.

A class of cloud computing services known as "Function as a Service" (FaaS) offers a platform that lets users create, use, and manage application capabilities without having to worry about creating and maintaining the infrastructure that is generally needed to create and launch app. databases are effective. This paradigm enables database administrators to concentrate on the strategic operations while also serving as consultants for the organisation.

IaaS, PaaS, and SaaS should be centralised across an organisation for two key reasons: to save resources and to provide a common user experience. It takes a lot of resources to purchase new equipment and software. By establishing a centralised team that is in charge of tool research, acquisition, and training for all teams, agencies may save overall costs. Moreover, by reducing maintenance and compliance activities, centralising services (such as those in the following categories) may help save money.

By providing shared tools to many teams within an agency, centralization enables agencies to create a single experience and promotes cooperation. Knowledge may spread outside of a team thanks to centralised documentation. When a cloud service is down, teams may cooperate by using the same ticketing, pager, and monitoring. Also, it makes onboarding easier when staff members change teams within an organisation.

To help other teams benefit from their experience, teams who have worked on cloud-based projects may also share their best practises and problem areas. In order to share their knowledge and expertise with younger teams, members of more seasoned teams may serve

as mentors, boosting the total investment in people. Two compelling arguments in favour of centralising cloud technologies are the need to save resources and eliminate organisational silos.

A Platform for Agencies:By purchasing cloud infrastructure in bulk and allocating access to various teams as required, agencies may consolidate access to current IaaS technologies. The right amount of access will be allowed, and newer teams will be able to use the infrastructure right away thanks to this. By providing options to standardise operating systems, software libraries, and logging, an agency's cloud team may behave more like a PaaS. Combined, these ideas will hasten the creation of cloud-based digital services and reduce resource use.

A centralised IaaS may set normative practises, enforce compliance, and start to lighten the load development team's bear from security paperwork like ATOs. Major IaaS systems include compliance checks, such alerting teams when a storage container is public or not encrypted so teams may swiftly resolve the problem. By using a single language in their software support programme (SSP) agreements and inheriting NIST SP 800-5334 controls from the organisational account when all teams utilise the same platform, speedier paperwork is possible. Agencies may set up artefact repositories and centralise "gold image" VMs so that teams can share IaC containers. The logging requirements described in OMB M-21-31 may also be set up in the VMs and containers. Although agencies may add security monitoring to the basic pictures, it's crucial to also maintain the images performant by not overloading the systems with too much additional processing. Here is where the conflict between usability and security can be seen. Enforcing frequent patching throughout the environment may increase security further.

Development teams are getting closer to the concept of "immutable workloads" with the help of artefact repositories and IaC; once cloud infrastructure and code are deployed, they are not manually updated or altered. The systems would be redeployed when any modifications were made using the CI/CD pipelines. By supplying certificates for the programme on the web server, encryption services may guarantee that the application utilises secure communication channels (like TLS). Moreover, when appropriate, these services may natively encrypt data at rest or use managed services to do so. Applications may cycle keys and passwords on a set schedule without being interrupted thanks to key and password management. Moreover, this service must be able to revoke keys if they are hacked.

Development Services and Tools:For constructing and sustaining applications fast and effectively, development tools and services are essential. While not complete, this section lists typical services and tools used in application development. The software development lifecycle and DevSecOps heavily rely on requirement tracking, documentation, and communication tools to exchange current status evaluations both inside and across teams. Cultures of sharing and cooperation are created through agency-wide collaboration and documentation procedures.

A CI/CD pipeline's basis is its source control system, which determines which tools may be used for developing, testing, and delivering code. Another crucial component of CI/CD that can be standardised throughout the agency is the execution of code quality control using "linters," which can examine code for problems that would prevent execution or create difficult-to-read code, and checking for coding "anti-patterns," which can prevent sloppy coding conventions that create insecure or under-optimized code. The ability to centralise and standardise security testing throughout an organisation is crucial, since the consistent

application of security results in improved overall procedures. Security testing that is both static and dynamic may provide an early line of protection against unintentionally releasing issues into production.

Services aimed towards the public:Centralization might be advantageous for certain elements of the digital services that organisations provide to the general public.Obtaining a domain, setting up DNS records for the website, and installing certificates for secure hypertext transfer protocol are routine steps in launching a new website (HTTPS). The agency is able to keep an accurate inventory of their online presence because to the centralization of these operations.Internet-accessible applications and APIs need security against malicious traffic. WAFs, API gateways, and content delivery networks (CDNs) that also serve as DDoS protection are a few examples of security measures. In addition to inspecting the requests sent to the web server to search for typical website threats, WAFs may regulate access to the network generally. A single API gateway may safeguard several APIs by controlling access to them for certain users.

CDNs may typically absorb excess traffic in denial of service (DoS) attacks or restrict network traffic through firewalls in addition to offering a mechanism to store cached material closer to customers for quicker delivery. This security will be necessary for all of an agency's online assets, and buying in bulk may reduce costs.To the maximum degree practicable, security services agencies should implement centrally integrated security services throughout the organisation. A company's attack surface is less when there are fewer distinct instances of the same service. Application safeguards including logging, authentication, authorisation, encryption, and key management are provided by security services.

The secret to effective incident response is centralised logging. It eliminates the need for locally stored logs and lessens the effect of their erasure. The amount of time required for an incident responder to investigate is also decreased by centralised logging. The sharing of logs between components and their parent Department is also provided for in OMB M-21-31, and centralization makes this procedure more efficient. Moreover, threat hunting across CSPs and on-premises systems is made easier by centralising logging.ICAM via single sign-on is a great place to start for agency services since the CIO probably already has the ability to allow workers to log in to services like email. Employees won't have to remember yet another password thanks to Lightweight Directory Access Protocol (LDAP), which allows broker access to cloud services even on-premises. A potential architecture with centralised identification and logging is shown in Figure 8.1.

LDAP considerations:Agencies are urged to cooperate with CSPs to make sure that federated identity services are protected with the proper logging enabled, given that the majority of agencies depend on on-premises LDAP services, such as Active Directory, to access cloud services and resources. Agencies should continue to collaborate with their suppliers and cloud providers to move ICAM services to the cloud as the main identity provider since CSP LDAP services are expanding quickly.

When integrating a new SaaS product into the centralised ICAM system, agencies should aim for the least amount of friction possible. Prior to full acceptance and integration, they should look into chances to prototype or trial linking such services. Agencies may be able to resolve their original performance and security-related issues as a result, and it may also provide more information for better integration. ICAM services often contain authorization (i.e., the process of figuring out whether a user has authority to do an activity). Nonetheless, there are situations when apps need to impose authorisation. The ability to continuously seek

authorization-related information is necessary for development teams. Agencies will transition from role-based to attribute-based permission as they continue their zero trust journey, and authorisation will need to include data from many pillars. Authorization gains from centralization as well.

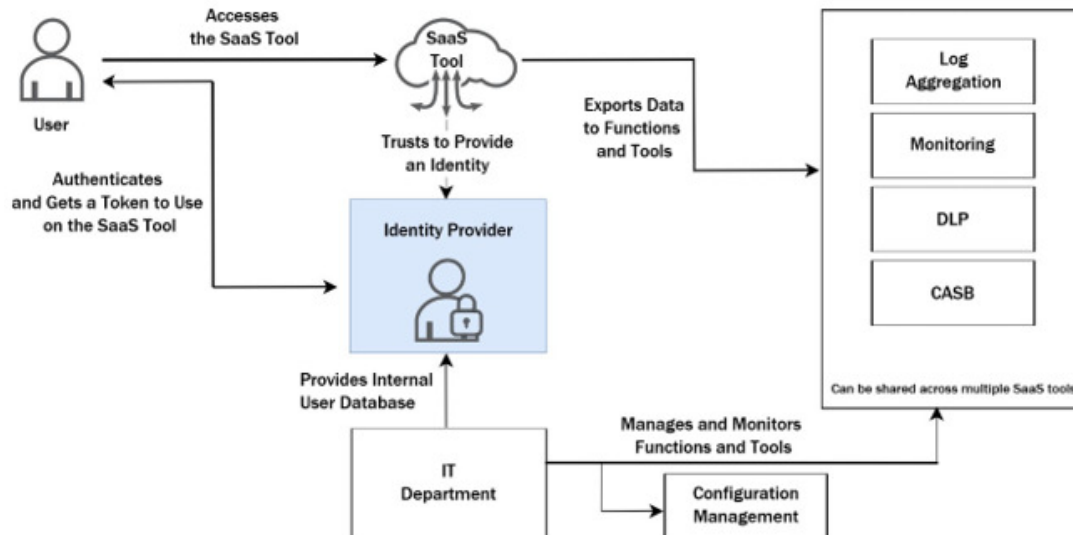


Figure 8.1: Reference Architecture on Centralized Security Services

The Human Element: Changes to processes and procedures are necessary for the personnel working on deploying tools and applications, as well as for the stakeholders and users of these tools, in order to build scalable, repeatable architectures using CSPs. For initiatives using the cloud, agencies will need to make personnel investments. Process reform, employee education, and dependable access facilitation will also be required.

Spend money on people: A project's success depends on investing in the best personnel to execute cloud-based solutions. This is divided into three parts: hiring, training, and procurement. Federal personnel who have been working in conventional software development settings may be retrained in cloud technology, but doing so requires agencies to invest in their staff via outside trainings, workshops, certifications, and the utilisation of off-the-clock study time. This may also include instruction on contemporary project management techniques. Agencies must provide their staff with opportunity to exercise their newly acquired abilities in order to reinforce the trainings (e.g., through access to sandbox environments that allow for experimentation with these new technologies). Experimentation, iteration, and liberty to "fail fast" can assist workers that are new to cloud technologies grow their abilities and produce stronger digital services.

Another method to invest in people is to hire new government workers who have prior experience working on cloud-based initiatives. Since there is often a small skill pool, it might be difficult to fill a new post. The Subject Matter Expert Qualification Assessments (SME-QA) method from the Office of Personnel Management (OPM) is one approach for more effectively employing technical applicants. This enables organisations that need comparable personnel, such as designers or product managers, to exchange job postings, screen applicants using technical tests, and build a pool of competent individuals from which they may individually choose. SME-QA has decreased the time required for competitive recruiting while increasing the number of positions filled via it. It might be challenging to draw in seasoned software developers from the commercial sector. Salary is one of many causes of

this, but it is also one of them. With the help of OPM, agencies may develop methods to increase pay using the General Schedule (GS) scale while still providing signing incentives and high-quality training opportunities.

Finally, it is possible to hire contractors to create digital services and introduce CSP goods. TechFAR Hub36 is a tool that procurement experts may use to understand how to make it easier to get IT services, including cloud services and companies that can make software for them, to be purchased. To assist procurement professionals in learning more, TechFAR Hub offers a project called Digital IT Acquisition Professional Training (DITAP).

Support Staff Agencies should work to help federal personnel by providing extra training and timely access via onboarding and other defined processes. In order to utilise the new CSP technologies and understand how the usage of cloud tools alters the security paradigm, all staff members will need some extra training. Anti-phishing instruction and correct data management are two examples of security training.

Both individuals actively engaged in the development of digital services (such as those working in DevSecOps) and those supporting digital services from a non-technical standpoint will need to adapt to this new security paradigm. Silos that may have developed over time may be broken down through better communication between development teams and stakeholders. Information silos will be broken down and cooperation will be facilitated by combining development, security, and operations. Giving workers timely access also lessens the chance that they may create "shadow IT" services that evade IT or security team scrutiny and deteriorate the Federal Government's overall cybersecurity posture.

Management of cloud security posture: The cloud security posture management (CSPM) and associated security capabilities and results are introduced in this section. This section also discusses organisational requirements for designing cloud services and reducing cloud risks, as well as some important factors to keep in mind while moving to the cloud. The background of CSPM is also explained in terms of how such capabilities might make it easier to create zero trust systems. The part that follows will:

1. Describe the word "CSPM" and how it is used in this paper in contrast to other reliable sources.
2. Describe the needs for implementation: To achieve desired security results, emphasise organisational demands and factors that should be taken into account when CSPM and zero trust are implemented.
3. Align the objectives of the executive order: Explain how the CSPM supports the zero-trust objectives.

CSPM Definition: The adoption and usage of cloud computing sometimes include the use of several networking and cybersecurity buzzwords. Several of these words have definitions and meanings that are standardised or generally accepted. Yet, a lot of these expressions have varying meanings and are used differently by various parties (e.g., within a given organization, across the Federal Government, within industry, etc.). The phrase "Cloud Security Posture Management" is a relatively new invention, and different organisations have given it varying definitions. Although many of these definitions are identical, they are stated enough differently from one another to provide some room for interpretation. To establish agreement on the meaning of such disparities in this term's definition and others, more clarification among stakeholders may be necessary. In this paper, "CSPM" refers to a continuous process of monitoring a cloud environment that identifies, alerts on, and mitigates cloud vulnerabilities as well as lowers risk and enhances cloud security.

The following activity outcomes are intended to be supported by the CSPM capabilities described in this document: Governance and Compliance, Standards and Policies, Privilege and Identity Access Management, Data Protections, Infrastructure and Application Protections, System Health and Resource Monitoring, and Incident Response and Recovery. These skills include Identity, Credential, and Access Management (ICAM), Continuous Monitoring and Alerting, Security and Risk Assessments, DevSecOps Integration, and Security-Based Artificial Intelligence (AI) and Machine Learning (ML) Capabilities.

Moreover, although the link between cloud adoption and zero trust migration is emphasised in this article, it is not implied that switching to cloud services automatically results in a zero trust architecture. The dispersed nature of cloud requires extra configuration and management tools in order to accomplish the kind of security and visibility over assets, people, and data that a zero trust architecture would demand. This is one of the reasons why cloud services allow zero trust.

Need CSPM: Agencies can access and manage cloud resources, apps, and data thanks to CSPM. By migrating data and apps to the cloud, organisations modify how they handle governance and compliance requirements for their data and programmes as well as outsource physical access to these distributed resources. As cloud installations become more advanced, they are more complicated and often include many suppliers and technologies. Moreover, recent cyberattacks have had far-reaching repercussions; these attacks demonstrate the need of proactive management and monitoring provided by cloud services for protecting the Federal Government from cyberthreats. Agencies should continuously control their risk by. In a rapidly changing world with emerging threats and where CSPs are continually updating their product and service offerings, monitoring and strengthening their entire cybersecurity capabilities is essential.

When organisations go to the cloud, there are chances to incorporate granular controls and safeguards as well as to manage cloud security by employing automated technologies for watching over every part of the cloud, spotting dangers, and sending out alerts when anything seems out of the ordinary. The CSPM encourages ongoing development of an agency's cybersecurity skills and posture, allowing them to stay ahead of new threats, guard against configuration errors, and lower the likelihood of a security incident or data breach. Even though some Agencies might be better positioned to benefit from these capabilities, antecedent actions, like creating a warm site, could provide all Agencies with immediate security benefits, operational resilience, and a foundation for adopting additional capabilities.

CSPM support Zero Trust: According to Executive Order 14028, organisations moving to the cloud should embrace zero trust principles and, as soon as is practical and consistent with their risk tolerance, convert their environments to zero trust architectures³⁸. To do this, agencies should concentrate on enhancing core cybersecurity capabilities that are integrated across on-premises and cloud settings, such as identity management, asset management, network security, application security, and data safeguards. Agencies should also implement oversight, visibility, and automation and orchestration across these domains. Agencies may build a zero trust architecture in a number of ways, including via improved identity governance, logical micro-segmentation, and network-based segmentation technique, of NIST SP 800-207. Yet, a complete zero trust solution will have components from each of those three strategies. For designing a plan to implement zero trust, organisations may also utilise the CISA Zero Trust Maturity Model³⁹.

Agencies should focus their efforts on developing an identity management system that offers enterprise-wide identity awareness across cloud and on-premises systems. Agency users will

have identities across a number of suppliers when agencies move services to the cloud. Agencies will need to link their on-premises identities with those in the cloud environments in order to manage these identities and associated credentials efficiently and to coordinate security safeguards comprehensively.

Agencies may utilise CSPM capabilities to enable monitoring, analysis, and automated configuration of access controls for deployed services at scale and across environments, including for service, network, and workload identities. Asset and vulnerability management should be integrated into all agency settings, using automation wherever feasible. Agencies will need to make sure that the devices used to access services and data, particularly those stored in cloud settings, are secure. Tools for CSPM may be used to ensure compliance and collect data on vulnerabilities. Agencies should segment their networks in a zero trust architecture in both on-premises and cloud settings to minimise lateral movement, restrict rights, and manage attack vectors.

Applications will need to be designed for cloud deployment, and cloud-native products will need to be taken into account for application delivery. In their design, agencies should give top priority to data and access demands. In order to guarantee the safeguards have the visibility and fidelity required to deliver effective security, agencies should align application security measures based on zero trust principles and integrate their security controls more closely with their application operations. For all cloud-deployed apps and services, agencies should do continuous and dynamic application health and security monitoring. Application deployment settings may be managed and monitored using CSPM capabilities.

Finally, a zero trust architecture necessitates a review of how organisations safeguard their cloud data. Always safeguard data when it is in the cloud at rest, in transit to and from the cloud, and while it is really using the cloud. The use of CSPM technologies enables continuous monitoring, alerting on unusual behaviour in access logs, and identification and prevention of misconfigurations that may result in data loss and leakage.

CSPM Results: A selection of the cybersecurity objectives that are supported by the use of CSPM are described in this section. These results are roughly divided into a number of categories that relate to various security procedures that agencies should take care of. Agencies may build solid foundations for the security of their cloud deployments by attaining these various goals. Protections can then be implemented throughout deployment, operations, and post-incident reaction and recovery.

The following outcomes will be covered in this section: governance and compliance, standards and policies, privilege and identity access management, data protection, application and infrastructure protection, system health and resource monitoring, and incident response and recovery.

Administration and Compliance: Agencies must adhere to internal rules and procedures as well as regulatory and governance requirements when establishing and implementing cloud governance to direct operations and deployment. As a result, agencies should identify the relevant laws, rules, and binding guidelines that apply to the whole government, as well as establish internal procedures and tools for determining compliance. Agencies should make sure that compliance covers all facets of their cloud services, not only deployment and operations, including procurement requirements, invoicing and contract renewal, and service termination. CSPs often provide native services that go by several of these specifications, providing a base degree of compliance. Several solutions additionally include with tools for regularly evaluating cloud deployments and settings in comparison to these specifications. Agencies should think about how service provider modifications and updated terms of service

may naturally assist compliance with relevant regulations or may result in non-compliance and need further remediation.

To assist guarantee that cloud installations and services deliver a minimum degree of operability, agencies should take into account industry standards and best practises in addition to governance and compliance. The variety of deployment needs, which might vary in terms of things like physical security, operational continuity, and data controls, are assisted by standards and best practises. Again, many of these goals are supported natively by the cloud, so agencies can evaluate whether metrics meet their particular needs and take any necessary steps to close any gaps that may exist.

This evaluation against standards and best practises should also take into account policies specific to the cloud service, policies governing non-cloud aspects of an agency enterprise that would intersect with the cloud deployment, policies for pertinent on-premises services, and other policies in addition to reviewing cloud deployment policies. Agencies should continue to review and revise policies as necessary when service providers make modifications and upgrades.

This result aids agencies in defining and setting policies to match their unique needs, along with governance and compliance. How agencies implement and enforce these regulations for cloud services is equally crucial. CSPM capabilities may guarantee that these rules are followed in a number of different ways throughout the deployment stage. When configuring cloud infrastructure and services, methods like IaC or policy-as-code may allow monitoring, remediation, and automated enforcement of rules. Otherwise, other results like ICAM, data protection, and others show more clearly how these regulations are implemented and enforced.

Management of identity, credentials, and access

The management of identities, credentials, and access restrictions is a crucial aspect of policy enforcement. ICAM controls may be integrated across the full identity life cycle with the aid of CSPM solutions, which also provide continuous monitoring and analysis. Anomaly-producing activity that can point to a compromise or other possible problems can be found via behavioural pattern analysis and account activity log monitoring. ICAM controls are continuously managed and defined by CSPM capabilities, ensuring that services immediately inherit the correct settings. This fixes flaws including, but not limited to, excessively liberal access rules and unlimited code execution rights.

While zero trust offers a more all-encompassing approach to ICAM with granular account access controls, directory services, application and resource authorization, and policy compliance, CSPs are also working towards natively integrating those features. With these built-in features, the CSP may then assist agencies in implementing a zero-trust strategy by using commercially approved standards that are built on top of scalable infrastructure. Common cloud solutions also provide federated access while fostering interoperability, efficiency, and reuse.

Agencies must make sure that their own on-premises ICAM controls are current and in line with those used by their CSPs. Best practises may prevent illegal access and privilege escalation in the network, directory services, and apps by enabling phishing-resistant MFA and granting privileged accounts more granular levels of access and permissions. In addition to audits and reporting, ICAM also uses analytics for monitoring in order to assist

compliance. Agencies may consult additional architectural and governance guidelines, such as those from the GSA-managed Federal Identity, Credential, and Access Management (FICAM) program.

Data Protection: Determining who is responsible for data management and protection must be done in collaboration with data protection agencies' CSPs. For agency data stored in the cloud, data security is required at all stages (creation, storage, access, roaming, sharing, and retirement), for all data types (unstructured, structured, and semi-structured), and for all states (at rest, in transit, and in use). Several types of data protection may be offered via CSPM capabilities that assist policy enforcement.

Data loss and data leakage are two of the biggest issues in data security. To lessen the potential effect of these threats, agencies must adopt and enforce data security as they transfer data to, from, and within their cloud environments. Agencies must choose what needs to be safeguarded, how much protection to use, and (3) who will be in charge of overseeing requests for data exchange and access for each individual CSP. Agencies should also have procedures in place for handling data in the event that cloud services are discontinued in any manner. The appropriate sanitization and inaccessibility of deleted data, accounts, and machine images depend on access management. The security teams of an organisation should do assessments, implement controls, and provide analytics for data security and monitoring. Knowing about current or possible threats, agency policies and management choices, risk assessments, and vulnerability findings will be necessary for this.

In order to develop the proper governance for data protection, agencies should take into account the policies, rules, regulations, and standards. Agencies should set up cryptographic services like key management, PKI, or symmetric encryption appropriately as well as use digital rights management to safeguard intellectual property when necessary. The contract with the CSP should be examined as part of re-evaluations as well. It could be necessary to update both new and current features or modify their service level agreements (SLAs). The agency wouldn't need to re-encrypt data at rest, for instance, if the CSP currently completely encrypts data at rest and compliance and privacy requirements are still satisfied. Nevertheless, the service will need to be reevaluated if the service providing changes.

A crucial data security strategy is encryption. To prevent data breaches and to safeguard data in case other security measures are unsuccessful, data at rest is often encrypted. Data may be sent over networks with the use of encryption while being inaccessible to unauthorised parties. Agencies may also choose to utilise client-side or server-side encryption techniques based on their requirements. Client-side encryption prevents the CSP from seeing the data being saved since the agency generates their own key and does not share it. By using server-side encryption, the data are encrypted as they reach the cloud. To guarantee that encrypted data may only be viewed by authorised individuals, agencies should adhere to secure key management procedures (see Section 5.3.10). Managing account access, regularly testing backups to prevent accidental leaks, segregating resources to prevent them, and monitoring cloud regions, including unused and unsupported cloud regions, are other examples of data security techniques to take into mind. Data protection may benefit from emerging technology for secure computing and operations.

Protection of the Infrastructure and Applications

For various levels of cloud use, security may be provided through infrastructure and application protection. They include offering security for the apps, resources, and network connected to a company's cloud computing services. In order to scan their infrastructure, which includes VMs, virtual networks, apps, containers, and other services that may be

scanned, agencies should implement vulnerability management methods and tools. In order to enable agencies to issue warnings, CSPs have made initiatives to natively incorporate workload security and posture management logs into management dashboards. Certain warnings, including those that activate to return changed settings to a predetermined security baseline, may be handled automatically. Third-party tools may also be utilised for the following purposes: To prioritise proactive actions, align resources and applications, provide context for service settings, and provide dashboards for reviewing and assessing cloud security posture. CSPs may assess decision efficacy and produce reports using such dashboards and tools for risk and security management.

These dashboards, tools, and reports might aid agencies in making choices that are more effective, consistent, or appropriate. Several of these services are regarded as proactive and may increase defence effectiveness against various popular and conventional attack vectors. The migration of agencies to zero trust architectures is also supported by all of these infrastructure safeguards, which provide automation and orchestration along with visibility and analytics into people, devices, network environments, application workloads, and data.

Network Security: The right setup of network defences guarantees that networking permissions, segmentation, firewall, proxy, certificates, etc., are set up to facilitate secure usage. The ability to isolate systems depending on factors like location, application, environment (such as development or production), or resource type, should be supported by host, firewall, and other controls based on characteristics like location, application, environment (such as production or development), or resource kind. It improves the overall security posture to be able to implement multiple security rules for various service types. Agencies should also exercise prudence when it comes to network access and security settings, such as employing multi-factor authentication that is resistant to phishing attacks and encrypting connections. This is in line with the zero trust principle, according to which every communication is encrypted.

Protection of Resources: Infrastructure protection also includes resource protection, which includes protection for CSP service configuration. Every resource or service provided by a CSP and used by a tenant are considered resources in this context. By putting into practise some of the elements from the aforementioned categories, including data protection, CSP SLA clauses provide protection and may hold CSPs accountable for safeguarding a part of the resources made available. Additional SLA features might include infrastructure placement considerations to fulfil requirements as well as physical access protection and monitoring. In a zero trust architecture, enabling safe access to resources with automated policy enforcement is a core security feature.

Application Security: Application security includes: Short-lived containers may not need scanning while they are being executed, however container images should be inspected for vulnerabilities prior to deployment. Application layer firewalls are used, mitigations for distributed denial of service attacks are put into place, platforms or middleware are scanned, containers are scanned, and applications are scanned before being released for production or when containers are uploaded to a container repository. The extent to which apps may be accessible by other agency resources and vice versa should be assessed by agencies and, if necessary, limited. This involves being careful with how application accounts are handled, with what access or permissions they have, with who has access to the accounts, and with how they are secured. Agencies may safeguard the apps and allow speedier reactions to their possible exploitation by carrying out this identification and mitigation. Good application security is yet another essential zero trust design tenet.

Vulnerability Control: Versions, fixes, and vulnerability management are all interrelated. Agencies may make sure that vulnerabilities are systematically found and fixed by executing scans on a regular basis. By doing so, systems will be maintained up to date, patched to the necessary versions, and help find and get rid of outdated software. Updates may differ depending on the architecture being utilised to administer the systems. Some updates could take effect immediately, while others happen when a vulnerable resource is replaced with a resource that has just been fixed. No access will be lost during these changes. In a zero trust architecture, vulnerability management is essential to protecting all resources. The Binding Operational Directive 20-0144 of CISA states that it must also have a vulnerability disclosure policy (VDP), which is an element that is visible to the outside world.

Monitoring of System Health and Resources: In addition to addressing risks posed by malicious actors and activities to cloud service installations, CSP technologies also provide insight into how the service is generally used to guarantee optimal utilisation and system health. For instance, symptoms like excessive central processing unit (CPU) consumption or memory shortages cannot be a sign of hostile actors instead pointing to faulty setup or a less than ideal state of services and systems. These technologies keep an eye out for security-related events and may send users alerts or automate responses to problems. These automated processes aid in ensuring that the service is more reliable and that there are enough and available resources. This monitoring can include broader indicators of the health of the cloud services, such as checking billing and payment status, comprehending utilisation metrics, and tracking the number of users and their level of activity, in addition to directly handling resource requirements, such as using load balancers to adjust the number of active instances. To allow ongoing visibility into assets and applications, several provider technologies offer curated dashboards for displaying the most obvious or immediately critical areas of concern. An essential component of a zero trust architecture is monitoring the integrity and security posture of all cloud installations. Also, an agency's security posture should be regularly improved using this information.

There are various difficulties brought on by the variety of tools accessible. In addition to taking into consideration multi-cloud deployments that may utilise different data to signal system health, agencies must avoid possible fragmentation or a lack of interaction across many solutions from multiple providers, especially from third party vendors.

CSPs and other parties provide a variety of reaction options, including setting off warnings and automating responses to possible dangers, via their management consoles and CSPM capabilities. These measures allow for quick correction and stop the escalation of serious threats. Also, they enable human security operations to respond to less serious threats with more considered measures. The backend cloud architecture facilitates recovery in a similar way by substituting fresh resources for compromised ones in order to maintain service. During the post-incident investigation, the immediate shutting of possibly compromised instances may also enable untainted forensic analysis.

Plans for incident response and recovery are essential to reducing risks, guaranteeing service continuity, and preserving artefacts for forensic investigation after an incident. These designs should use cloud capabilities and take into consideration native CSPM tools. Provide adequate automatic response setup, streamline access to archived cloud instances, and collaborate with the CSP's incident response plans are a few examples of what this may include. Agencies need to be aware of the variations and difficulties involved in incident response and recovery in the cloud. For instance, it's doubtful that agencies will ever have access to the actual hardware that houses their resources. This also entails becoming ready to analyse exploited cloud resources using digital forensics. Moreover, just because an agency

uses cloud services does not mean that its data, apps, and infrastructure are automatically backed up. In order to assist reaction and recovery, agencies should preposition capabilities, such as solid backup policies and procedures, and should routinely carry out audits and inspections as part of maintaining the efficacy of their response plans.

Implementing CSPM Skills: Agencies may desire to transition their current on-premises infrastructure, data, and operations to one or more clouds. While theoretically simple, the methods used by an agency to migrate are complex and subtle. Older systems may not be the best fit for redesigned cloud-centric solutions or cloud environments in general. Agencies must choose the alternatives that work best for their cloud infrastructures. To create a strong security posture, capabilities including monitoring, scanning, reporting, mitigation, and other solutions should be assessed. Using CSPM capabilities to accomplish the results mentioned in section the broad CSPM capabilities that are accessible to agencies and their main responsibilities are described in the section that follows. Nevertheless, when agencies transition to CSPs and implement these capabilities, there are certain situations specific to each agency that will need to be taken into consideration. In order for agencies to retain situational awareness over the security of their linked services, difficulties with integrating capabilities across many CSPs should be addressed using the shared responsibilities paradigm. In order to assist the delivery of CSPM capabilities, this section also examines the manner in which security tools may be deployed separately or as a component of an integrated deployment.

This comprises: The following terms are used to describe various types of capabilities: CSPM Capabilities, Independent and Integrated Capabilities, CSP Account Management Hierarchies, Identity, Credential, and Access Management, Evolution of the Perimeter, Visibility and Sensor Placement, Monitoring, Application Programming Interfaces, Telemetry and Logs, and Deployment, Automation, and Orchestration.

Capabilities of CSPM: Via both native services and outside vendors, CSPs provide CSPM capabilities. Although some CSPs may permit the integration of third-party products into their services, other CSPs may prohibit or limit the integration of external third parties into their services. CSPs provide organisations the ability to incorporate CSPM capabilities at scale in addition to the conventional infrastructure- and service-level configurations and conventional intrusion detection and prevention systems (IDS/IPS). Examples comprise:

Security and Risk Assessments: Security assessment tools quantify policy effectiveness, posture, and compliance while also continuously monitoring identities and their permission sets in an automated risk-based framework. Similar to conventional on-premises port, service, and configuration scans, CSPs include features that may be used to manage accounts, examine traffic, find service vulnerabilities, and examine code repositories. Moreover, CSPs may have built-in tools for enhancing continuous visibility, automating security, and keeping track of compliance.

Continuous Monitoring and Alerting: CSPs may provide monitoring capabilities that let organisations record events and other forensic evidence in order to give agencies insight into system resources and data (e.g., supplementing continuous monitoring and alerting with the generation and analysis of network metadata, sourced pervasively from inside cloud environments). To increase visibility, logging services should be created for continuous diagnostic reporting. Alerts may be created as part of the monitoring services depending on metrics or unusual activity. Moreover, depending on the logs supplied by a CSP, third-party security information and event management (SIEM) systems may be utilised to gather, watch for, and issue alerts.

ICAM Capabilities: CSPs provide the capacity to carry out crucial operations in order to connect to or execute authentication on behalf of third-party authentication brokerages. These duties include generating, setting, and monitoring privilege escalation and resource access as well as maintaining and rotating keys, credentials, and certificates. While the majority of CSPs provide these capabilities, agencies should be aware of the specifics and restrictions for each of these services before implementing. Agencies shouldn't take for granted that these services are safe or adhere to all regulations.

Integration with DevSecOps: Centralized controls may automate CI/CD via the security integration into each part of the DevSecOps pipeline. With regionally focused deployments to react to current situations, the pipeline may also be enhanced. This involves: Performing real-time asset health and performance monitoring; addressing configuration errors caused by both human and automated deployments;

Monitoring and traffic rerouting through CDNs to extend data visibility and access controls outside the traditional network perimeter; Segmenting networks and provisioning segmentation for workloads, containers, and cloud objects; and Using infrastructure as code (IaC) to include practises and procedures that reduce environmental drift.

Security-based AI and ML capabilities: For the purpose of automating processes, enhancing performance, and conducting analytics on data streams and data storage, CSPs may provide AI and/or ML integration to existing security capabilities. With insight-based prioritising, these skills may enhance analytics and automation, but iterative review should be used to lower the danger of both natural bias and hostile machine learning.

Although many of the aforementioned services may be offered by CSPs, agencies may also seek to third-party solutions to enhance, replace, or extend the native CSP capabilities. Agencies may incorporate external third-party solutions through features like cloud access security brokers (CASBs), secure access service edge (SASE), and SECaaS options in addition to specialised third-party products offered in a CSP's marketplace. With the use of automated, managed security solutions, these services may provide agencies the option to outsource some of their security and monitoring duties to a CSP or other party.⁴⁸ Although combining conventional capabilities may result in the results mentioned in Section 5.2, CSPM capabilities, whether provided directly by a CSP or via a third party, make these outcomes more likely to occur.

Although certain features may encourage vendor lock-in, agencies should think about how integrated they want to be with each CSP they utilise. Using integrated services from a CSP may be advantageous for designing and implementing services as well as for managing and safeguarding the cloud environment. The internal testing of a CSP and better integration with its other features may help native CSP capabilities. Nonetheless, there may be instances when a CSP's tools are insufficient for an agency's requirements. The agency should assess third-party tools from the CSP market or commercial off-the-shelf (COTS) solutions in certain circumstances to fill the gaps. For a comprehensive, integrated strategy, organisations operating in multi-cloud environments may wish to employ capabilities that cross their CSP accounts. Both deployment and security operations may benefit from this.

Capabilities that are Independent and Integrated: To more effectively identify current vulnerabilities and ongoing compromises and to thwart future breaches, agencies' security postures may either be developed around the use of independent capabilities (i.e., stand-alone, or non-integrated capabilities) or integrated capabilities across cloud deployments. Both kinds of capabilities might contribute to the long-term security of every element of a service deployment. A pipeline component can also be modified by these capabilities

depending on the level of control an agency has over them (see Section 3.1 for the Shared Responsibilities Model). As a result, agencies typically maintain the same level of control over their deployment pipeline with integrated capabilities but can change control as necessary with independent capabilities.

There is minimal interaction between separate capabilities when they are used in the pipeline. This is shown by the division of the capabilities in the deployment pipeline at the top between Vulnerability Scanning and Assessment (VS/A), CASB, and IDS. Under this strategy, agencies are free to choose and use capabilities as they see appropriate. Therefore, agencies will need to leverage a third-party tool or create their own solutions in order to have a comprehensive picture of all of their deployed capabilities.

In addition, installations could entirely rely on integrated tools for managing unified coordination across services, such as SIEMs. While this strategy may provide agencies less discretion to deploy the capabilities of their choice, it may offer improved visibility across the deployed capabilities and across numerous deployment pipelines. An integrated set of scanning, authentication, and logging capabilities are applied to various parts of an agency's cloud deployment via the hypothetical deployment pipeline shown at Figure 12's bottom.

Figure 8.2 shows two distinct applications of independent and integrated capabilities to hypothetical deployment processes. Using the hypothetical action flows, this graphic conceptually illustrates the transfer of control from security capabilities to security capabilities:

There are three categories of data: validated (caught by the capacity), unvalidated (not yet collected by the capability), and unmonitored (not captured by a capability).

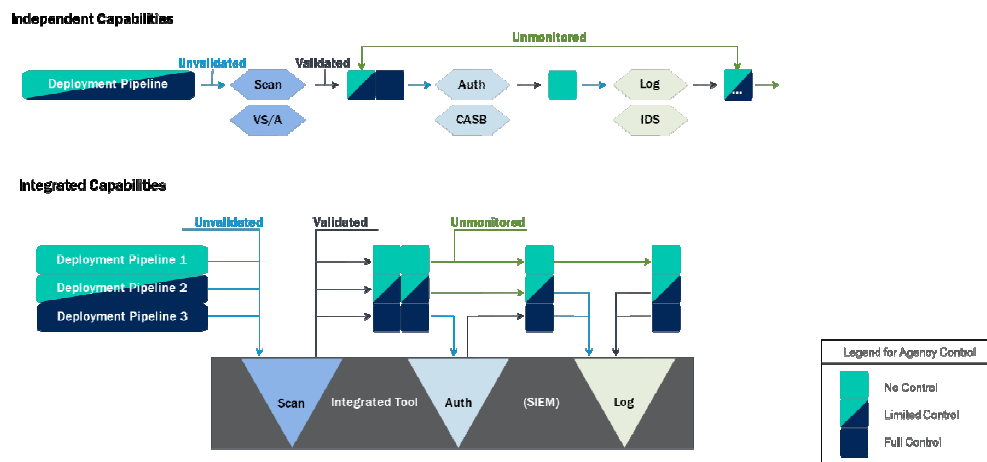


Figure 8.2: Service Deployments and Integrated Solutions

Hierarchies for CSP Account Management: Although some CSPs provide on-demand account creation, others necessitate agreement-related coordination in advance of account formation. Different CSPs have different policies regarding whether and how several accounts may be connected together inside a CSP. A main account that may keep an eye on other accounts owned by a tenant is permitted by certain CSPs. Some CSPs use a main account organisational structure that enables distinct companies to function with their own subscriptions, users, and responsibilities. No matter how a CSP manages account hierarchy and linkage, having the ability to monitor numerous accounts from a single account may provide you a comprehensive overview of all your accounts as well as a close-up view of any

specific account or organisational unit. Consequently, ICAM features, such as the ability to audit activities, develop and enforce security rules, and impose expirations on account validity, are ideally suited to managing CSP Accounts. The auditing process may also include security and risk assessments, continuous monitoring, and alerting capabilities for both monitoring real-time user behaviour and assuring the security of security policies.

Several factors that affect account structure and organisation may apply to agencies. Government and commercial cloud service providers (CSPs) may not be able to transfer data internally between the two spheres, thus agencies need to be aware of this. The agency would have to pay to have the log data leave one of the accounts (either commercial or government) and be transmitted to the monitoring account if the agency has both commercial and government accounts with a CSP and wanted all log information for accounts in a single place. If an organisation wishes to pay for that capability and move all security data on-premises for study and monitoring, many CSPs may also construct direct network connections to on-premises settings.

To limit access to assets inside a particular account, agencies should think about utilising a CSP to create several accounts or to divide entities within their organisation using the account hierarchy features that are already incorporated into the CSP. Next, agencies should create standards for establishing the hierarchical structure for accounts and allowing access. A strategy for how agencies will set up accounts with CSPs for development and testing should also be developed. For instance, IaC enables fast replication of production settings, enabling developers to test code with assurance before publication.

Management of identity, credentials, and access

Management of Identity and Credentials: How and where authentication will be carried out is one of the first architectural choices that organisations must make when transferring to the cloud. CSPs allow integration with identity providers in addition to native (e.g., isolated) authentication. To help with governance and compliance and to give advice on policies and procedures for identity and access management systems, agencies should consult papers and resources including NIST SP 800-63, OMB M- 22-0950, and the FICAM Playbooks. When accessing various CSPs, such as email housed in SaaS and an application hosted by an agency in their IaaS, federated identity providers are often used to allow users to authenticate to a single identity provider.

Users logging onto on-premises resources may also get authentication services via a federated identity provider. The Federated Identity scenario in Appendix A has further information about federated identification. Certain authentication services support single sign-on and/or MFA integration. While many authentication service providers may provide MFA, the MFA might not be PIV-enabled or phishing-resistant, which are required for government systems. In certain cases, third-party MFA programmes may be added to an authentication service, but they'll cost extra and some of them can need you to buy real hardware tokens or utilise virtual ones.

ICAM solutions from CSPs may enhance current deployments by integrating user management across on-premises and various cloud infrastructures. To improve Monitoring and Alerting capabilities, the Security Risk and Assessment and ICAM logs may be combined using a SIEM, and alerts may be established to activate when certain users carry out particular actions. Agencies can: By enabling DevSecOps Integration and AI/ML capabilities to respond to these alerts:

1. Incorporate access control into their service pipeline and utilise behaviour analysis to restrict users' activities.
2. Automatically rectify aberrant user behaviour in the CI/CD pipeline and through rollbacks in IaC.

The many authentication realms that agencies will utilise in their settings should be properly managed. Any special kind of authentication that permits a person, process, or system to access another process or system is referred to as an authentication realm. A web server, a database server, and a file server are just a few of the resources in a hypothetical IaaS cloud architecture. They're all hosted on virtual machines. These resources are accessible to the cloud administrator through a federated identity provider. This provider may be located on-site, in that cloud, or in another cloud.

A username and password are required to access the server and database that the VM server and database administrators control, respectively. A certificate is used by the webserver and the server administrator to access the respective web server and virtual machine server. With a username and password, end users may access the web server. The oval outlines in Figure 8.3 show the four unique authentication realms used in this case. Due to resource overlap, malicious behaviour in resources outside of one authentication realm might result from an exploitation in that authentication realm.

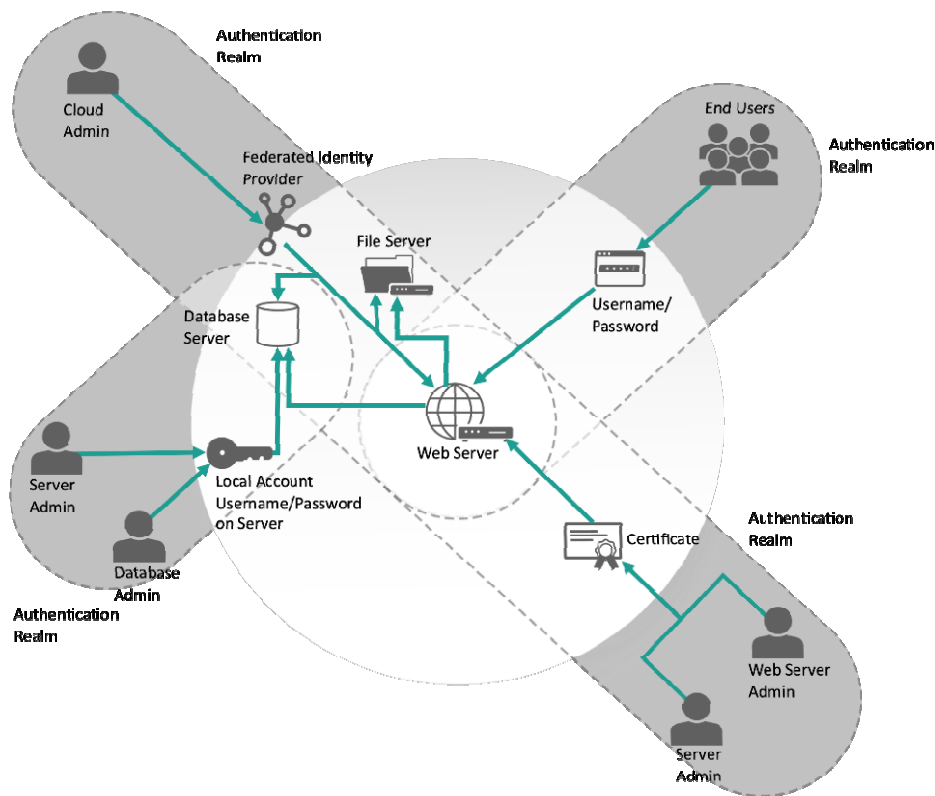


Figure 8.3: Representing the Authentication Realms.

Use of PaaS architecture, which does away with the underlying servers that house the database and web server, enables organisations to restrict the number of authentication realms, as hypothetically seen in Figure 8.4. Instead of using the methods shown in Figure 8.4 for system resource authentication, the web server and database administrators would authenticate using the federated identity provider.

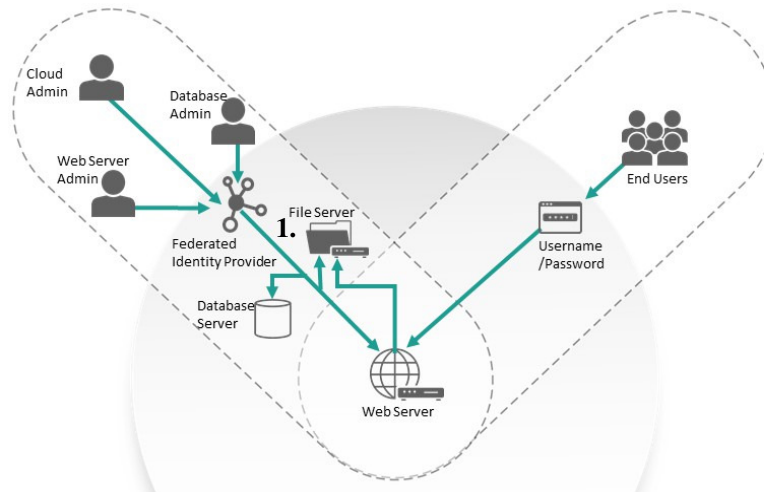


Figure 8.4 for system resource authentication, the web server

Access Control: By implementing a zero trust security strategy, agencies should impose least privileges within each authentication realm. This should involve allocating access to agency resources (such as computers, networks, administration, and data) and to individual accounts in a manner that keeps the amount of access granted to each to a level that is strictly essential for each to carry out its duties. In Security and Risk Assessments, agencies may employ auditing processes to find too privileged accounts and other account misconfigurations.

Agencies should think about establishing cloud infrastructure entitlement management in order to enforce least privilege and keep track of identity authorizations due to the complex and dynamic nature of the cloud.

The Development of the Periphery: On-premises systems have always relied on robust perimeter defences like firewalls and access control lists. This castle and moat metaphor cannot secure an agency's assets when it migrates to the cloud. In a multi-cloud scenario where they have varying degrees of control over perimeters, agencies will probably operate. Agencies will presumably be able to mimic conventional network defences in IaaS systems and add action-based defences that provide adversary identification via deception and redirection; these safeguards are probably absent in a SaaS context.

Agencies may not be able to apply allow lists or deny lists to IPs or ports to the CSP console since administrative access to a CSP's console (for example, through the web or command line) is available to the public internet. Nonetheless, organisations may apply security measures for assets located inside virtual networks built by a CSP.

Agencies should think about adopting a data-centric strategy as part of the development of their perimeter, where security measures put data and asset protection ahead of the security of applications and services. Moreover, it is possible to segregate data, applications, and services inside the network to impose more specific security rules for access to these resources.

When organisations migrate to the cloud, agencies may assess both their new cloud networks and test older networks for developing vulnerabilities by employing security scans inside Security and Risk Assessment capabilities and audits within ICAM capabilities. The addition of CDNs to DevSecOps Integration expands the perimeter beyond the systems owned, operated, or used by agencies, increasing the attack surface for data in transit and enhancing

regional visibility of activities. AI/ML skills may be utilised to better determine the categorization and sensitivity of information to detect, forecast, and monitor data exfiltration from agency and CSP infrastructure in order to analyse data in transit over an expanded perimeter. AI/ML may also use adaptive load balancing and automatic firewall management to safeguard the availability of services. To increase awareness of current and emerging risks in both their current and CSP-based infrastructure, agencies may combine all previously described capabilities into their Continuous Monitoring and Alerting capabilities.

Location of Sensors and Visibility: Agencies must be aware of the restrictions of sensor location when moving to the cloud and how these limitations may affect their ability to see log data, events, assaults, and other problems. Network taps and logs produced by equipment like firewalls are both examples of sensors. Using networking that is under CSP management, traffic is sent to a tenant on a CSP. CSPs go about their business while traffic moves along this corridor own research and may reduce or get rid of prospective dangers, such a DDoS assault. By restricting tenant sight, these CSP-internal defences may affect how well an agency understands the dangers to its cloud resources. Several CSPs include the ability to mirror network traffic, however this feature is only available for traffic that reaches a virtual network within the CSP. An agency won't have any view of such traffic if the CSP drops it for whatever reason. Moreover, even if CSPs provide safeguards, it's possible that they won't let tenants know when such protections are being used to stop harmful or suspected malicious activity. By differing in their service models, offers, and configurations, limitations on sensor placement and situational awareness might affect an agency's capacity to respond to threat actors and satisfy its visibility requirements.

In addition to monitoring and controlling traffic between services inside a network, agencies should take into account the location of sensors for both incoming and outgoing traffic. It is important to keep an eye on all incoming and outgoing traffic as it enters and leaves the property. Agencies should also think about keeping an eye on traffic between services, particularly traffic that travels over peer networks. Traffic that is unauthorised or suspicious should trigger alarms and may be diverted to secure surveillance areas. With infrastructure at scale, agencies may be better able to comprehend their cloud-based infrastructure by using the virtualization of interface monitors, network traffic sensors, and system logs Continuous Monitoring and Alerting capabilities. As part of the DevSecOps Integration capability, maintaining control over the CDN enables agencies to expand visibility with sensors deployed well beyond the agency boundary, obtaining more insight over legitimate user and malicious traffic.

By combining CSPM skills like Continuous Monitoring and Alerting and DevSecOps, Monitoring Agencies may develop a strong monitoring capacity. These capabilities may search for vulnerabilities, check system availability during regular operations and in a simulated environment, find misconfigurations, and aid in repair. More precisely, these monitoring features may assist analysts and engineers investigate breaches and maintain uptime during events, as well as detect and enumerate service uptime, quality of service, prevent malicious activity by assuring content integrity, and more. For further information about logging.

Agencies may enhance their monitoring solution(s), such as a SIEM, with Continuous Monitoring and Alerting features, which also offer scalability for both new and old log aggregation systems. To enhance quality of service, content integrity, and service uptime needs inside monitoring apps, agencies may combine automatic performance enhancements and analytics from AI/ML with their preferred authentication deployment provider, from ICAM. CDNs within DevSecOps Integration allow organizations to better understand

underlying network constraints, such as regional availability, quality of service, and demand that may make it more difficult to provide services to users, similar to how the perimeter, visibility, and sensor placement have evolved. By examining IaC, service configuration, and ICAM permissions, CSPs and other parties provide dashboards and other tools that enable agencies to find misconfigurations across their infrastructures.

Agencies should utilise monitoring to track the impact of the services they use. This may accomplish two key goals: Identify unlawful usage of services (such as shadow IT) that may occur by employees who run their own accounts with CSPs and/or operate in unapproved areas. Preserve and maintain an inventory of CSPs, CSP region operations, services, apps, accounts, and other assets.

In order to successfully monitor cloud resources, agencies will also need to take into account the threat models and geographical deployments that are suitable for their operations. Agencies may utilise a unified CSP interface or combine various areas in a third-party service if such monitoring capabilities are implemented across many geographical regions. Monitoring services should make sure that reported monitoring data and operational cycles like updates and patches are in sync once they are set. Monitoring services must make sure patches are applied correctly and report on the status of cloud deployments in order to achieve this. While existing locked-in services are frequently the only ones that can be monitored via their respective CSP monitoring services, CSP specific monitoring, like integrated capabilities, may promote vendor lock-in. The ability of agencies to complete any of the monitoring outcomes (such as compliance verification, vulnerability scanning, misconfiguration identification, and incident remediation) may be hampered by this. Particularly in a multi-cloud scenario, third party monitoring services can improve situational awareness across cloud resources. However, these same third-party monitoring services might not have the same level of depth of visibility into a specific cloud environment because CSPs might not make all pertinent monitoring data available to their users or third parties.

Interfaces for Application Programming: The abundance of APIs in cloud environments is a notable departure from on-premises environments. The ability to access numerous cloud services and functionalities is improved thanks to APIs. In order to implement automation and effective controls, use best practises that minimise environmental drift, and enable the use of their services by third parties, agencies can also adopt an API-centric and/or microservices approach to their cloud deployments. Agencies will need to manage the complexity of scaling services from both user and backend infrastructure responsibly as networks expand. The recommendations in this subsection were influenced by NIST's special publications NIST SP 800-204 and parts A, B, and C, 51, 52, 53, and 54, thus agencies are advised to review these.

Agencies may get real-time data about their use of CSP-based APIs and other service use data by integrating CSP-based Continuous Monitoring and Alerting. In order to restrict and manage access and guarantee the implementation of agency least privilege regulations, ICAM capabilities may also be embedded into API services. To guarantee that agencies utilise APIs for the administration of their cloud-based infrastructure while keeping compliance with privilege monitoring and vulnerability assessments, CSPM-based Security and Risk Assessment capabilities may be deployed. The CI/CD pipeline and API applications under DevSecOps Integration work together to benefit both parties. Development inside the CI/CD pipeline guarantees correct API use validation is completed, while utilising APIs within the CI/CD pipeline speeds up service patches.

Using APIs might expand an agency's attack surface since they include code created by outside parties over which the agency has neither visibility nor control. Thus, suitable security policies should be designed to reduce any possible cybersecurity concerns brought on by their use. Versioning of APIs should be used by agencies to track API changes and manage them over time. Versioning should be used by CSPs for their APIs, and agencies should confirm that these security precautions are in place for those API services. As a result, CSPs should provide API versioning and give tenants enough time to upgrade between releases. Agencies may use one or more of the following strategies to strengthen the security of their API-related operations.

1. Encrypt data in transit to safeguard the output and input secrecy of APIs.
2. Use API access keys as identifiers; you can track which users make certain API requests using them.
3. Agencies should create and correctly implement an API key revocation policy in the event that an API key is compromised, as well as a matching key reissuance policy, to supplement this strategy. Keys should be kept secret but also be easily disposed of when needed.
4. Use API authorisation to impose user restrictions on API requests.

Architecture focused on APIs: To establish APIs as a component of a cloud-based system, there may be a number of options. CSPs make considerable use of APIs and provide tenants with access to API families as building blocks for tasks like administration, logging and monitoring, service architecture, and interface development. For the cloud services they install, agencies should consider how they may both use CSP APIs and create their own API families. To guarantee the right permissions and access to both CSP APIs and APIs created and implemented by an agency, great effort should be taken to provide sufficient security.

Agencies should consider how and where they will gather telemetry and how they will make their logs accessible when developing apps or APIs that will be made available to consumers. Considerations for telemetry include security, speed, faults, connectivity, etc. The data structures of the logs that agencies employ as well as the APIs they utilise to acquire logs should both include versioning.

Microservices Agencies have the option to base their development and production around microservices. Using this architectural strategy, cloud-native applications are implemented as a group of lightweight, independent services. The idea of service-oriented architecture (SOA), which in turn developed from the deployment of monolithic services, gave rise to microservices. While the application code for monolithic servers is often simpler and requires less time to write than that for SOA or microservices, it does not scale well due to the close coupling between its activities and offers substantially less resilience when managing application problems. Microservices therefore provide a suitable answer to the requirements of cloud-based infrastructure. For further information, see the Microservices scenario in Appendix A.

Due to their lightweight resource use and simple deployment, container technologies and the deployment of microservices enable lower resource utilisation than when adding extra full-featured VMs or hardware (e.g., monolithic deployments). With organisational expansion, an agency may employ a container management system to monitor its microservices. Moreover, since each service may be expanded in accordance with its own demand, rather than the whole programme being scaled around a single bottlenecked service in a monolithic paradigm, microservices provide for more granular application scalability.

For inter-service communication, microservices use APIs. An API gateway is a layer that gives an agency a consistent interface to handle security, deployment, analytics, and other service use. The agency may often decide to combine various communications into one. When additional microservices are implemented, using an API gateway becomes more advantageous since each microservice is individually deployed and developed.

This microservices strategy compliments the DevSecOps capabilities of CSPM since one of the main advantages of using microservices is to minimise overall time and effort in the application development process. Each microservice may be built independently and adhere to the development, deployment, testing, security, and automation steps of the CI/CD pipeline. Also, an organisation may use microservices to execute cloud security monitoring and other tasks and scale them in accordance with operational requirements.

The manner that each agency creates software applications will likely need to alter both technologically and culturally if agencies embrace a microservices design. Changes in code structure are necessary, but the underlying process also necessitates fresh perspectives on software development lifecycles, especially when an agency switches services from a monolithic distribution. Moreover, agencies should avoid overcompensating by over-functionalizing microservices, since this might increase costs and reduce the return on investment of this model. Agencies should be aware of the amount of visibility their monitoring services have with regard to their microservices setups since microservices will also increase operational complexity because several independent services will support a single application.

Authentication and Authorization in the Cloud: Using a service mesh is a popular strategy for implementing a microservices architecture and controlling complexity. A service mesh is a specific kind of infrastructure layer that allows network regulations, traffic flow, and service-to-service communication to be configured independently of application code. "Sidecar proxies" that are installed specifically for each application and operate simultaneously with each application are used to offer service meshes. These sidecars are often built as separate containers from the actual microservices apps. Traffic must be able to be routed to and from the application via the service mesh's ingress and egress gateways or the sidecar proxies themselves. These sidecars may then be used at runtime to implement security rules. This architecture guarantees that all service communication is encrypted in transit, a critical security operational step. The service mesh must function as a certificate authority (CA) for its sidecars and support an X.509 certificate infrastructure. Despite the fact that certain service mesh solutions have this functionality, an agency should not utilise one to encrypt communication in a production cloud environment.

An attribute-based access control (ABAC) framework may be used in cloud-native deployments that make use of a service mesh. Several functional modules that are arranged into clearly defined architectures are used to construct ABAC, which defines and enacts access restrictions between a user and a protected resource. These modules' fundamental function is to establish properties based on user-object connections and set limitations on the user's ability to interact with the object.

ABAC enables the creation and execution of granular and strong authorization and authentication frameworks for data flowing inside applications when used in combination with a service mesh. Defined security rules for the microservices architecture should be encoded and distributed via a decoupled control plane in the service mesh. Then, using the control plane, an agency may support authorisation and authentication.

At the service or end-user level, or based on access control models on the control plane of the service mesh, authorization rules may be established. The sidecar proxies that carry out their enforcement are subsequently given these policies. Based on request metadata, these rules define the terms under which access may be permitted or restricted. Examples include metadata that is either source- or destination-based, such as IP addresses or ports or HTTP request parameters or characteristics. The authorization framework must also support the three tenets that make up the reference monitor concept:

Every access attempt must first use the authorization mechanism, which is supplied by the ingress/egress gateways and sidecar proxies. Modification protection is also given by the immutable and distinct ABAC modules (through independent testing and verification of each module in both shadow IT and production).

Authentication may also be provided at the level of the service or the end-user. End-user authentication is supplied via the provision of credentials, while service-based authentication is carried out using service identity profiles. In a service mesh, the sidecar proxy must enforce end-user authentication.

Agencies should assess which software stack is most suited at each tier for their particular needs when evaluating access control solutions in terms of performance, flexibility, extensibility, scalability, and process isolation. Agencies may better manage their authentication and authorization requirements as their users and resources grow in the cloud by integrating ABAC rules with service meshes.

Monitoring and Logging: While using cloud services, agencies must be aware of the logs and telemetry that are accessible to them. Setting up the groundwork for monitoring and alerting requires a thorough examination of log management procedures. Agencies should be aware of:

What data fields are in the collected logs, which log kinds are accessible, when logs are sent, and how collected logs will be processed, stored, and retrieved? In order for security teams to get the logs they need to carry out their activities more rapidly, this may assist agencies manage log creation more effectively. Moreover, organisations should take measures to evaluate and ensure that the logs they collect are correct and preserved properly (e.g., in warm storage for on-hand analysis versus cold storage for longer term retention).

Agencies may check their log use and obtain insight into their log statistics using the Continuous Monitoring and Alerting features to make sure they are recording the right information. Moreover, agencies may use these monitoring tools to make sure that the number of incoming logs does not overload log ingestion systems and to develop unique triggers for unusual occurrences. To filter log and telemetry data, agencies may use cloud AI/ML skills to remove noise. They can also use these capabilities to detect abnormal traffic based on behaviours and historical data. Using data from the CSP, such as traffic patterns and threat detection, cloud-based AI algorithms may be trained to help agencies' logging functions and reaction plans adapt to changes in telemetry (that would otherwise be unachievable). With end-user service in CDNs, agencies may employ DevSecOps Integration capabilities to gather logs from pre-deployment in the CI/CD pipeline, greatly expanding the scope of telemetry and logs.

Agencies shall follow the logging guidelines established by OMB M-21-31 in accordance with Section 8 of Executive Order 14028 when collecting logs from SaaS, PaaS, or IaaS cloud instances. The Federal Bureau of Investigation (FBI), CISA, and federal agencies are

given a set of standards to help them find risks and vulnerabilities on government cloud installations. Agencies may adhere to the following broad rules to achieve this goal:

1. Ensure that identity services are appropriately monitored for atypical authentication and login attempts, particularly in relation to "break glass" accounts, changes to privileged management or roles, and alterations to keys or secret vaults.
2. Keep an eye on alert rules and access regulations for unauthorised modifications, as well as API activity logs and service metrics for unusual behaviour.
3. Carry out basic system administration tasks, such as data loss prevention, log upkeep, and monitoring for unforeseen logging policy changes.
4. Use detection and prevention services, access managers, firewalls, web application firewalls, flow, and DNS records to identify changes in the cloud environment affecting production applications, data/log storage, and cloud network.

Time Coordination: Agencies should make sure that all gathered logs are accurate, recorded in the same time zone, and have synchronised clocks. This will make it possible to correlate all of an agency's logs notwithstanding provider or geographical variances. The latency of the logs gathered and made accessible by the CSP should be known to and understood by Agencies. As an example, many CSPs have a delay of up to 15 minutes, which prevents real-time monitoring and may exacerbate already-existing security issues related to latency. Moreover, certain telemetry and log collecting require an agency to take action in order to obtain it, including installing logging agents on virtual machines (VMs). Agencies need to be aware of how the time-related fields on each log function when collecting logs from different areas and time zones. While collecting logs as well as when they are in use, agencies should confirm the time zone in which each log is recorded. It may be necessary to configure things such that all log timestamps use the same default time zone. If that is not practicable, it may be possible to do log data normalisation during intake to guarantee proper querying of events. Agencies should also engage their CSPs to learn how they maintain correct timestamps of records and verify for clock drift in time-creating or -reporting devices.

Combining and Centralizing: Agencies should keep track of the version numbers attached to the logs and telemetry they gather so that, in the event that new versions are released, they may compare the differences and make any required modifications. It is recommended to set up a lot of logs to be automatically gathered and transmitted to storage or integrated monitoring systems (either CSP provided or third party). Whatever the method of collecting, independent of regional or provider variations, logs must finally be gathered in one central place. Some CSPs further provide the delivery of logs from numerous accounts to a central account, enabling the monitoring of logs from all accounts from a single place. Agencies should carefully prepare for how they will handle logs acquired from several regions or from numerous CSPs since some of these integration services that span regions may result in extra expenses.

On-Premises using Cloud Telemetry, Logs, and Forensics: Data gathered via the cloud and data collected on-premises may vary greatly. Latencies for log delivery from CSPs often exceed 15 minutes or more before being accessible. Some of the telemetry that agencies had accessible for their on-premises operations may not be provided by CSPs. Agencies may not have access to forensic artefacts, such as memory snapshots of computers suspected of being hacked, in certain circumstances or won't have access to them in other cases. Agencies need to understand these sortsof changes and how they affect their present threat hunting, incident response, security operations centre (SOC), and processes.

Factors to Consider While Provisioning API: Agencies should consider how and where they will gather telemetry and how they will make their logs accessible when developing apps or APIs that will be made available to consumers. Considerations for telemetry include security, speed, faults, connectivity, etc. Agencies should also enforce rate limitations to stop DoS attacks, version both the APIs used to gather logs and the data structures of the logs they utilise, and keep an eye on API activity for future measurements and reporting. For event-based infrastructure, agencies may want to think about developing webhooks to assist lighten the burden on API requests.

SaaS considerations: There are several methods for log gathering for SaaS companies. An linked IaaS or PaaS account, API requests to gather logs, third-party collection tools, and log exports are all ways to make logs accessible. If at all feasible, avoid manually exporting logs in favour of an automated, scalable collecting system. Tenants are not able to gather extra log data for security reasons beyond what the service provider gives since the service provider is in charge of the technological stack and the SaaS offering. Logs are frequently sorted by API families in SaaS settings and are typically produced by API calls that the service provider utilised to develop the SaaS product. The most common method of accessing logs is via APIs created by the service provider, although some of them could also include security dashboards or log viewers in their administrative panel. Several SaaS providers construct their services on top of those provided by other CSPs. This can restrict the amount of data that the SaaS provider and, by extension, the tenant, can access.

Taking into account IaaS and PaaS: The CSP makes a large number of logs accessible for IaaS and PaaS implementations that may be collected to improve situational awareness of the environment. They might contain health logs, access and identity logs, network flow logs, API call logs, and service event logs. The majority of IaaS and PaaS providers provide built-in solutions for gathering logs and storing them in a centralised place. In order for one account inside a CSP to monitor numerous accounts utilised by an agency, mechanisms for gathering and sharing logs across linked accounts may also be offered. This enables the creation of accounts depending on roles or functions.

Orchestration, Automation, and Deployment: Agencies may combine services and automate deployment together in ways that are not possible on-premises because to the cloud's dynamic nature. By integrating DevSecOps into their development processes, agencies may automate the deployment of new applications. This paradigm encourages a security-first attitude, which is crucial for managing the difficulties brought on by CSPs' frequent modifications to cloud services.

Include DevSecOps in order to build and deliver code that has security built-in from the start rather than put on afterwards, DevSecOps brings together development, security, and operations teams as one cohesive entity. While though DevSecOps is often focused on production cloud deployments, this security-first attitude can be applied to any cloud environment.

To create and test their deployments, developers utilise CI. Operation engineers use CD mechanisms to plan their deployments and keep an eye on them to make sure they're healthy and available. To ensure that the new deployments adhere to security requirements, security engineers collaborate with developers to write tests that are executed as part of unit integration and/or system testing. The DevSecOps team's security professionals also make sure automated tests are set up to check for typical application vulnerabilities before deployment. In order to ensure that the right security practises are used, security people collaborate with developers throughout the design phase. They also cooperate with operations

staff to guarantee that the deployment is secure, adequately monitored, and patched on schedule. Security specialists keep an eye out for security problems throughout the iterative DevSecOps process. For further information on DevSecOps.

Management of Deployments: Agency cloud installations may be swiftly and easily modified because to the virtual environment of the cloud. Patches for vulnerabilities and in-place operating system (OS) and application upgrades are common in on-premises settings. This procedure is often carried out outside of regular business hours, which causes some downtime. Tools provided by several CSPs and third-party providers alter this paradigm by allowing "zero-downtime" updates (i.e., deploying upgrades without halting current operations). Incorporating appropriate ICAM capabilities with adaptive AI/ML skills may enable agencies to deploy and orchestrate IaC with greater accuracy and faster reaction times.

Agencies may produce basic or "golden" VM images as well as container images to do this. Security apps are installed, security rules are set up, and necessary patches and updates are deployed to these images. The outputs of the operation are then validated by scans to see whether a picture satisfies all necessary security requirements. After being created, these pictures may be saved in repositories and subsequently utilised to replace active production images. The process of creating fresh images may be done on a regular basis such that they are issued every month, week, hour, or even in reaction to newly identified vulnerabilities. For instance, CI/CD pipelines should take action to inform about and fix codebases that employ insecure settings, packages, and libraries. In order to prevent regressions caused by new upgrades to applications (OS) or services (container) on golden images, system and integration tests should also be revalidated.

A container that is constructed nightly to include the most recent libraries it required for functioning may serve as an example of this form of deployment. Once the container passes a series of security checks and testing, it may be sent via deployment. The new container may then be pointed to by all incoming connections. The prior container is terminated as soon as connections to it stop working. Depending on whatever pre-deployment tests the new container fails, the relevant engineers are notified, and they may fix the highlighted problem (s). Agencies should utilise a vetting solution to make sure library dependent versions are "secure/updated" throughout this deployment process and be mindful of supply chain issues with open source technologies.

Also, the cloud enables organisations to assign a large number of maintenance jobs to the CSP, which provides IaaS, PaaS, and SaaS computing alternatives. Agencies may be able to concentrate on their mission demands as a result. In the aforementioned container scenario, an agency might deploy its containers using a serverless platform provided by a CSP. In this instance, the agency does not have to worry about a variety of deployment-related issues, including server installation, licencing, patching, monitoring, and upgrading, as well as the licencing and installation of container orchestration software. The agency could yet need to configure the container software orchestration programme in some way.

Agencies should adopt an IaC approach while developing, configuring, and deploying. IaC enables organisations to deliver configuration settings for anything from VMs and networks to CSP-managed services. There are several CSPs that provide IaC management and scripting tools, as well as third-party vendor software that may be compatible with different clouds. Agencies should follow recommended procedures when using configuration management tools code management and storage for IaC code. For instance, confidential data like keys, emails, and passwords shouldn't be stored in code repositories. One method of asynchronously managing settings across computers is using version control systems.

Key Administration: Modern cloud-first key management techniques may provide seamless encryption across a company's cloud deployment. Agencies may use the server-side encryption (SSE) offered by CSP or use a third-party key management provider. It is not recommended for agencies to create their own encryption software. Therefore, agencies should make sure the service complies with their threat model's needs before choosing any key management provider. Agencies may try to adopt a different key management technique if they discover that a CSP or third-party supplier does not satisfy their needs.

For instance, an agency could want to make sure that the information gathered by their application is protected so that only the agency can access it and examine it, and the CSP cannot. To make sure that no one has simultaneous access to encrypted material, keys, rules, and monitoring, agencies should think about instituting separation of tasks. Secrets needed for services (such databases, network file sharing, APIs, etc.) should also be rotated periodically in addition to keys. Agencies may try to employ CSP and outside vendor products that provide rotating passwords, certificates, and keys. Agencies should also decide how secrets will be kept, whether in a software (such as a time-based one-time password (TOTP) authenticator programme) or hardware (such as a hardware security module (HSM)) arrangement, and consider choices in line with their threat model.

Implementation Management: Agencies should keep an eye out for drift, or unintentional configuration changes, in their settings now that quick deployment is possible on the cloud. Little, gradual modifications may easily go undetected, but a significant configuration change is likely to be recognised and discovered fast. These drifts may eventually add up and cause a major shift in the environment, rendering it incompatible with the security plans and ATO for which it was first certified. To make sure that rogue or unintentional modifications may be found and fixed, planned changes must be authorised. For further details on configuration management.

As federal agencies continue to utilise cloud technology, this Cloud Security Technical Reference Architecture demonstrates suggested methods for cloud migration and data protection. These methods will help the federal government better secure the.gov business as well as identify, detect, defend, react to, and recover from cyber events. Also, when organisations' network infrastructures change, these methods educate organisations on the benefits and inherent hazards of adopting cloud-based services. By an emphasis on cloud modernization, this Cloud Security Technical Reference Architecture helps federal agencies continue to develop within a quickly changing technological environment.

CHAPTER 9

FEDERATED IDENTITY MANAGEMENT

Dr. Suchithra R

Professor, Department of Computer Science and Information Technology,
Jain (Deemed to be University) Bangalore, Karnataka, India
Email Id-r.suchithra@jainuniversity.ac.in

Further information on federated identity management, microservices, and a warm standby site in the cloud is provided by the following three scenarios. They don't aim to address every implementation scenario and are purposefully limited in scope. Identity Management in the Cloud: The key to business security is identity management. Agencies must make choices about how to manage identities across the many domains, services, and apps they employ when they migrate to the cloud.

In the past, software was deployed in a typical business setting after being acquired from suppliers. Agencies are increasingly purchasing services from cloud service providers or from vendors that operate their software as a service outside of an agency's environment, going beyond the conventional on-premises setting. Without an integrated authentication solution, identity providers would be needed for each unique service environment, resulting in the need for several identities for each user inside an organisation.

Federated identity management is a way to ease the load of handling these various identities. Applications and services may utilise a single identity provider as the source of authentication for identities across domains by employing authentication standards like the most recent iterations of SAML and OpenID. Yet, just because there is only one identity supplier doesn't imply that an organisation has to utilise it or even ought to. To estimate how many identity providers are required to satisfy an agency's system needs, a number of criteria should be taken into account. The authentication standards create a trusting connection between each domain or service provider and the identity provider.

A user asks access to a service in Figure 9.1. The service and an identity manager who controls identities have a trust relationship. The user may submit credentials at the service and the service will send them to the identity provider, or the user may be diverted to the identity provider and then sent back to the service provider, depending on how the authentication is performed. Because of their mutual trust, the identity provider and service provider may approve the user's login once the identity provider has validated their credentials.

Factors for Implementation: Single sign-on may also be used to lessen the friction experienced by employees who often switch between apps and services as part of their work. Federated identity management systems may include multi-factor authentication that is resistant to phishing. For the business, identity management takes place in a single, centralised location rather than across domains or apps. This makes it simpler to manage identities when onboarding new employees or denying access to leaving employees.

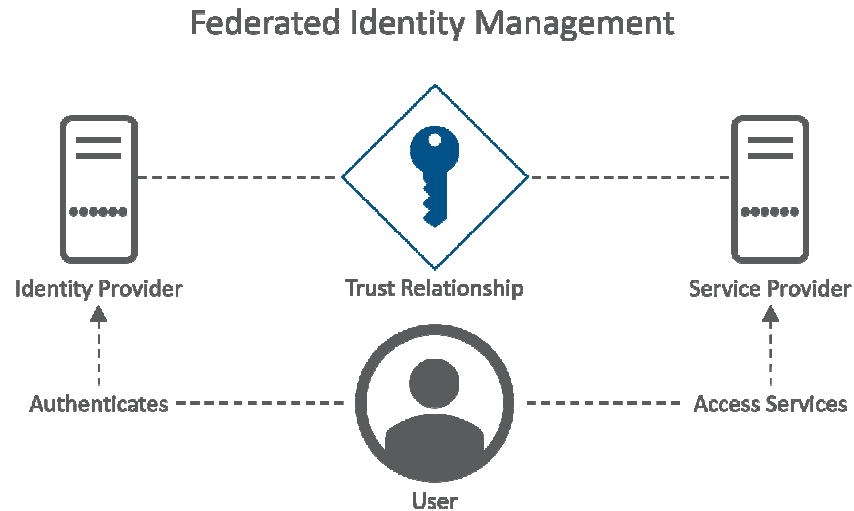


Figure 9.1: Federated Identity Management

Instead of having services in place to gather authentication logs across domains or services for analysis, centralised authentication logs enable quick analysis of user behaviour to discover questionable logins or login attempts. Threat actors may utilise a user's compromised credential to gain access to other services if a federated identity management system has been hacked. A service or application does not necessarily need to use federated identity management. In certain situations, it could be preferable to establish a distinct authentication realm for services, applications, or data with high levels of security.

Microservices: In order to better safeguard its assets, a well-established agency with experienced development and DevOps teams wishes to build a zero-trust architecture (ZTA) as part of its migration to the cloud. The organisation aims to combine this technology with its present infrastructure to save costs while still managing rising service demand and being adaptable to new needs.

Such an organisation would often use a monolithic design, requiring any new services to alter a centrally located codebase where, typically, modifications are difficult to scale. Moreover, network rules would have strict definitions and manual on-site setup. Also, setting up services would be done on a device-by-device basis, which may lead to consistency flaws and management issues with policies. A service mesh and other associated elements of a microservices architecture enable configuration to be sent to networked devices universally or granularly, depending on agency needs. To handle many services, the agency chooses to use a service mesh with a secure authentication and authorisation architecture. The service mesh with sidecar proxy offers additional security advantages that improve independent development and deployment since each microservice is untrusted:

The mesh in this situation enables the deployment of DevSecOps pipelines for IaC and policy-as-code, including security from the outset. Each microservice carries out a single, well defined business function; they are atomic in nature and autonomous. As a result, the creation of microservices is carried out in a decentralised manner, generally with small teams providing separate code for a service. A microservices-based architecture is complemented by the mesh by compartmentalising different cross-cutting phases of data analysis pipelines. This skill aids in the large-scale collecting and analysis of very diverse and unstructured data. The design may be used in a variety of specialised fields, including SCADA and IoT systems with limited resources. An example of one such solution is shown in Figure 9.2, where the

service mesh is built using sidecar proxies that are deployed per-service (depicted via circles containing opposing arrows) proxies in a sidecar applications that isolate certain elements from the primary architecture, such as inter-service communication, monitoring, and security, to make managing and maintaining the application as a whole simpler. This is the method for granularly pushing network and security settings to microservices. Each microservice makes use of a separate dedicated data store and may be created by a different development team. Agencies may use the API gateway, which controls interfaces for all microservices, to access business operations.

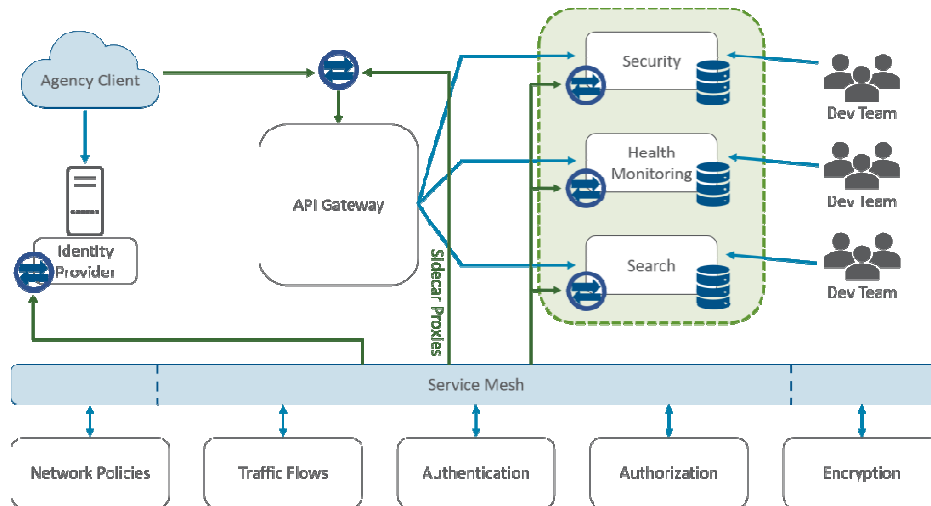


Figure 9.2: Represent the Microservices

Factors for Implementation: An agency should take into account the needs of each security service and be aware of the risk involved before integrating a security feature into a microservice. For instance, using a reverse proxy to provide TLS encryption across containers in the microservices-based architecture described above may result in single points of failure. This has to be compared to the dangers posed by unencrypted data in transit.

As each microservice must have its own autonomous development, data consistency may be a problem. Agencies should assess the trade-offs between data consistency and availability and choose a method that is suitable for their individual requirements. In order to ensure consistency, distributed data stores may be examined and updated by a separate function using a rolling data update technique.

Warm Standby: In times of crisis and heavy consumption, an agency would prefer to effortlessly move workloads to a warm location in the cloud. Regular updates must be made to this warm site, and if it is practical after failovers, updates must be made to the local live systems as well.

The cloud warm site should be synchronized to replicate security management, network access, service gateway, and data storage functions; however, the function of data manipulation should not occur in a warm site. A hot site is often built up by agencies seeking high availability for their operations to supplement performance and traffic that fail over from their principal systems. In these "hot sites," infrastructure replicas are synchronised right away (see Figure 9.3 Agency Main Site for an example), and traffic is distributed equally across all replicas to improve network performance via load balancing. By shifting target protection strategies and reducing denial of service attacks, this load balancing also improves security.




Figure 9.3: Cloud Warm Site Synchronization and Failover Movement

A warm site is deployed as opposed to cold sites, which have long-term storage with rare access, and hot sites, which have complete, instantaneous, mirror tools, to allow operation continuation under poor availability or very hostile situations. In contrast to normal operations, these warm sites simply handle traffic and simple read-only requests while recovering systems and creating new copies. Warm sites synchronise security management components like firewalls, network access components like routers, service gateways like web servers, and service data with their primary counterparts at progressively less frequent intervals. For instance, security management systems must be updated immediately to ensure proper configuration, whereas service data systems should not be updated immediately to stop the spread of adversarial data corruption. Figure 17 illustrates how, in addition to a conventional hot site fail over and load balancing system, a cloud-based warm site may be employed to handle fail overs. Keep in mind that data manipulation and processing shouldn't be enabled in warm clouds. The warm site in this case uses IaaS, while alternative service implementations might be employed.

Factors for Implementation: In the worst-case situation, a warm site guarantees activities will continue while maintaining original settings and enabling the affected environment to stay undamaged, assisting in reaction and recovery. As writes may further taint sensitive data and since working in several settings without real-time synchronisation may result in discrepancies in data storage across cloud and conventional environments, data should largely be accessed in a read-only form.

Agencies may use CSPM capabilities, such as Security and Risk Assessments and DevSecOps, to increase security without making changes to the current infrastructure as the agency's network is extended towards cloud-based warm sites by starting with the implementation of secondary fail over measures in cloud environments. An agency will need staff experienced with synchronization, access control, capability implementation, and the general vulnerabilities and constraints of CSP environments to guarantee correct configuration and administration before to and during crises. These employees will support agencies as they install new cloud-based systems in the future.

Application Programming Interface (API): A system access point or library function that has a clear syntax and may be used from user code or application programmes to give functionality that is clear and well-defined. Any special kind of authentication that permits a person, process, or system to access another process or system is known as an authentication realm. An official management decision to authorise the operation of an information system and to explicitly accept the risk to organisational operations (including mission, functions, image, or reputation), organisational assets, people, other organisations, and the Nation based on the implementation of an agreed-upon set of security controls is known as a "authority to operate" (ATO).

An authorising authority must provide permission for the functioning of every component of an information system, with the exception of any independently approved systems to which the information system is linked. **Cloud Access Security Brokers (CASBs):** A computer programme that controls access to protected data and keeps records using up-to-date encryption keys and log files. **Cloud Security Posture Management (CSPM)** is a continuous process for keeping an eye on a cloud environment, finding and alerting on cloud vulnerabilities, and reducing them as well as enhancing cloud security.

A third-party business that offers its customers a platform, infrastructure, applications, and/or storage services is known as a cloud service provider (CSP). In order to provide safe, quick, and effective data delivery, a linked network called a content delivery network (CDN) delivers caches of files or services across several locations. **Continuous Integration (CI)** is the practise of automating and incorporating changes to code made by many teams when developing software. **Continuous Delivery (CD)** is the practise of automating the release of new software into production. **Continuous Monitoring (ConMon):** A procedure that guarantees CSPs continually maintain the security of their FedRAMP-authorized systems by giving the Joint Authorization Board (JAB) and Authorizing Officials (AOs) visibility into the system's security posture on a monthly basis.

Desktop-as-a-Service (DaaS): DaaS is a cloud computing service in which a service provider uses the Internet to provide virtual desktops to end users who have paid a per-user subscription. A software development strategy known as Development, Security, and Operations (DevSecOps) closely combines developing code with testing, securing, and deploying that code. Applications and services that offer digital information, such as data or content, as well as those that provide transactional services are together referred to as "digital services" (e.g., online forms, benefits applications) spanning several platforms, gadgets, and distribution methods (e.g., websites, mobile applications, and social media) similar to CSP services.

A subset of U.S. government departments and agencies that excludes the Department of Defense and organisations in the Intelligence Community is known as the Federal Civilian Executive Branch (FCEB). In order to meet federal business goals, identity, credential, and access management (ICAM), a core and crucial cybersecurity capability, guarantees that the appropriate individuals and objects have access to the appropriate resources at the appropriate time for the appropriate purpose. The capacity to supply processing, storage, networks, and other essential computing resources is known as infrastructure-as-a-service (IaaS). With this capability, the consumer is able to create and run their own software, which may include operating systems and applications. While the user has no management or control over the underlying cloud infrastructure, they do have some limited influence over certain networking components, operating systems, storage, and installed applications (e.g., host firewalls).

IT infrastructure management and provisioning utilising machine-readable configuration files rather than physical hardware configuration or interactive setup tools is known as infrastructure as code (IaC). Intrusion Detection and Prevention Systems (IDS/IPS): Software that automates the process of keeping track of the activities taking place on a computer system or network, examining them for any indications of potential incidents, and making attempts to thwart any potential incidents that are identified. A design concept known as least privilege states that each entity should only be given the system resources and permissions that are absolutely necessary for it to carry out its purpose.

Multi-Factor Authentication (MFA) is an authentication method that necessitates the use of several unique authentication factors. A multi-factor authenticator or a group of authenticators that each give a separate factor may be used to achieve multi-factor authentication.

Platform-as-a-Service (PaaS): The capacity to install consumer-generated or purchased applications produced using programming languages, libraries, services, and tools supported by the provider into the cloud infrastructure is made available to the customer. The customer has control over the deployed apps and perhaps the configuration options for the application-hosting environment but does not manage or control the underlying cloud infrastructure, including the network, servers, operating systems, or storage.

A certificate-based public key cryptography system's development and operation are supported by the architecture, organisation, methods, and processes known as the public key infrastructure (PKI). Public key certificates may be issued, updated, and revoked using a framework. SCADA is a control system architecture that uses computers, networked data transmission, and graphical user interfaces to provide high-level monitoring of equipment and processes.

Service Level Agreement (SLA): A contract that specifies the particular duties of the service provider and establishes the expectations of the client.

Service Mesh: A specialised infrastructure layer that enables network rules, traffic flow, and inter-service communication without relying on application code. A microservices architecture is supported by a service mesh.

Software-as-a-Service (SaaS): The ability to utilise the provider's applications that are operating on a cloud infrastructure is made available to the customer. Via a programme interface or a thin client interface like a web browser (for example, web-based email), the programmes may be accessed from a variety of client devices. With the potential exception of a small number of user-specific application configuration choices, the customer does not manage or control the underlying cloud infrastructure, which includes the network, servers, operating systems, storage, or even particular application capabilities.

Telemetry: Artifacts produced by security systems that show the state of security.

Technical visibility includes things like assets, users, systems, data, logs, etc. Operational visibility includes things like usage, criticality, risks, etc., and organisational visibility includes things like mission functions, operations, priorities, etc. Though one aspect may be specified, the three together are frequently of concern.

Zero Trust: A group of notions and theories intended to reduce uncertainty in enforcing precise, least privilege per-request access choices in information systems and services in the face of a network that is thought to be hacked.

Zero Trust Architecture: A cybersecurity strategy used by an organisation that includes component interactions, workflow design, and access controls. As a result, a zero trust enterprise is an organization's network infrastructure (both physical and virtual) and operational rules as a result of a zero trust architectural plan.

CHAPTER 10

INFRASTRUCTURE SECURITY

Dr. Murugan R

Associate Professor, Department of Computer Science and Information Technology,
Jain (Deemed to be University) Bangalore, Karnataka, India
Email Id-murugan@jainuniversity.ac.in

An instance of an issue related to this first risk element is a security hole in Amazon Web Services (AWS) that was revealed in December 2008. Users who chose to use HTTP instead of HTTPS had a higher chance that their data may have been changed in transit without their knowledge, even though utilising HTTPS would have reduced the integrity risk.

Adequate Access Control is ensured: A company employing a public cloud now risks a considerable increase in risk to its data due to some portion of these resources (or maybe even all of them) being exposed to the Internet. Even after the event, it is probably impossible to audit the network operations of your cloud provider (much alone do any real-time monitoring, as on your own network). Your capacity to conduct in-depth investigations and obtain forensic evidence will be constrained, and you will have less access to pertinent network-level logs and data.

The issue of reused (reassigned) IP addresses is one illustration of the issues connected to this second risk element. Typically speaking, when IP addresses are no longer required for one client, cloud providers do not "age" them properly. Once addresses become available, other customers often reassign and reuse them. This makes sense from the standpoint of a cloud service provider. IP addresses are both a billable asset and a limited resource. Yet, the persistence of IP addresses that are no longer in use might be problematic from a customer's security standpoint. A client cannot trust that the termination of network access to its resources occurs when its IP address is made public. Between changing an IP address in DNS and having that address cleared from DNS caches, there is a delay that is unavoidable. When physical (i.e., MAC) addresses are updated in ARP tables, there is a comparable delay until the previous addresses are purged from cache; an old address remains in ARP caches until they are cleaned. As a result, users may still access these ostensibly nonexistent sites even if addresses may have changed since the (now) previous addresses are still stored in cache. Several allegations of "non-aged" IP address issues at one of the biggest cloud providers have recently surfaced; this is likely what prompted AWS to announce the availability of Amazon Elastic IP in March 2008. Customers that use elastic IP addresses get a block of five routable IP addresses that they may assign different addresses to. In addition, Simson Garfinkel claims:

Yet, routable IP addresses are not the only ones affected by the problem of "non-aged" IP addresses and illicit network access to resources (i.e., resources intended to be reachable directly from the Internet). The problem also affects non-routable IP address assignments and internal networks used by cloud providers for client usage. Even though they may not be immediately accessible from the Internet, your resources must be available through private addressing inside the cloud provider's network for management reasons. Every resource with a public or Internet-facing address also has a private address. It's possible that less than honest clients of your cloud provider might use their networks to access your resources

inside. The Washington Post stated that AWS has had issues with users abusing its resources in ways that harm the general public and other customers.

IP address reuse is an issue that some new solutions on the market* will assist to fix, but unless cloud providers provide these products as managed services, consumers are paying for yet another third-party product to tackle a problem that their cloud provider's policies caused for them.

Ensure Internet-Facing Resources Are Available

Since more data or more employees of a business rely on externally hosted devices to maintain the availability of cloud-provided services, network security has become more important. As a result, your company must agree to the three risk factors listed in the previous section. This third risk aspect is well shown by BGP prefix hijacking (i.e., the fabrication of Network Layer Reachability Information). Announcing an autonomous system's address space that belongs to someone else without her consent is known as prefix hijacking. Such announcements often result from setup errors, but such errors may still have an impact on the accessibility of your cloud-based services. Some hundreds of these misconfigurations happen each month, according to a research that was presented to the North American Network Operators Group (NANOG) in February 2006. § Most people are probably most familiar with the misconfiguration blunder that Pakistan Telecom committed in February 2008 when it announced a fake route for YouTube to its own Hong Kong-based telecoms partner, PCCW. In Pakistan, YouTube was intended to be blocked due to several allegedly blasphemous films that were uploaded there. As a consequence, YouTube was down for two hours everywhere.

Misconfigurations are not the only thing that happens; purposeful attacks sometimes occur. Prefix hijacking as a result of malicious assaults is less often than misconfigurations, but it nevertheless happens and may prevent access to data. Less than 100 assaults take place each month, according to the same report that was presented at NANOG. While prefix hijackings are not new, the number of attacks will undoubtedly increase and perhaps considerably as cloud computing grows. The value of cloud-based resources for clients improves as cloud computing use rises. Customer value has risen, which raises the possibility that malevolent behaviour may try to undermine that availability.

DNS# assaults are one more issue connected to this third danger aspect. In reality, there are a number of DNS attack types to be concerned about in relation to cloud computing. The problem with DNS and cloud computing is an increase in an organization's risk at the network level due to increased external DNS querying (reducing the effectiveness of "split horizon" DNS configurations) along with some increased number of organisational personnel being more dependent on network security to ensure the availability of cloud-provided resources being used, even though DNS attacks are not new and are not directly related to the use of cloud computing.

While the "Kaminsky Bug" (CVE-2008-1447, "DNS Inadequate Socket Entropy Vulnerability") received the majority of the attention in 2008 regarding network security, other DNS issues also have an influence on cloud computing. Not only are the DNS protocol and implementations vulnerable, but DNS cache poisoning attacks, which deceive DNS servers into accepting false information, are also extremely common. Despite the widespread belief that DNS cache poisoning attacks were no longer a threat, particularly in the context of cloud computing, this is untrue. These attacks are still a major concern.

The target domain's name server (NS), the NS record, and replying before the actual NS are examples of variations on this fundamental cache poisoning attack (called DNS

forgery). Denial of service (DoS) and distributed denial of service (DDoS) assaults are a last illustration of issues related to this third risk element. Although while DoS/DDoS assaults are nothing new and have nothing to do with the usage of cloud computing, the problem with these attacks with cloud computing is an increase in an organization's network-level risk due to some greater use of resources outside of your organization's network. For instance, there are still claims of ongoing DDoS assaults on Amazon, which would prevent customers from using the services for hours at a time. § (Amazon has not accepted that DDoS assaults are the real cause of service outages.)

The danger of a DDoS attack while adopting IaaS, however, is not only external (i.e., Internet-facing). The IaaS provider's network, which is independent from its corporate network and is utilised by its clients, is likewise susceptible to internal DDoS attacks. Customers utilise this internal (non-routable) network to access their private instances (such as Amazon Machine Images or AMIs), while the provider uses it to manage its own network and resources (such as physical servers). If were a malicious customer, nothing would stop me from using my customer access to this internal network to locate and attack other customers or the infrastructure of the IaaS provider—and the provider most likely wouldn't have any detective controls in place to even be alerted of such an attack. Some customers' main preventative measures would be the degree of hardening applied to their instances (such as AMIs) and if they make use of a provider's ability to firewall off groups of instances (e.g., AWS).

Domains to Replace the Existing Network Zones and Tiers Model

With the public IaaS and PaaS clouds, the traditional isolation concept of network zones and tiers no longer applies. For years, network security has depended on zones to separate network traffic for better protection, such as intranet against extranet and development versus production. This approach was built on exclusion; only people and systems with certain responsibilities could enter certain zones. Similar to this, systems inside a given tier often only have specialised access to resources within or across that tier. For instance, systems within a presentation tier are only permitted to connect with another approved system within the application zone; they are not permitted to contact directly with systems in the database tier. Similar properties apply to SaaS clouds based on open IaaS or PaaS clouds. Nonetheless, a public SaaS based on a private IaaS (such as Salesforce.com) may adhere to the classic isolation approach, but topological details are often not disclosed to clients.

"Security groups," "security domains," or "virtual data centres," which have logical separation between tiers but are less exact and provide less protection than the earlier established architecture, have taken the role of the conventional model of network zones and tiers in public cloud computing. For instance, the security groups feature in Amazon enables communication between your virtual machines (VMs) via a virtual firewall that may filter traffic based on IP address (a particular address or a subnet), packet type (TCP, UDP, or ICMP), and ports (or a range of ports). Based on DNS, domain names are utilised in a variety of networking scenarios as well as for application-specific naming and addressing requirements.

Development systems and production systems were physically segregated from one another at the host level in the established concept of network zones and tiers in addition to their logical separation at the network level (i.e., they ran on physically separated servers in logically separated network zones). Nevertheless, with cloud computing, this division is gone. Logical separation is offered by the cloud computing concept of domain separation for

addressing-only needs. A test domain and a production domain may very well be on the same physical server, hence there is no longer any "necessary" physical separation.

Moreover, the previous logical network separation is no longer present; rather, logical separation now takes place at the host level, with both domains operating on the same physical server and being only logically isolated by VM monitors (hypervisors).

Mitigation at the Network Level

What can you do to lessen these higher risk factors, given the reasons covered in the sections above? First, be aware that network-level dangers remain irrespective of the "cloud computing" services being used (e.g., software-as-a-service, platform-as-a-service, or infrastructure-as-a-service). Because of this, the major factor in determining risk level is not the IaaS being utilised, but rather whether your firm plans to use or already uses a public, hybrid cloud, or private cloud. Virtual network zoning is a feature that some IaaS clouds provide, however they may not be compatible with an internal private cloud environment that uses stateful inspection and other network security measures.

Your risks will go down if your company is big enough to afford private cloud resources, providing you have a real private cloud that is internal to your network. Depending on the capabilities and maturity of the provider, a private cloud housed within the premises of a cloud provider could in certain circumstances help you achieve your security needs.

By implementing encryption, especially certified implementations of cryptography for data-in-transit, you may lessen the dangers to your secrecy. Data integrity is ensured by secure digital signatures, which make it very difficult, if not impossible, for someone to alter your data.

Unless your company uses a private cloud that is inside to your network architecture, it is far more difficult to minimise availability issues at the network level using cloud computing. You will experience greater risk at the network level even if your private cloud is a private (i.e., non-shared) external network at a cloud provider's facility. A public cloud is significantly more vulnerable. Yet, let's keep things in perspective. Greater than what?

Even huge businesses with plenty of resources struggle to maintain infrastructure security at the network level. Are the hazards connected to cloud computing really worse than the threats businesses are now facing? While establishing such a comparison, take into consideration partner connections and existing private and public extranets. Is the danger of utilising public clouds—assuming that such organisations lack the resources required for private clouds—really greater for major corporations without considerable resources or for small to medium-sized businesses (SMBs) than the hazards associated with their present infrastructures? The majority of the time, the answer is probably no there is no increased danger.

Security of Infrastructure at the Host Level

You should take into account the context of cloud service delivery methods (SaaS, PaaS, and IaaS) and deployment models when considering host security and evaluating risks (public, private, and hybrid). VM escape, system configuration drift, and insider threats via lax access control to the hypervisor are some virtualization security threats that carry over into the public cloud computing environment, despite the fact that there are no known new threats to hosts that are specific to cloud computing. From the standpoint of security management, the dynamic (elasticity) character of cloud computing may provide new operational issues. The operational approach encourages quick provisioning and transient VM instances. As the pace

of change is far greater than in a typical data centre, managing vulnerabilities and fixes is significantly more difficult than just conducting a scan.

Consider the "velocity of attack" element in the cloud. In addition, the fact that clouds harness the power of thousands of compute nodes mixed with the homogeneity of the operating system used by hosts implies that threats may be magnified fast and cheaply. More significantly, you need to comprehend the trust boundary and your obligations to safeguard the host infrastructure that you are in charge of. The same should be compared to providers' obligations to secure the portion of the host infrastructure that the CSP is in charge of.

Security for SaaS and PaaS Hosts

As hackers might use such knowledge to their advantage when attempting to get into the cloud service, CSPs generally avoid disclosing information about their host platforms, host operating systems, and the security measures that are in place to protect the hosts. Because of this, host security in the context of SaaS (such as Salesforce.com, Workday.com) or PaaS (such as Google App Engine, Salesforce.com's Force.com) cloud services is opaque to consumers, and the CSP is in charge of keeping the hosts secure. You should request that the vendor disclose information under a non-disclosure agreement (NDA) or simply insist that the CSP release the information through a controls assessment framework like SysTrust or ISO 27002 in order to get assurance from the CSP on the security hygiene of its hosts. From the standpoint of controls assurance, the CSP must guarantee that the proper preventive and detective controls are in place and will have to do so via a third-party evaluation or an assessment methodology similar to ISO 27002.

It is typical for CSPs to include virtualization platforms, such as Xen and VMware hypervisors, in their host computing platform architecture since virtualization is a major enabling technology that enhances host hardware usage in addition to having other advantages. You should be aware of the provider's method for safeguarding the virtualization layer as well as how the company is using virtualization technology.

A host abstraction layer is used by the PaaS and SaaS platforms to abstract and conceal the host operating system from end users. The accessibility of the abstraction layer that conceals the operating system services that programmes use is a significant distinction between PaaS and SaaS. SaaS users do not have direct access to the host abstraction layer; instead, only developers and the CSP's operations staff have access to it. In contrast, PaaS users have indirect access to the host abstraction layer through a PaaS application programming interface (API), which in turn communicates with the host abstraction layer. In summary, if you use SaaS or PaaS, you depend on the CSP to provide a secure host platform on which the CSP and you, respectively, create and install the SaaS or PaaS application.

In conclusion, the CSP now has responsibility for host security in SaaS and PaaS services. A significant advantage from a security management and cost perspective is the fact that you do not have to worry about shielding hosts from host-based security threats. Yet, maintaining information housed by cloud services is still a risk that belongs to you as a client. Your duty is to get the required degree of assurance on the CSP's management of host security hygiene.

IaaS Host Protection

IaaS clients, in contrast to PaaS and SaaS users, are principally in charge of protecting the cloud servers they have been given access to. Considering that practically all IaaS services now offered use virtualization at the host layer, the following categories should be used to describe host security in IaaS:

Security software for virtualization

The layer of software that is built on top of bare metal that allows users to build and delete virtual instances. Any virtualization approach, such as OS-level virtualization (Solaris containers, BSD jails, and Linux-VServer), paravirtualization (a mix of the hardware version and versions of Xen and VMware), or hardware-based virtualization, may be used to implement virtualization at the host level (Xen, VMware, Microsoft Hyper-V). This layer of software, which is present between the hardware and the virtual servers, must be protected. Customers do not have access to this software layer in a public IaaS service; the CSP alone is responsible for its management.

Security of a virtual server or guest OS for customers

The virtualized version of an operating system, such as different Linux distributions, Microsoft, or Solaris, that is installed on top of the virtualization layer and made available to users through the Internet. Virtual servers are completely accessible to customers.

Security Software for Virtualization

Customers will not have access to or insight into the virtualization software that runs on top of the hardware since the CSP oversees it. Hardware or OS virtualization allows the safe concurrent use of many operating systems and programmes on a single machine by allowing the sharing of hardware resources across numerous guest VMs without interfering with one another. We assumed that IaaS services are using "bare metal hypervisor" technologies, sometimes referred to as type 1 hypervisors, such as VMware ESX, Xen, Oracle VM, and Microsoft's Hyper-V, for the sake of simplicity. These hypervisors support a wide range of guest operating systems, including Sun's OpenSolaris, Microsoft Windows, and numerous Linux "flavours".

It is crucial to safeguard the hypervisors from unauthorised users since hypervisor virtualization is a crucial component that ensures the separation and isolation of client VMs from one another in a multitenant environment. There is already a new arms race between hackers and defenders (CSP) in the area of virtualization security. The IaaS cloud design depends heavily on virtualization, therefore any assault that may jeopardise the integrity of the compartments would be disastrous for the whole client base on that cloud. A recent incident at the little UK-based business Vaserv.com serves as an example of the danger to hypervisor security. 100,000 websites hosted by Vaserv.com were destroyed by hackers who used a zero-day flaw in HyperVM, a virtualization programme created by Lxllabs. Due to the zero-day vulnerability, the attackers had access to critical Unix commands like `rm -rf`, which demands a complete deletion of all files. Seemingly, a few days before to the breach, an unidentified person uploaded a lengthy list of unpatched vulnerabilities in Kloxo, a hosting control panel that interfaces with HyperVM, on the hacker website milw0rm. Over 50% of Vaserv's clients choose unmanaged service, which excludes data backup, making their condition worse. It's yet unknown whether the proprietors of those websites will ever be able to recover their lost data.

CSPs should implement the required security measures, such as limiting physical and logical access to the hypervisor and other virtualization layers that are being used. IaaS users should be aware of the technological and procedural security measures put in place by the CSP to safeguard the hypervisor. With reference to your host security standard, rules, and regulatory compliances, this will assist you in understanding the compliance and gaps. You could be forced to take a leap of faith and trust CSPs to deliver an "isolated and protected virtualized guest OS" since CSPs often lack transparency in this area.

Hypervisor-related threats

The integrity and availability of the hypervisor are crucial to ensuring the integrity and availability of a public cloud based on a virtualized environment and are of the highest significance.

All user domains might be accessible to malevolent insiders with a weak hypervisor. Moreover, subversion attacks may be able to affect hypervisors. Several members of the security research community presented a "Blue Pill" attack on a hypervisor to show the vulnerability of the virtualization layer. Joanna Rutkowska, Alexander Tereshkin, and Rafal Wojtczuk from the Invisible Things Lab presented a variety of techniques to exploit Xen's virtualization during Black Hat 2008 and Black Hat DC 2009 even so Rutkowska and her group have pointed out issues with Xen implementations, but overall they seem to be quite supportive of the Xen strategy. Nonetheless, their example does highlight the difficulty of protecting virtualized systems and the want for fresh strategies to defend hypervisors from such assaults.

As most virtualization layers in public clouds are proprietary and closed source (although some may utilise a variant of open source virtualization software like Xen), security researchers are unable to examine the source code of the software that CSPs use.

CHAPTER 11

SECURITY OF VIRTUAL SERVERS

Dr. Ananta Charan Ojha
Professor, Department of Computer Science and Information Technology,
Jain (Deemed to be University) Bangalore, Karnataka, India
Email Id-oc.ananta@jainuniversity.ac.in

IaaS users have full access to the virtualized guest VMs that are hosted by hypervisor technology and kept separate from one another. Customers are thus in charge of the guest VM's security and continuous security monitoring. An open IaaS, like Amazon's Elastic Compute Cloud (EC2), provides a web services API for managing IaaS platform resources including provisioning, decommissioning, and replicating virtual servers. These system management operations, when organised effectively, may give flexibility for resources to increase or shrink in accordance with workload need. If the process to manage the virtual servers is not automated with suitable processes, the dynamic life cycle of virtual servers may lead to complexity. The virtual server (Windows, Solaris, or Linux) may be available to everyone on the Internet from an attack surface standpoint, hence adequate network access mitigation measures should be adopted to prevent access to virtual instances. The CSP typically disables all port access to virtual servers and advises clients to manage virtual server instances using Secure Shell (SSH) on port 22. The cloud administration API increases the attack surface and has to be taken into account when discussing how to secure virtual servers in a public cloud. New host security risks in the public IaaS include some of the following:

1. Stealing host-access and host-management keys (e.g., SSH private keys).
2. Targeting vulnerable, unpatched services that are listening on common ports (e.g., FTP, NetBIOS, SSH).
3. Attacking systems that are not effectively protected by host firewalls; hijacking accounts that are not properly secured (i.e., using weak or no passwords for standard accounts); deploying Trojans embedded in the software component in the VM or inside the VM image (the OS).

Virtual server security

In an IaaS platform, self-provisioning new virtual servers is so straightforward that there is a chance that vulnerable virtual servers may be generated. It is necessary to provide secure-by-default setup by adhering to or surpassing existing industry baselines.

Strong operational security processes and procedure automation are required for cloud virtual server security. Here are a few suggestions:

1. Choose a setup that is secure by default. While creating VMs (the guest OS) in a public cloud, harden your image and utilise a typical hardened image. Building bespoke VM images with just the capabilities and services required to run the application stack is a recommended practise for cloud-based apps. Reducing the capabilities of the underlying application stack decreases the number of patches required to maintain the security of that application stack, as well as the host's total attack surface.
2. Keep track of the OS and VM versions that have been readied for cloud hosting. Some of these VM images are provided by the IaaS provider. For hosts within the

company, a virtual image from the IaaS provider should go through the same degree of security testing and hardening. Offering your own image that complies with the same security criteria as internal trusted hosts is the best solution.

3. Prevent unwanted access to the hardened image's integrity.
4. Keep your private keys safe so you can access hosts in the public cloud.
5. In general, keep the decryption keys separate from the cloud where the data is stored, unless they are absolutely essential for decryption and even then, only temporarily. As the key will be colocated with the programme, it may not be able to safeguard it if your application needs a key to encrypt and decrypt data continuously.
6. Except for a key to decrypt the filesystem key, your virtualized images should not include any login credentials.
7. Disallow password-based shell access authentication.
8. Demand passwords for role-based access or sudo* (e.g., Solaris, SELinux).
9. Use a host firewall and only allow access to the fewest ports required to provide an instance's services.
10. Just activate the necessary services, and disable the unnecessary ones (e.g., turn off FTP, print services, network file services, and database services if they are not required).
11. Put in place a host-based IDS like OSSEC or Samhain.
12. Activate event and system auditing, and report the security events to a specific log server. Separate the log server and add access limits and better security protections.
13. Shut down the instance, take a snapshot of your block volumes, and back up the root filesystem if you think there has been a breach. Afterwards, you may do forensics on an unharmed system.
14. Provide a procedure for patching both offline and instantiated images in the cloud.
15. Regularly check logs for ominous activity.

Application-Level Security for Infrastructure

A crucial component of any security programme should include application or software security. The majority of businesses with information security programmes have yet to implement a programme to handle application security. Existing application security programmes will need to reassess current procedures and standards in order to design and develop apps that are intended for deployment on a cloud platform. From solitary, single-user programmes to complex, multiuser, e-commerce systems used by millions of users, the application security spectrum includes them all. Both small and big enterprises employ web applications including content management systems (CMSs), wikis, portals, bulletin boards, and discussion forums. Moreover, a sizable number of enterprises create and maintain bespoke online applications for their operations utilising a variety of web frameworks (PHP, .NET, J2EE, Ruby on Rails, Python, etc.). Before 2007, few criminals, according to SANS, targeted weak websites since other attack paths were more likely to result in an advantage in unlawful access to financial or informational resources. Cross-site scripting (XSS) and other assaults have shown, however, that criminals seeking financial gain may take advantage of vulnerabilities brought about by poor web development as new means to infiltrate significant businesses. We will restrict our discussion to web application security in this section: Users of common Internet browsers, such as Firefox, Internet Explorer, or Safari, may access web applications in the cloud from any computer with an Internet connection.

It is crucial for application security programmes to include browser security within the purview of application security given that the browser has established itself as the end user client for accessing in-cloud apps. Together, they assess the effectiveness of end-to-end cloud

security, which contributes to preserving the availability, confidentiality, and integrity of the data handled by cloud services.

Risks to Application-Level Security

Web application vulnerabilities in open source and bespoke programmes made for roughly half of all vulnerabilities found between November 2006 and October 2007, according to SANS.

The current threats take use of well-known application defects such as cross-site scripting (XSS), SQL injection, malicious file execution, and other problems brought on by coding faults and design flaws (e.g., the OWASP Top 10; see http://www.owasp.org/index.php/Top_10_2007). Hackers are continually searching the Internet for application vulnerabilities using their expertise and tools to scan online apps. Next, they are turning reputable websites into criminal servers hosting client-side exploits and phishing schemes, among other unlawful activities, by taking advantage of the vulnerabilities they have found. All web frameworks and all kinds of online applications are susceptible to web application security flaws, which might range from inadequate validation to logical faults in the programme.

Web applications deployed in tightly regulated environments, such as corporate intranets and private clouds, have traditionally been protected from outside hackers by a mix of perimeter security rules and network- and host-based access restrictions. Online applications created and deployed on a public cloud platform are vulnerable to attack, hacking, and possible exploitation by criminals for fraudulent and unlawful purposes. Security must be integrated into the Software Development Life Cycle (SDLC) and web applications deployed in a public cloud (the SPI model) must be built for an Internet threat model in that threat model; see Figure 11.1.




Figure 11.1. Represent the sdlc.

The DoS and EDoS

You should also be aware of application-level DoS and DDoS assaults, which have the ability to severely impair cloud services. These assaults often come from infected computer systems connected to the Internet (hackers frequently take over and operate PCs infected with viruses,

worms, and malware, as well as sometimes very effective unsecured servers). High-volume reloads of web pages, XML web services queries sent over HTTP or HTTPS, or protocol-specific requests sent via a cloud service are all examples of application-level DoS attacks. It is quite challenging to selectively filter the malicious traffic without affecting the service as a whole since these malicious requests mix in with the normal traffic. For instance, a DDoS assault on Twitter on August 6, 2009, resulted in the service being unavailable for a number of hours.

DoS attacks may rapidly deplete your company's money for cloud services in addition to interrupting cloud services, which has a negative effect on user experience and service-level implications. Your cloud utility bill will dramatically rise as a consequence of DoS assaults on pay-as-you-go cloud services since more network bandwidth, CPU, and storage will be used. Another name for this kind of assault is "economic denial of sustainability" (EDoS).

Hackers now have an equal playing field because to the minimal obstacles that small and medium-sized businesses must overcome to embrace cloud computing for lawful purposes. Hackers will be able to join together computer resources to accomplish large quantities of computation without paying for any of the capital infrastructure expenditures by using cloud accounts that have been taken over or compromised. You could see DoS assaults conducted from IaaS or PaaS clouds against other cloud services in the not-too-distant future (such hostile and offensive cloud models are being characterised as dark clouds).

User Protection

As a user of a cloud service, you are in charge of end user security chores, such as security precautions to safeguard your Internet-connected Computer, as well as "safe browsing." Using security software on your Internet-connected computer, such as anti-malware, antivirus, personal firewalls, security updates, and IPS-type software, is one way to take precautions. The new catchphrase "the browser is your operating system" effectively communicates the idea that browsers have replaced operating systems as the de facto standard method of accessing cloud services. All Internet browsers often have software flaws that leave them open to end user security intrusions. Thus, we advise cloud customers to take the necessary precautions to shield browsers against threats. Customers must practise proper browser hygiene in order to ensure end-to-end security in a cloud. To reduce risks associated to browser vulnerabilities, the maintaining patched and updated browsers (such as Internet Explorer, Firefox, and Safari). Users are advised to often check their browser vendor's website for security updates, utilise the auto-update option, and install patches on a timely manner to preserve end user security even if browser security add-ons are currently not commercially accessible.

Charge of Cloud Web Application Security:

The extent of security obligations will be shared by the client and the cloud provider, depending on the service-level agreement (SLA) and delivery model for cloud services (SPI). The key is to be aware of how your security obligations differ from those of the CSP. Several security studies have underlined the fact that a barrier to cloud adoption is the lack of openness in security procedures and controls used by CSPs in this context. To begin with, cloud users lack the openness needed to understand software flaws in cloud services. Customers are unable to control the operational risk that can be associated with the vulnerabilities as a result. Additionally, CSPs are preventing security researchers from checking their software for bugs and security issues by designating it as proprietary. (Cloud service companies using open source software are an exception.) Customers have little option but to rely on their CSPs to report any new vulnerabilities that may compromise the

confidentiality, integrity, or availability of their application as a result of this lack of transparency. For instance, no well-known IaaS, PaaS, or SaaS companies are involved in the Common Vulnerability and Exposures (CVE) initiative as of March 2009. As an example, Amazon fixed a vulnerability that Colin Percival reported in May 2007 after waiting 7.5 months. This vulnerability, which impacted Amazon's database API (SimpleDB) and EC2 API services, was a cryptographic flaw in the request signing code that was not made public until after it had been addressed in December 2008. Colin admits that Amazon treated this problem seriously at all times; the extended schedule was merely a result of the significant amount of effort required to send out a fix to the impacted services. Enterprise clients should be aware of cloud service providers' vulnerability disclosure policies and take them into consideration when evaluating their risk exposure. The web application security is discussed in the parts that follow in relation to the SPI cloud service delivery paradigm.

Security for SaaS applications

According to the SaaS business model, the supplier must oversee the whole set of apps sent to consumers. As a result, SaaS providers bear the bulk of the responsibility for protecting the customer-facing software and components they deliver. Operational security services, such as user and access management enabled by the provider, are often the responsibility of the customer. Prospective clients often ask for details about the provider's security procedures, generally under an NDA. Design, architecture, development, testing for both white-box and black-box application security, and release management should all be included in this data.

In order to independently get confidence, some clients would even go as far as to hire outside security firms to do penetration testing (sometimes known as "black-box security testing") of SaaS apps. Penetration testing, however, may be expensive, and not all suppliers consent to this kind of verification.

The authentication and access control mechanisms provided by SaaS CSPs need special consideration. Often, it is the sole security measure available to control information risk. The majority of services, including those from Google and Salesforce.com, include a web-based administrative user interface tool to handle application authentication and access management. Certain SaaS programmes, like Google Apps, include built-in tools that customers may utilise to provide other users read and write access. The privilege management tools, however, may not be sophisticated or have fine-grained access, and they could have flaws that make them incompatible with the access control requirements of your company. The way Google Docs handles photos that are embedded in documents and access rights to earlier versions of a document is one illustration that illustrates this problem. Clearly, the sharing settings that safeguard a document do not protect embedded photos that are kept in Google Docs. This means that even when you stop sharing the document, the recipient will always be able to see any embedded photos you have shared. This access control anomaly was found by a blogger, who alerted Google about it. Google has recognised the problem, but in its reaction, it makes it clear that it doesn't think the issues constitute a serious security danger to its consumers.

Another issue using Google Documents had a privacy flaw* that improperly shared access to a tiny portion of word processing and presentation files kept in the Google Apps cloud service (Google asserts that just 0.05% of the files were impacted). Because the papers were shared only with those whom the Google Docs users had previously shared documents, rather than with the world at large, the situation underscores the necessity to review and understand cloud-specific access control techniques.

Customers of the cloud should make an effort to comprehend the access control methods that are unique to the cloud—including support for strong authentication and privilege management based on user roles and functions—and take the appropriate precautions to safeguard data stored there. To safeguard the application from insider threats, additional controls should be built to govern privileged access to the SaaS administration tool and ensure division of tasks. Customers should develop a strong password policy that requires users to choose strong passwords when authenticating to an application in accordance with security best practises. †

SaaS companies often combine their customers' structured and unstructured data in a single virtual data store and depend on data tagging to guarantee data separation. Data is tagged and stored with a distinct customer identification in that multitenant data store paradigm, where encryption may not be possible owing to key management and other design constraints. As the data is processed, the application layer's business logic may ensure customer separation thanks to this distinctive data tag. It is possible that during software updates by the CSP, the application layer maintaining this isolation might become weak. Customers should be aware of the virtual data store architecture and the safeguards SaaS providers use to provide the isolation and compartmentalization necessary in a virtual multitenant environment.

Salesforce.com, Microsoft, and Google are well-known SaaS companies that have a reputation for investing in software security and incorporating security assurance into their SDLC. Yet, as there is no industry standard for evaluating software security, it is very difficult to compare service providers to a standard.

Application Security in PaaS

Vendors of PaaS may be largely divided into the following two categories:

PaaS software may be used by businesses to provide a solution for internal consumption while they evaluate private clouds. While Eucalyptus does provide a small experimental pilot cloud for developers at Eucalyptus.com, it is not yet known that any significant public clouds are employing commercial off-the-shelf or open source PaaS technologies. Nonetheless, it is advised that businesses considering PaaS software do a risk analysis and use the software security standard just as they would if they were buying any other piece of business software.

A PaaS cloud (public or private) is a platform as a service that provides an integrated environment for the design, development, testing, deployment, and maintenance of unique applications written in the languages that the platform supports. Security for PaaS applications is divided into two software layers:

Security of client applications installed on a PaaS platform, including security of the runtime engine of the PaaS platform.

Broadly speaking, the platform software stack, which includes the runtime engine that powers client applications, must be secured by PaaS CSPs (such as Google, Microsoft, and Force.com). The third-party application provider may be in charge of protecting their services as PaaS applications may employ third-party apps, components, or web services. Customers should thus be aware of how dependent their applications are on these services and evaluate the risks associated with using third-party service providers. CSPs have up to now been hesitant to provide platform security-related information, citing the possibility that doing so may give hackers an edge. Enterprise clients should, however, demand openness from CSPs and look for the data they need to do risk analysis and continuous security management.

Application container for PaaS

The confinement and isolation of multitenant apps from one another are the fundamental security principles of the multitenant PaaS service delivery paradigm. According to that paradigm, only your corporate users and the apps you control and administer should have access to your data. The CSP owns the intellectual property for the PaaS platform runtime engine's security model, which is necessary to supply the "sandbox" architecture in a multitenant computing paradigm. As a result, the platform runtime engine's sandbox feature is crucial to protecting the privacy and integrity of your application when it is deployed via PaaS. CSPs are in charge of keeping an eye out for new flaws and problems that might be utilised to attack the PaaS platform and circumvent the sandbox design. The privacy implications for customer-sensitive information are undesired and might have a significant negative impact on your organisation in this sort of circumstance, which is the worst case scenario for a PaaS service. In order to learn more about the confinement and isolation architecture of the PaaS service, business users should contact the CSP.

The PaaS cloud provider is also in charge of maintaining network and host security outside of the PaaS platform (i.e., monitoring of a shared network and system infrastructure hosting customer applications). Customers of PaaS providers should be aware of how PaaS CSPs maintain their platform, including runtime engine updates and change, release, and patch management.

Security for Applications Deployed by Customers

For the deployment and management of software modules that impose security restrictions, PaaS developers must get acquainted with a set of particular APIs. Developers must also get acquainted with platform-specific security features since the API is special to PaaS cloud services. These capabilities are accessible to them in the form of security objects and web services for defining authentication and authorisation controls inside the application. As there is presently no standard for PaaS API design and no organised effort by CSPs to provide a universal and uniform API across clouds, porting an application between PaaS clouds is a laborious process. Today, only Python and Java are supported by the Google App Engine, while Apex is the sole language supported by Salesforce.com's Force.com. Apex is different from programming languages like C++, Java, and .NET. Apex, in contrast to other languages, is significantly more constrained in scope and is tailored to the development of commercial applications on the Force.com platform. Cloud services have a greater chance than conventional software licencing of retaining clients in this aspect. The absence of an API standard has implications for cloud application portability and security management.

Developers should anticipate that CSPs will provide a set of security capabilities, including SSL or TLS support, user authentication, single sign-on (SSO) via federation, authorization (permission management), and SSO. There isn't yet a standard for PaaS security management since every CSP has a different security model and every provider has different security features. While utilising the Google App Engine, a developer may customise the user profile and choose HTTPS as the transport protocol by using Python or Java objects. Similar to this, Force.com provides an Apex API that can be used to allocate specific TCP ports for application-to-application connection-type interactions, change different runtime options, and define security parameters.

According to our analysis of the main PaaS CSPs, the security capabilities accessible to PaaS apps are restricted to the bare minimum: basic privilege management, SSL setup, and user authentication using the provider's identity database. User federation is very seldom supported by the Security Assertion Markup Language (SAML).

Application Security for IaaS

IaaS cloud providers (like Amazon EC2, GoGrid, and Joyent) perceive the programmes running on client virtual instances as a "black box," and as a result, have no control over how those apps are run or managed. The complete stack, including client applications and the runtime application platform (Java,.NET, PHP, Ruby on Rails, etc.), is installed and maintained by customers and runs on their virtual servers. Thus, customers are solely responsible for the security of any applications they install to the IaaS cloud. As a result, customers shouldn't count on CSPs to provide any support with application security beyond some basic advice and firewall policy features that may impact an application's contacts with other apps, users, or services either within or outside the cloud.

Web applications that are delivered in a public cloud must be built with standard security countermeasures against typical online vulnerabilities and for an Internet threat model (e.g., the OWASP Top 10). They should also be routinely evaluated for vulnerabilities in accordance with best practises for security development, but most significantly, security should be integrated into the SDLC. Customers are completely responsible for ensuring that their runtime platform and apps are patched to safeguard the system from malware and hackers searching for weaknesses to access customer data in the cloud. It is strongly advised that you create and put into practise apps utilising a "least-privileged" runtime paradigm (e.g., configure the application to run using a lower privileged account).

IaaS cloud application developers are need to provide their own functionality for handling authentication and authorisation. Cloud apps should be created to take use of delegated authentication service capabilities enabled by an enterprise identity provider (such as OpenSSO, Oracle IAM, IBM, CA) or third-party identity service provider, in accordance with corporate identity management policies (e.g., Ping Identity, Symplified, TriCipher). Custom Authentication, Authorization, and Accounting (AAA) feature implementations should be avoided wherever feasible since they have the potential to become a weak point if not handled correctly.

In conclusion, business online applications with an n-tier distributed design are most similar to the architecture for IaaS hosted apps. Distributed applications are operated in an organisation with a variety of security measures in place to protect the host and the network that connects the dispersed hosts. Similar controls must be implemented via a network, user access, or as application-level controls to an IaaS platform since they are not present by default. IaaS cloud users are in charge of all facets of their application security and must take the required precautions in a multitenant, hostile Internet environment to safeguard their applications.

Clients who are considering using the public cloud should be aware of its limits in terms of supporting unique security features. A public SaaS, PaaS, or IaaS cloud does not meet security needs such an application firewall, SSL accelerator, cryptography, or rights management utilising a device that supports PKCS 12. Depending on client demand, IaaS and PaaS providers may eventually provide some of these more advanced security capabilities. In general, mitigating measures that ask for the installation of a device or locally connected peripherals in a public IaaS/PaaS cloud are not yet practical.

CHAPTER 12

DATA SECURITY AND STORAGE

Dr. A Rengarajan

Professor, Department of Computer Science and Information Technology,
Jain (Deemed to be University) Bangalore, Karnataka, India
Email Id-a.rengarajan@jainuniversity.ac.in

Data security becomes increasingly crucial when employing cloud computing at all "levels": infrastructure-as-a-service (IaaS), platform-as-a-service (PaaS), and software-as-a-service in today's world of network, host, and application-level infrastructure security (SaaS). Many facets of data security are including:

Data processing, includes multitenancy, data-in-transit, data-at-rest, data lineage, data provenance, and data remanence.

Users in assessing their data security situations and determining the risk to their companies. Not all of these factors of data protection are equally crucial in all topologies, much as other aspects of cloud computing and security, e.g., the use of a public cloud versus a private cloud, or non-sensitive data versus sensitive data.

Many Data Security Issues

The main danger when it comes to data-in-transit is not employing a tested encryption technique. It is usual for people not to comprehend this necessity while utilising a public cloud, regardless of whether it is IaaS, PaaS, or SaaS, despite the fact that this is evident to information security specialists. In particular if the protocol is used for data transmission over the Internet, it is crucial to make sure that it offers secrecy in addition to integrity (e.g., FTP over SSL [FTPS], Hypertext Transfer Protocol Secure [HTTPS], and Secure Copy Program [SCP]). A non-secured protocol (such as "vanilla" or "straight" FTP or HTTP) and data encryption alone may offer secrecy, but they cannot guarantee the integrity of the data (e.g., with the use of symmetric streaming ciphers).

While it may seem apparent to use encryption to safeguard data while it is in transit, this is not always the case. Encrypting data-at-rest is feasible and highly advised if you use an IaaS cloud service (public or private) for basic storage, such as Amazon's Simple Storage Service or S3. Nevertheless, it is not always possible to encrypt data-at-rest that is used by a PaaS or SaaS cloud-based application (for example, Google Apps, Salesforce.com). A cloud-based application typically does not encrypt data that is at rest since doing so would make it impossible to index or search the data.

The economics of cloud computing often dictate that SaaS and PaaS-based applications adopt a multitenancy design when dealing with data that is at rest. In other words, data gets mixed with the data of other users when it is processed by a cloud-based application or saved for use by a cloud-based application (i.e., it is often kept in a huge data store, like Google's BigTable). Unauthorized access to commingled data is frequently prevented by features like data tagging in applications, although it is still feasible if an application vulnerability is exploited. Data is not on a platform that is only used by one entity, even if some cloud

providers have their apps assessed by other parties or validated using outside application security technologies.

Although an organization's data-in-transit may be encrypted while being transferred to and from a cloud provider and its data-at-rest may be encrypted if using simple storage (that is, if it is not connected to a specification application), an organization's data is undoubtedly not encrypted if it is processed in the cloud (public or private). Each programme must be able to handle unencrypted data. There was no known way to completely process encrypted data up to June 2009. Hence, the data will be unencrypted for at least a portion of its life cycle in the cloud—processing at the very least—unless it is just being stored there.

In June 2009, IBM said that one of its researchers has created a completely homomorphic encryption system that enables data to be processed without being decoded in collaboration with a doctoral student from Stanford University. This is a tremendous advancement in cryptography, and as cloud computing enters the implementation phase, it will significantly benefit from it. At Stanford University, earlier work on completely homomorphic encryption (such as 2-DNF) was also carried out, but IBM's revelation surpassed even that promising work. While the homomorphic approach has theoretically overcome the obstacle of completely homomorphic encryption, it took a significant amount of computing work. Ronald Rivest, an MIT professor and the creator of the well-known RSA encryption technique, claims that the procedures necessary to make it work practically won't take long. Predicate encryption is one of the other forms of cryptographic research that is being done to reduce the quantity of data that would need to be decrypted for processing in the cloud.

Regardless of whether an organisation has encrypted the data it has placed in the cloud, knowing the precise location and time the data was precisely stored there is helpful and may be necessary (for audit or compliance reasons). For instance, the data could have been transferred to a cloud service provider, like Amazon Web Services (AWS), on date x_1 at time y_1 and stored in a bucket on Amazon's S3 in `example1.s3.amazonaws.com`, then processed on date x_2 at time y_2 on an instance being used by an organisation on Amazon's Elastic Compute Cloud (EC2) in `ec2-67-202-51-223.compute-1.amazonaws.com`. Data lineage, sometimes referred to as mapping application data flows or data route visualisation, is crucial for an auditor's assurance (internal, external, and regulatory). Even when a company has perfect control over the environment, giving data lineage to auditors or management takes time. It is essentially impossible to attempt to give reliable reporting on data lineage for a public cloud service. What physical system, and precisely where was it situated, is that bucket on `example1.s3.amazonaws.com` from the previous example? What was the physical system's condition at that time, and how might a client or auditor confirm that information?

Even if data lineage can be verified in a public cloud, some clients have an even more difficult need and issue: demonstrating data provenance not simply demonstrating the data's integrity, but also the data's more precise provenance. There is an essential distinction between the two words. Data that has not been altered inadvertently or by an unauthorised party is said to have integrity. According to provenance, the data not only has integrity but is also computationally accurate, meaning that it was computed correctly. The anticipated result of such equation is \$2.00. If the response were different, there would be a concern with integrity. The \$2.00 is obviously assumed to be in U.S. dollars, however the following linked assumptions might result in the assumption being incorrect:

1. The formula is only valid for the Australian, Bahamian, Barbadian, Belizean, Bermudan, Brunei, Canadian, Cayman Islands, Cook Islands, East Caribbean, Fijian, Guyanese, Hong Kong, Jamaican, Kiribati, Liberian, Namibian, New Zealand,

- Samoan, Singaporean, Solomon Islands, Surinamese, New Taiwan, Trinidad and Tobago, Tuvaluan, or Zimbabwean dollars.
2. The dollar is designed to be changed from money from other nations into money from the United States.
 3. The conversion is computed accurately and can be verified, and the appropriate exchange rate is utilised.

In this case, the equation has integrity but lacks provenance if it meets those presumptions. There are several cases from the actual world when data integrity alone is inadequate and additionally calls for data provenance. Calculations used in science and finance are two apparent examples. When employing shared resources in a cloud computing environment, how can data provenance be demonstrated? Even if you are aware of certain identifying information about the systems (such as their IP addresses) and the "generic" location, you are not under physical or even logical control of those resources, and you most likely have no way to trace the systems utilised or their condition at the times you used them (e.g., a country, and not even a specific data center).

Data persistence is the last component of data security. The remaining representation of data that has been officially wiped or eliminated is known as data remanence. This residue might be caused by data that was not completely removed during a normal delete operation or by the physical characteristics of the storage media. If the storage media are released into an uncontrolled environment (such as being tossed in the trash or handed to a third party), data remanence may make accidental exposure of sensitive information feasible. Regardless of the cloud service you choose, data remanence poses the danger that an organization's data may be accidentally exposed to an unwanted party (SaaS, PaaS, or IaaS). The danger of adopting SaaS or PaaS is almost likely accidental or unintended exposure. After an unauthorised disclosure, however, it is not comforting, and prospective clients should inquire about what third-party methods or evaluations are used to aid verify the security of the provider's apps or platform.

While data security has become more crucial, cloud service providers (CSPs) pay startlingly little attention to data persistence. Many services don't even bring up data persistence. Yet if the question of data security is brought up, many CSPs rather glibly mention compliance with DoD 5220.22-M of the United States Department of Defense (the National Industrial Security Program Operating Manual). The reason we used the word "glibly" is because it seems that the providers (and other sellers of information technology) have not truly read this document. DoD 5220.22-M lists the two acceptable data (destruction) security techniques, however it doesn't specify how these two ways are to be carried out in detail or provide any guidelines on how to do so. Three paragraphs of DoD 5220.22-M, a 141-page document, are all that are necessary for understanding data remanence:

Clearing "Clearing is the process of removing the data from the media before reusing the media in a setting that offers an appropriate degree of security for the data that was on the media prior to clearing. To properly prevent access to previously saved information, all internal memory, buffers, and other reusable memory must be deleted. Sanitization is b.

Sanitization is the process of removing the data from the media before reusing it in an environment that does not provide an adequate degree of security for the data that was on the media prior to sanitization. Before to being released from controlled access to sensitive material or released for usage at a lower classification level, IS resources must be sanitised.

The National Institute of Standards and Technology (NIST) Special Publication, 800-88, "Guidelines for Media Sanitization," should be consulted by providers for detailed information on how data security should be accomplished.

Many businesses, particularly those in regulated sectors, voluntarily conform to NIST rules and standards, despite the fact that this NIST book merely offers recommendations and is intended only for government civilian departments and agencies. It's crucial to follow these NIST recommendations as there isn't a single industry standard for data persistence.

Mitigation of Data Security

Since a portion of a customer's infrastructure security moves outside of its control and a provider's infrastructure security may (for many enterprises) or may not (for small to medium-sized businesses, or SMBs) be less robust than expectations, you will be dissatisfied if prospective customers of cloud computing services expect that data security will serve as compensating controls for potentially weaker infrastructure security. Data in transit may and should be encrypted, but any use of that data in the cloud that goes beyond basic archiving necessitates its decryption. As a result, it is nearly a given that data will not be encrypted on the cloud. Customer-unencrypted data will also most likely be housed in a multitenancy environment if you are employing a PaaS-based application or SaaS. The risks of data security for customers are greatly increased when you consider the exposure of the data, the challenges in determining the data's lineage and provenance (where necessary), and even the widespread failure of many providers to adequately address such a fundamental security concern as data remanence.

What therefore should you do to lessen these threats to data security? Making ensuring that no regulated or sensitive data is stored in a public cloud is the only practical method of mitigation (or that you encrypt data placed into the cloud for simple storage only). CSPs are not providing strong enough controls around data security given the current economics of cloud computing and the current limitations of encryption. It's possible that the economics shift and providers continue to provide their present services alongside a "regulatory cloud environment" (i.e., an environment where customers are willing to pay more for enhanced security controls to properly handle sensitive and regulated data). Making ensuring that no regulated or sensitive data is stored in a public cloud at this time is the only practical mitigating strategy.

Security of Provider Data

Customers should be worried about the provider's collection of data and the CSP's methods for protecting it in addition to the security of their own customer data. In particular, what information does the provider have about your client data, how is it protected, and how much access do you, the customer, have to that metadata? The value of the metadata rises as your data volume with a certain supplier climbs.

Also, a significant quantity of security-related data is collected by and kept secure by your provider. For instance, your supplier should gather, monitor, and safeguard data from your router flow, firewall, intrusion prevention system (IPS), security incident and event management (SIEM), and other network-level systems. At the host level, your provider should gather system log files, and SaaS providers should gather application log data, which includes login and authorization details.

For the purposes of its own audit, the provider cares about the data your CSP gathers and how it monitors and secures that data. The possibility that it may be required for incident response and any digital forensics necessary for incident analysis makes this information relevant to both providers and customers.

Storage

We are referring to data stored in the cloud (i.e., storage-as-a-service) and not data connected to an application operating in the cloud on PaaS or SaaS. These data saved in the cloud (like Amazon's S3) are subject to the same three information security issues as data kept elsewhere: confidentiality, integrity, and availability.

Confidentiality

There are two things you could worry about in terms of the privacy of data kept in a public cloud. To start, what access controls are in place to safeguard the data? Both authentication and authorisation are part of access control. CSPs often utilise weak authentication techniques (such as username + password), and the authorization ("access") controls that users have access to are frequently fairly broad and imprecise. This broad authorisation poses serious security issues for big enterprises.

Cloud suppliers often just offer administrator authorisation (i.e., the account owner) and user authorization (i.e., all other authorised users), with no tiers in between (e.g., business unit administrators, who are authorised to approve access for their own business unit personnel). The second possible issue, however—how is the data that is actually kept in the cloud protected—is unquestionably pertinent to this part. In all actuality, using encryption is required to secure data kept on the cloud.

So, when data is kept on the cloud, is it genuinely encrypted? If so, what encryption technique and key strength are used? It varies, and in particular, it depends on which CSP you're using. Data of a client is encrypted, for instance, via EMC's MozyEnterprise. Amazon S3 does not, however, encrypt client data. S3 does not provide encryption, although customers may do it themselves before uploading their data. The second factor to think about is what encryption algorithm a CSP will employ if it encrypts a customer's data. There are variations among encryption algorithms. Several algorithms used in cryptography don't provide enough security.

The cryptography community should only utilise algorithms that have been openly reviewed by a formal standards organisation (like NIST). It is essential to avoid using any proprietary algorithms. Please take note that we are discussing symmetric encryption techniques here. A single secret key is used in symmetric encryption (see Figure 12-1) to encrypt and decode data simultaneously.

Because symmetric encryption has the processing power and speed to handle data encryption for huge amounts of data. An asymmetric method wouldn't be used in this encryption use case very often. The notion (i.e., a single shared, secret key) is employed in data storage encryption even though the example in Figure 12.1 relates to email.

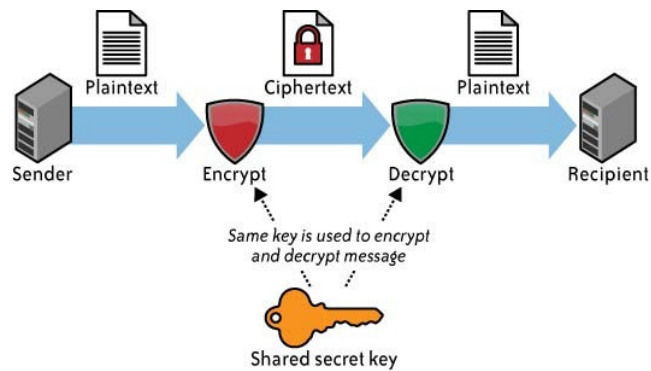


Figure 12-1: Symmetric encryption

A public key and a private key are not used in data storage encryption, despite the fact that the example in Figure 12.2 relates to email.

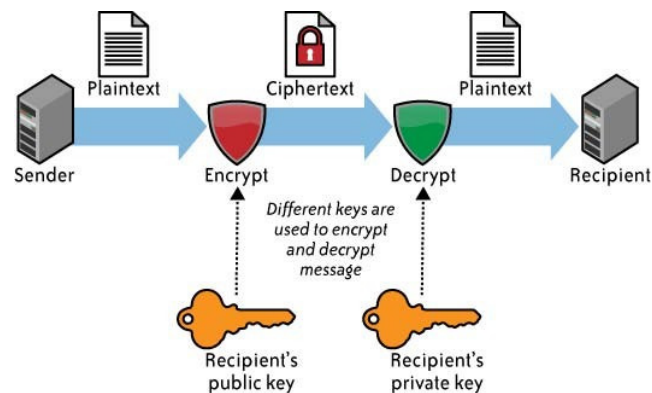


Figure 12.2. Asymmetric encryption

What key length is employed is the next thing you should think about. In symmetric encryption, the better the encryption, the longer the key length is (i.e., the more bits in the key). Long key lengths are more secure, but they are also more computationally expensive demanding, and might place a burden on computer CPUs. What can be mentioned is that for Triple DES (Data Encryption Standard) and AES (Advanced Encryption Standard), both NIST-approved algorithms, key lengths should be at least 112 bits and 128 bits, respectively.

Key management is another aspect of encryption that affects secrecy. Who will control how and how the utilised encryption keys are managed? Will you control your own keys? Hopefully, you have the knowledge to handle your own keys if the answer is yes. It is not advised that you let a cloud provider—at least not the same service that manages your data—to handle your keys. This indicates that more tools and resources are required. Nonetheless, maintaining good key management is a challenging and complicated undertaking. A client should at the very least review all three sections of NIST's 800-57, "Recommendation for Key Management":

Key management is complicated and challenging for a single client, and it is significantly more complicated and challenging for CSPs to attempt to manage the keys for several customers. Because of this, some CSPs manage clients' keys ineffectively. For instance, it is

typical for a provider to use a single key to encrypt all of a customer's data. Worst still, we are aware of one cloud Storage Company that employs a solitary encryption key for each and every one of its clients! Such problems are being attempted to be addressed by the Key Management Interoperability Protocol (KMIP) of the Organization for the Advancement of Structured Information Standards (OASIS).

Integrity

You must be concerned not just with the integrity of your data but also with its secrecy. Data may be encrypted for secrecy reasons, but you may not be able to check the integrity of such data since confidentiality does not guarantee integrity. For secrecy, encryption alone is adequate, but message authentication codes are also necessary for integrity (MACs). The easiest method of using MACs on encrypted data is to employ a block symmetric algorithm in cypher block chaining (CBC) mode, together with a one-way hash function, as opposed to a streaming symmetric algorithm. This is not for those who are not familiar with cryptography, and it is one of the challenges of efficient key management. Cloud users should at the very least inquire about these issues from providers. This is crucial for maintaining the integrity of a customer's data as well as revealing the level of sophistication—or lack thereof—of a provider's security operation. Nevertheless, keep in mind that not all service providers encrypt client data, particularly for PaaS and SaaS services.

Another issue of data integrity is crucial, particularly when utilising IaaS for bulk storage. How can a client verify the accuracy of the data saved in the cloud if they have several gigabytes (or more) of it there for storage? IaaS transfer fees and network usage (bandwidth) issues for the customer's own network are linked with transferring data into and out of the cloud*. Instead of downloading and reuploading the data, what a customer actually wants to do is verify the accuracy of the data while it is still in the cloud.

Since this activity must be completed on the cloud without explicit access to the whole data set, it is significantly more challenging. Consumers often don't know where such systems are situated or which physical computers are used to store their data. Moreover, that data set presumably changes regularly and is dynamic. The usefulness of conventional integrity insurance measures is negated by these frequent changes. Instead, a proof of retrievability—a mathematical method to confirm the accuracy of the data as it is dynamically stored in the cloud—is required.

Availability

You should be worried about your data's accessibility even if a customer's info has remained secret and accurate. There are now three main concerns in this area; none of them are novel to computing, but due to increasing risk, they all become more significant in cloud computing.

There have been some prominent cloud provider disruptions. For instance, Amazon's S3 had outages of 2.5 hours in February 2008 and 8 hours in July 2008. Imagine the challenges that other, smaller, or less experienced cloud companies are facing as Amazon is one of the more established cloud service providers. Since S3 serves a sizable number of users, many of whom are heavily (if not entirely) dependent on S3's availability for their own operations, these Amazon disruptions were all the more noticeable.

In certain instances, data stored in the cloud has actually been lost in addition to service interruptions. For instance, "cloud-based storage service provider Carbonite Inc. filed a lawsuit in March 2009 alleging that defective gear from two hardware vendors caused backup failures that led the firm to lose data for 7,500 clients two years earlier."

Whether cloud storage providers will continue to operate is a more important subject for cloud clients to think about. Coghead, a supplier of cloud services, abruptly went out of business in February 2009, leaving its clients less than 90 days (nine weeks) to get their data off its servers or risk losing it forever.

Lastly, potential users of cloud storage should be sure to find out exactly what services their supplier is truly providing. While the data is saved on the cloud, it is not always backed up. Several cloud storage companies provide storage in addition to backing up client data. Yet, a lot of cloud storage companies simply back up client data on occasion or solely as an extra feature that costs extra. Data saved in services like Amazon S3, Amazon SimpleDB, or Amazon Elastic Block Storage, for instance, "are redundantly kept in many physical locations as a standard feature of such services and at no extra price." AWS does not make backups because "data that is retained inside operating instances on Amazon EC2, or within Amazon S3 and Amazon SimpleDB, is all client data." § This is a seemingly simple yet important inquiry that users should pose to cloud storage providers about availability.

The service-level agreement (SLA) that a CSP offers to its clients should cover all three of these factors: confidentiality, integrity, and availability. The CSP SLAs, on the other hand, are now quite weak and, in reality, are basically useless. Even in cases when a CSP seems to have at least a minimally adequate SLA, it is challenging to really quantify such SLA. Due to all of these factors, users should pay close attention to data security issues and how data is truly kept in the cloud.

CHAPTER 13

IDENTITY AND ACCESS MANAGEMENT

Dr. Pawan Kumar

Assistant Professor, Department of Computer Science and Information Technology,
Jain (Deemed to be University) Bangalore, Karnataka, India
Email Id-pawan.bhuphd@gmail.com

The "trust border" is mostly static in a normal business where applications are installed within the organization's perimeter and is tracked and managed by the IT department. Under that conventional paradigm, the network, systems, and applications housed in a private data centre overseen by the IT department are included inside the trust boundary (sometimes third-party providers under IT supervision). Moreover, network security measures such as virtual private networks (VPNs), intrusion detection systems (IDSs), intrusion prevention systems (IPSs), and multifactor authentication are used to safeguard access to the network, systems, and applications.

The organization's trust boundary will become dynamic and leave IT control with the adoption of cloud services. With cloud computing, an organization's network, system, and application boundaries will extend into the domain of the service provider. (This could already be the situation for the majority of big businesses involved in e-commerce, supply chain management, outsourcing, and community and partner engagement.) If not handled appropriately, this loss of control will obstruct the adoption of cloud services inside an organisation by challenging the existing trusted governance and control paradigm (including the trusted source of information for workers and contractors).

Organizations will be compelled to depend on additional higher-level software controls, such as application security and user access controls, to make up for the loss of network management and to increase risk assurance. Strong authentication, authorisation based on role or claims, reliable sources, single sign-on (SSO), user activity monitoring, and auditing are examples of these measures. Organizations should pay close attention to the identity federation architecture and procedures in particular because they may improve governance and build confidence with cloud service providers (CSPs).

For managing the varied, dynamic, loosely tied trust connections that define an organization's internal and external supply chains and cooperation model, one emerging industry best practise is identity federation. A salesperson engaging with Salesforce.com from a corporate network, for example, is an example of how federation facilitates the interaction of systems and applications separated by an organization's trust boundary. Federation will be crucial in speeding the adoption of cloud computing inside enterprises because it may offer robust authentication via delegation, online single sign-on, and entitlement management through centralised access control services.

In certain instances, a lack of central governance and identity information architecture may negatively affect IAM practises inside an organisation. Many administrators often maintain identity storage manually, and user provisioning procedures are not effectively coordinated. This technique will spread current poor practises to cloud services in addition to being inefficient. In these circumstances, the weak access model will provide unauthorised users extra rights to cloud services.

IAM works both ways: For clients to benefit from and expand their practise to maintain compliance with internal rules and standards, CSPs need to enable IAM standards (like SAML) and practises like federation. Traditional IT applications' transition from trusted corporate networks to a trusted cloud service model will be sped up by cloud services that integrate IAM capabilities like federation. For clients, properly established user IAM policies and procedures will aid in safeguarding the integrity and security of information stored in the cloud as well as managing compliance. The adoption of new cloud services and the movement of IT applications from trusted corporate networks into a trusted cloud service model may both be accelerated by cloud services that implement IAM standards like SAML.

IAM: IAM practises have historically been an investment for businesses to increase operational effectiveness and to meet regulatory, privacy, and data protection requirements:

Boost operational effectiveness

By automating user onboarding and other repetitive operations, properly designed IAM technology and processes may increase productivity (e.g., self-service for users requesting password resets that otherwise will require the intervention of system administrators using a help desk ticketing system).

Management of regulatory compliance

Organizations implement "IT general and application-level controls" derived from industry standard frameworks like ISO 27002 and Information Technology Infrastructure Library to safeguard systems, applications, and data from internal and external threats (such as disgruntled employees deleting sensitive files) and to adhere to various regulatory, privacy, and data protection requirements (such as HIPAA, SOX) (ITIL). Organizations may achieve access control and operational security goals by enforcing compliance obligations like "segregation of responsibilities" and assigning staff members with just the credentials they need to carry out their jobs, among other IAM procedures and practises. In order to assist the management of regulatory compliance procedures, such as Payment Card Industry (PCI) Data Security Standards (DSSs) and the Sarbanes-Oxley Act of 2003, auditors often map internal controls to IT controls (SOX).

IAM can allow new IT delivery and deployment methods in addition to increasing operational effectiveness and compliance management (i.e., cloud services). For instance, federated identity, a crucial IAM component, makes it possible to connect and transmit identity data across trust boundaries. As a result, it makes it possible for businesses and cloud service providers to connect different security domains via federated user provisioning and online single sign-on.

The following are a few examples of cloud use cases that call for IAM support from the CSP:

1. IT administrators accessing the CSP management console to provision resources and access for users using a corporate identity (e.g., IT administrators of Newco.com provisioning virtual machines or VMs in Amazon's EC2 service, configured with identities, entitlements, and so on). Employees and on-site contractors of an organisation accessing a SaaS service using identity federation
2. Developers setting up partner user accounts on a PaaS platform (e.g., developers from Newco.com provisioning accounts in Force.com for Partnerco.com employees contracted to perform business process tasks for Newco.com)

3. End users employing access policy management tools to share files and objects with users within and outside of a domain while using cloud storage services (like Amazon S3)
4. A cloud service provider's (such as Amazon EC2) application using another cloud service's storage (e.g., Mosso).

Due to the ability of IAM capabilities like SSO to externalise authentication functions, organisations may quickly acquire *aaS services (Salesforce.com is an example) by cutting down on the time needed to interact with service providers. IAM capabilities can also assist a company in outsourcing a task or service to partners while minimising the impact on the company's security and privacy. For instance, order fulfilment company staff members can use their federated identities to access real-time data stored in a merchant application to manage orders procedure for product fulfilment. In brief, your organization's user access management practises and procedures may be expanded to the cloud by extending your IAM strategy, practise, and architecture. As a result, businesses with established IAM procedures may quickly embrace cloud services while keeping their security policies effective and efficient.

IAM Obstacles

Managing access for various user groups (workers, contractors, partners, etc.) to internally and externally hosted services is a key difficulty for IAM. IT has ongoing challenges in providing quick access to users whose roles and responsibilities often change according to business needs. The organization's user churn is another problem. Industry- and function-specific turnover may result from company developments including mergers and acquisitions, new product and service launches, business process outsourcing, and shifting responsibilities. Seasonal personnel swings in finance departments are one example of a change in turnover. As a consequence, maintaining IAM procedures may become an ongoing problem.

Information access regulations are seldom centrally and uniformly administered. Organizations may have several directories, which may result in intricate webs of user identities, access privileges, and processes. This has made user and access control methods ineffective and exposed these companies to serious security, legal compliance, and reputation problems.

Several businesses have looked for technological solutions to offer centralized and automated user access control in order to handle these difficulties and hazards. Given that the issue is often significant and intricate, it is not unexpected that many of these projects are undertaken with great expectations. Usually, IAM improvement programmes take many years and cost a lot of money. In order to address the root causes of inefficiency while maintaining the effectiveness of the control, businesses should approach their IAM strategy and architecture with both business and IT drivers (related to access control). The organizations' chances of success and return on investment will be better only then.

Definitions of IAM

We'll begin by outlining the fundamental ideas and definitions of IAM roles for all services:

Authentication

Verifying a user's or system's identity via authentication (e.g., Lightweight Directory Access Protocol [LDAP] verifying the credentials presented by the user, where the identifier is the corporate user ID that is unique and assigned to an employee or contractor). Typically, authentication implies a more reliable type of identification. In some use situations, such

service-to-service communication, authentication entails confirming the network service asking for access to data provided by another service (such as a travel agency).

Online service that connects to a credit card gateway on the user's behalf to validate the credit card.

Authorization

After an identity has been established, the user or system's privileges are determined via the authorization procedure. To put it another way, authorisation is the act of enforcing regulations. In the context of digital services, authorization often comes after the authentication phase and is used to assess whether the user or service has the required rights to execute certain tasks.

Auditing

IAM auditing involves reviewing and examining authentication, authorization records and activities to assess the effectiveness of IAM system controls, confirm adherence to established security policies and procedures (such as separation of duties), find security service breaches (such as privilege escalation), and suggest any modifications that are necessary for countermeasures.

IAM Practice and Architecture

IAM is not a one-size-fits-all solution that can be quickly implemented to add capabilities. It is both a component of architecture (see Figure 5-1) and a collection of technological elements, operational procedures, and best practises. The typical business IAM architecture includes a number of layers of services, technology, and procedures. A directory service (like LDAP or Active Directory) that serves as a repository for the identity, credentials, and user characteristics of the user pool for the organisation forms the basis of the deployment architecture. The directory communicates with IAM technology elements including provisioning, user management, authentication, and federation services that enable organisational standards for IAM practises. Organizations often employ several directories, whether they were added via company mergers and acquisitions or were installed for environment-specific purposes (e.g., Windows systems using Active Directory, Unix systems using LDAP).

The following general categories may be used to group the IAM processes that serve the business:

User administration

Actions for successful identity life cycle governance and management.

Authentication control

Actions to ensure that the process for assessing whether an entity is who or what it purports to be is effectively governed and managed.

Authorization control

Actions required for the administration and management of the entitlement rights determination process, which determines which resources an entity is allowed to access in line with the organization's rules.

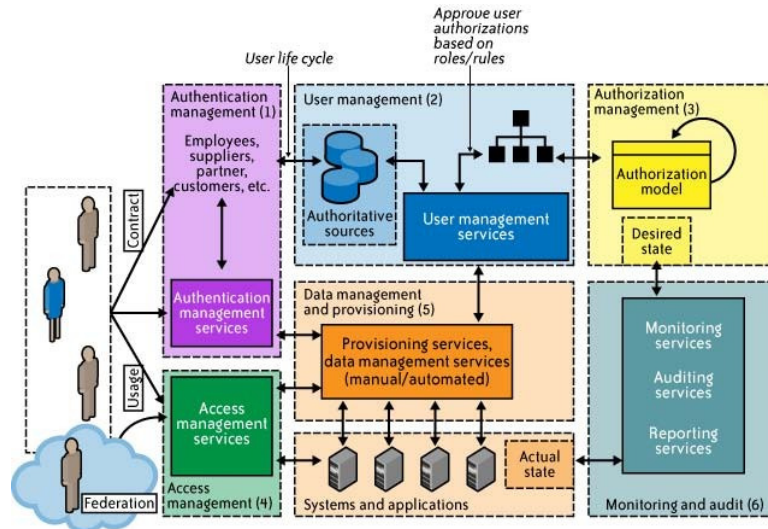
Access control

Response to a request from an entity (user, service) wishing to access an IT resource inside the company by enforcing access control restrictions.

Data provisioning and management

Identity and data dissemination for IT resource authorisation through automated or manual operations

Monitoring, auditing, and reporting user compliance with relation to using Observing and evaluating: organization resources in accordance with the set policies.



Figur1 13.1: IAM Practice and Architecture

IAM is not a one-size-fits-all solution that can be quickly implemented to add capabilities. It is both a component of architecture (see Figure 13.1) and a collection of technological elements, operational procedures, and best practises. The typical business IAM architecture includes a number of layers of services, technology, and procedures. A directory service (like LDAP or Active Directory) that serves as a repository for the identity, credentials, and user characteristics of the user pool for the organisation forms the basis of the deployment architecture. The directory communicates with IAM technology elements including provisioning, user management, authentication, and federation services that enable organisational standards for IAM practises. Organizations often employ several directories, whether they were added via company mergers and acquisitions or were installed for environment-specific purposes (e.g., Windows systems using Active Directory, Unix systems using LDAP).

The following general categories may be used to group the IAM processes that serve the business:

1. User administration
2. Actions for successful identity life cycle governance and management
3. Authentication control
4. Actions to ensure that the process for assessing whether an entity is who or what it purports to be is effectively governed and managed
5. authorization control

6. Actions required for the administration and management of the entitlement rights determination process, which determines which resources an entity is allowed to access in line with the organization's rules
7. Access control: Response to a request from an entity (user, service) wishing to access an IT resource inside the company by enforcing access control restrictions.
8. Data provisioning and management: Identity and data dissemination for IT resource authorisation through automated or manual operations
9. Observing and evaluating: Monitoring, auditing, and reporting user compliance with relation to using organisation resources in accordance with the set policies (Figure 13.2).




Figure 13.2: Life cycle of an identity.

Preparing for the Cloud

Organizations preparing for cloud services must first consider fundamental user administration tasks like user account provisioning and continuing account management, which includes prompt deprovisioning of users when they no longer need access to the cloud service.

The infrastructure and architecture of organisations that have invested in identity and access management methods should be able to be used to their advantage. Businesses that don't yet have identity and access management policies in place may employ cloud-based solutions from a variety of identity management service providers (examples include Symplified, Ping Identity, Conformity, and TriCipher). Multiprotocol gateways offered by manufacturers like Symplified and Vordel are probably going to be extensively utilised for a while to come since there are several IAM standards available, all of which are at different stages of development and acceptance. These vendor-based solutions provide cloud service gateways that support federation. In "Identity management-as-a-service".

Companies should invest in the fundamental technological components that allow user management and federation and begin with an IAM strategy and architecture. As SSO is supported by federation, users won't need to sign in more than once or remember user authentication details particular to each cloud service provider (e.g., one user ID and

password per provider), reducing risks to enterprises while also ensuring a consistent user experience.

Organizations will be able to support an identity provider (IdP), also known as an SSO provider (using an existing directory service), by architecting an identity federation architecture or service for cloud-based identity management). With such architecture, businesses may exchange IDs with reputable CSPs without disclosing confidential user information or user credentials. The definition, description, and administration of obligatory, non-mandatory, and important characteristics are essential stages to be ready for federation. Maintenance of identity attributes also plays a part in federation. Using a common federation architecture to federate identities and provide single or reduced sign-on to cloud services, this strategy may assist enterprises in extending IAM operations and practises.

The foundation of federation technology is often a centralised identity management architecture that makes use of industry-standard identity management protocols like Liberty Alliance, WS Federation, and Security Assertion Markup Language (SAML). Out of the three main federation protocol families, SAML seems to be accepted as the de facto standard for enterprise-controlled federation.

The Organization for the Advancement of Structured Information Standards (OASIS), the Liberty Alliance, and the Shibboleth Project all contributed to these federation standards, which merged their efforts to improve SAML 1.0 to develop SAML 2.0. As of today, manufacturers and organisations from all over the globe support SAML 2.0 as the de facto industry standard for delivering and administering open identity-based systems. SAML 2.0 was adopted as an official OASIS industry standard in March 2005. The U.S. Federal E-Authentication profile, the globally accepted Liberty eGov profile, the higher education Shibboleth and Eduserv federations, as well as many other industry federations, all need SAML.

In a private community cloud, SAML may be used to provide federation among community members, which is necessary for safe information exchange. This was further supported in 2007 by the industry research company Gartner, which referred to SAML 2.0 as the facto federation standard across sectors.

Organizations must implement an Internet-facing IdP in order to establish a user federation model for their users. These steps include establishing an authoritative source for the identity, determining the crucial user profile attributes, and planning and implementing an IdP that supports an SSO service and is reachable by CSPs. Using federation technology components that communicate with your directory, internet-facing IdPs may be set up. The foundation of a business architecture's access management features is a directory, such LDAP or Active Directory. Companies that intentionally or accidentally replicate directories that are accessible through a DMZ network may be able to hasten the adoption of federation. Similar to this, federation may be performed in businesses with federation-friendly architecture where directories may be available to authorised third-party providers through network access restrictions (such firewalls, site-to-site VPNs, or proxies). To allow delegated authentication or the SSO functionality, businesses often implement an identity federation solution that easily interfaces with their directory service (examples include Sun's OpenSSO, Oracle's Federation Manager, and CA's Federation Manager).

IAM Procedures and Standards That Are Applicable to Cloud Services

We outlined the prerequisites and advantages of implementing common IAM principles and practises to cloud services in the preceding sections. The pertinent IAM standards that

encourage enterprises to embrace cloud services will be covered in this section. Businesses should additionally assess the CSP's commitment to and support for IAM standards when selecting cloud services based on business and operational factors.

Organizational IAM Standards and Specifications

Organizations may create effective and efficient user access management procedures and practises in the cloud by using the IAM standards and specifications listed below. The four main difficulties in user and access control that cloud users encounter are used to arrange these sections:

1. How can I provide my users a single sign-on experience and prevent the duplication of identities, characteristics, and credentials? SAML.
2. How can the provisioning and deprovisioning processes for user accounts in cloud services be automated? SPML.
3. How can I manage my users' entitlements and create user accounts with the proper permissions? XACML.
4. How can I give cloud service X permission to view my data in cloud service Y without giving away my login information? OAuth.

Language for Security Assertion assertions (SAML)

The most developed, comprehensive, and frequently used set of standards for browser-based federated sign-on for cloud users is SAML. After establishing her identity with the identity service, the user may freely access provisioned cloud services that are part of the trusted domain without having to go through the sign-on procedure particular to the cloud. As SAML permits delegation (SSO), users may choose to use strong authentication (multifactor authentication) for certain cloud services by employing risk-based authentication rules. Using the organization's IdP, which enables delegated authentication and strong authentication, makes this task simple to do. Users are less susceptible to phishing attempts, which have been continuously increasing online, by using robust authentication methods like dual-factor authentication. Strong

It's also recommended to authenticate to cloud services to safeguard user credentials against man-in-the-middle attacks, which happen when computers or browsers are attacked by trojans and botnets. The CSP may transfer the authentication rules to the client organisation by implementing a SAML standard that offers a delegated authentication paradigm for cloud customers. Simply said, SAML enables CSPs to become independent of client authentication needs.

A browser-based SSO into Google Apps (Figure 13.3). The graphic shows the steps that a user who is federated to Google must do in order to complete SSO:

1. A user from your company tries to access a Google application that is hosted, such Gmail, Start Pages, or another Google service.
2. A SAML authentication request is produced by Google. The URL for the IdP providing the SSO service in your company contains the encoded and embedded SAML request. The SSO URL also includes the Relay State parameter, which contains the encoded URL of the Google application the user is attempting to access. It is intended for this Relay State argument to be an opaque identifier that is returned without being altered or checked.

3. Google gives the user's browser a redirect. The encoded SAML authentication request for your organization's IdP service is included in the redirect URL.

4. Your IdP decodes the SAML request and gets the URL for the user's destination URL as well as Google's Assertion Consumer Service (ACS) (the Relay State parameter). The user is then authenticated by your IdP. Your IdP may request legitimate login information or look for valid session cookies to verify the user's identity.

5. Your IdP creates a SAML response that includes the username of the verified user. This answer is digitally signed using the partner's public and private DSA/RSA keys in line with the SAML 2.0 standard.

6. Your IdP encrypts the Relay State parameter and the SAML answer before sending them to the user's browser. Your IdP offers a method for the browser to transmit that data to Google's ACS. For instance, your IdP may allow you to embed the destination URL and SAML answer in a form, along with a button that users can click to send the form to Google. Moreover, your IdP could add JavaScript to the website that sends the form to Google automatically.

7. Google's ACS validates the SAML answer using the public key of your IdP. If the verification of the answer is successful, ACS connects the user to the target URL.

The user has signed into Google Apps and has been sent to the target URL.

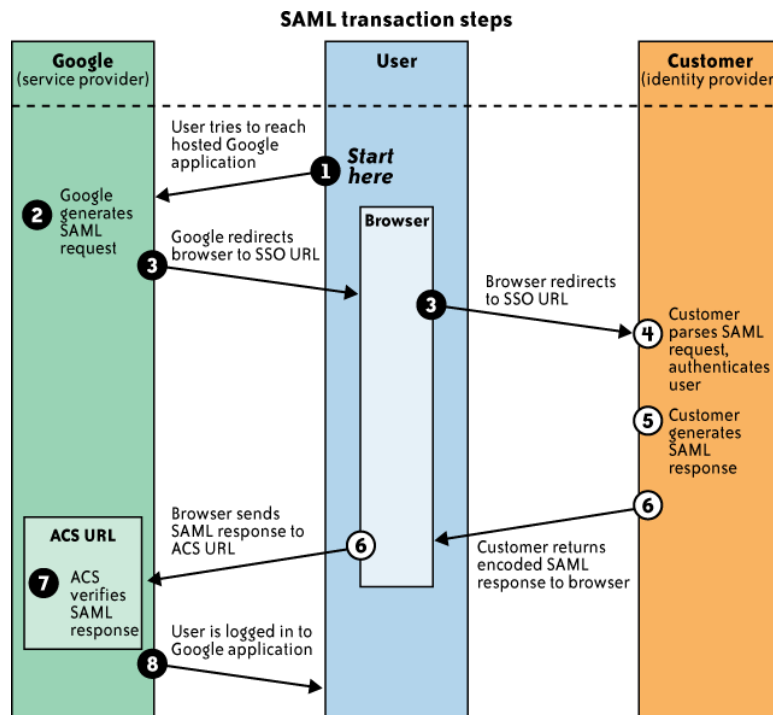


Figure 13.3. SSO transaction steps using SAML

Language for Service Provisioning (SPML)

OASIS is creating SPML, an XML-based framework enabling collaborating organisations to exchange user, resource, and service provisioning information. An emerging standard called SPML may assist businesses in automating the supply of user IDs for cloud services (e.g., an application or service running at a customer site requesting Salesforce.com for new

accounts). Organizations should utilise SPML as it becomes available to supply user profiles and accounts with the cloud service. Software as a service (SaaS) providers may allow "just-in-time provisioning" to instantly generate accounts for new users if SPML is enabled (as opposed to preregistering users). According to that paradigm, the cloud service provider (CSP) collects characteristics from a new user's SAML token, generates an SPML message on the spot, and sends the request to a provisioning service, which then adds the user identity to the cloud user database.

By adopting SPML, user or system access and entitlement rights to cloud services may be standardised and automated, preventing consumers from being tied to proprietary solutions.

A cloud-based provisioning system is requested by an HR system via an SPML request, as shown in Figure 13. 4. The HR System of Record (requesting authority) is a client of SPML web services engaging with the SPML provisioning service provider at the cloud service provider, which is in charge of providing user accounts on the cloud services, in the figure (provisioning service target).




Figure 13.4. SPML use case

Language for eXensible Access Control (XACML)

For managing policies and making access choices, XACML is an all-purpose, XML-based access control language that has received OASIS certification. It offers an XML schema for a broad policy language that is used to safeguard any sort of resource and control who has access to it. In addition to providing a model of the policy language, the XACML standard also suggests a processing environment model to control the policies and reach access choices. The request/response protocol that the application environment may use to interact with the decision point is also specified in the XACML context. An access request answer is also expressed using XML.

The majority of apps, whether web-based or not, come with an integrated authorization module that allows or prohibits access to certain application features or resources depending on the user's allocated entitlements. With an IAM system that is centrally controlled, application-specific authorisation models, it is challenging to identify each user's access privileges across all programmes because of (silos). As a result, XACML's objective is to

offer a standardised language, a means of access control, and policy enforcement for use by any applications that implement a single authorization standard. Several authorization policies and guidelines depending on the user role and job function are used to make these authorisation choices. XACML enables uniform authorisation policies, to put it briefly (i.e., the use of one consistent XACML policy for multiple services).

Figure 13.5 depicts the interaction between a variety of healthcare professionals who each have certain responsibilities (permission privileges) and access to private patient information kept in a healthcare application.

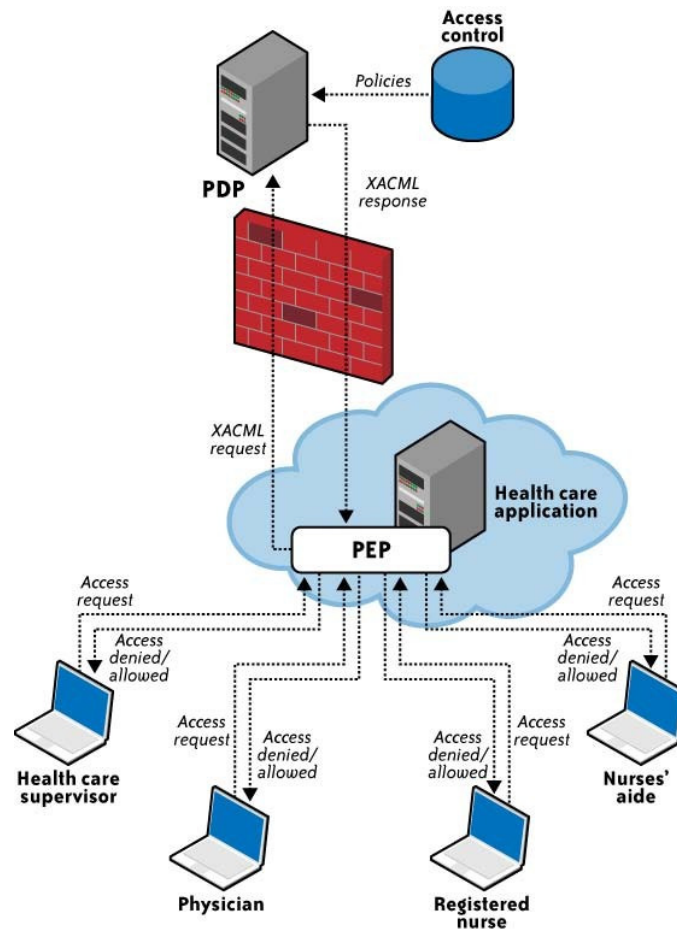


Figure 13.5: Use scenario for XACML

The figure depicts the stages of the XACML process as follows:

1. The health care application controls how different hospital staff members—including the doctor, registered nurse, nurses' assistant, and health care supervisor—access different parts of the patient file. This application sends the request to the PEP in reliance on the policy enforcement point (PEP).
2. The PEP serves as the application environment's interface. It receives the requests for access and assesses them using the policy decision point (PDP). The resource is then either made accessible or not (the health care record).
3. The request is then sent to the PDP by the PEP. The PDP serves as the primary arbitrator for access requests. It gathers all the essential data from the information sources that are

accessible and comes to a conclusion about what access to allow in the end. The PDP should be situated in a secure network with strict access control regulations, such as a corporate secure network that is fortified by a corporate firewall.

4. The PDP provides the PEP the XACML answer after assessment.

5. By upholding the PDP's authorisation decision, the PEP satisfies its duties.

The communication takes occur through a request-response protocol with the payload being a XACML message. In this method, the assessment of rules against access decision requests is sent via XACML.

Accessible Authentication (OAuth)

OAuth is a new authentication standard that enables users to exchange their private resources—such as pictures, videos, contact lists, and bank accounts—between CSPs without having to reveal the authentication details (e.g., username and password). OAuth is an open protocol that was designed to make it possible to authorise users using a secure application programming interface (API), which is a straightforward and common technique for desktop, mobile, and online apps. OAuth is a way for publishing and interacting with protected data for application developers. OAuth gives consumers access to their data housed by another provider while safeguarding their login credentials for CSPs.

When using a web services SSO paradigm inside a business, OAuth may be used to allow SSO with a trusted service provider. Without needing a specific federation architecture, OAuth makes it easier to authorise two services to communicate with one another. OAuth, like OpenID, was developed with the goal of enabling consumer services to access customer data stored across providers. In order to improve usability, Google has introduced a hybrid OpenID and OAuth protocol that integrates the permission and authentication processes in fewer stages. OAuth compatibility was just introduced for Google's GData API. SAML is supported by GData for browser SSO.

The interactions between a client's or partner's web application, Google services, and the end user are shown in Figure 13.6:

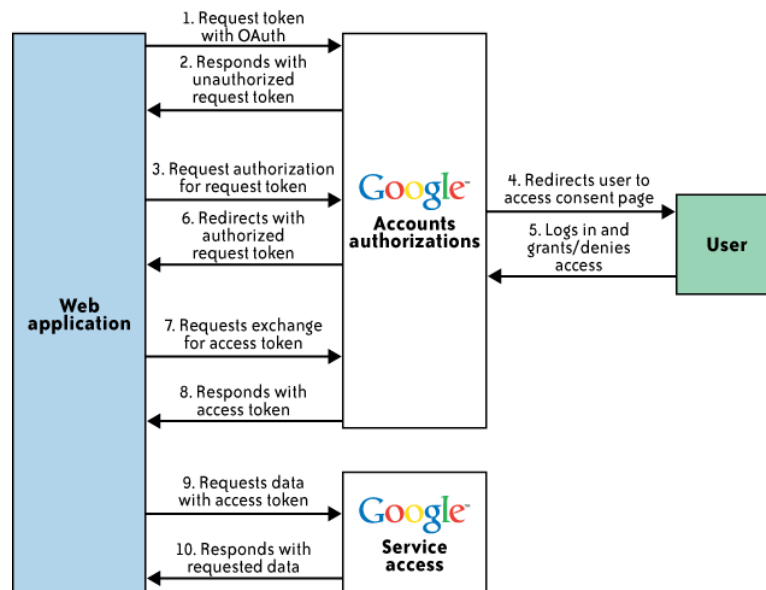


Figure 13.6: Representing the OAuth usage

1. Customer web application contacts the Google Authorization service, requesting for a request token for one or more Google service.
2. Google checks to see whether the web application is registered and then returns an unauthorised request token.
3. Using the request token, the web application sends the user to a Google authorisation page.
4. The user is asked to connect into his account on the Google authorization page (for verification), after which he may choose to give or refuse the web application restricted access to his Google service data.
5. Whether to allow or refuse access to the web application is up to the user. If the user declines access, he is not sent back to the online application but rather to a Google page.
6. Should the user consent to access, the Authorization service sends him back to the web page that was defined by the web application that was registered with Google. The now-approved request token is included in the redirect.
7. To exchange the allowed request token for an access token, the web application makes a request to the Google Authorization service.
8. After confirming the request, Google provides a legitimate access token.
9. A request is sent to the relevant Google service by the web application. The access token is included in the request, which is signed.
10. The Google service provides the requested information if it recognises the token.

CHAPTER 14

IAM STANDARDS, PROTOCOLS, AND SPECIFICATIONS FOR CONSUMERS

Dr. Taskeen Zaidi

Department of Computer Science and Information Technology,
Jain (Deemed to be University) Bangalore, Karnataka, India
Email Id-t.zaidi@jainuniversity.ac.in

The protocols and requirements listed below are intended for consumer cloud services and are thus not applicable to business cloud computing. OpenID is a free, open-source standard for user authentication and access management that enables users to sign on to a variety of services using only one digital identity, or a single sign-on for users of services that implement OpenID. By enabling a user to log in once and access the resources of various software systems, it therefore substitutes the standard logon procedure, which requires a logon username and password. OpenID is mainly intended for use with consumer services provided by online businesses like Google, eBay, Yahoo!, Microsoft, AOL, BBC, and others. Due to trust difficulties, adoption of OpenID for corporate usage (i.e., non-consumer use) is almost non-existent; some studies have demonstrated that OpenID might speed up phishing assaults that could lead to the compromise of user credentials.

Detailed cards: Another open standard for identification on the Web is information cards. The Information Card Foundation, whose steering committee includes officials from Google, Microsoft, PayPal, Oracle Novell, and Equifax, is in charge of directing the standard itself. According to the Foundation, its goal is "to decrease the occurrence of identity theft by safeguarding digital identities in lieu of conventional logons and passwords." This standard aims to provide consumers a secure, reliable, and phishing-resistant user interface without the need for a login and password.

For convenience, users may utilise their information card digital identity on several websites without risking the security of their login credentials (similar to using an OpenID identity across multiple sites). The Information Cards Protocol is intended for usage in high-value contexts, like banking, where phishing resistance and support for secure authentication methods like smart cards are essential operational needs.

Information cards may be implemented, issued, and accepted by any service provider (also called i-cards). The requirements are much more complex than OpenID since information cards are created using WS specifications rather than HTTP redirect. This method provides excellent security against identity theft and phishing, but it still has a few problems that keep it from accomplishing its goal. The system's biggest flaw is that it only functions if the website is up.

Participating and accepting information cards is utilised by the customer. The information card is worthless without this association.

The technique will be more and more helpful as websites embrace information cards, although they are now of limited utility. For instance, a managed information card from Microsoft Windows Live ID may be used to enable single sign-on for the majority of Microsoft websites, such as MSDN, TechNet, Live, and Connect.

Accessible Authentication (OATH)

OATH is an initiative by professionals in the IT sector to provide a reference architecture for robust universal authentication across all users, devices, and networks. This programme aims to address the three main forms of authentication:

1. SIM-based authentication utilising the Global System for Mobile Communications/ General Packet Radio Service (GSM/GPRS) Subscriber Identity Module (SIM)
2. Authentication via Public Key Infrastructure (PKI) (using an X.509v3 certificate)
3. Authentication using a One-Time Password (OTP)

This authentication protocol makes use of federated identity protocols as well as well-known infrastructure elements like a directory server and a Remote Authentication Dial-In User Service (RADIUS) server.

API for Open Authentication (OpenAuth)

Via the usage of the AOL-exclusive OpenAuth API, users of AOL and AOL Instant Messaging (AIM) may be authenticated on other websites and services. An AIM or AOL registered user may use this authentication mechanism to sign in to a third-party website or application and access AOL services or new services that are developed on top of AOL services. AOL claims that the OpenAuth API has the following capabilities:

1. A safe way to sign in. The websites or apps that a user logs into are never made aware of their user credentials.
2. A safe way to manage which websites may access private or protected material.
3. Permissions are only automatically granted if the user chooses Allow Always on the Consent screen.
4. Requesting user permission before reading any private or protected material on the website or application (e.g., separate consent requests to allow Buddy List information, to send IMs, to read albums).
5. Access to other websites that do not belong to AOL without having to register for a new account on each one that accepts AOL OpenAuth APIs.
6. Since OpenAuth is a proprietary protocol, it is not accepted outside of the AOL network and is not regarded as an open standard by the cloud computing industry.

Standards and Protocols for Consumer and Enterprise Authentication

IAM Methodologies in the Cloud

IAM procedures in the cloud are still developing in comparison to the conventional business application deployment paradigm.

Standards support by CSPs (SaaS, PaaS, and IaaS) is not uniform among providers given the present level of IAM technology. Large providers like Google, Microsoft, and Salesforce.com seem to have some fundamental IAM capabilities, but in our opinion, they still fall short of what is needed for handling corporate IAM needs for regulatory, privacy, and data protection standards. Based on the authors' evaluation and applied to all SPI service delivery models.

The maturity model tackles the four primary elements of the IAM automation process and takes into consideration the dynamic nature of IAM users, systems, and applications in the cloud:

1. User Administration, New Users

2. User Modifications and Management
3. Handling of Authentication and Authorization

While standard corporate IAM procedures and processes may be used to cloud services, they must be modified for the cloud context in order to reap their advantages. User management features in the cloud may be broadly divided into the following categories:

Administration of cloud identities, Federation or SSO, Authorization management, and Compliance management.

Cloud Identity Management

The life cycle management of user identities in the cloud should be the main emphasis of cloud identity administration tasks. This includes provisioning, deprovisioning, identity federation, SSO, password or credential management, profile management, and administrative management.

Companies should look towards cloud-based identity management services if they cannot support federation. This new breed of services often serves as the organization's proxy IdP and synchronises the internal directories with the directory (which is typically multitenant).

Organizations may prevent duplicate identities and attributes being stored with the CSP by federating identities using either a private internal IdP or a cloud identity management service provider.

Customers may need to develop unique strategies to handle user management tasks in the cloud due to CSPs' uneven and scant support for identity standards. When federation is not enabled, provisioning users might be difficult and complicated. The use of manual procedures, web-based administration, outsourced (delegated) administration including the uploading of spreadsheets, and execution of customised scripts at both the client and CSP sites is commonplace in enterprises.

The latter architecture is undesirable since it cannot be scaled across different CSPs and will be expensive to administer over time.

Identity Federation (SSO)

Companies wanting to create identity federation that allows SSO for users may choose one of the following two approaches (architectures):

1. Establish an enterprise IdP inside the boundaries of the organisation.
2. Connect your system with a reputable cloud-based identity management solution.
Each architecture has benefits and drawbacks.

Organizational identity provider

With this design, an organization's IdP will handle authentication for cloud services. Inside a trusted ring of CSP domains, the organisation federates identities in this delegated authentication architecture. With all the domains that are permitted to delegate authentication to the IdP, a circle of trust may be established. More control may be taken over user identities, characteristics, credentials, and rules for authenticating and authorising users to a cloud service under this deployment architecture, where the company will offer and maintain an IdP. The deployment architecture of the IdP is shown in Figure 14.1.

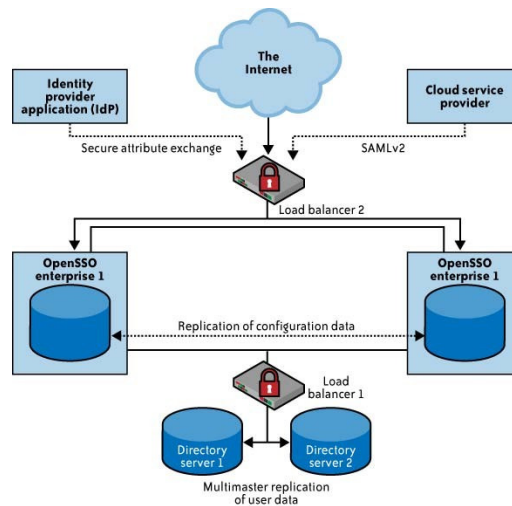


Figure 14.1: Identity provider deployment architecture

The following are this strategy's distinct benefits and drawbacks:

Pros

Businesses may expand their IAM processes to the cloud by using the current investment they have made in their IAM infrastructure. For instance, businesses who have used SSO for their data center's applications show the following advantages:

Cons

They have direct control over the service-level agreement (SLA) and security of the IdP. They incrementally invest in improving the current identity architecture to accommodate federation. They are compatible with internal rules, procedures, and access management frameworks.

Due to the inclusion of life cycle management for non-employees like clients, additional inefficiencies might arise if the infrastructure supporting federation is not changed.

The majority of firms will probably keep employing naturally established IAM infrastructures and methods to handle employee and long-term contractor IDs. Yet, they seem to prefer to hire a reputable cloud-based identity provider as a service partner to handle partner and customer IDs.

Online identity management

Under this architecture, identity management as a service (IDaaS) providers are able to handle authentication on behalf of cloud services. Under this strategy, businesses contract with companies like Ping Identity, TriCipher's Myonelogin.com, or Symplified.com to handle user management and federated identity management technologies.

Organizations may need to use their IAM system and procedures to manage the identity life cycle when federating IDs to the cloud. Nevertheless, if the company has to communicate with several different partners and cloud service federation schemes, it could profit from using an external multiprotocol federation gateway (identity federation service). For instance, as of the time of writing, Google Apps and Salesforce.com both support SAML 2.0. An identity management CSP like Symplified or TriCipher may host a multiprotocol federation gateway that is useful for businesses using Google Apps and Salesforce.com.

An organisation may also outsource credential issuance (and background checks) to a service provider, such as the GSA Managed Service Organization (MSO), which issues personal identity verification (PIV) cards and, optionally, the certificates on the cards, in circumstances where credentialing is challenging and expensive. Federal civilian agencies may use of the USAccess management end-to-end solution as a shared service from the GSA MSO.

This is essentially a SaaS paradigm for identity management, where the SaaS IdP serves as a proxy for users of the company accessing cloud services and keeps identities in a "trusted identity store," as shown in Figure 14.2.

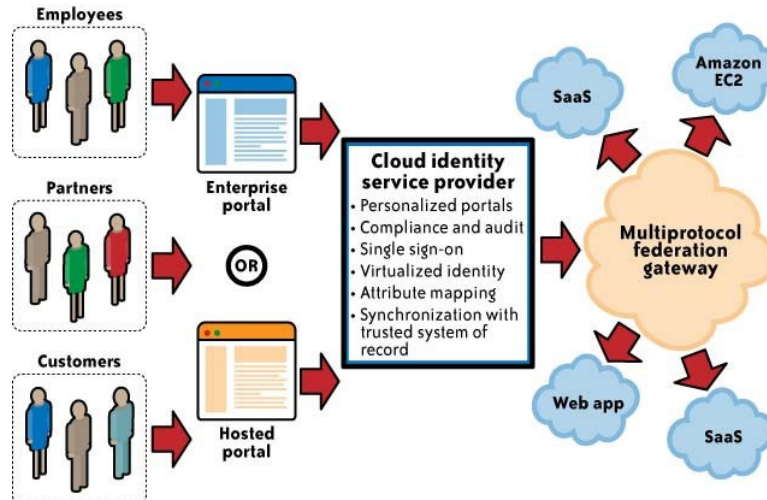


Figure 14.2: Identity management as a service, (IDaaS)

By a provider-proprietary method (e.g., agents operating on the customer's premises syncing a portion of an organization's identity store to the identity store in the cloud using SSL VPNs), the identity store in the cloud is maintained in sync with the corporate directory.

The company should collaborate with the CSP to assign authentication to the cloud identity service provider after the IdP has been set up in the cloud. Before cloud users may access any cloud services, the cloud IdP will verify them (this is done via browser SSO techniques that involve standard HTTP redirection techniques).

The following are this strategy's distinct benefits and drawbacks:

The difficulty of working with several CSPs supporting various federation standards is hidden by assigning certain authentication use cases to the cloud identity management service. Salesforce.com and Google are two examples of services that offer delegated authentication using SAML. Nevertheless, as of the time of this writing, they only support SAML 1.1 and 2.0 respectively for Salesforce.com and Google Apps. Organizations embracing cloud services may mask this integration difficulty with cloud-based identity management systems that accept both SAML standards (multiprotocol federation gateways).

The minimal amount of architectural modifications required to accommodate this paradigm is another advantage. Users may access cloud services after identity synchronisation between the organisation directories or trusted system of record and the identity service directory is set up.

Services that make use of corporate identification, credentials (both static and dynamic), and authentication guidelines. If you depend on a third party to provide your identity management service, you can have limited access to information about the service's implementation and architecture. As a result, the SLA, performance management, and availability of the identity management service provider determine the availability and authentication performance of cloud apps. It's critical to comprehend the identity management service provider's service level, architecture, service redundancy, and performance guarantees.

This method may not be able to provide bespoke reports to satisfy internal compliance needs, which is another disadvantage.

However, when identity attributes are improperly created and connected to identities, identity attribute administration may become complicated (e.g., definitions of attributes, both mandatory and optional). To control user attributes that cross the organization's trust border, new governance procedures may need to be approved for different actions (add, alter, or delete attributes). Over the identity's life cycle, its qualities will evolve and may get out of sync.

While these methods allow users to be identified and authenticated for cloud services, as we will explore in the next section, other capabilities and integration details are unique to the SaaS, PaaS, and IaaS service delivery models.

Management of Cloud Authorizations

For their cloud users, medium-sized and big enterprises often have special needs for authorisation tools (i.e., assignment of privileges, or entitlements, to users based on their job functions). Role-based access control (RBAC), in this case authorization is built to meet the needs of the organisational functional roles, may be necessary in specific situations for a business application. As of this writing, the capabilities for managing and enforcing permission for cloud services are limited, and even when they are, they are often quite coarse-grained. The services on offer may not satisfy the needs of your business.

The majority of cloud services support an administrator and end user with at least dual roles (privileges). The assignment of administrative privileges to the administrator position is standard procedure across CSPs. Administrators with these rights are able to supply and deprovision identities, basic attribute profiles, and, in certain situations, access control rules like password complexity and approved trusted networks.

The ideal standard for defining and implementing authorization and user authentication requirements is XACML, as we have explained. We are not aware of any cloud services that allow XACML to express authorization rules for users as of this writing.

Assistance from IAM for Compliance Control

In addition to having a significant influence on the effectiveness of internal IT operations, cloud IAM design and practises also have a significant impact on managing compliance inside the company. The efficacy of the controls defined by compliance frameworks may be increased with the aid of properly implemented IAM procedures and processes. For instance, enterprises may lower the risk of unauthorised users accessing cloud services and satisfy your privacy and compliance needs by automating the timely provisioning and deprovisioning of users and entitlements. Moreover, identity and attribute management will be a major area of concern for regulatory and privacy problems in terms of compliance; appropriate IAM governance mechanisms should be implemented to deal with these challenges.

A consolidated view of corporate activities is provided by IAM practises and procedures, and an automated procedure that may thwart insider threats before they materialise is also provided. However, due to the CSP's limited support for IAM standards like SAML (federation), SPML (provisioning), and XACML (authorization), you should individually evaluate the CSP's capabilities and set up procedures for handling compliance with regard to identity (including attribute) and access management.

IAM Practice for Cloud Service Providers

IAM features need to be included in the design criteria for the cloud service from the viewpoint of the CSP (SaaS, PaaS, or IaaS), with the intention of handing over user authentication and authorization to the client utilising user management and federation standards. Integration issues arise for both consumers (such as single sign-on and user provisioning) and CSPs as a result of support for IAM functionalities (e.g., billing, accounting resource utilization). IAM integration considerations made early in the service design process will assist the client and CSP avoid costly retrofits. Hence, IAM features should be included into cloud service architecture and platform application development at different phases of the product life cycle, including architecture, design, and implementation (e.g., externalise the authentication from the application using the federation feature).

From the standpoint of a cloud user, the governance, integration, and user experience of the cloud service will be impacted by the application's IAM capabilities (or lack thereof), such as identity federation (e.g., barriers to adopt the cloud service). Hence, IAM requirements for cloud applications should be understood by architects, designers, and developers, and baked into the RFP or CSP assessment criteria.

1. Provisioning of cloud service accounts to users, including administrators, is one of the prerequisites for enterprise IAM.
2. Cloud service provisioning for service-to-service integration (e.g., private [internal] cloud integration with a public cloud).
3. SSO support based on federation standards for users (e.g., SAML support).
4. Support for internal and regulatory policy compliance demands, such as role-based access control, rules-based access control, or claims-based authentication techniques for the separation of roles. RBAC features encourage the usage of a least-privilege-based access model, in which a user is only given the rights necessary to complete a task. Since claims-based technique only permits the user's entitlements—not her real identity—to flow with communications, it supports several crucial privacy use cases and provides fine-grained permission without the need to actually embed the user's identity into messages.
5. User activity tracking, reporting, and monitoring as required by organisational rules and legal requirements including SOX, PCI, and HIPAA.

SaaS

Security management is one of the main issues that IT and business decision-makers have with software-as-a-service applications. Even though the majority of SaaS vendors have been able to show that their cloud-based applications are secure from an operational perspective, organisations still need to address access control issues to make sure that their corporate data is completely secure from the perspective of corporate policies and procedures.

Since SaaS apps have a low adoption barrier and a pay-as-you-go service model, they are becoming more and more popular and need to be addressed. In certain circumstances, business divisions are avoiding IT and working directly with SaaS suppliers, which might

cause more IT problems. IT must mitigate risks that may result from this loss of visibility and control and be able to guarantee that the appropriate users have access to data housed by SaaS suppliers at the appropriate level.

Two key difficulties for identity management should be taken into account by organisations thinking about integrating with SaaS services:

- a. Is the business prepared to provide and manage the user life cycle by extending its current IAM practise to the SaaS service?
- b. Can user provisioning and life cycle management be automated by the SaaS provider without the usage of a customised SaaS service solution?

Customers' obligations

Customers have less accountability and accessible controls to safeguard information in SaaS services. SaaS solutions are often multitenant and provided to the client via a web browser. Only IAM controls, such as identity provisioning, authentication rules (such as password strength), profile configuration, and fundamental authorization policies that appear as user profiles, are exposed to the client. From an IAM viewpoint, clients have the following obligations:

Provisioning of users

The SaaS provider often uses exclusive techniques for user provisioning. Consumers must comprehend the preferred technique, the time it takes to activate customers, and the user characteristics that the SaaS service supports. Most commonly the provisioning procedure is manual and may entail uploading spreadsheets or documents in XML format. As provisioning users in bulk is the most frequent use case, almost all SaaS providers offer it. When a user clicks on a URL that is specific to their user identity, for example, certain SaaS providers may allow just-in-time provisioning, where user identities are generated instantly via a provisioning request (which may sometimes utilise SPML).

Profile control

Customers could have the option to build user profiles that are used in user authorisation as part of the provisioning process. Assigning rights to users inside a SaaS application may be done using user profiles like manager and user. It's true that these are not complicated capabilities, so users will need to familiarise themselves with the administration and flexibility of the profiles.

Evaluation of SaaS IAM capabilities

Customers are in charge of assessing how well CSPs offer IAM capabilities like SSO (using identity federation). SAML is currently supported by major SaaS providers and is the de facto standard for federating IDs (among them Google and Salesforce.com).

SAML 2.0 isn't supported by every provider, and some may only offer SAML 1.1. For instance, Google Apps supports SAML 2.0 but Salesforce.com only supports SAML 1.1. So, it's crucial to comprehend which federation protocols are supported by specific providers as well as the integration needs for federation and SSO.

Investigation assistance

Investigating occurrences often also necessitates the use of logs and audit trails. For instance, if the service provider has a breach, PCI DSS mandates that the company "provide for rapid forensic investigation". Monitoring (let alone investigation) is challenging since the SaaS

provider's logs are internal and not always available to the outside world or by customers. Be careful to negotiate access to the provider's logs as part of any service agreement since it is necessary for PCI compliance and may be needed by auditors and regulators.

Compliance supervision

Although if firms outsourcing their data with SaaS have the same security issues as those that exist inside their own networks—securing the network, hardware, apps, and data—trust and transparency make the problem worse with cloud computing. Since the extent of the data housed in SaaS includes compliance with governmental rules like SOX, the Gramm-Leach-Bliley Act (GLBA), and HIPAA as well as with industry standards like PCI DSS, it may be difficult to fulfil those requirements. Customers of SaaS services are often in charge of compliance management, but the data are hosted by the supplier. Examine if the SaaS providers' access control, logging, reporting, and auditing capabilities are sufficient to satisfy compliance management needs.

Tasks of CSP

Certain IAM-related duties are within the purview of the CSP, while others are the customer's. The following are CSP duties:

Services for authentication

Unless the SaaS provider enables delegated authentication via federation, it normally authenticates the SaaS users using a user identification and static password through a web form supplied over SSL. It is up to the SaaS provider to authenticate customers depending on the network trust level since users may access the service from anywhere on the Internet. For instance, certain CSPs may preregister the IP address or IP range of a user's location (home, workplace, etc.) to safeguard data from hackers who are employing keyboard loggers (perhaps placed covertly on the user's computer) to steal the user's identity and login information. It is essential for the CSP to supply and maintain a constantly accessible authentication service since authentication activity comes before real SaaS service consumption.

Policies for account management

The account management rules, such as account lock-outs (after several failed login attempts), account provisioning techniques, and privilege account administration responsibilities, should be communicated by CSPs.

Federation

Customers that want to utilise this function and enable SSO for their users should be provided with the information by CSPs that support identity federation using standards like SAML. Such information includes the version (SAML 1.1, SAML 2.0), a use case implementation example, and implementation details of the federation using the API (e.g., support for SAML using REST and SOAP).

PaaS

There aren't many solutions available to organisations thinking about extending their current IAM policies to PaaS cloud providers. Typically, PaaS CSPs use federation to assign authentication tasks to the IdP of the PaaS provider (for example, the Google App Engine assigns authentication to Google's authentication service). Delegated authentication has limited support in certain circumstances, such as Salesforce.com's Force.com, and is often

carried out without the use of SAML assertions (e.g., it is proprietary to each PaaS provider implementation). Moreover, CSPs provide software components that may be used for limited authorisation and authentication using programming languages (often PaaS-specific).

The SAML 2.0 compatible "Geneva" Claims-Based Access Platform was just released by Microsoft (and is still in beta as of this writing). The project's objective is to aid businesses in federating users utilising Microsoft's federation service, Security Token Service, and to assist developers in externalising authentication, authorization, and customization from .NET applications (STS). Potentially, STS implemented in an organisation might be used by Microsoft developers to interact with apps running on the Azure platform. Customers of Microsoft who are interested in federating their Active Directory directory into a cloud are the target audience for this service. As a result, it is unclear if the Geneva Claims-Based Access Platform will work with current SAML 2.0 compliant SaaS and PaaS providers.

IaaS

As of this writing, businesses looking to expand their current IAM procedures to IaaS cloud service providers (computing and storage) have few, if any, alternatives. IaaS providers do not have access to applications that are hosted on the IaaS platform since they provide compute or storage-as-a-service. Few providers, like Amazon Web Services (AWS) EC2, provide a web portal to supply users, manage user keys, and assign users to security groups that pertain to the administrative duties of IaaS. Almost all IaaS providers employ Secure Shell (SSH) to log on and administer users and credentials.

The following are some of the duties and difficulties associated with managing users in IaaS services:

Provisioning of users

User provisioning for administrators and developers on dynamic IaaS systems. Considering that hundreds of computers are set up for workload management, user provisioning must be automated and based on policies at the time an image is created. To prevent duplicate identities on systems, systems should ideally depend on corporate directories (LDAP, Active Directory) for user management. While it should be evaluated on a per-CSP basis, the virtual network architecture in cloud and network security rules may conflict with directory-based authentication systems.

Privileged user administration

Keeping track of system administrators' private keys and securing them when they depart the firm (e.g., SSH host keys).

Assigning keys to customers

Distributing the IDs and keys needed to use the service. These keys are used for both service user authentication and controlling billing-related access to client accounts. For instance, each EC2 client of Amazon is given an Access Key ID, a Secret Access Key, an X.509 certificate, and a matching private key.

Either the X.509 certificate and private key or the Access Key ID and Secret Access Key are used to authenticate to an Amazon request. As a result, it is the customers' responsibility to furnish and secure these keys.

User management for developers

IaaS instances are provisioned with developers and testers, and those instances are deprovisioned when access is no longer needed.

End-user administration

Granting access to those who need it to apps hosted by IaaS.

To prevent a duplicated user database at each of the IaaS clouds, there is currently no automatic mechanism to synchronise an organisational LDAP or Active Directory directory with IaaS providers. Nevertheless, several third-party identity management service providers claim to have created adapters for managing and deploying EC2 users.

Advice One of the biggest barriers to organisational adoption of cloud services is the management of identity and access in the cloud. Secure cooperation with international partners and safe access for international personnel who are consuming critical information from any place and any device at any time are just two examples of how IAM supports business demands.

Although the fundamental technology building blocks for IAM already exist (trusted identity stores, provisioning procedures, authorization and authentication methods, federation), extending or migrating those technologies into cloud services won't produce the alleged IAM benefits of efficiency, efficacy, and business agility in their current form. It will be difficult to scale and automate processes to manage users in a dynamic environment—both users and applications in the cloud—given the sheer volume of dynamic cloud compute resources (compute nodes, storage, and network policies) and the number of users and services accessing those resources.

The issue will be made worse by the enterprise's older IAM solutions. IAM architecture and solutions are costly to expand to cloud services and need substantial customisation in their present state. Reliable sources of identities in the cloud continue to be a problem that requires attention. On the other hand, CSP support for IAM standards and practises is patchy and insufficient for the majority of businesses. Many SaaS cloud services are beginning to accept federation standards like SAML, although PaaS and IaaS providers mostly ignore them.

A limited number of CSPs (primarily significant SaaS providers like Salesforce.com, Google, and Microsoft) are starting to take business IAM needs seriously, including support for standards like SAML that enable SSO utilising identity federation methods. Nevertheless, from an enterprise standpoint, the IAM capabilities are at best basic given the major organisations' early adoption cycles. Clients should keep pushing for IAM capabilities from their CSPs, such as support for SAML, SPML-based user provisioning, XACML-based authorisation, and an open API to enable different user and access automation procedures.

A new generation of cloud-based identity management services that relocate your identity trust boundary outside of your perimeter and into the cloud have been created in response to this IAM capabilities gap.

Microsoft's Azure STS and other identity services and frameworks enable user SSO from an on-premises Active Directory to Microsoft's cloud services and provide federation between Active Directory and those services. In addition, a variety of SSO access control, use monitoring, and centralised management strategies are being offered by start-ups including Symplified, Ping Identity, and TriCipher. While these cloud-based identity management systems are making it easier for small and medium-sized companies (SMBs) to get started,

certain corporations may still find them to be insufficient for meeting strict requirements for bespoke reporting and compliance monitoring.

As most businesses are unwilling to keep their trusted sources of identification outside of their regulated organisational boundaries, trusting cloud service providers and user data management are further entrance obstacles. Use scenarios where attribute data related to identities is either duplicated or stored in the cloud service further aggravate this problem. Enterprises continue to face significant difficulties in synchronising several identity stores. These obstacles will be lessened by using cloud-based services and tackling synchronisation problems via federation, virtual directories, and an open API.

Organizations wanting to implement cloud-based services should include the IAM strategy into their cloud services strategy road plan to prevent expensive retrofits and integration with after-market solutions. Hence, using an optimised internal IAM strategy and practise in the cloud should benefit organisations who have been investing in directories, IAM capabilities, and IAM processes. The presence of a strong directory and federated identity management capability within the organisation (internal or cloud-based identity management service) is crucial for an enterprise to successfully manage identities and access control in the cloud. This capability includes architecture and systems, user and access life cycle management processes, and audit and compliance capabilities.

Organizations must focus on usability and simplicity in addition to risk-based authentication techniques when authenticating users and services to the cloud (e.g., log when sensitive data is accessed). All clouds are not created equal, therefore businesses must have a plan for using risk-based IAM techniques like strong authentication, automated provisioning, deprovisioning, auditing, and monitoring to handle risks particular to a CSP.

With a well-architected IAM architecture and IT procedures, identity and authentication problems can be solved (when the CSP makes such capabilities accessible); nevertheless, authorization services in the cloud are still relatively primitive and developing.

Cloud service users should be aware that granular application authorisation is still in its infancy. When it does, it is often done utilising the CSP's exclusive profiles and simple roles; these roles frequently include "user" and "administrator." Customers should be lobbying for more help as a long-term strategy even if XACML hasn't been used internally, of XACML-compliant entitlement management on the side of cloud service providers. Any applications that implement a single authorization standard may impose access control and policy using a defined language and technique thanks to XACML. Chief information security officers (CISOs) should, at the very least, consider authorization criteria and resist the urge to tailor solutions to their provider's capabilities.

The enterprise's business and IT stakeholders should promote the standardisation of roles that are linked to higher-level business functions at a coarse granular level (e.g., accounts payable manager, HR manager, or purchase order approver). In certain companies, restricted usage of the role characteristic may be helpful for access control, while it is not practical in others. Nevertheless, rule-based access control (which makes use of several features, groups, etc.) is more adaptable since it enables rapid responses to alterations in your organisation.

Future cloud service roles or profiles supplied by CSPs should be aligned to clearly define corporate roles. We anticipate that SPML and XACML will contribute in this area. As of right now, we are not aware of any initiatives to standardise enterprise role naming practises.

While this may speed up the adoption of cloud-based services like cloud identity management services, policy-based authentication, centralised logging, and auditing, IT architects should promote externalisation of authentication and authorization components from applications (loosely connected). The Geneva claims-based authentication framework from Microsoft and OpenSSO from Sun Microsystems, for instance, may aid in externalising authentication.

organisations with less developed IAM procedures, capabilities, and infrastructure should work to standardise IAM features across all of their applications, whether they are cloud-based or not. Easy administration should be a key advantage of centralised management, enabling businesses to embrace the cloud more quickly. Moreover, self-service options, password management, auditing, and reporting tools may all assist increase productivity and effectiveness in meeting compliance requirements. Attributes may be used more rapidly and with common security models using standard identity repositories for cloud apps.

Companies considering cloud services should include CSP support for identity and access management, including support for federation, in the assessment criteria in addition to the basic CSP service features. IAM capabilities like identity federation and SSO are currently only weakly and inconsistently supported by CSPs. Consumers will need to consider each CSP's identity integration option individually. If at all possible, customers should resist the urge to tailor their user access management solution to a particular CSP since doing so would ultimately decrease the effectiveness of the IAM process and raise administration expenses. IAM process standardisation will reduce risks associated with illegal access to cloud services and profit from unified user and access management across CSPs (different clouds).

CHAPTER 15

SECURITY MANAGEMENT IN THE CLOUD

Karthikeyan Palniswamy

Assistant Professor, Department of Computer Science and Information Technology,
Jain (Deemed to be University) Bangalore, Karnataka, India
Email Id-karthikeyan.mp@jainuniversity.ac.in

A large part of your network, system, applications, and data will move under third-party provider control with the adoption of public cloud services. In addition to creating islands (clouds) of virtual perimeters and a security architecture with shared responsibility between the client and the cloud service provider, the cloud services delivery methodology (CSP). The organization's IT operations team will face additional security management problems as a result of this shared responsibility paradigm. In light of this, the first query a chief information security officer (CISO) must address is whether she has sufficient visibility from cloud services to manage the governance (shared responsibilities) and implementation of security management processes (preventive and detective controls) to reassure the company that the data in the cloud is adequately protected. In order to properly manage security in the cloud, an enterprise's security management tools and procedures must be adjusted. The solution to this issue consists of two parts: what security controls the client must give in addition to those built into the cloud platform. Based on the sensitivity of the data and changing service levels, both responses need to be reviewed often.

You should begin your exercise in understanding the trust boundary of your cloud services as a cloud client. You should be familiar with every layer of the cloud service you use, touch, or own, including the network, host, application, database, storage, and web services, including identity services (see Figure 15.1). Also, you must be aware of the extent of your IT system administration and monitoring duties, which include managing access, modification, configuration, patches, and vulnerabilities.

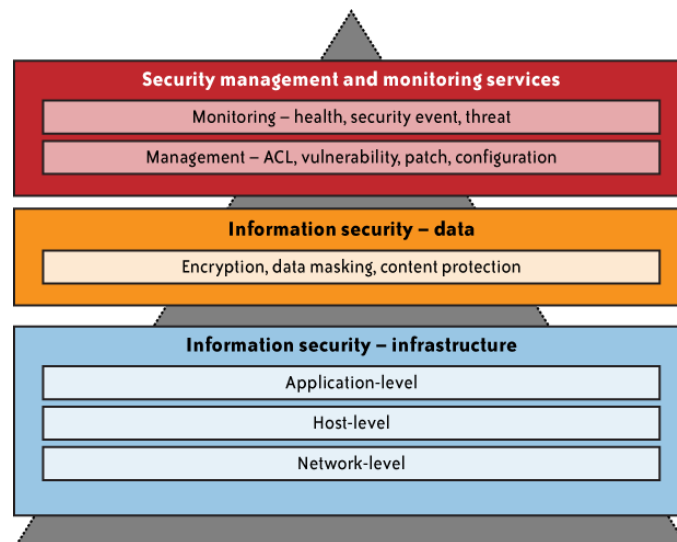


Figure 15.1. Security management and monitoring scope

Even though you might be handing over some operational duties to the provider, the extent of your operational obligations will differ and depend on a number of variables, such as the service delivery model (SPI), provider service-level agreement (SLA), and provider-specific capabilities to support the expansion of your internal security management processes and tools.

It is common knowledge that established IT firms use security management frameworks like ISO/IEC 27000 and the ITIL service management framework. These widely accepted management frameworks provide instructions for organising and putting in place a governance programme with ongoing management procedures that safeguard information assets. For instance, ITIL provides a thorough explanation of many crucial IT practises with precise checklists, tasks, and processes that can be customised for each IT firm. Cloud computing is compatible with a fundamental ITIL principle that states that organisations (people, processes, and information systems) change over time. In order to align and realign IT services to changing business demands, management frameworks like ITIL will aid in continual service improvement. Finding and applying changes to the IT services that support business operations, such as sales force automation via a cloud service provider, is known as continuous service improvement. The actions involved within the security management procedures must be continuously updated to be current and efficient given the dynamic nature of cloud computing services.

The cloud security management will be heavily reliant on security management, which is a continuous activity. The ITIL Security Management framework has two main objectives:

Understanding security needs

The SLA and other external standards which are outlined in supporting contracts, laws, and internally or externally enforced policies—generally establish security needs.

Realizing a fundamental degree of security

This is required to ensure the organization's security and continuity as well as to achieve a more straightforward service-level management for information security management.

The objective of well-established security management procedures is to safeguard the confidentiality, integrity, and availability of information. These processes are also in line with an organization's IT policies and standards. A business's ITIL life cycle is shown in Figure 15-2. Relevant ISO and ITIL functions reflect security management disciplines.

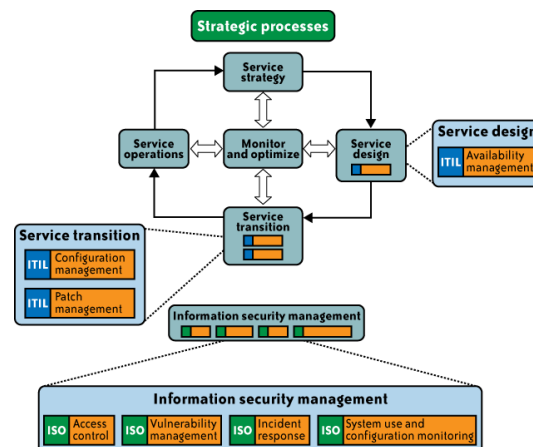


Figure 15.2: The enterprise's ITIL life cycle

So, how do cloud services implement security management? What security management tasks might clients anticipate CSPs performing? In the context of the SPI (SaaS, PaaS, IaaS) service delivery paradigm, er addresses the situation of security management in public clouds, your duty scope, and recommendations.

Standards for Security Management

The standards that are pertinent to security management practises in the cloud, according to the authors' judgement, include ITIL and ISO/IEC 27001 and 27002.

ITIL

A collection of best practises and recommendations known as the Information Technology Infrastructure Library (ITIL) provide an integrated, process-based approach to managing information technology services. Almost all types of IT environments, including cloud operating environments, may use ITIL. Effective information security measures should be implemented at the strategic, tactical, and operational levels, according to ITIL. Information security is seen as an ongoing activity that has to be managed, planned, carried out, assessed, and maintained.

Information security is divided up into:

Policies

The broad goals that an organisation is aiming to accomplish.

Processes

What must to place in order to accomplish the goals?

Procedures

Who does what task when to meet the goals?

Work directives

Directions for doing certain tasks.

The information security management code of practise, or ISO/IEC 17799:2005, is the foundation for the ITIL-process security management. Almost every other ITIL procedure is connected to the security management process. Yet, as they have a significant impact on the system's degree of security, the linkages with the service-level management process, incident management process, and change management process will be the most visible (server, network, or application). Given that ISO/IEC 20000 was the first worldwide standard for IT Service Management, ITIL is also tied to that standard (ITSM). It is based on the previous British standard, BS 15000, and is intended to replace it.

It is not possible to certify organisations or management systems as "ITIL-compliant." Nonetheless, a company that has integrated ITIL guidelines into ITSM is still able to comply with ISO/IEC 20000 and apply for certification under that standard.

The essential specifications for an Information Security Management System are clearly stated in ISO/IEC 27001 (ISMS). It also serves as a certification standard and identifies appropriate information security controls inside the ISMS using ISO/IEC 27002. Organizations are able to choose and apply controls as they see appropriate because ISO/IEC 27002 is just a code of practice/guideline and not a certification standard.

Information security professionals agree that the ITIL security management best practises should be revised in order to strengthen the application and logical security in the Information and Communication Technology (ICT) infrastructure domain given the current trend of organisations moving towards ISO/IEC 27001 for information security management.

In essence, the frameworks for ITIL, ISO/IEC 20000, and ISO/IEC 27001/27002 aid IT firms in internalising and responding to fundamental queries like:

- a. How can I verify that the present security levels are enough for your needs?
- b. How can I establish a baseline for security throughout your operation?

In doing so, they assist you in finding an answer to the query: How can I make sure that my services are secure?

Managing security in the cloud

We (the authors) picked the following pertinent processes as the advised security management emphasis areas for safeguarding cloud services after examining the management process disciplines across the ITIL and ISO frameworks:

1. Availability control (ITIL)
2. Vulnerability management (ISO/IEC 27002)
3. Access control (ISO/IEC 27002, ITIL)

Patch management, configuration management, incident response, and system usage and access monitoring (ISO/IEC 27002) are all examples of ITIL-compliant practices and their potential to reduce overall risk to the enterprise, we chose these security management methods. The subset of ITIL management domains that have the greatest influence on organizations in managing security and operational risk. Other ITIL management domains, such as problem management and service continuity management, may be more pertinent to your company in the context of security management, we'll talk about the procedures for security management that apply to cloud services. Also, we have made an effort to show how current cloud service deployment and delivery models assist security management procedures (private, public, and hybrid). This is undoubtedly a developing field, therefore we advise you to routinely review cloud service capabilities and adjust your security management procedures as necessary.

In the context of deployment models, the applicability of key security management tasks that are accessible to you for each SPI cloud delivery type (private and public). The chart shows that security management technique is applicable to all delivery and deployment modes. Your cloud security operations model has to take these tasks into account. As a result, businesses wishing to expand their usage of the public cloud for specific use cases may apply and expand internal security management procedures already established for internal private cloud services.

Availability Control

Cloud services are not impervious to outages, and depending on the circumstances of the outage, the extent and intensity of the damage on the client might change. The criticality of the cloud application and its connection to internal business operations will determine the business effect of a service outage, just as it does for any internal IT-supported application. Even a brief service outage may have a significant effect on your company's productivity, revenue, customer happiness, and service-level compliance in the case of business-critical applications where organizations depend on the continuous availability of service.

Major CSPs have experienced downtime ranging from a few minutes to a few hours, according to the Cloud Computing Incidents Database (CCID), which analyses cloud service disruptions. A service outage once went on for more than 24 hours! Moreover, disruptions may impact all or a subset of customers depending on the severity of the event and the extent of the impacted infrastructure. Affected customers will not be able to use the cloud service during an interruption, and in certain situations, their performance or user experience may be negatively impacted. For instance, when a storage service is interrupted, it will have an impact on the availability and functionality of a computer service that relies on it.

Elements Affecting Availability

The CSP's data centre architecture (load balancers, networks, systems), application architecture, hosting location redundancy, diversity of Internet service providers (ISPs), and data storage architecture are just a few of the variables that affect the resilience and availability of the cloud services. The following is a summary of the key elements:

- a. Redundancy in SaaS and PaaS application architecture.
- b. The data centre architecture for cloud services, as well as network and system architecture, including fault-tolerance and geographically diversified design.
- c. The consistency and redundancy of the customer's and CSP's Internet connection.
- d. The capacity of the customer to react swiftly and rely on internal apps and other processes, including manual ones.
- e. The fault's visibility to the customer. It may be difficult to understand the entire effect of certain downtime events if the impact only impacts a small fraction of users, which may make troubleshooting more challenging.
- f. The dependability of the software and hardware utilised to supply the cloud service.
- g. The ability of the network and security architecture to protect the cloud service against a distributed denial of service (DDoS) assault.
- h. The effectiveness of security procedures and controls that reduce human error and safeguard infrastructure from hostile internal and external threats, such as privileged users abusing their position of power.

Availability Management in SaaS

SaaS service providers are accountable for business continuity, application, and infrastructure security management practises due to the nature of the service delivery and business model. This indicates that the CSP will now manage the duties that your IT group previously performed. When they attempt to translate internal service-level categories to a CSP, certain established businesses that are in alignment with industry standards, such as ITIL, may confront new issues with governance of SaaS services. As an example, if a marketing application is important and has a high

Service-level requirement: Based on the SaaS provider's SLA, how can the IT department or business unit satisfy the internal marketing department's availability expectation? SaaS companies may choose to deal with service terms via terms and conditions rather than offering SLAs. Salesforce.com, for instance, does not provide a comprehensive SLA that outlines and details performance standards and service obligations. Nevertheless, NetSuite, a different CRM SaaS provider, provides the following SLA terms:

Uptime Objective: NetSuite guarantees a 99.5% uptime for the NetSuite application, excluding routine maintenance windows.

Both planned and unscheduled maintenance are excluded from the downtime calculation. If a maintenance window is announced at least two complete business days in advance, it is considered routinely planned. Regularly planned maintenance time normally takes less than 10-15 hours per quarter, is notified at least a week in advance, and is scheduled to happen at night on the weekend.

NetSuite thus notifies you that every Saturday night from 10:00 p.m. to 10:20 p.m. Pacific Time is set aside for necessary regular maintenance.

Another SLA example is as follows:

The Google Apps Covered Services web interface shall be operational and accessible to Customer at least 99.9% of the time in any calendar month during the Term of the relevant Google Apps Agreement (the "Google Apps SLA"). Customer will be qualified to obtain the Service Credits detailed below if Google does not adhere to the Google Apps SLA and Customer complies with its duties under this Google Apps SLA. The Customer's Exclusive and Exclusive Remedy for Any Failure by Google to Provide the Service is outlined in this Google Apps SLA. Client Must Seek Credit for Services. Within thirty days of the moment the customer becomes qualified to obtain a service credit, the customer must inform Google in order to get any of the service credits mentioned above. Customer's eligibility for a Service Credit will be lost if this condition is not met.

Service Credit maximum. A total of fifteen days of service, added to the end of the customer's term for the service, may not be granted by Google to a customer for any and all Downtime Periods that happen in a single calendar month. Service Credits cannot be converted to cash or traded for cash. Exclusions from the Google Apps SLA. The Google Apps SLA does not apply to any services that explicitly state in their documentation that they are not covered by it, nor does it cover performance issues that are either (i) brought on by circumstances beyond Google's reasonable control, (ii) caused by the customer's equipment, third-party equipment, or both, or (iii) both (not within the primary control of Google).

Cloud service providers do not use standardised service level agreements (SLAs). Each provider will have a different uptime guarantee, service credits, and service exclusions clause.

Consumer Obligation

To learn about service disruptions, customers should understand the SLA and available communication channels, such as email, RSS feeds, and website URLs. Customers should utilise automatic tools like Nagios or Siteuptime.com whenever they can to check the SaaS service's accessibility.

Customers of a SaaS service currently only have a few alternatives to enable availability management. So, when a disruption occurs, consumers should try to understand the availability management aspects, including the service's SLA, and explain any gaps with the CSP about SLA exclusions and service credits. The effectiveness of SaaS SLAs was examined in the context of software suppliers switching to a SaaS delivery model in a recent white paper by the American Software & Information Industry Association (SIIA). According to the paper's findings, the SLA must contain a number of essential components in order to be effective. These components include: clear expectations and communication between the service provider and its clients in order to determine what is important and reasonable in terms of standards and expectations.

Consumers of cloud services should be aware that a "one size fits all" operational philosophy is frequently used when designing multitenant service delivery models, which implies CSPs typically provide a single SLA for all clients. This means that if the normal SLA does not satisfy your service-level needs, CSPs may not be willing to provide bespoke SLAs. A bespoke SLA may still be possible if you are a medium or big business with a substantial budget, however.

Customers should be aware of how resource democratisation works inside the CSP as the majority of SaaS providers employ virtualization technology to offer a multitenant service. This will help them estimate the possibility of system availability and performance during business volatility. It is possible that a very demanding tenant may starve other tenants if the resources (network, CPU, memory, and storage) are not distributed fairly among the tenants to fulfil the task. This may lead to reduced service levels or a bad user experience.

Health SaaS Monitoring

Customers have the following alternatives for monitoring the status of their service:

1. The CSP's dashboard on service health. On their websites, SaaS providers like Salesforce.com often post information about the service's current condition, any outages that may have an effect on users, and impending planned maintenance services.
2. The Database of Cloud Computing Incidents (CCID). (This database is often community-supported; it may not accurately represent all CSPs and all occurrences.)
3. A mailing list of consumers that informs them of current and previous disruptions.
4. Internal or external service monitoring solutions that systematically assess the state of SaaS providers and notify clients when services are disrupted (e.g., Nagios monitoring tool).
5. An RSS feed that is hosted by the SaaS provider.

Availability Management in PaaS

Customers (developers) construct and deploy PaaS applications using the PaaS platform provided by the CSP in a typical PaaS service. The PaaS platform is generally constructed using a network, servers, operating systems, storage infrastructure, and application components that are owned and controlled by the CSP (web services). The availability management of customer PaaS applications can be challenging because they are built using CSP-supplied application components and, occasionally, third-party web services components (mash-up applications). Take, for instance, a social network application on Google App Engine that depends on a Facebook application for a contact management service. The responsibility for availability management is split between the client and the CSP in that muddled software deployment architecture. Although the PaaS CSP is in charge of the PaaS platform and any other services it provides, the client is in charge of controlling the availability of applications they have built and services from third parties. Customers are in charge of managing the apps created and deployed on the AppExchange platform, for instance, while Force.com is in charge of administering the platform itself.

Understanding your application's dependence on third-party services, including those provided by the PaaS vendor, is crucial since PaaS applications may be built to rely on additional third-party web services components that are not included in the PaaS service offerings (e.g., your web 2.0 application using Google Maps for geo mapping). Your application may be dependent on the availability of these service components, such as the message queue service, identity and authentication service, and database service that PaaS

providers may also supply (an example is Google's BigTable). As a result, the reliability of your PaaS application, the PaaS platform on which it was created, and third-party web services components all play a role in its availability.

Clients are urged to examine and comprehend the PaaS platform service levels, including any quota triggers that can restrict the amount of resources available for their application (usually outlined in the SLA, or in the terms and conditions of the PaaS service). In situations where the PaaS platform sets compute resource limits (CPU, memory, network I/O), the application may not be able to react within the usual latency expectations and may finally become inaccessible if the thresholds are reached. For instance, the Google App Engine uses a quota system in which each App Engine resource is compared to either a fixed or a billable quota.

Billable quotas are resource ceilings that you, the program's administrator, have established to keep the cost of the application from going above your predetermined spending limit. Every application is given a free portion of each billable quota. By enabling billing, establishing a daily budget, and finally assigning the budget to the quotas, you may enhance the billable limits for your application. Only resources that your programme really utilises will be paid for, and only those that are utilised in excess of the free limit.

Fixed quotas are upper limits on resources defined by App Engine to protect the stability of the system. All applications are required to operate within the restrictions set forth by these resources, which specify the bounds of the architecture. They make sure that the performance of your app won't be impacted by another programme that is using excessive amounts of resources.

The Apex governor function of Force.com is another example. The Apex runtime engine carefully imposes a variety of restrictions since the Apex application operates in a multitenant environment to prevent runaway scripts from monopolising common resources. Based on a common policy with consumers, governors keep track of and enforce the limitations. The connected governor produces a runtime exception that cannot be handled if a script ever goes above a limit.

Consumer Obligation: The PaaS application customer should carefully examine the dependencies of the application on the third-party web services (components), taking into account all the variable parameters in availability management, and then outline a comprehensive management strategy to manage and monitor all the dependencies.

For PaaS clients, the following matters:

Platform as a Service service levels

Consumers should thoroughly read the SLAs of the CSP and comprehend the limitations on availability.

PaaS Health Inspection

Applications hosted on the PaaS CSP platform are often web-based in nature (e.g., your Java or Python application hosted on the Google App Engine). Because of this, most methods and procedures used to monitor SaaS applications also work with PaaS apps. Customers should keep an eye on both their application and the services provided by other parties that make up PaaS apps. Understanding the web services protocol (HTTP, HTTPS) and the necessary protocol parameters (e.g., URI) to validate the service's availability) is necessary when configuring your management tools to monitor the health of web services.

Monitoring your application may involve a standard web services protocol, such as Representational State Transfer (REST), Simple Object Access Protocol (SOAP), eXtensible Markup Language/ Hypertext Transfer Protocol (XML/HTTP), and in some circumstances, proprietary protocols, when CSPs support monitoring via application programming interfaces (APIs).

Customers may choose from the alternatives listed below to check on the status of their service:

The CSP's Service Health Dashboard, available at <http://status.zoho.com>, for instance.

CCID (this database is primarily community-supported, and may not include all CSPs and all events that have happened) (this database is generally community-supported, and may not reflect all CSPs and all incidents that have occurred).

Availability Management for IaaS

A compute and storage (persistent and ephemeral) infrastructure in the cloud should both be taken into account when determining the availability of the IaaS delivery model. Moreover, IaaS providers could provide additional services like account administration, message queue, identification and authentication, database, invoicing, and monitoring. So, availability management has to take into account all the services that you rely on for your business and IT requirements. Consumers are in charge of all availability-related issues, as they are in charge of supplying and maintaining the virtual server life cycle.

Five things affect how you manage your IaaS virtual infrastructure in the cloud:

1. The accessibility of an infrastructure for the CSP network, host, storage, and support applications. This variable is determined by:
2. CSP data centre design, which includes a fault-tolerant and geographically varied architecture.
3. The customer's and the CSP's Internet access is dependable, diverse, and redundant.
4. The architecture for ensuring the dependability and redundancy of the hardware and software parts required to supply computing and storage services.
5. The processes and procedures for availability management, including those implemented by the CSP for business continuity.
6. Accessibility of web console or API services. To control the virtual server life cycle, the web console and API are necessary. Customers are unable to setup, start, stop, and deprovision virtual servers while such services are unavailable.
7. SLA. This element differs amongst CSPs, thus the SLA should be examined and reconciled, taking exclusion clauses into account.
8. The accessibility of your virtual servers as well as the permanent and ephemeral storage that is associated for compute services (such Amazon Web Services' S3 and Amazon Elastic Block Storage, for example).

The accessibility of the virtual storage on which your users and virtual server rely for storage. This covers use cases for both synchronous and asynchronous storage access. Asynchronous use cases are more forgiving of latency and availability, while synchronous storage access use cases require minimal data access latency and constant availability. User authentication, video streaming, and database transactions are a few examples of synchronous storage use cases. Inconsistency or interruptions to storage in synchronous storage has a significant influence on overall server and application availability. A prominent example of an

asynchronous use case is a cloud-based storage solution for backing up your computer via the Internet.

Availability of your network connectivity to the Internet or virtual network connectivity to IaaS services. In some cases, this can involve virtual private network (VPN) connectivity between your internal private data centre and the public IaaS cloud (e.g., hybrid clouds) (e.g., hybrid clouds). Availability of network services, including a DNS, routing services, and authentication services required to connect to the IaaS service.

IaaS Health Monitoring

The following options are available to IaaS customers for managing the health of their service:

The CSP's dashboard on service health.

CCID (this database is generally community-supported, and may not reflect all CSPs and all incidents that have occurred) (this database is generally community-supported, and may not reflect all CSPs and all incidents that have occurred).

CSP customer mailing list that notifies customers of occurring and recently occurred outages.

Internal or third-party-based service monitoring tools (e.g., Nagios) that periodically check the health of your IaaS virtual server. For example, Amazon Web Services (AWS) is offering a cloud monitoring service called CloudWatch. This web service provides monitoring for AWS cloud resources, including Amazon's Elastic Compute Cloud (EC2) (EC2). It also provides customers with visibility into resource utilisation, operational performance, and overall demand patterns, including metrics such as CPU utilisation, disc reads and writes, and network traffic.

Web console or API that publishes the current health status of your virtual servers and network.

Similar to SaaS service monitoring, customers who are hosting applications on an IaaS platform should take additional steps to monitor the health of the hosted application. For example, if you are hosting an e-commerce application on your Amazon EC2 virtual cloud, you should monitor the health of both the e-commerce application and the virtual server instances.

CHAPTER 16

ACCESS CONTROL

Raghavendra R.

Assistant Professor, Department of Computer Science and Information Technology,
Jain (Deemed to be University) Bangalore, Karnataka, India
Email Id-r.raghavendra@jainuniversity.ac.in

Access needs for your users and system administrators (privileged users) who access network, system, and application resources are often covered by access control management, which is a comprehensive role. The following issues should be handled by the access control management functions: organization's access rules and standards should include the aforementioned facets of the access control domain and be in line with the roles and responsibilities of all users, including privileged system administrators and end users.

The Cloud Access Control System

Network access control will become less important in a consumption model for cloud computing because customers may access cloud services from any host that is connected to the Internet. The rationale is that conventional network-based access controls concentrate on preventing illegal access to resources using host-based characteristics, which are often insufficient, are not consistent between users, and may result in erroneous accounting. Network access control in the cloud takes the form of cloud firewall rules that enforce host-based access control at the points of entrance and exit into the cloud as well as the logical grouping of instances inside the cloud. This is often accomplished via protocols (rules) that make use of basic TCP/IP parameters, such as source IP, source port, destination IP, and destination port.

User access control should be prioritised above network-based access control in the cloud because it can tightly link a person's identity to the resources there and supports fine-grained access control, user accounting, compliance assistance, and data security. The security and integrity of your data in the cloud are significantly protected by user access management rules, such as robust authentication, single sign-on (SSO), privilege management, and logging and monitoring of cloud resources.

Six access control goals are outlined in ISO/IEC 27002, including end user, privileged user, and network, application, and information access control. We advise readers to evaluate cloud services and comprehend the pertinent ISO/IEC 27002 control goals that reduce risk to the organisation the most. With relation to cloud services, the following user access management control statement from ISO 27002 is very important:

The goal is to protect information systems from illegal access and to assure authorised user access. To regulate the distribution of access rights to information systems and services, formal mechanisms need be in place.

All phases of the user access lifecycle, from initial registration of new users to ultimate de-registration of users who no longer need access to information systems and services, should be covered by the protocols. When necessary, special consideration should be given to the need to regulate the distribution of privileged access privileges, which enable users to bypass system restrictions.

The six control statements are as follows:

- a. Limit who has access to what data.
- b. Control user access privileges.
- c. Promote sensible access procedures.
- d. Limit who has access to network resources.

Control access to operating systems, as well as to software and hardware.

ITIL mandates an access management function that was included as a new procedure to ITIL v3, much like ISO 27002. IT security was the driving force behind the choice to incorporate this specific procedure.

For the following reasons, it should be of utmost significance that only authorised users be allowed access to IT services and applications from a security standpoint. This function's goal is to allow authorised users to use a service while restricting access for unauthorised users. The procedures for access management fundamentally carry out IT security management rules.

SaaS access control

The CSP is in charge of overseeing every element of the network, server, and application infrastructure in the SaaS delivery paradigm. Under that architecture, network-based restrictions are supplemented or replaced by user access controls, such as one-time password authentication, as the programme is provided to end users as a service, often via a web browser. To safeguard the data stored by SaaS, clients should concentrate on user access controls (authentication, federation, privilege management, deprovisioning, etc.). Some SaaS systems, like Salesforce.com, combine user access control with network access control (e.g., source IP address/network-based control), giving clients the choice of enforcing access based on user and network policy criteria.

Providers' support for user access control varies, as do their capabilities. A limited number of CSPs (primarily significant SaaS providers like Salesforce.com, Google, and Microsoft) are starting to take business IAM needs seriously, including support for standards like SAML that enable SSO utilising identity federation methods. Nevertheless, from an enterprise standpoint, the IAM capabilities are at best basic given the major organisations' early adoption cycles. Clients should keep pressing their CSPs to include IAM functionality, such as SAML support, SPML-based user provisioning, and an open API to accommodate different user and access automation procedures. To facilitate user access management and federation, organisations should make use of their current identity management techniques, procedures, and architecture (such as IdP).

Access Management: PaaS

The CSP is in charge of overseeing access control to the network, servers, and application platform infrastructure in the PaaS delivery model. Nevertheless, access control to the apps installed on a PaaS platform is the customer's responsibility. End user access management, which includes user provisioning and authentication, appears as access control to apps.

Providers' support for user access control varies, as do their capabilities. Most PaaS providers provide basic user access control capability as of this writing, with the exception of Force.com and Microsoft Azure (currently in testing). Businesses that use their internal identity provider (IdP) must be aware of PaaS features, such as federation support. It is possible for a PaaS CSP to provide an industry-standard API like OAuth to regulate

application access and handle authentication. To improve usability, Google, for instance, implements a hybrid OpenID and OAuth protocol that integrates the permission and authentication processes in fewer stages. If the CSP supports federation standards like the Security Assertion Markup Language, you might potentially delegate authentication to your IdP. (SAML) pertaining to information on identity and access management in a PaaS delivery paradigm.

Access Control: IaaS IaaS clients are solely in charge of controlling all facets of access to their cloud-based services. The client will be responsible for designing and managing access to the virtual servers, virtual network, virtual storage, and applications hosted on an IaaS platform. Access control management fits into one of the following two groups in an IaaS delivery model:

Control of CSP infrastructure access

Access control administration for the host, network, and management software that the CSP owns and maintains

Customer control of access to virtual infrastructure

Virtual storage, virtual networks, and applications housed on virtual servers are all within your control when it comes to access control management.

Control of CSP infrastructure access: The CSP is in charge of controlling access to the network utilised by administrators to carry out their duties. Access control for administrative tasks like backups, network and host (hypervisor) maintenance, router and firewall policy management, and system monitoring and administration are all included in this. Role-based access control and strong authentication should be used to secure access to administrative operations. The implementation of robust operational processes should enable the granting and removal of administrative credentials. To verify least privileges and division of tasks, periodic access control audits and administrative user certifications should be put into place. The core idea of least privilege serves as a framework for AWS's information security policies, according to the aforementioned AWS security white paper. By demanding that no person, programme, or system be given greater access rights than are required to complete the job, least privilege safeguards customer information assets. Any employee who is discovered to have violated this rule may face punishment, up to and including termination.

Access control for customer virtual infrastructure

IaaS users must first comprehend the virtual resources (network, host, firewall, load balancers, management panel, etc.) and the available security measures to prevent unauthorised access.

Only authorized users may access. Customers often get complete root access and administrative control over leased virtual servers from CSPs. Customers may also be given the ability to control network access restrictions for both the entry and exit of their virtual servers and network. As a result, it is the customer's responsibility to take all required precautions to secure access to virtual resources.

IaaS CSPs often provide APIs (REST, SOAP, or HTTP with XML/JavaScript Object Notation [JSON]) to enable users to execute the majority of administration tasks, such access control, from a distance. Some vendors additionally provide a web-based panel where users may activate access control features. The design and implementation of access management

procedures with access request or approval and a gatekeeper, as well as the maintenance of a list of privileged users with access to IaaS resources, should be done by organisations using IaaS services.

While managing access control for your cloud infrastructure, keep the following things in mind:

Network access management

Ask the provider what the default settings are for the firewall that the CSP normally uses to enforce network access. By default (factory settings), CSPs often block all access to your virtual servers, which prevents any incoming traffic from reaching them. In order to access your cloud-based virtual servers, you must add new rules explicitly. For instance, you might allow access to IP 10.0.0.1 from 192.168.0.1 to port 22 (Secure Shell, or SSH), where 10.0.0.1 is the IP address of the virtual server and 192.168.0.1 is the trusted IP address from which 10.0.0.1 can be accessed using SSH. In order to implement various ingress restrictions as necessary, Amazon EC2 provides network group capabilities that enable the construction of several security groups. According to Amazon, a client may limit traffic to each EC2 instance by protocol, service port, or originating IP address and manage each security group using a PEM-encoded X.509 certificate.

Control of virtual server access: Access restrictions, such as operating system authentication techniques, should be used to secure virtual servers running your choice OS (Linux, Solaris, or Windows). Configuring Unix servers with SSH-based logins and robust authentication is common practise. Robust authentication guards against a variety of security risks (e.g., IP spoofing, fake routes, man-in-the-middle, and DNS spoofing). The authentication techniques include Kerberos authentication, pure RSA authentication, one-time passwords using S/Key, and host authentication based on the Rivest-Shamir-Adleman (RSA) encryption algorithm. While employing RSA keys, it is advised to keep the keys on a secure medium and to protect them using a password. These precautions aid in preventing unwanted access to your keys.

Station for cloud management: The majority of the time, client apps that use a CSP-exclusive API (REST, SOAP, or HTTP with XML/JSON) to control distant resources are used to manage your cloud-based virtual resources. Installed on the management station is a client management toolkit (provided by the CSP), which communicates with the CSP management service through the public API. The cloud management station should be thought of as a command and control centre for the cloud infrastructure since it houses sensitive data such as host and user keys and firewall rules. Thus, access to the management station should be secured using reliable access provisioning practises and rigorous authentication.

Online console: Some CSPs add a web-based console capability to the cloud management station so that clients may control access to their virtual infrastructure. The console provides an alternate method for administering the cloud infrastructure to the cloud management station. The console, which serves as a management station for your cloud infrastructure, provides quick access to private data similar to the management station, including your host keys and firewall settings. You should sufficiently secure console access since the web console is a strong tool that can manage your virtual network and virtual server instances. For instance, only SSL protocol should be used to access the web console.

Access Control Overview: When it comes to traditional deployment methods and the SPI (SaaS, PaaS, and IaaS) cloud delivery paradigm, access control is a crucial security management role (public, private, and hybrid). In the absence of encryption and other data protections, access management may be the main method of security control for your data stored in SPI clouds. Access control capabilities in public clouds are still developing and inconsistent as of this writing. Due to the following factors, access control capabilities provided by CSPs may not be sufficient for corporate clients in their present state:

Access control procedures, procedures, and practises vary across CSPs. Customers must put in additional effort to comprehend CSP-specific access control capabilities and tailor them on a CSP basis in order to manage access control to their virtual cloud architecture.

It is very challenging to control access across different clouds due to the absence of a common API across CSPs. For instance, none of the major CSPs, including Amazon, provide SAML support.

The majority of cloud resource access constraints are lax. Granular network-level access management is often supported by access controls from CSPs, but coarse user access management is not. User access restrictions are primarily concerned with the authentication components.

Handling user permission to the cloud infrastructure is at best basic. The concepts of least privilege and division of tasks should be supported by CSPs by roles that provide granular privilege access (e.g., console manager, network access manager, zone manager, host manager).

Access management is a crucial security procedure to maintain the confidentiality, integrity, and availability (CIA) of data stored in the cloud from the viewpoint of a corporate client. Procedures for provisioning, prompt deprovisioning, flexible authentication, privilege management, accounting, auditing, and assistance for compliance management should all be included in a complete access management programme. The access control features for networks, systems, and applications that are special to CSPs should be understood by cloud clients, and access should be managed accordingly.

Patch, Configuration, and Security Vulnerability Management

Cloud services continue to face a serious danger from malware's (or a hacker's) ability to remotely attack apps, network services, and infrastructure weaknesses. A public PaaS and IaaS delivery architecture where the client is still responsible for managing vulnerabilities, patches, and configuration poses an even higher risk. Clients should keep in mind that all tenants in a multitenant virtual environment have the lowest or highest common denominator of security in cloud computing environments. So, it is the responsibility of the clients to comprehend the extent of their security management duties.

To assist customers in comprehending and preparing for complementary security management duties, consumers should insist that CSPs increase transparency about their cloud security operations.

Generally speaking, CSPs are in charge of managing the infrastructure (networks, hosts, apps, and storage) that is maintained and administered by CSPs as well as any third-party services that they may use. This is known as vulnerability, patch, and configuration (VPC) management. Nonetheless, consumers are not exempted from their VPC tasks and should grasp the VPC components for which they are liable. End-to-end security concerns and customer-managed systems and applications that interact with cloud services should be

covered by a VPC management scope. CSPs may have implemented these programmes as regular procedure inside their security management domain, but often the procedure is internal to the CSP and is not made clear to clients. CSPs should use ISO/IEC 27002-style control and assurance frameworks to reassure their clients about their technical vulnerability management programme.

The obligations for CSPs and their clients with regard to VPC are outlined in the sections that follow and are discussed in the context of the SPI delivery model.

Handling of Security Vulnerabilities: To assist defend hosts, network devices, and applications against attacks against known vulnerabilities, vulnerability management is a crucial threat management component. Mature businesses have put in place a vulnerability management process that includes routinely scanning the systems linked to their network, evaluating the risks of vulnerabilities to the company, and remediating the risks (often by feeding the results into a patch management programme). Technical vulnerability management control goal, which reads: Objective: To decrease risks arising from exploitation of publicised technical vulnerabilities, is known to be used by organisations utilising ISO/IEC 27002.

Technological vulnerability management should be done in a fashion that is efficient, organised, and repeatable, and measurements should be conducted to verify this. Operating systems and any other apps in use should be taken into account. Depending on the SPI service used, either the client or the CSP is in charge of managing cloud infrastructure vulnerabilities.

Security Patch Control: Security patch management, like vulnerability management, is an essential component of threat management that guards hosts, network devices, and applications against unauthorised users abusing a known vulnerability. The activities commanded by your vulnerability management software immediately feed into the patch management procedures, which adhere to a change management framework. The administration of security patches reduces the danger that both internal and external attacks pose to your company. As a result, SaaS providers should regularly check for new vulnerabilities and patch the firmware and software on all devices that are used to supply clients with the aaS service.

Customers are relieved of patch management responsibilities in a SaaS environment, whereas they are accountable for managing patches for the entire stack of software (operating system, applications, and database) installed and used on the IaaS platform. The scope of patch management responsibility for customers will have a low-to-high relevance in the order of SaaS, PaaS, and IaaS services. Patching apps installed on the PaaS platform belongs to the customers as well.

Administration of Security Configurations: Another important threat management technique to protect hosts and network devices from unauthorised users exploiting configuration flaws is security configuration management. A component of overall IT configuration management, security configuration management is closely tied to the vulnerability management programme. Monitoring and access control to crucial system and database configuration files, including OS configuration, firewall rules, and application configuration, are necessary for protecting the network, host, and application configuration.

An access control management database, locally and remotely connected storage, and network zone setup. SaaS and PaaS service providers are accountable for managing the configuration of their platforms, whereas IaaS customers are accountable for managing the configuration of the operating system, application, and database hosted on the IaaS platform.

According to the SPI service delivery model, configuration management from a customer responsibility perspective has a low to high relevance in the order of SaaS, PaaS, and IaaS services. Clients are also in charge of managing the settings of the apps they install on the PaaS platform.

SaaS VPC Administration: Vulnerabilities, security patching, and system configuration in the CSP-managed infrastructure as well as the client infrastructure interacting with the SaaS service are the main areas of concern for SaaS VPC management. It is crucial to protect the endpoints from which the cloud is accessible since the SaaS delivery model is founded on the idea that the application service is given over the Internet to a web browser operating on any computing device (personal computer, virtual desktop, or mobile device). As a result, a VPC management programme has to be customised for the corporate context and should contain endpoint VPC management requirements. The majority of businesses have made it normal procedure to implement an OS image for personal computers that includes security features like firewalls, antivirus, and anti-malware software.

1. SaaS provider obligations
2. SaaS VPC scope is represented by the list below:
3. Systems, networks, hosts, applications, and storage that are managed by outside companies in addition to those that are owned and controlled by the CSP
4. Smartphones and personal PCs held by SaaS contractors and staff
5. Customer obligations for SaaS

The customer has minimal duties for VPC administration of the cloud infrastructure since SaaS services are often given to web browsers and, in certain circumstances, are connected with client applications (through an XML interface). But, SaaS users are in charge of managing the VPCs on their systems that connect to the SaaS service. A SaaS users own computers, applications, and services that communicate with the SaaS service are among the obligations.

SaaS service security testing. While SaaS providers are in charge of managing software delivered as a service's vulnerabilities, some enterprise. Customers have the option to evaluate the security of an application on their own. Clients considering this independent verification option should get the CSP's approval first since SaaS security testing can only be done with the vendor's assistance. In order to find application vulnerabilities, this sort of application testing, which is often carried out by a third party tester, may comprise both an active examination of the programme and a simulation of actual attack scenarios. Since that this is a qualitative approach, the testing's parameters may change depending on the discovered vulnerability. So, it is advisable to confirm and settle on the scope before the exercise. The top web application vulnerabilities, as listed in the OWASP Top 10, may be found with this kind of testing. Throughout the cycle of application vulnerability testing, popular forms of vulnerabilities include SQL injection, parameter manipulation, cookie poisoning, and cross-site scripting (XSS).

PaaS VPC Management

PaaS VPC management concentrates on managing VPCs in the infrastructure that is managed by the CSP as well as the client infrastructure that interfaces with the PaaS service. The programme should include endpoint VPC management scope because applications deployed on a PaaS platform are accessed from a web browser running on an endpoint device (a personal computer, virtual desktop, or mobile device).

PaaS provider obligations: Similar to a SaaS model, the PaaS CSP is in charge of managing VPCs for both its own infrastructure and any third-party services it might use. For a list of responsibilities, see "SaaS provider responsibilities".

Customer obligations for PaaS: PaaS customers are responsible for VPC management of the applications implemented and deployed on the PaaS platform, in addition to the duties described in "SaaS customer duties". Applications running on a PaaS platform should be treated similarly to a standard application running in your data centre in terms of vulnerabilities or configuration weaknesses (e.g., private cloud).

Software vulnerabilities can be caused by coding mistakes or design flaws. By misconfiguring an application in the areas of authentication and privilege management, configuration weaknesses can be introduced. Furthermore, third-party web service vulnerabilities may cause PaaS applications that rely on them to become weak and vulnerable, and you have no control over that. However, you

You must cooperate with the PaaS vendor or outside service providers to address any flaws or vulnerabilities in their services if you do not have the ability to fix vulnerabilities in the source code of your PaaS application. Customers should be aware of the SLAs, PaaS policies, and vulnerability disclosure methods used by third-party service providers. Customers of PaaS should adhere to the best practises incorporated into the Software Development Life Cycle (SDLC), which reduces the vulnerability of software applications. The following are some examples of accepted procedures:

White-box application testing: Utilize testing tools, such as Ounce Labs and Fortify source code analysis tools, to analyse the source code for flaws like buffer overflows. Black-box application testing: Testing professionals who perform this type of testing must be familiar with the functionality of the application. In most cases, source code access is not necessary. SQL injection, parameter manipulation, cookie poisoning, and XSS are just a few of the OWASP Top application vulnerabilities that can be discovered through this kind of testing. As an illustration, consider companies like Cigital and Veracode.

Penetration testing of applications: Although PaaS providers are in charge of managing software platform vulnerabilities when they deliver their services, some enterprise customers can opt to independently assess the security of the application platform. Platform testing can only be done with the permission and cooperation of the PaaS vendor, so customers considering this independent verification option should first check with their PaaS CSPs. In order to find application vulnerabilities, this type of application testing, which is typically carried out by a third party tester, actively analyses the application and simulates actual attack scenarios. Given that this is a qualitative approach, the testing's parameters may change depending on the discovered vulnerability. Therefore, it is wise to confirm and settle on the scope before the exercise.

Vulnerabilities warnings: Customers should be aware of how PaaS providers, organisations, or communities that support the PaaS programming language inform them of vulnerabilities. PaaS providers have a few options for reaching out to their clients, including email, RSS, and web portals. You should also pick the right communication channels to stay updated on any new platform or third-party service provider vulnerabilities.

Customers of PaaS are also in charge of managing the VPCs on their systems that connect to the PaaS service. These include the user's personal computer and the browsers used to access the PaaS service.

Programs installed on the client's property that connect to the PaaS service

IaaS VPC Administration: IaaS VPC management focuses on both the customer infrastructure that interfaces with the IaaS service as well as the CSP-managed infrastructure. IaaS VPC management differs from SaaS and PaaS in that the boundaries between the infrastructure, customer networks, and CSP infrastructure are hazy. The customer and CSP are responsible for managing VPC in each of the infrastructure layers (network, host, and storage) from their respective perspectives (i.e., the CSP is responsible for the common CSP infrastructure available to all customers, and the customer is responsible for the virtual infrastructure available to the customer for the duration of use). Consequently, a VPC management programme should take into account both the shared and common infrastructures.

IaaS provider obligations

In general, an IaaS CSP is in charge of VPC management for the infrastructure that the CSP owns and operates as well as any third-party services and infrastructure they might use. The systems, networks, hosts (hypervisors), storage, and applications that are owned and run by CSP should be covered by the VPC management scope.

Third-party-managed systems, networks, hosts, storage, and applications

1. The management station or web console that customers use to manage their virtual infrastructure
2. Personal computers used by IaaS contractors and employees
3. IaaS customer obligations

The virtual infrastructure that an IaaS CSP allots to a customer for use must be managed by the customer's VPC. The management area for VPCs should cover:

Online servers: Both active and inactive VMs fall under this category. The OSs of the virtual servers must be taken into account when managing VPCs for VMs, and the programme must be adjusted accordingly (e.g., Fedora Linux, Solaris 10, Windows 2003). Customers are advised to manage virtual machines (VMs) according to best practises, which include:

Image standardisation through the use of security by default

After the image has been sufficiently hardened using the security-by-default method, customers are advised to standardise it. In the early stages of cloud services, until experience and best practises catch up, loss of security by default is more obvious. The concept of "security by default" refers to the implicit security present in daily operations.

Configuration guidelines: To lessen their overall attack surface, the OS, applications server, database, and web server must be installed and configured in accordance with least-privilege and security hardening principles. For instance, based on accepted best practises for the deployment, configuration, and operation of networked systems, the Center for Internet Security publishes Internet security benchmarks for major OS, databases, and application servers. All three factors that contribute to Internet-based attacks and disruptions—technology (software and hardware), process (system and network administration), and human—are taken into account by the center's security-enhancing benchmarks (end user and management behavior).

Configuration control

This is an example of centralized configuration management where a large number of nodes and zones in a public IaaS cloud need to be managed and the proper configuration data is required. Numerous configuration management tools are available, including open source tools (e.g., Puppet) and tools from commercial vendors such as BMC, Configure soft, HP, Microsoft, and IBM. However, due to the distinctiveness of the management API specific to each CSP, configuration management of virtual servers hosted in the cloud will require customization per CSP.

Network access regulations

The security architecture heavily relies on network zoning and firewalling to create security zones for applications hosted in an IaaS cloud. To reduce risk from improper configuration, network policies that allow traffic into and out of a customer infrastructure should be carefully configured. Network access policies that have been improperly configured may have left vulnerable services open to Internet crackers.

Typically, policies are divided into the following trust categories:

Internet etiquette

Permit Internet traffic between customers virtual servers and hosts (e.g., allow only ports 22, 80, and 443 to servers). Deny all outgoing traffic coming from virtual servers used by customers. Zone policy Permits communication between cloud-based virtual servers (e.g., allow port 3306 [MySQL] from server zone A to server zone B).

IaaS administrators are also in charge of managing the VPCs for their systems that connect to IaaS services. These include the following:

Cloud management station, which is the host that the customer manages for managing the virtual infrastructure in an IaaS cloud:

Internet-accessing browsers for IaaS services

To manage the deployment of their public and private IaaS clouds, IaaS customers have the option to use third-party services like RightScale, Enomaly, Elastra, and 3tera. To include security management functions in your SLA, you will need to work with your provider as the nature of security management services varies between providers.

Incident response and intrusion detection

Significant incident response and intrusion management challenges are brought on by the multitenant delivery model of a large-scale cloud provider offering SaaS, PaaS, and IaaS services. These challenges affect both customers and CSPs. A corporate information security management domain's core operations for managing risks like intellectual property loss, regulatory non-compliance, brand erosion, and fraud include intrusion and incident management. The ability of organisations to respond to intrusions and data breaches is facilitated by these crucial functions that support security management. Additionally, organisations are required by law to address privacy data breaches. The custodian of personal and regulated data must notify individuals whose data may have been compromised during a security breach under laws that have been adopted by more than 44 U.S. states. Since shared infrastructure resources and services are used to deliver public cloud computing, which is multitenant by definition, to customers, both the customer and the CSP are in charge of

managing intrusion and incident response. It will be necessary for both parties to be ready to manage and respond to security breaches.

The following control recommendations are provided by ISO 27002 for incident response and notification:

Management of security-related incidents and upgrades. A consistent and efficient approach must be used to manage information security incidents, according to the goal.

Once information security events and weaknesses have been reported, there should be responsibilities and procedures in place to effectively handle them. The handling of information security incidents, including their response, monitoring, assessment, and overall management, should be done through a process of continuous improvement. In order to ensure compliance with legal requirements, evidence should be gathered where it is necessary.

In the past, medium-sized and large enterprise customers would either use an internal security operations centre (SOC) or a third-party managed service to manage security and incident monitoring processes. Today's SOC keeps track of activity from firewalls and intrusion detection systems and uses the CERT process to respond to incidents. Because monitored firewalls and IDS will no longer be able to protect cloud applications, the traditional network security-monitoring model will be put to the test. The SPI (SaaS, PaaS, IaaS) delivery model, CSP-specific SLA, incident disclosure policy, and data governance model within the CSP will all have an impact on the responsibility scope for intrusion monitoring and incident response in the cloud. The scale of operation will present various challenges to CSPs because they may host hundreds of thousands of virtual servers (IaaS), application instances (PaaS), and commingled customer data (SaaS).

However, unlike current incident management procedures used by a SOC or CERT team, incident notification in the cloud is not as straightforward. The notification and corrective action for all applications under the control of the organization's IT department are handled by one internal group in the traditional model, which groups those processes into a single governance and incident response model. In the case of a cloud where thousands of application owners have a stake, the notification process is more complex and will not follow traditional methods. New incident response tools may need to emerge to manage the complexity—e.g., an application registry implemented by the CSPs, with the contact details of the application owners and an automated notification system to handle a large number of customers (tenants) (tenants).

Best practices from a privacy perspective dictate the isolation of application data. In the traditional architecture, the breach management process will focus on one entity and not several. Unfortunately, in the cloud, the data separation will blur quickly and an incident procedure will have to be very specific to handling a commingled data environment and identify the dependencies so that the incident notification can be delivered to all parties in a line of data custody.

Customer versus CSP Responsibilities

Given the shared infrastructure and responsibilities, both the customer and CSP should have in place a security incident response plan to address any kind of security breach thoroughly and expeditiously. The team should promptly disclose to other tenants the existence of a vulnerability that affects its operation to prevent further ripple effects (e.g., cascading infections within or outside the cloud) (e.g., cascading infections within or outside the cloud).

The serviced customer may have to inform its own customers or employees of the occurrence of the breach. The cloud service provider may also have to inform its other tenants that the breach has occurred.

In the case of an IaaS or a PaaS environment, the system and application trust boundary interlaces both the CSP and customer environment, and as a result, both parties share responsibilities for security monitoring and incident response domains. Those responsibilities should be clearly identified and documented. For example, a PaaS CSP should be responsible for intrusion detection and incident response for the shared network and system infrastructure, for the PaaS platform runtime engine software, and for supported service components; the customer is responsible for their deployed applications and hosted data.

The provider collects and must protect a huge amount of security-related data. For example, at the network level, the provider should be collecting, monitoring, and protecting firewall, intrusion prevention system (IPS), security incident and event management (SIEM), and router flow data. At the host level, the provider should be collecting system logfiles, and at the application level, SaaS providers should be collecting application log data, including authentication and authorization information. They should have in place a security monitoring and incident response plan to address the security breach thoroughly and expeditiously.

Caveats: Prior to designing a VPC programme, customers are advised to read and understand the terms and conditions and user agreements with their CSP, because there may be potential restrictions for scanning network services, brute force testing, and penetration testing of applications deployed on that CSP's IaaS platform. Furthermore, network port scanning, application security scanning, and active penetration testing can trigger a CSP's intrusion detection system/ intrusion prevention system (IDS/IPS) alarms, which in turn can result in suspension or deactivation of your service temporarily or permanently. For example, Amazon AWS, as a matter of policy, prohibits port scanning of your virtual servers.

QUESTIONNAIRE

1. Explain Cloud Security?
2. Explain Cloud Computing?
3. What are Cloud Computing Threats?
4. Name a Few Cloud Computing Attacks.
5. Who takes Responsibility for Application-level control in Platform as Service Cloud?
6. What are the Different Types of Cloud Services?
7. Example for Software as Service.
8. Example for Infrastructure as Service
9. Example for Platform as Service
10. Name a Few Cloud Security Tools

REFERENCE BOOKS FOR FURTHER READING

1. CSA Guide to Cloud Computing: Implementing Cloud Privacy and Security by Raj Samani.
2. Cloud Computing Security: Foundations and Challenges by John R. Vacca.
3. Cybersecurity for Executives in the Age of Cloud by Teri Radichel.
4. Cyber-Vigilance and Digital Trust: Cyber Security in the Era of Cloud Computing and IoT by Wiem Tounsi.
5. Enterprise Cloud Security and Governance: Efficiently Set Data Protection and Privacy Principles by Zeal Vora.
6. Mastering AWS Security by Albert Anthony.
