CENTRE FOR LAND WARFARE STUDIES

ISSUE BRIEF

No. 120 December 2017

Identity and Access Management in the Cloud

Debashish Bose joined CLAWS as a senior fellow in 2016. Prior to coming to CLAWS, he was heading the Army HQ Computer Centre at Sena Bhawan. He was responsible for the launch of the 'Army Cloud' in November 2015. His current area of work is "Cloud Computing: Security Implications on Military Networks". Other areas of passionate interest are Cyber Warfare, Influence Operations (Nation State / TBA & CI / CT) and Cyber Security. He is always eager to engage with like-minded people on the areas of interest.

In the present day digital environment, the Identity and Access Management (IAM) has become one of the fundamental pillars for providing a safe and secure digital ecosystem. This area is one of the most deeply impacted, with the migration to cloud computing. The deep impact does not imply that IAM fundamentals or methods of implementation have changed, but now IAM implementation has become a deeper and all pervasive issue. It has become a zero or one kind of game, wherein there are no shades of grey in IAM deployments in the cloud environment. This is primarily because the stakes are very high in the case of a compromise of the IAM system in the cloud.

The term "IAM" is not universal and is often referred to as Identity Management (IdM). Gartner defines IAM as "the security discipline that enables the right individuals to access the right resources at the right times for the right reasons." Rest assured that every word in the definition is loaded. Hereafter, in this brief, we will refer to the identity management infrastructure and associated software and hardware as an Identity Management System (IDMS).

Key Points

1. The Identity Management System is the first and most important line of defence in a cloud computing ecosystem.
2. It has the facility of a single log-in to access all the applications hosted in the cloud.
3. Multi-Factor Authentication is an essential requisite in all cloud deployments.
4. Due diligence is required for implementation of the federation of identity directories in the cloud.
5. Attribute-Based Access Control Vs Role-Based Access Control.

The Centre for Land Warfare Studies (CLAWS), New Delhi, is an autonomous think-tank dealing with national security and conceptual aspects of land warfare, including conventional and sub-conventional conflict and terrorism. CLAWS conducts research that is futuristic in outlook and policy-oriented in approach.

Website: www.claws.in Contact us: landwarfare@gmail.com

Identity and Access ...

All the three cloud delivery models of Infrastructure as a Service (IaaS), Software as a Service (SaaS) and Platform as a Service (PaaS) require that the user organisations' Information Technology (IT) department and the Cloud Service Provider (CSP) work jointly to extend the organisation's IDMS practices, processes, and procedures to cloud services in ways that are scalable, effective, and efficient for both the provider and the user organisations and departments. Clouds, because of the inherent architecture, tend to change faster and are more distributed. This not only adds to the complexity of the management plane but also requires fast network communications for all its activities, which opens up core infrastructure administration to network attacks.

The major area of concern arises from the fact (as applicable to all other areas of cloud computing), that cloud deployment, management and daily operations are basically a trust relationship between two parties, the CSP and the user. This trust relationship gains great significance when we look at the IDMS deployment. In the case of the IDMS, the trust relationship gets further refined with a clear division of responsibilities and the technical means used to implement the same. In the commercial domain, this issue is compounded by the fact that any company may use cloud services from multiple providers, thus, the central IDMS requirement also needs to be standardised across the spectrum of different providers. This complication, however, may not exist in the defence forces ecosystem (in our context), where the cloud services are generally provided by a single entity, but this ecosystem may have other kinds of complications like single vendor lock in, etc. However, this may also become applicable for the defence forces in case we have different cloud deployments among the three components.

IDMS, in addition to being an important functionality, is also extremely complex. Primarily, we are mapping some form of an ***entity*** (a person, system, piece of code, etc.) to a verifiable ***identity*** associated with various ***attributes*** (which can change based on current circumstances), and then making a decision on what they can or can not do, based on ***entitlements***. Determining and enforcing attributes and entitlements across disparate systems and technologies bring both processes and technical issues to the fore. The jargon, and a little bit of inline glossary would be helpful at this stage:

- **Entity**: It is described as the person or thing that will have an identity. Refers to an individual, a collective system, a specific device, or application code.

- **Identity**: It is the unique expression of an entity within a given directory. One particular entity can be represented in the digital world by multiple digital identities, for example, the same individual can have a work identity, a social media identity and a personal identity.

- **Attribute**: These are facets of an identity, and they can be relatively static or highly dynamic [like the IP address, location information, etc.]

- **Entitlement**: Entitlement is what an identity is allowed to do, and for the purpose of documentation, the details are stored in an entitlement matrix. We will see this in detail later.

Even when one single administrative entity is controlling the entire chain of the above processes, managing it across disparate systems (multiple cloud platforms / multiple application platforms) and technologies in a secure and verifiable manner, especially at scale, it comprises a challenge.

As far as cloud identities are concerned, both users and providers of the cloud infrastructure need to take basic decisions on how they are going to manage their identities:

- The cloud provider will always have to support internal identities, identifiers and attributes for users who directly access the service, and, at the same time, support federation so that the using organisations don't have to manually provision and

administer the users created in the cloud provider's system and issue everyone separate credentials.

- In the same breath, cloud users need to decide where they want to manage their identities, and, in particular, which architectural models and technologies suit their existing organisational policies and processes, and, more importantly, make the process of integration with the cloud provider smooth.

As a user of the cloud, one can also create all the required identities on the CSP infrastructure. This is not a very comfortable arrangement because it is not scalable in most use cases. Thus, in such situations, people turn to federated architectures. In spite of federation, there are specific cases where users like to keep all or some of their identities isolated with the cloud provider, such as the privileged administrator accounts, to help debug any problems with federation, in case they occur. When we are thinking about implementing identity management through federation, we first need to clearly identify the authoritative source of the identity. In most cases, this is an internal directory / identity server. Once this is identified, the next decision is to directly use the authoritative source as the identity provider, or use a different identity source that feeds from the authoritative source, as a proxy / identity broker. The identity broker who is handling federation between the identity provider and the user organisations (departments), may be located on the network edge (as far as the department is concerned), or in the cloud, since this facilitates web Single Sign On (SSO). There are two possible architectures for implementing these options:

- **Hub and Spoke Model:** The internal identity source connects with a central broker / proxy / repository; this then serves as an identity provider for federation to cloud providers.

- **Free Form Model:** In this case, the internal identity source connects directly to the cloud provider. This model has a few major security related problems, such as the directory server needs internet (i.e. outside the home network) access. This can be a problem in terms of implementation with respect to the existing architecture, and inadvertently, it may affect the security policies of the organisation. Additionally, it may require the users to connect back to the home (organisational / corporate) network, before one is able to access cloud services. Another major practical issue from the point of view of implementation is the situation of the existence of multiple directory servers, in different departments of the organisation. In such a situation, federating to a cloud provider who is typically outside the department, is complex and technically difficult. This is more applicable to the defence forces, that have been implementing IT applications for decades now, as a result of which a large number of departments have their own directory servers running in silos.

Identity providers need not be located only on the premises (within the department). The cloud provider can also provide cloud-based directory servers that support federation internally and with other cloud services.

Another useful area that requires looking into is the identity provisioning process that already exists in the organisation and how to integrate that in the cloud. There may be multiple processes for different use cases. However, the focus of the architect should be towards having a unified process as much as possible. If the existing process is efficient, efforts should be made to extend the same to the cloud. In case the existing processes are weak and inefficient, then migration to the cloud should be used as an excuse for developing new processes.

The entire thought process discussed till now is heavily biased towards users accessing services, but we also need to give thought to managing identities for the application code, systems, devices and other services. Thus, we will have to pay attention to services talking to services, systems or devices.

Steps for Adding Cloud Providers into Existing Departmental IDMS Infrastructures.

- Map attributes (including roles) between the cloud identity provider and user organisation / department being integrated.
- Enable the required monitoring / logging (preferably as per existing departmental process), including identity related security monitoring such as behavioural analytics.
- Build an entitlement matrix.
- Document any break / fix scenarios in case there is a technical failure of any of the federations (or other techniques) used in the relationship.
- Ensure Standard Operating Procedures (SOPs) are in place for incident response for potential account takeovers, with special emphasis on breach of privileged accounts.
- Implement de-provisioning or entitlement change processes (as per existing departmental processes) for identities and the cloud provider. Proper scrutiny needs to be carried out to ensure that there is no break in the existing departmental processes when they are migrated to the cloud.
- The cloud providers need to decide which identity management standards they will be supporting as per organisational requirements.

The important points to be kept in mind while choosing an identity protocol are as follows:

- No protocol is an ultimate solution that will solve all identity and access control issues.
- The identity protocols which are being considered must be scrutinised in the context of their use case. As in the case of browser based SSO, API (Application Programme Interface) keys, mobile to cloud authentication, etc—each will lead the user to a different approach.
- The most important aspect to be kept in mind is that the identity is a perimeter in and of itself. Thus, it should be so selected that it can withstand attacks and manipulations.

Authentication and Credentials

Authentication is the process of proving or confirming an identity. The common user understanding applies to the process of a user logging into his account / network. In general, we will assume authentication anytime an entity proves who he / she is, and assumes an identity. Authentication is assumed to be the responsibility of the identity provider. Authentication commonly relies on at least one of the factors mentioned below:

- **Something You Know**: This refers to authentication using some information known only to you, such as a password or Personal Identification Number (PIN). This information is shared between the user and the authentication service provider. Ideally, it is supposed to be both secret and hard to guess. However, in the dynamic cloud scenario, this mechanism is considered to be the least secure authentication mechanism and is susceptible to replay attack and identity theft since the user password / PIN can be easily stolen. Another common practical implementation problem in the defence environment, where levels of trust are generally very high among members, is the issue of sharing of passwords and usage of overly simple passwords among many users.
- **Something You Have**: This generally refers to something that we physically possess, such as a token. This is commonly considered as a stronger authentication mechanism. The secret user authentication credentials are encrypted and stored in hardware or software tokens. This information is checked against the credentials

stored with the IDMS before providing access to the cloud resources / services. However, these authentication tokens are also prone to identity theft attacks because these cards can be lost / stolen and thereafter the hacker can use them for his own malicious purposes. However, the strength of this mechanism can be enhanced by using it in conjunction with Something You Know.

- **Something You Are**: In this category of authentication, the IDMS authenticates the user based on biometrics. In this, user verification is performed on the basis of some natural characteristic such as fingerprint, voice pattern, iris characteristic. These are unique to every individual. To make the system stronger, this can also be combined with the above two.

The most major impact that cloud computing has had on the process of authentication is the implementation of Multi-Factor Authentication (MFA), so as to ensure the availability of strong authentication. This has happened because cloud services are always accessed over the network which is generally not under the user's (organisation's / department's) control. Thus, in such a situation, the loss of credentials could result in an account takeover by a hacker who does not have to be on the local network. The singular aspect which makes credential loss so sensitive in the cloud environment is the feature of SSO; the loss of one set of credentials could potentially lead to the compromise of a large number of cloud services. Till this day, MFA offers one of the strongest methods to prevent the compromise of a user's account. As a result, use of a single factor (such as a password only) poses a very high risk, particularly in the cloud environment. In the cloud environment, when we use MFA with federation, the identity provider should pass the MFA status as an attribute to the user party.

The commonly available options for MFA in the current eco system are as follows:

- **Hard Tokens:** These are physical devices which generate a One Time Password (OTP) which have to be manually entered or can also be plugged into a reader. This is considered to be the best option when a very high level of security is required.

- **Soft Tokens:** In functionality, they are similar to hard tokens, but since they are software applications, they require a phone or computer to run on. Soft tokens are also considered to be an excellent option, but they can become an issue if the system (mobile / personal computer) on which they run gets compromised.

- **Out of Band Passwords:** These usually comprise a text or message sent to the user on his phone (mobile or land line) and are thereafter entered like any other OTP generated by a token. However, message interception in the selected band (for OTP transmission) should also be considered when threat modelling is being carried out.

- **Biometrics:** In the present day, this is another commonly used option, because of easy availability of biometric readers in smart phones or as cheap stand-alone devices. As far as cloud services are considered, the biometric input is a local protection and the exact biometric information is not sent to the cloud provider and is instead an attribute that can be sent to the provider. As in the case of soft tokens and out of band passwords, the security and ownership of the local device is an important consideration.

Entitlement / Authorisation / Access Control

The meanings and implications of all these three terms overlap slightly, so we will define them before going ahead.

- **Authorisation:** Authorisation is permission to do some activity, such as accessing a file or network or performing certain functions like an API call on a particular resource.

... Management in the Cloud

- **Access Control:** This allows or denies the expression of the above defined authorisation. It involves issues like assuring that the user is authenticated before allowing access.

- **Entitlement:** Entitlement maps identities to authorisations and any required attributes. Thus, a user is allowed access to a specific resource, when the given attributes have designated values. A map of these entitlements is created which is known as an entitlement matrix. Entitlements are generally converted as technical policies for distribution and enforcement.

Cloud Impact of Entitlement / Authorisation / Access Control

- The **cloud provider** will have his own set of potential authorisations specific to him. The user will generally need to configure entitlements within the cloud platform directly.

- The **cloud provider** is responsible for enforcing authorisations and access controls.

- The **cloud user** is responsible for defining entitlements and properly configuring them within the cloud platform.

- Cloud platforms tend to have greater support for Attribute-Based Access Control (ABAC), than the Role-Based Access Control (RBAC) model. ABAC offers greater flexibility and security than RBAC. RBAC is the traditional model for enforcing authorisations and generally relies on a single attribute or defined role. ABAC, on the other hand, gives granular control and context aware decisions by including multiple attributes such as role, location, authentication method, and many more. ABAC is the recommended model for cloud access management.

- In the federated architecture, the cloud user is responsible to map the attributes, including roles and groups (groupings of users), to the cloud provider and ensure that these are properly communicated during authentication.

- Cloud providers are responsible for supporting granular attributes and authorisations to enable ABAC and effective security for cloud users.

Privileged User Management

From the risk perspective, this is a very important activity. All the requirements discussed earlier for strong authentication, need to be implemented in totality for privileged users. Additionally, account and session recording should be religiously implemented to ensure full accountability and visibility of privileged users. For certain privileged users, it may also be prudent to log in through a separate tightly controlled system which implements higher levels of assurance for credential control, digital certificates, physically and logically separate access points, and / or jump hosts.

IDMS Attack Patterns

Since the IDMS is the primary security feature in the cloud environment, its study is not complete without a look at the current known attacks against these systems. Presently, the study of security of cloud-based identity management systems is still in its early stages and requires further observation and research. The existing systems face various security and performance-based issues, which tend to inhibit their adoption as a viable solution for a dynamic cloud environment. The attacks listed below are all either against the IDMS or use identity as an attack vector. These attacks give a fairly good idea about the features required / lacking for achieving a good IDMS, which can keep the identities of the users secure:

- **Brute Force Attack**: This kind of attack generally allows the attacker to gain access to the identity credentials, stored in an identity management server, using different possible combinations of

the user ID and password. A dictionary attack is a possible example of a brute force attack.

- **Cookie Replay Attack**: The attacker steals a cookie containing valid session information along with the user's ID credentials and thereafter reuses it to trick the identity management server into assuming that a previously authenticated session is still continuing. The attacker gets access to the victim's confidential information as well as to cloud services and resources to which the victim was authorised access.

- **Data Tampering Attack**: This type of attack relates to unauthorised change of data relating to identification of an user in an identity store in the cloud. As a result of these changes, the attacker may be able to affect the cloud services and resources. This attack basically challenges the integrity of the data stored in the cloud.

- **Denial of Service (DoS) Attack**: DoS attacks can be launched against those identity management systems which have poor processes for logging user activities. As a result of this, the attacker is able to overwhelm the identity management server in the cloud with fraudulent authentication or authorisation requests, and, as a result, use all available resources so that genuine requests for authentication/authorisation cannot be processed.

- **Eavesdropping**: This attack happens when the user is communicating with the cloud—when the user and the identity management server are exchanging user credentials for authentication/authorisation. It happens through real time reading/stealing of user credentials by the attacker through listening or reading unencrypted data from the network.

- **Elevation of Privilege**: This involves getting in as a legitimate user, thereafter illegitimately escalating his rights and, thus, impersonating a user with higher privileges, and gaining access to resources/services the user is not supposed to access.

- **Identity Forgery/Cloning/Spoofing Attack**: This refers to unauthorised copying or manipulation of identity credentials obtained from the trusted source. The aim of doing this is to deceive the investigator investigating the reported breach. To avoid this kind of attack, the cloud-based IDMS should strictly employ the two factor authentication.

- **Identity Theft**: This refers to stealing someone's identity as referred to in the name / PII (including Aadhar number, PAN number, etc.) / credit card information, without the permission / knowledge of the owner. The aim of doing this is to acquire cloud resources or for financial fraud in the original owner's name. In many cases, the original owner may be prosecuted for the actions of the hacker. This is the first step for committing other crimes such as fraud, forgery, etc.

- **Luring Attack**: It can be understood as a type of privilege escalation attack. In this type of attack, the genuine user unknowingly executes the hacker's code, while operating in the privileged security mode. To be more specific, the hacker lures the original user to perform the illegal activities on his behalf. This kind of attack mostly happens in those IDMS which don't provide robust logging and reporting functionalities.

- **Phishing Attack**: Phishing is the act of acquiring the user's information such as name, password, Aadhar / PAN number, bank account numbers, credit card details, by fooling the user into entering his details into a duplicate website whose appearance very closely matches that of the actual website. The attacker will manage the whole session in such a manner that the user feels he is providing his details to the trusted authority.

- **Replay Attack**: This type of attack occurs when the IDMS is unable to ensure the confidentiality of the identity credentials during their transmission. During the attack, the hacker captures the valid identification information, and impersonates the original user by retransmitting the same. Unless properly addressed, the IDMS will end up processing the request assuming that it is servicing the genuine user. Accordingly, all authorised services and resources will be made available to the attacker.

- **Repudiation**: This kind of attack occurs when the cloud service user denies any action, which as per the IDMS was legally done by the user. This attack may also be facilitated by the fact that the IDMS may not be implementing proper / detailed user activity logs, so as to be able to carry out proper forensics after the breach. In this kind of situation, the user is at liberty to deny any kind of malicious activity carried out by him on the IDMS or any of the cloud resources / services.

- **Side Channel Attack**: The IDMS may fall victim to this kind of attack if it does not have stringent procedures and protocols for implementation of federation and access control. In this particular type of attack, the hacker may steal / accurately access information like session identifiers, timing information, OAuth tokens, electromagnetic leaks, etc, from the physical implementation of the IDMS system. To obviate this kind of attack, it is recommended to store sensitive identity information distributed at multiple locations across a federated system. Thereby, making it technically and physically difficult for the attacker to analyse the extracted side channel information as one single whole.

- **Skimming Attack**: In this type of attack, the hacker steals credential information from smart cards / credit cards / etc. To obviate this attack, the IDMS should implement strong encryption and distribute identity credentials across multiple servers.

- **Snooping**: This type of attack enables illegitimate collection of sensitive information such as identity credentials, network architecture, available services, etc from an identity server running in the cloud. Snooping differs from eavesdropping in the level of sophistication. This may involve highly technical and specialised techniques to intercept secure communications through key loggers, remote desktop captures, etc.

Way Forward

- There is no ideal or preferred security protocol. There is a need to check applicable use cases and organisational structures and processes to decide on the correct protocol.

- Organisations / departments should have a comprehensive and approved plan, along with applicable processes for managing identities and authorisations with cloud services.

- While connecting to the cloud provider, use federation if possible to extend existing identity management. Reduce silos of identities in cloud providers that are not tied to internal identities. This reduces risk.

- Use identity brokers where appropriate.

- Cloud users are responsible for maintaining the identity provider and defining identities and attributes. These should be based on an authoritative source.

- Cloud users should use MFA for all cloud accounts and send MFA status as an attribute when using federated authentication. Privileged identities should always use MFA.

- Develop, test, fine-tune and ratify an entitlement matrix for each cloud provider and project. Translate the same into technical policies, which should be supported by the cloud provider and the cloud platform.

- ABAC is to be preferred over RBAC for cloud computing.

- The cloud provider is to provide both internal identities and federation using open standards.

Conclusion

Historically, organisations have been using on-premises IDMS software to manage identity and access policies, but the present day trend is for organisations to add more cloud services to their environments, hence, the process of managing identities is getting more complex. Therefore, adopting cloud-based Identity-as-a-Service (IDaaS) and cloud IDMS solutions seems to be the future.

References

1. https://www.networkworld.com/article/2163744/infrastructure-management/identity-and-access-management-as-a-cloud-based-service-eliminates-time-p.html
2. Cloud Security Alliance: Security Guidance v4
3. https://securityintelligence.com/identity-management-cloud-tips-secure-identities-iam/
4. https://auth0.com/learn/cloud-identity-access-management/
5. http://www.computerweekly.com/tip/Identity-and-access-management-IAM-in-the-cloud-Challenges-galore
6. https://casmodeling.springeropen.com/track/pdf/10.1186/s40294-014-0005-9?site=casmodeling.springeropen.com

CENTRE FOR LAND WARFARE STUDIES (CLAWS)
RPSO Complex, Parade Road, Delhi Cantt, New Delhi 110010
Tel.: +91-11-25691308, Fax: +91-11-25692347, Email: landwarfare@gmail.com
Website: www.claws.in
CLAWS Army No. 33098