cloud security alliance℠

SecaaS Implementation Guidance

Category 1 // Identity and Access Management

September 2012

Contents

Foreword 6
Letter from the Co-Chairs 7
Acknowledgments 8
1.0 Introduction 9
1.1 Intended Audience 9
1.2 Scope 10
2.0 Requirements Addressed 11
2.1 Authentication 11
2.1.1 Strong Authentication 11
2.1.2 Risk-Based Authentication 12
2.2 Identity Federation Services 12
2.2.1 Federated Identity Management 12
2.2.2 Federated Single Sign-On 12
2.3 Identity Management Services 13
2.3.1 Provisioning and Deprovisioning 13
2.3.2 Centralized Directory Services 13
2.3.3 Privileged User Management 13
2.4 Authorization and Access Management 13
2.4.1 Authorization Management 14
2.4.2 Access Policy Management 14
2.4.3 Audit and Reporting 14
3.0 Implementation Considerations and Concerns 16
3.1 Considerations 16
3.1.1 Control 16
3.1.2 Visibility and Transparency 17
3.1.3 Portability 17
3.1.4 Interoperability 17
3.1.5 Costs and Investment Considerations 18
3.1.6 Multi-Layer Management 19

3.1.7 Performance/Availability Considerations 19
3.1.8 Service Level Agreements 19
3.1.9 Hybrid Cloud/Non-Cloud Services Integration 19
3.1.10 Unwanted Access 20
3.1.11 Scalability 20
3.2 Concerns 20
3.2.1 Standards/Openness/Vendor Lock-in 20
3.2.2 Identity Theft 21
3.2.3 Unauthorized Access, Insider Threats, Fraud and Accidental Access 21
3.2.4 Elevated Privilege Control 21
3.2.5 Non-Repudiation 22
3.2.6 Least Privilege & Excess Privileges/Excessive Access 22
3.2.7 Performance/Availability 22
3.2.8 Features/Functionality Gaps/Weaknesses 22
3.2.9 Disaster Recovery and Business Continuity Management 22
4.0 Implementation 24
4.1 Architecture Overview 26
4.1.1 Conceptual Architecture 26
4.2 Guidance and Implementation Steps 31
4.2.1 Provisioning and Deprovisioning of Accounts 31
4.2.2 Authentication 31
4.2.3 Directory Services 33
4.2.4 Directory Synchronization (Multi-Lateral as Required) 33
4.2.5 Single Sign-On 34
4.2.6 Federated Single Sign-On 34
4.2.7 Web Single Sign-On 36
4.2.8 Authorization (Both User and Application/System) 36
4.2.9 Support for Policy and Regulatory Compliance Monitoring and/or Reporting 38
4.2.10 Access and Activity/Session Monitoring 39
4.2.11 Tamper-Proof Audit 39
4.2.12 Policy Management 39
4.2.13 Role-Based Access Controls (RBAC) 39

4.2.14 Centralized Directory Service 40
4.2.15 Privileged Accounts 41
5.0 References and Useful Links 43
5.1 References 43
5.2 Useful Links 43

Foreword

Cloud Computing represents one of the most significant shifts in information technology many of us are likely to see in our lifetimes. We are reaching the point where computing functions as a utility, promising innovations yet unimagined. The major roadblock to full adoption of Cloud Computing has been concern regarding the security and privacy of information.

Much work has been done regarding the security of the cloud and data within it, but until now, there have been no best practices to follow when developing or assessing security services in an elastic cloud model—a model that scales as client requirements change.

One mission of the Cloud Security Alliance is to provide education on the uses of Cloud Computing to help secure all other forms of computing. To aid both cloud customers and cloud providers, the CSA SecaaS Working Group is providing Implementation Guidance for each category of Security as a Service, as delineated in the CSA's SecaaS Defined Categories of Service. Security as a Service was added, as Domain 14, to version 3 of the CSA Guidance.

Cloud Security Alliance SecaaS Implementation Guidance documents are available at https://cloudsecurityalliance.org/research/working-groups/security-as-a-service/.

We encourage you to download and review all of our flagship research at http://www.cloudsecurityalliance.org.

Best regards,

Jerry Archer	Alan Boehme	Dave Cullinane
Nils Puhlmann	Paul Kurtz	Jim Reavis

The Cloud Security Alliance Board of Directors

Letter from the Co-Chairs

Security as a Service is a specialized area categorized two years ago as growing rapidly and in unbound patterns. Vendors were struggling. Consumers were struggling. Each offering had its own path. We felt it was urgent to address the needs and concerns common to the implementation of Security as a Service in its many forms.

The Defined Categories of Service helped clarify the functionalities expected from each Category. In this series, we hope to better define best practices in the design, development, assessment and implementation of today's offerings.

We want to thank all of the many contributors worldwide who have worked so hard to produce these papers providing guidance for best practices in Cloud Computing Security. Many have been with the Security as a Service Working Group since the beginning; many others joined in this effort. Each has spent countless hours considering, clarifying, writing and/or editing these papers. We hope they help move forward toward those unimagined innovations.

Sincerely,

Kevin Fielder and Cameron Smith
SecaaS Working Group Co-Chairs

Acknowledgments

Co-Chairs

Ulrich Lang, Object Security
Valmiki Mukherjee, Mycroft

Contributors

Sameer Anja, Innovative 21st Century Technology Minds
Kanchanna Balraj, Engineering Ingegneria Informatica SPA
Aradhna Chetal, The Boeing Company
Andrey Dulkin, Cyber-Ark Software
Karen Lu HongQian, Gemalto
Anish Mohammed, Accenture
Sundar Narayanswamy, Ernst & Young
Balaji Ramamoorthy, GE
Laura Robinson, Microsoft
Kent Welch, Courion

Peer Reviewers

Martin Antony, Alcatel-Lucent
Jim Brigham, EMC/RSA
Tim Brooks, Signify Solutions, Ltd.
John Linn, EMC/RSA
Sandeep Mittal
Jean Pawluk
Robert Polansky, EMC/RSA
Michael Roza, Bridgestone
Said Tabet, EMC

CSA Global Staff

Aaron Alva, Research Intern
Vicki Hahn, Technical Writer/Editor
Luciano JR Santos, Research Director
Kendall Scoboria, Graphic Designer
Evan Scoboria, Webmaster
John Yeoh, Research Analyst

Category Sponsors

1.0 Introduction

Identity and Access Management (IAM) includes people, processes, and systems that are used to manage access to enterprise resources by assuring that the identity of an entity is verified, then granting the correct level of access based on the protected resource, this assured identity, and other context information.

This guidance discusses the significant business and technical decisions that need to be considered by an organization seeking to implement the IAM component of Security as a Service (SecaaS) as part of the cloud environment, or an organization that is looking for guidance as to how to assess an IAM offering. This document is intended to assist with the planning, design, implementation and assessment of SecaaS offerings in the area of Identity and Access Management. It is meant to serve as a source of reference for best practices in the industry today.

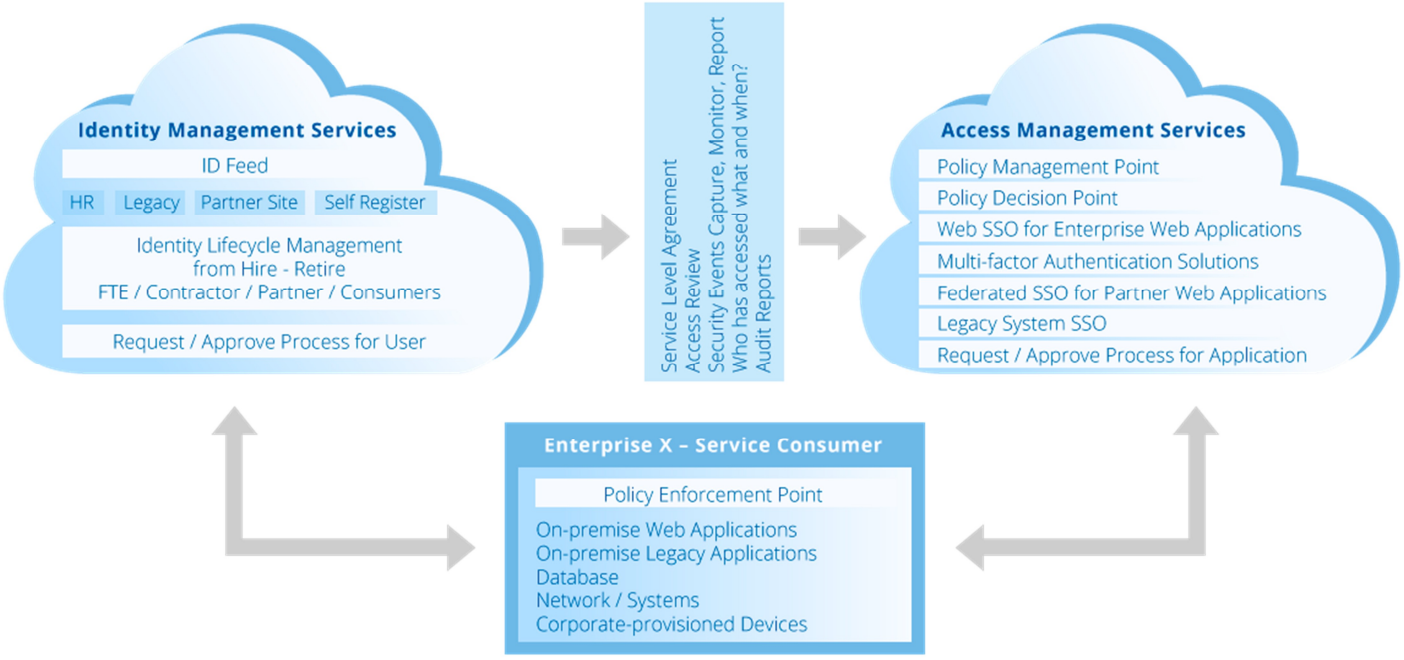

Figure 1: Security as a Service – IAM Components

1.1 Intended Audience

This document addresses personnel involved in the identification and implementation of the IAM solution in the cloud. It will be of particular interest to those with the responsibility of designing, implementing and integrating the consumption of services of the IAM function within any cloud application of SecaaS. Business processes are intended to be shared with stakeholders who have responsibility for ensuring that the solution has full functionality to support the demands of their business. This paper also provides direction for enterprise security stakeholders responsible for ensuring the security of IAM solutions in a corporate IT environment.

Section 2 offers a high-level overview of Identity and Access Management as it is applied to Cloud Computing development and implementation. The material is written for executive level discussion, and it indicates a baseline for best practices in implementation and design of IAM services in the cloud.

Section 3 details the considerations and concerns that should be part of the decision making conversation, whether by a design team or within the context of a purchasing decision. Section 3 is written for those who are implementing or evaluating IAM services.

Section 4 is a highly technical discussion of the architecture and implementation of IAM Security as a Service (SecaaS). This material is written for systems architects, designers and developers. It also provides best practice guidelines that should help purchasers better assess IAM offerings under consideration for purchase.

Section 5 supplies references and useful links to trusted sources of information regarding IAM Security as a Service in Cloud Computing.

1.2 Scope

This Implementation Guidance documents best practices for the design, implementation and assessment of Identity and Access Management services, especially as they are applied within Cloud Computing. IAM components discussed include:

- Centralized Directory Services,
- Access Management Services,
- Identity Management Services,
- Identity Federation Services,
- Role-Based Access Control Services,
- User Access Certification Services,
- Privileged User and Access Management,
- Separation of Duties Services, and
- Identity and Access Reporting Services.

2.0 Requirements Addressed

Data is an asset to any business, and may be the most valuable asset a business owns. Data must be treated with the same degree of concern required to protect any significant asset. Allowing unauthorized access could lead to various avenues of risk exposure, including identity theft, privilege escalation, loss of intellectual property, and fraud. Identity and Access Management functions are critical parts of any data protection mechanism.

The Principle of Least Privilege states that entities should be allowed access only to that data they have a need to know in a particular, often dynamically changing context. Identity and Access Management is critical to enforcing that principle.

Identity Management includes the creation, management and removal (deletion) of a digital identity. Access Management includes the authorization of access to only the data an entity needs to access to perform required duties efficiently and effectively. This section discusses the requirements of secure Identity and Access Management, and the tools in use to provide IAM security in the cloud.

2.1 Authentication

Authentication is the process of verifying the credentials of an entity trying to access a protected resource. Authentication must be done in a secure, trustworthy, and manageable manner. For accounts that require higher levels of security, multiple factor authentications may be required. Authentication systems should have the capability to use business transaction risk definition as guidance, and provide adaptive authentication based on the level of risk of the transaction.

Single Sign-On (SSO) is the functionality within access management where user is authenticated once and the credentials for the session are trusted across different applications within a security domain. This is typically done within one security or risk domain. SSO is a critical requirement within organizations operating all applications within a specific cloud infrastructure.

2.1.1 Strong Authentication

Strong authentication typically requires the use of two or more of the three types of authentication factors. In the cloud platform, authentication services should include strong authentication mechanisms for validating the credentials and determining the authenticity of the user. Functionalities such as a One Time Password should be supported as a standard feature. The multi-factor and the risk-based approaches described above should be available as options in the IAM SecaaS service offering. The use of a onetime password is strongly encouraged during provisioning and communicating first-login passwords to users.

2.1.2 Risk-Based Authentication

Risk-Based Authentication (RBA) is a dynamic response to the changing conditions of *the risk rating attached to a particular user (agent) at a particular time.* Risk ratings are assigned to both the transaction and the agent. The higher the risk rating of the transaction, the higher the authentication level required of the agent, in order to complete the transaction.

RBA determines the agent's risk level by examining behavioral contexts, such as a history of similar requests, location of the agent/transaction, and timing of the request. Any behavioral anomalies can trigger a reaction ranging from a request for further authentication to a denial of the transaction.

2.2 Identity Federation Services

Federated identity services allow an organization to manage both the identity and access of its users to resources of partner organizations providing services authorized for those specific users. While Federated Identity Management processes help manage the lifecycle of the users' identities and accounts in the partner systems, Federated Single Sign-On assists with the authentication of users internally, then relays that identity to its cloud services provider, as a trusted token. This enables the organization to maintain control of the authentication process.

2.2.1 Federated Identity Management

Federated Identity Management provides the policies, processes and mechanisms to manage identity and trusted access to systems across organizations. This allows for reuse of users' identities across organizational boundaries, and ensures efficient user lifecycle management, compliance, and congruence of relevant user information between two partner organizations without excessive administrative overhead. The primary objective of federated identity management is to provide the users of one security domain the ability to access the systems of another domain in a seamless manner, thus enabling Federated Single Sign-On.

2.2.2 Federated Single Sign-On

Federated Single Sign-On (SSO) further enables a user's authentication in one domain to be trusted across different domains (e.g., different service providers). This provides convenience to users and better security, if the authentication domain maintains a strong security posture. Federated Single Sign-On is required as a standard functionality for facilitating inter-organizational and inter-security domain access to resources leveraging federated identity management.

2.3 Identity Management Services

2.3.1 Provisioning and Deprovisioning

Provisioning and deprovisioning are critical aspects of Access Management. Provisioning is the process of creating accounts to allow users to access appropriate systems and resources in the cloud. The goal of user provisioning is to streamline account creation and provide a consistent framework for providing access to end users.

Deprovisioning is the process whereby a user account is disabled when the user no longer requires access. This may be due to users leaving an organization, transferring within an organization, a change in role, etc. In the cloud computing environment, deprovisioning refers to the termination or disabling of user accounts in cloud platforms, or those managed by the cloud-based IAM service.

2.3.2 Centralized Directory Services

Directory service is one of the basic building blocks of security in an enterprise and in the cloud. A directory service provides an organized repository of information stored and identified by a unique identifier and location.

Lightweight Directory Access Protocol (LDAP), based on the X.500 standard, is a primary protocol for directory service. Each entry in an LDAP directory server is identified through a Distinguished Name (DN). In a cloud environment, directory services are heavily utilized by the Identity and Access Management framework as a security repository of identity and access information. Access to Directory Services should be part of the Identity and Access Management solution and should be as robust as the core authentication modes used. The use of Privileged Identity Management features is strongly encouraged for managing access of the administrators of the directory. If these are hosted locally, rather than in the cloud, the IAM service will require connectivity to the local Lightweight Directory Access Protocol (LDAP) servers, in addition to any applications and services for which it is managing access.

2.3.3 Privileged User Management

Privileged User Management provides the special requirements to manage the lifecycle of user accounts with highest privileges in a system. This also should fulfill requirements to authenticate, authorize, log, monitor and audit, and manage the password of the privileged users. While services offered for these accounts likely will be similar to “normal” user accounts, the policies and procedures that must be adhered to in relation to the management of these accounts may be considerably more stringent.

2.4 Authorization and Access Management

Authorization and access management forms a broad category of services that is required to be fulfilled in order to ensure security in the cloud. While authorization determines the user’s right to access a certain resource, access management in general has the responsibility to enforce that users’ access to a certain resource is managed with respect to the user’s credentials and attributes associated with the identity.

2.4.1 Authorization Management

Authorization management in the cloud should ensure that users have appropriate rights to access cloud as well as enterprise managed resources. Both policy definition and enforcement functions need to be available. User access needs to be approved or disapproved in real time with respect to the authorization policies in place. Completely trusted and anonymous authorization should be restricted, and detailed user authorization should be implemented. This needs to work seamlessly across on-premise systems and the enterprise in the cloud, offering real-time synchronization for both provisioning and deprovisioning.

2.4.2 Access Policy Management

Access Policy Management functions should be available at each layer of the cloud solution, including infrastructure, platform and software as a service. In a SecaaS environment, there is a requirement to have end-to-end management of an Access Policy so that security is not compromised at any level. Access policies usually emanate from a common set of security policies. Access policy management is critical because if the policy that is to be enforced is incorrect, unmanageable, or not verifiable (compliance), then the entire access control system is effectively rendered useless.

2.4.3 Audit and Reporting

The Identity and Access Management function of Security as a Service should be able to track and furnish information for all basic audit requirements as that in traditional IAM implementation which includes the information regarding:

- Who has access to what information,
- If the access is appropriate for the job being performed
- If the access is monitored, logged and reported properly

Monitoring cloud resources for user access is critical. An appropriate evaluation of the risk of exposure is not possible without having the right measures in place for monitoring and reporting, which could provide the metrics that will be used by audit. Monitoring cloud resources for user access is required to:

- Identify and prevent access violations
- Quantify risk exposure and residual risk
- Enforce segregation of duties
- Have a role-based access control mechanism

Audit logs are a critical part of the IAM process. Logs of activity, including all authentication and access attempts (both successful and failed) should be kept by the application/solution. Different types of reports need to be created and used within the scope of IAM as security information and events. Many of the reports that are typically created are used for operational purposes, such as reports of system performance activities, tasks and queue management functions, and reconciliation events. Audit reports include those that describe:

- User identities and their associated access

- Access approval information
- Administrative and privileged accounts and their associated owners
- User count and associated statistics for a particular resource
- Access Failures
- Privilege Access Failures

The processes and supporting systems should be able to provide reports that detail access approvals and reviews. Reporting on both the data and process is equally important and should be included.

3.0 Implementation Considerations and Concerns

3.1 Considerations

The use of a cloud service for IAM instead of an in-house service introduces numerous changes, some of which are not immediately obvious. This section discusses a number of those considerations when cloud IAM SecaaS is implemented.

The IAM SecaaS should provide the following features:

- Control over elevated privileges
- Reporting on access success and failures (including, but not limited to, privilege access)
- Least privilege/need-to-know
- Segregation of administrative (provider) vs. end user (client) interface and access
- Removal and archiving of identity information at the end of the lifecycle
- Disabling of access for the identity at the end of the lifecycle (Disabling and not removal is essential to establish audit and forensic trails for compliance requirements. Removal of the record also may give rise to inconsistencies in the access.)
- Real-time provisioning and deprovisioning
- Dynamic trust propagation and development of trusted relationships among service providers
- A user centric access control where user requests to service providers are bundled with their identity and entitlement information.
- Violation reporting during provisioning (both new and change control)
- Enforcement of segregation of duties
- Integration with Single Sign-On, dual mode authentication features
- Provisioning to add multi-authentication layers for specific roles

3.1.1 Control

Potential lack of control over data – Subscribers have a potential lack of control over identity and access policy data compared to an in-house installation, as the IAM SecaaS provider often holds all the identity and policy data. Identity and policy import/export features are critical and should be offered by IAM SecaaS providers.

Potential lack of control over features and functionality – While the adequacy of the functionality must be considered for any IAM product or service, features offered by an IAM SecaaS provider may potentially be more fluid than in an in-house installation (where, for example, feature update cycles can be controlled better).

Potential lack of control over IAM SecaaS provider operations – IAM cloud subscriber organizations may have specific requirements over the operations of the IAM SecaaS (e.g., for compliance purposes). While cloud services are not necessarily automatically less secure or less compliant than in-house installations, it is potentially harder to demonstrate compliance.

3.1.2 Visibility and Transparency

Potential lack of visibility for compliance purposes – Each subscriber organization is responsible for their own compliance (if applicable), no matter if IAM is used as SecaaS or in-house. However, as a result of effectively outsourcing the IAM operation when IAM SecaaS is used, there is a potential lack of visibility for compliance purposes. If compliance visibility is an issue, it is important to ensure the provider offers sufficiently detailed logs and reports on a sufficiently timely basis. Appropriate logging and reporting also should be able to facilitate and satisfy the requirements of independent audits where necessary.

Multi-jurisdictional regulatory requirements – Depending on where the IAM SecaaS is hosted, there may also be potential jurisdictional issues related to export of privacy related information. The geography of the IAM SecaaS, therefore, should be considered.

3.1.3 Portability

In order to avoid vendor lock-in, cloud environments and services should be portable. Without data portability, it would be impossible to switch cloud service providers. For IAM functions serviced in the cloud, there could be a great challenge to business function portability as much as for data portability. IAM functions are known to be complicated in terms of implementation and are usually tied to a particular implementation type. However, changing cloud vendors could pose a great challenge with this migration, as there are no specific standards for business function migration as there are for data transformation and migration. The use of standard protocols and interfaces for integration with cloud-based IAM services can simplify technical aspects of portability.

To ensure portability, the IAM vendor can use standardized formats for data exchange, and provide connectors to various cloud platforms which will enable the enterprise to switch cloud platforms if so required. The core engine in this case remains the same, while the connectors provide the ability to connect to different types of cloud platforms. At the very least, IAM solutions should have SDK/ API interface built in which can be used by the enterprise to determine portability strategies.

In terms of legal contracts, some outsourcing companies have a data-hostage clause in their contracts which is related to early termination and non-payment. In the case of a cloud environment, it becomes very important that this clause be reviewed and rewritten to accommodate the newer terms of cloud computing service offerings.

3.1.4 Interoperability

Interoperability with key business and technology platforms is a major concern for cloud computing platforms, especially any proprietary solutions, which could lock the customer into using other solutions from the same vendor. This would inhibit a customer from exploring other options and alternative products in the market.

Existing applications also might be affected, if, after migration to the cloud, a customer finds that the existing in-house applications are not compatible with the cloud-based applications to which they just migrated.

3.1.5 Costs and Investment Considerations

IAM SecaaS may not necessarily be cheaper than an in-house installation in the long run. This depends on the usage and charging pattern. Carefully evaluate all costs associated with each option.

IAM SecaaS may not immediately be able to turn all IAM costs from capital to operating expenditure, due to potentially significant integration, configuration, data management and other efforts related to using IAM SecaaS. This can be because the access control needs still occur on the protected resources, which may not support the access control mechanisms required (i.e. investment/engineering necessary). This can also be due to the complexities of implementing a SecaaS/in-house IAM mix. If the organization does not have identity/login/authentication/access policy data effectively managed, the cost of "doing things right" can be significant (with or without SecaaS).

3.1.5.1 Access Control

Access Control Integration – Identity and authentication are generally only a means to an end in information security, with access control being one of the core goals. The purpose of access control as a form of authorization is to limit any abuse a legitimate user has over a system. Access control constrains what a user can directly perform, or what a system, acting on behalf of a user, can perform. Access control (based on username/password logins, identity certificates, pre-authorized authorization tokens, fine-grained technical access control policies, ACLs etc.), must be managed by a protected resource when access is requested from a protected resource by a legitimate user within the system. As with in-house identity management deployments, this integration of cloud identity management SecaaS into the actual access control of cloud or in-house applications can be a considerable challenge, but is absolutely necessary. Access control needs to be tightly integrated with the enterprise application to which access is being requested. This will enable the mechanism to work seamlessly with single-sign on features and federated identity requirements and will use the core principles of the enterprise application for authorization.

Lack of Interoperable Representation of Authorization/Entitlement Information – Support of standards should be considered. For example, OASIS XACML can be used as a technical interchange format for authorization policies. Alternatively, authorization token standards (e.g., OAuth) can be used to exchange authorization information.

Access Control Granularity – In many enterprise scenarios today, IAM's access control policies must be fine-grained, contextual and feature-rich to support particular regulatory compliance and business requirements. This raises significant access policy management questions, which may be answered using newer and emerging approaches such as attribute-based access control (ABAC), potentially paired with authorization-based access control (ZBAC) model-driven security approaches for manageability.

Resource-Based Access Control – The assumption that access control is always (human) user-based does not hold any longer in many environments (e.g., interconnected IT landscapes such as Service Oriented Architectures, SOAs). Access control may need to be machine-to-machine or application-to-application-based,

and may only be easily enforceable if it is expressed with the protected resource in mind ("what is allowed on this system") rather than user-centric ("what user xyz is allowed to do").

Delegation of Authorizations/Entitlements – Delegation of authorizations/entitlements may work only if the same standards are followed by all involved. Checking for common standards should be considered if authorization token delegation is needed.

3.1.6 Multi-Layer Management

IAM is not necessarily only about users logging into a service to get coarse-grained access. In today's interconnected IT landscapes, where applications are often orchestrations of individual connected modules (e.g., web services in Service Oriented Architecture, SOA), granular service-to-service IAM may be required, and user and/or machine privileges may need to be delegated multiple hops across such orchestrations. If this is needed, the IAM deployment must support such integration at the best layer(s) in the software stack.

3.1.7 Performance/Availability Considerations

Depending on the particular deployment scenario, it may be of concern that staff cannot log into a service if the IAM SecaaS is not available due to network connectivity issues. While this may not be an issue if the IAM SecaaS is used for IAM of cloud services (because these would then quite likely also be down), this could be an issue if IAM SecaaS is used to control access to in-house services (e.g., in a hybrid deployment). While most likely not a problem for user sign-on, the use of IAM SecaaS versus in-house IAM may have general performance consideration in high-performance deployments that require IAM, for example, for machine-to-machine authentication.

3.1.8 Service Level Agreements

Service Level Agreements (SLAs) should ensure the service meets the intended requirements. This may often be easier said than done; no SLA will include all eventualities. Today, SLAs typically offer very limited guarantees to subscribers, as well as limited compensation in case of failures.

Cloud subscribers are responsible for their own compliance and liabilities, even if SLAs are in place. SLAs may allow cloud subscribers potentially to claim damages from cloud providers.

3.1.9 Hybrid Cloud/Non-Cloud Services Integration

Enterprises today have IAM deployments in place, in which much effort and resources have been invested. It is unlikely that enterprises will simply "switch off" their existing investments (or in fact any other in-house IT) and move to IAM SecaaS. As a consequence, integration between in-house IAM and IAM SecaaS need to be considered. Considerations include functional integration, technical integration, data compatibility, dealing with data duplication/redundancy, standards incompatibilities etc. The maintenance of such a hybrid solution must also be considered. Ideally, an IAM SecaaS solution should offer interoperability with existing IT systems and existing solutions with minimum changes

3.1.10 Unwanted Access

3.1.10.1 Contractual Access

A subscriber's data should no longer be accessed at a provider once it has been deleted by the subscriber or once the subscription relationship has concluded. The subscriber might either request an archival clause or an export clause, which specifies that subscriber related data is transferred to the subscriber once the contract expires or is terminated.

3.1.10.2 Government Access

If the government has a search warrant for any of the data collocated on a cloud IAM service, and the provider cannot convincingly separate out different customers' data, then the government will most likely request access to the entire storage/database, etc. This means that the government could potentially have (for unintended technical reasons) access to the data of many subscribers even without a search warrant.

3.1.10.3 Unauthorized Access by IAM SecaaS Staff

Proof of best security practices being followed should be demanded from the IAM SecaaS provider.

3.1.11 Scalability

IAM SecaaS providers should be able to dynamically scale up and down based on the requirements of the service consumer. This should be driven by consumer requirements. The demands on the cloud provider should be SLA driven and appropriate agreements should be in place to seek response and fulfillment based on the level of changes.

3.1.11.1 Access to Subscriber Data

The cloud provider should not have access to the subscriber data after termination. During the subscription period, the access by the provider needs to be determined by the subscriber via the SLA agreement and has to be subject to the relevant compliance requirements.

3.2 Concerns

This section discusses concerns such as points where data could be unencrypted, security of access to the data, separation of duties, and separation of logs when in multi-tenancy environments.

3.2.1 Standards/Openness/Vendor Lock-in

Lack of standards, and the resulting potential vendor lock-in, should be considered. If a technology area is still quite new, there may be no established standards. In those cases, openness is critical (availability of specifications/data formats/protocols), so that converters can be developed in the future.

Vendor lock-in is a general problem with using cloud services. Once IAM information is managed in one provider's service, it will be hard to export all information and import it into another provider's service. Standards may help somewhat, but most likely each provider will maintain meta-information that is vendor specific and/or non-exportable.

3.2.2 Identity Theft

Adequate security (encryption, authentication, access control, monitoring, etc.) should be in place for accessing identity and administration interfaces. This should be aligned with the requirement to protect not only the information assets but also identity specific information available within the services.

3.2.3 Unauthorized Access, Insider Threats, Fraud and Accidental Access

Identity information (or usage and other sensitive information) could be accessed by unauthorized IAM SecaaS provider staff or malicious intruders. Similarly, unless adequate security (encryption, authentication access control, monitoring, etc.) are in place for accessing identity and administration interfaces, unauthorized access over the network (eavesdropping) would be a risk factor.

Unintended government access may occur if the IAM SecaaS provider co-locates customer data in a way that cannot be easily disentangled. Cross-jurisdictional issues could occur due to such unauthorized access, esp. if provider and subscriber reside in different countries with strict privacy regulations, such as the EU.

Privacy across multiple tenants may potentially not be reliably preserved. Appropriate encryption and access measures should be explored to ensure proper separation of tenants' information and access.

Unintended access can be both fraudulent/malicious or accidental – assuming good intentions is therefore not good security practice.

3.2.4 Elevated Privilege Control

One of the most critical components of the IAM feature is Privilege Access Management. This applies to access by administrators, end-users with higher privileges and generally anyone who has or tries to gain access to privileged functions. Apart from the regular audit and logging requirements, privilege access needs to be a governance mechanism as well. Here, the provider and the subscriber need to define very clearly the management of access, roles, and what is expected to be retained by the provider even after the subscription has been terminated. Most enterprise applications suffer from mismanagement of roles, violation of segregation of duties and therefore reporting on this aspect becomes crucial. Apart from reporting, both success and failure events may need to be sent as an alert to defined roles/people within the enterprise. These also could be sent to the provider, depending on the type of contract existing between the subscriber and the provider in terms of managed services.

3.2.5 Non-Repudiation

It is always challenging to ensure true non-repudiation, outsourcing IAM to a SecaaS provider may make this even more difficult due to the trust boundaries between provider and subscriber. Examples of scenarios that might trigger non-repudiation concerns include:

- Login from multiple systems (smartphone, desktop/laptop)
- Known Impersonation during trouble-shooting
- Access from home desktop/laptop OR internet cafés which do not have a static IP

3.2.6 Least Privilege & Excess Privileges/Excessive Access

The IAM system should subscribe to the principle of Least Privilege and work forward using a workflow mechanism for additional approvals. True least privilege means that a user should be granted the least amount of privileges required to perform the current task. This can be established as a process and an IAM governance mechanism.

Because least privilege not only involves static roles and resources, but also complex context, and dynamically changes (both technical, changes, and non-technical changes), many IAM deployments today vastly overprovision effective access rights, and provide excessive access to users not requiring them (e.g., based on RBAC role definitions). This applies to both IAM SecaaS and traditional IAM.

To improve the implementation of least privilege, new/additional approaches that build on top of IAM and promise improvements could be considered, including those that take application security, fine-grained attribute-based access control (ABAC), authorization-based access controls (ZBAC), and policy automation technologies (model-driven security) into account.

3.2.7 Performance/Availability

Attacks on Identity Services or network connectivity, such as DDoS attacks or resource hogging, could jeopardize the availability or degrade the performance of an IAM SecaaS service. If high availability and/or performance are required, redundancy and fail-over options should be considered.

3.2.8 Features/Functionality Gaps/Weaknesses

The subscriber should make a list of the features desired in the IAM SecaaS system. A gap analysis can be conducted as a part of the solution evaluation to identify and provision for mitigation requirements.

3.2.9 Disaster Recovery and Business Continuity Management

IAM SecaaS may provide inherent high-availability due to the nature of the service. However, customers should not assume that cloud services come with in-built disaster recovery (DR) or business continuity management (BCM). In order to provide the same or similar level of DR as an on-premise IAM deployment, the customer should explicitly verify that the cloud provider meets or exceeds their DR requirements. At times, this may be

available at extra cost. However, without full DR capability, migration to cloud services could introduce additional technology or business risks which must be mitigated.

4.0 Implementation

From an enterprise perspective, when deploying any private and/or public cloud-based computing system(s), the same data access policies and regulatory conformance requirements need to extend into the private and/or cloud-based operations.

In publicly deployed cloud-based systems, privileged IT users may come from both the enterprise and the cloud service provider. To maintain conformance with regulatory requirements, privileged user access and entitlements for cloud services must be managed to conform to established enterprise data access policies. It is critical that the established SLAs between the enterprise and the cloud provider meet or exceed the enterprise's general requirements.

The agreement should provide for additional privileged accounts which can be defined in the service or infrastructure, but not directly associated with a specific user. These might be application or system accounts, emergency accounts, etc. Privileged accounts can be used either by a customer privileged user or by a service provider privileged user. It is important to control privileged accounts and ensure they are properly managed.

While many aspects of providing Identity and Access Management remain the same, whether provided in-house or in the cloud, there are a significant number of differences that must be accounted for and appropriately managed throughout the data lifecycle. Table 1 presents an overview of the impact on IAM's core functionalities when delivered in a SecaaS environment.

CORE FUNCTIONALITIES	IMPACT WITH CLOUD: SECURITY-AS-A-SERVICE
Provisioning/deprovisioning	Requires significant modification with newer and more standard technologies and protocols.
Authentication	Advanced forms authentication is preferred. Stronger authentication forms such as multi-factor, risk-based, knowledge-based or adaptive authentication, etc., may become a standard rather than just a "nice to have" requirement.
Directory services	With multiple repositories and authoritative sources, Directory Services would be comprised of a heterogeneous variety of information repositories, with Directory Services providing a unified view.
Federated SSO	Core; basic cloud functionality
Web SSO	Basic traditional functionality may need changes in a cloud-based solution.
Authorization	This will depend on the existing enterprise applications and their ability to inherit from the current module of authorization for the IAM system.
Authorization token management and provisioning	This will depend on the existing enterprise applications and their ability to inherit from the current module of authorization for the IAM system.
User profile & authorization management	This will depend on the existing enterprise applications and their ability to inherit from the current module of authorization for the IAM system.
Support for policy & regulatory compliance monitoring and/or reporting Federated Provisioning of Cloud Applications	Significant impact in a SecaaS environment
Self-Service request processing, like password reset, setting up challenge questions, request for role/resource etc.	Moderate changes required to the functionalities, however a number of enhanced features and technologies are available and need to be integrated to provide a better user experience.
Privileged user management/privileged user password management	Changes are required to the provider side of the service.
Tamper proof audit	Significant changes are required in a cloud SecaaS environment
Policy Management	Significant changes may be required because the IAM owner/operator is different than the subscriber/user. Control of policies may be very limited compared to traditional IAM deployments. Due to the more flexible nature of cloud mash-ups, policy management needs to be more flexible. Due to the changed organizational trust boundaries when using IAM SecaaS, policies may need to be much more fine-grained and contextual.
Role-Based Access Control	Changes to the RBAC model are required with cloud services.

Table 1: IAM Differences in the Cloud

4.1 Architecture Overview

Cloud-based Identity and Access Management architectures are similar but in many ways different from traditional on premise IAM. Characteristic similarities include the core principles of information security: confidentiality, integrity and availability. Serving the specific security needs of a business, and ensuring regulatory compliance continue to be central to IAM business objectives.

What really has changed with the cloud is the speed, flexibility and capability of IAM to be a business enabler, rather than a specific operational function. In terms of this new paradigm, IAM architecture spans across businesses, opening up a plethora of options to expand the portfolio of services that the business offers. This not only enhances the ability to provide services within the organization, but also the ability to collaborate with other businesses in more elegant and efficient ways. While a complete adoption of IAM SecaaS is in the early stages of adoption, it is expected to be a part of a standard enterprise architecture within a very short time.

The relationship between business applications and IAM SecaaS interaction can be defined in one of two ways: IAM in the cloud and IAM to the cloud. These two are complementary architectures.

The sample architectures discussed below are representative of the many possible ways to build an effective IAM system.

4.1.1 Conceptual Architecture

The conceptual architecture of IAM Security as a Service involves the combination of various traditional IAM capabilities into logical layers of services.

- **Access and Policy Services:** Various levels and depths of authentication, authorization and access policy management
- **User Services:** Provisioning, self-service password reset, delegated administration and centralized user administration services
- **Identity Services:** Identity data sync from authoritative sources, user data synchronization and correlation, password policy enforcement and sync, identity virtualization, impersonation and transformation (from a person level to an enterprise/group/corporate level)
- **Compliance Services:** Regulatory and policy compliance, maintaining audit trails, monitoring security events and reporting functionalities
- **Data Services:** Centralized identity and entitlements repository, providing business intelligence in terms of user identities, and providing access reporting data for operations as well as compliance

These layers of IAM SecaaS services interact with the entities from within and outside the organization using the services.

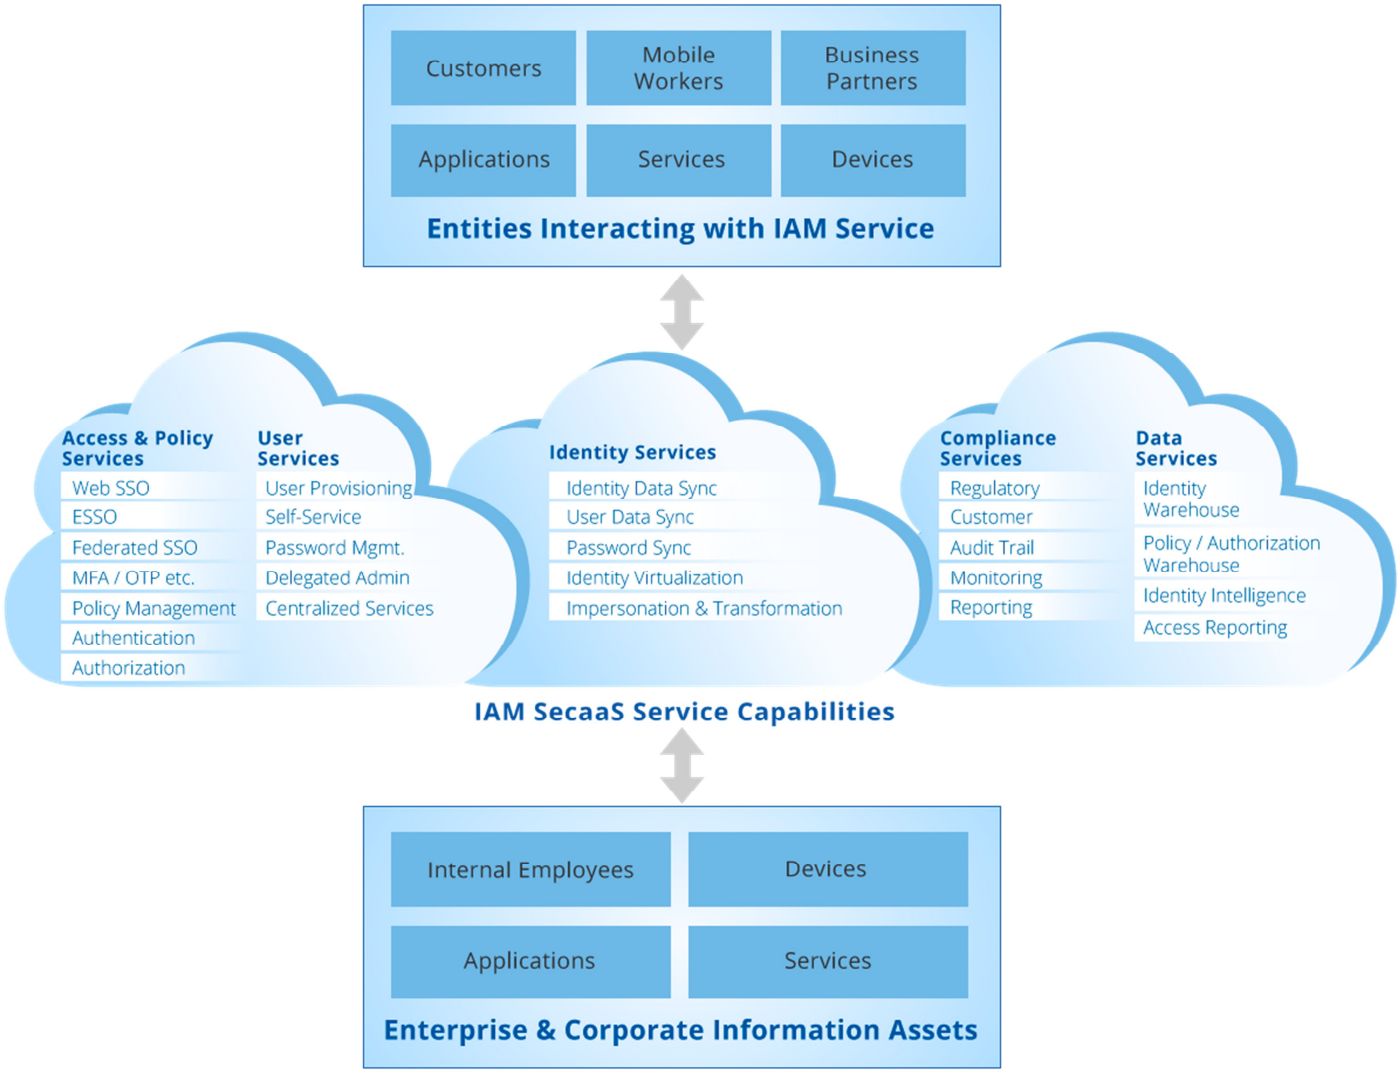

Figure 2: IAM Conceptual Capability Architecture

4.1.1.1 Cloud Implementation Architecture

IAM services can be provided by specialized service providers which cater to a specific layer or as a combination of one or more functions. Each layer can be abstracted and implemented along with the other complementary capabilities. These services work with each other in a cohesive manner by communicating through standards-based interfaces, such as SAML and WS-FED.

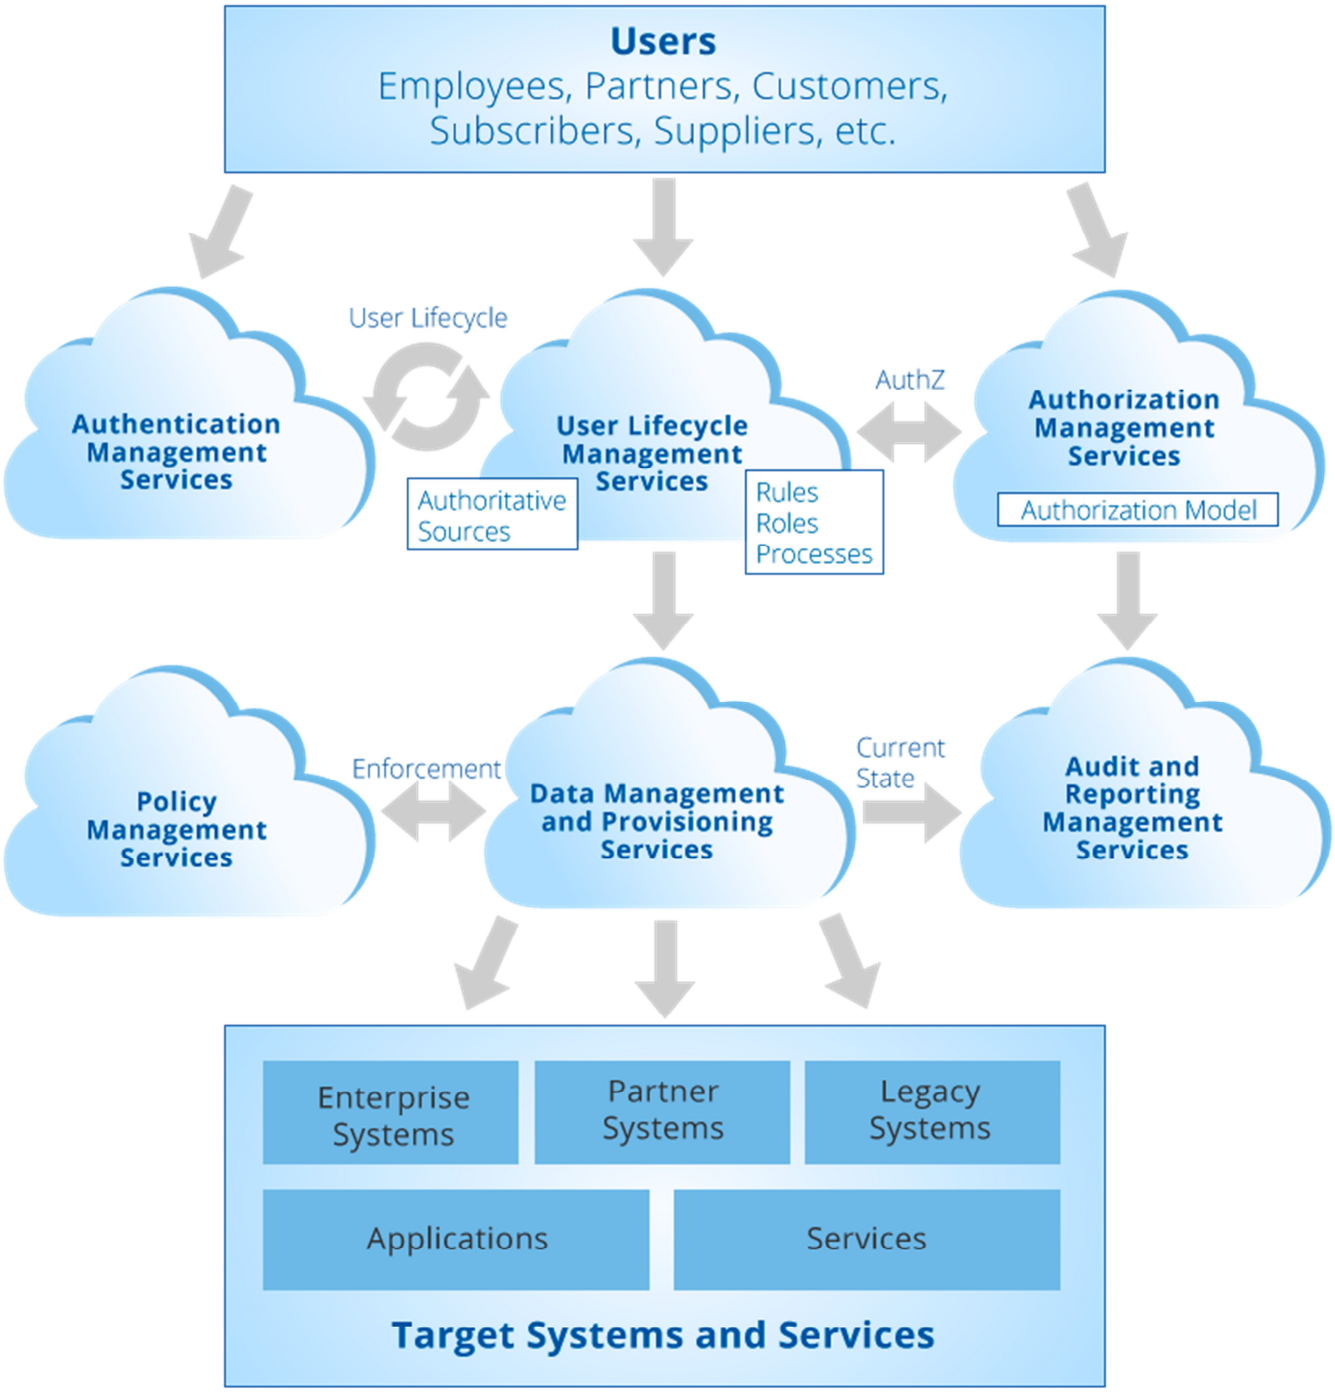

Figure 3: IAM SecaaS Independent Service Implementation

IAM Services also could be architected as consolidated and integrated service groups providing a single IAM service. These implementations use the security gateway architecture and typically are implemented by consumers using heterogeneous on-premise and cloud-based business applications.

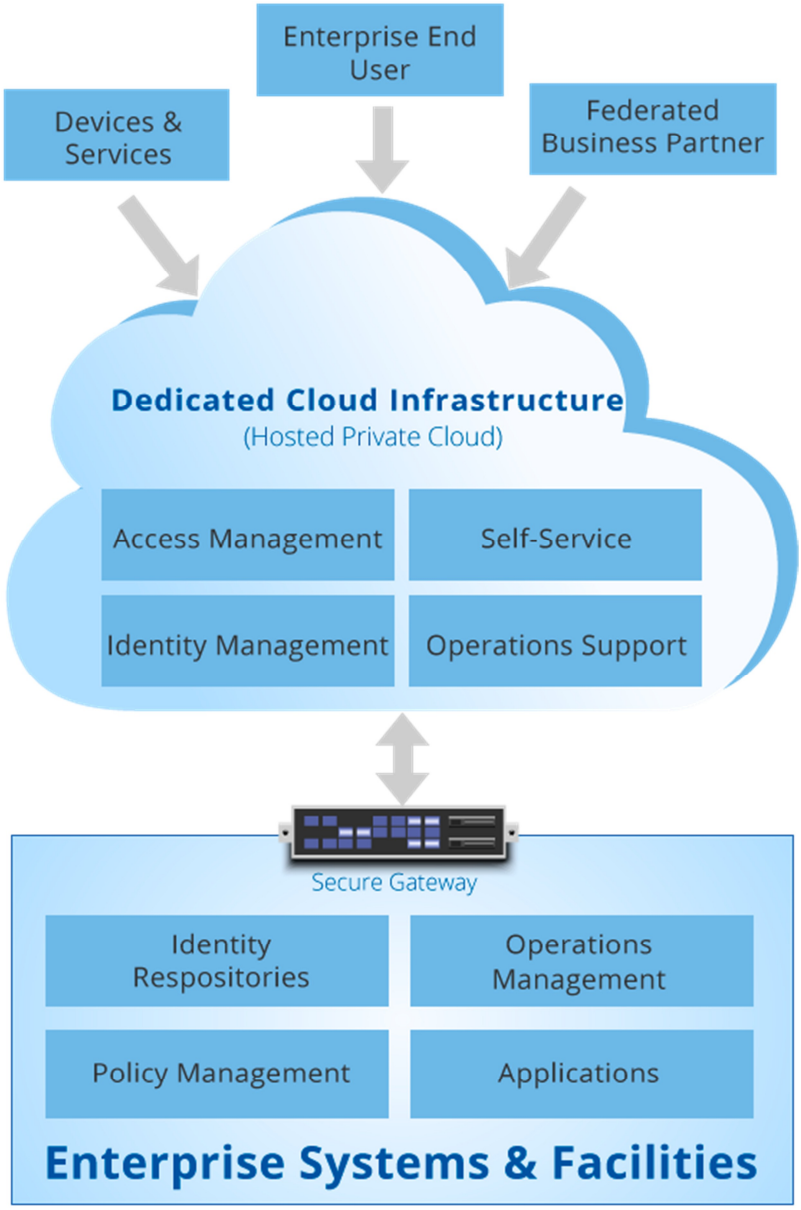

Figure 4: IAM SecaaS Consolidated Service Implementation

Apart from general user management, IAM SecaaS also considers privileged users and system administrator access to the IAM environment itself. There are several reasonable architectural paths; with the flexibility of IAM SecaaS, and the specific implementation chosen depends on what is appropriate to the business in question. Multi-factor authentication using One-Time Password (OTP) is a prevalent and accepted solution in the industry, which ensures that the level of protection for the IaaS and PaaS environments are dealt with extra security. Other options of strong authentication mechanism include Adaptive and Risk-Based Authentication and should be included in the consideration for providing authentication mechanisms for IAM SecaaS.

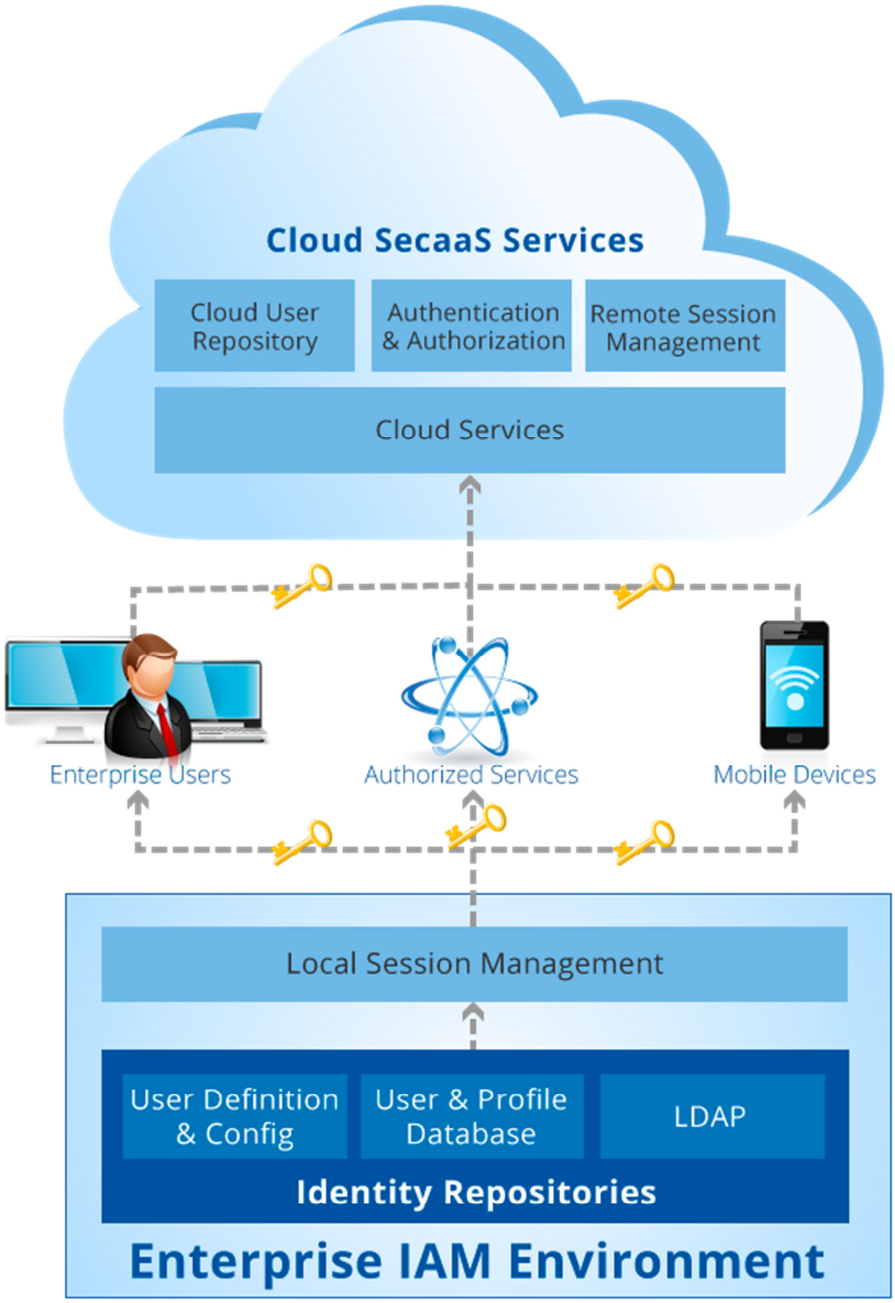

Figure 5: Remote Session Management of Cloud SecaaS Services

4.2 Guidance and Implementation Steps

4.2.1 Provisioning and Deprovisioning of Accounts

IAM services include provisioning and deprovisioning of accounts (of both cloud and on-premise applications and resources). Provisioning refers to the process of the creation of user accounts in cloud platforms so that users can access appropriate systems and resources in the cloud. The goal of user provisioning is to streamline the process of account creation and provide a consistent framework for providing access to end users.

In the cloud computing environment, provisioning as a service in the security domain means creating and managing the lifecycle of a user account according to the security requirements of the enterprise, providing appropriate access to computing and platform resources.

4.2.1.1 Provisioning

Provisioning of in the cloud refers to the management of the creation, modification, and revocation of user accounts according to defined security policies of a specific managed end point. The Provisioning component is also a means of propagating security policy, for example by setting access rights on managed cloud end points systems based on group memberships and/or role assignments. While provisioning in the enterprise accommodates for both automated and manual requirements and accommodates for the speedy or the lack of it within the enterprise, SecaaS IAM provisioning services need to maintain the agility of the cloud-based environment. This includes for accommodating the dependencies on the authoritative sources, identity silos feeding the user repository and provide a seamless experience providing user with timely access to the user.

4.2.1.2 Deprovisioning

Deprovisioning is the process whereby user accounts are disabled when data access is no longer applicable and appropriate for the specific user in question. This may be due to users leaving the organization, transferring within an organization, changing roles, etc. In the cloud computing environment, deprovisioning refers to the termination or disabling of user accounts on managed end points. This may at times necessitate the removal of data associated with the user based on business and regulatory requirements relevant to the business.

Services Required for Functionality – Automated removal of user accounts, workflow to check accidental removal of access, centralized user management, delegated user management, synchronization with authoritative source

4.2.2 Authentication

Authentication is the process of verifying the credentials of a user who is trying to access a protected resource. A user id password-based authentication is the predominantly implemented solution. While this serves as a sufficient mechanism for most business needs, it is susceptible to many known threats.

Authentication is critical to protect cloud resources. It verifies the identity of the entity (user or machine) who wants to access the cloud resources, and it must be done in a secure, trustworthy, and manageable manner.

The mechanisms through which a user can be authenticated fall into three categories. These factors have been widely adopted throughout the security implementations, and are the accepted standards for identity authentication:

- Something the user knows (e.g., password, PIN),
- Something the user possesses (e.g., ATM card, Smart card), and
- Something the user is (e.g., biometric characteristic such as a fingerprint or retinal pattern).

In a cloud platform, authentication services should include one of the multi-factor mechanisms for verifying the credentials of the user. Two factor authentication is predominantly becoming a minimum standard in the security implementations. SecaaS IAM Services should consider this and provide this as a minimum functionality with the option to extend to additional factors based on business requirements.

Authentication supports or works with other security features in the cloud. For accounts that require higher levels of security, multiple factor (strong) authentication may be a necessity. Strong authentication requires the use of two or more different authentication factors, e.g., a PIN and a fingerprint, or a password and an ATM card. For example, a client application may have (or can access) a secret key to use to sign the API requests, or a client application can authenticate to the cloud using a private key.

Risk-based authentication is an assessment of the risk involved in trusting the agent requesting a transaction, based on an assessment of behavioral, geographic, timing, and other factors.

At one end of a transaction, risk level typically is calculated for the transaction itself. If the calculated risk is low (a small-value purchase or a request for publicly available data), the level of authentication required is correspondingly low. If the risk level assigned to a transaction is high, additional authentication measures may be required.

On the other end of the transaction, the level of risk is calculated for *the user*. Geographic, behavioral and timing factors are among those calculated into the risk rating for a given user at a given time. If this user made a physical purchase in a different time zone within a few minutes of instigating the current transaction, the transaction may be declined, or the user may be asked for additional information. If an online transaction is attempted at a vendor from which the user typically purchases, the transaction likely will be approved. If a transaction is attempted that is outside the typical geographic location of the user, the transaction may require additional authentication, such as a password that has previously been established for the account.

A variety of behaviors, location information (including trending fraud patterns), timing, history, transaction risk level, and other factors are used to determine the overall risk for any given transaction.

SecaaS IAM implementations should provide mechanisms for strong authentication which include features such as multiple challenge response based on the context of the transaction. Including a multiple factor within the context of strong authentication is highly recommended accounting for the actual strength of the authentication mechanism with something that a user knows, user has or the user is. Using strong authentication provides for a higher level of protection of the target systems and the information contained in the cloud environment.

SecaaS IAM implementations also should seriously consider Risk-Based Authentication where appropriate and suitable, especially for financial transactions over the cloud.

4.2.3 Directory Services

A directory service provides an organized repository of information stored and identified by a unique identifier and location. Lightweight Directory Access Protocol (LDAP), based on the X.500 standard, is one of the primary protocols for directory services. Each entry in an LDAP directory server is identified through a Distinguished Name (DN).

In a cloud environment, directory services would continue to be heavily utilized by the Identity and Access Management framework as a security repository of identity and access information. Directory services can be used by identity and access management systems tightly integrated as a unified framework, or they can independently serve multiple IAM service consumers as a trusted source of user attributes for security functions.

SecaaS IAM implementations should provide mechanisms for strong authentication, which include features such as multiple challenge response based on the context of the transaction. Including a multiple factor within the context of strong authentication is highly recommended, accounting for the actual strength of the authentication mechanism with something a user knows, has or is. Using strong authentication provides for a higher level of protection of the target systems and the information contained in the cloud environment.

SecaaS IAM implementations should seriously consider Risk-Based Authentication where appropriate and suitable, especially for financial transactions over the cloud.

4.2.4 Directory Synchronization (Multi-Lateral as Required)

Directory synchronization is the synchronization of directory objects (users, groups, and attributes) from an on-premise User Directory (such as an Active Directory or LDAP) environment to the directory infrastructure of a cloud platform. These tools usually are installed on a dedicated computer in an on-premises environment.

Directory synchronization tools usually have an initial load of user data, and then accept additions, changes and deletions in regular feeds. As long as this synchronization process is in place, the cloud environment and applications can provide the same level of authentication and authorization that the on-premise applications can make.

During the implementation of directory synchronization, proper consideration of authoritative sources is essential. This, along with other IAM resources in the cloud ecosystem for a particular implementation, would ensure that the source of data in the cloud repository is accurate and appropriate. One-way synchronization from the authoritative source to the cloud end point is preferred, to ensure that the authoritative source data is not directly influenced by the cloud repository. The enterprise may consider having the cloud as the authoritative source of data. That way, even from their home, the authoritative data can be managed. This is where enterprises can leverage the advantage of the cloud and have a two-way sync established.

A synchronized directory on a SecaaS platform ensures several benefits such as enabling Single Sign-on into the cloud service. This also ensures that the round trip traffic for several security events can be avoided and carried out in the cloud. Synchronization makes it convenient for the cloud service provider to offer instant access to services for any existing user, and also change/modify the service components. Authorization can be carried out to the same level of granularity that would have been possible on-premise.

4.2.5 Single Sign-On

Traditionally, Single Sign-On has been a process that allows the user to access multiple applications requiring authentication by passing his credentials only once. The user first authenticates to some trusted authentication authority and then is granted access to all the applications trusting that authority. The applications only receive information about whether they may let the user in or not.

Since the user authenticates only once, exposure of sensitive information over the network is limited. An added benefit is that SSO systems usually redirect users to secure communication channels. SSO systems usually preserve the state of the user for some period of time, so the user may repeatedly access these applications without the need to authenticate each time.

Single Sign-On in the cloud is an extension of the existing web Single Sign-On in the on-premise application. As enterprises expand beyond their on-premise applications, there is expanding demand. The resulting implementation of single sign-on between the enterprise on-premise applications and cloud-based applications constitute an extended enterprise. A number of the use cases of Single Sign-On are provided by the federated model, with underlying standards such as SAML, OpenID, WS-Federation, etc. These implementations are predominantly federated SSO based on SAML 1.1/2.0 or OpenID. There are a number of Single Sign-On providers which provide out of the box connectors for major enterprise applications as well as cloud applications. While implementing Single Sign-On solutions, it is essential to understand the mechanism and keep it consistent across all integration.

Provisioning of identities in a store outside the enterprise poses several challenges for security and privacy. Most organizations find themselves in a hybrid environment, with both on-premise and hosted identity components. Identity bridges are on-premise appliances that enable identity services across a hybrid computing infrastructure. Identity services include directory synchronization, federation, mobile credential management, traditional provisioning, Web access management (WAM) and Extensible Access Control Markup Language (XACML) authorization.

4.2.6 Federated Single Sign-On

Federation provides the ability to share user identity and access information between multiple domains, which may be within the same or disparate IT infrastructures and organizations. Federated Single Sign-On allows multiple organizations to provide their services in a collaborative environment in a secure manner. Federated SSO, when properly implemented within a strong authentication domain, provides both security for the organization, and ease of use for the cloud consumer.

In a hybrid environment, Federated SSO can be deployed to secure configurations such as:

- Outbound SSO for users to access software as a service (SaaS) and business process, outsourcing (BPO) providers, and to connect with trading partners
- Inbound SSO for service providers, such as BPOs and managed services, to access the enterprise's resources
- Internal SSO for the enterprise and its acquisitions, affiliates, subsidiaries and joint ventures
- SSO to a third-party, hosted hub for users to share information among industry organizations

4.2.6.1 Standards

Federated Single Sign-On can be achieved using industry standards like SAML, WS-Security, Open ID and OAuth. When deploying Federated SSO, it is important to decide which standard to use, based on the use cases to be supported.

- SAML is the predominant protocol for browser-based identity federation, especially in B2B and employee-facing use cases.
- The composability of WS-* means that the profiles for interoperability are not as well-defined. If you wish to enable both identity-based Web services as well as browser-based interactions, then some combination of SAML and WS-Trust may be the right choice.
- OpenID is a relatively new SSO mechanism. OpenID provides a decentralized SSO model. OpenID requires the reliance on a third-party Identity Provider (IdP) to confirm the identity of the user requesting site access and this should be taken into consideration with respect to business and compliance requirements.
- OpenID Connect is a suite of lightweight specifications that provide a framework for identity interactions via RESTful APIs. The simplest deployment of OpenID Connect allows for clients of all types, including browser-based, mobile, and JavaScript clients, to request and receive information about identities and currently authenticated sessions. The specification suite is extensible, allowing participants (optionally) to support encryption of identity data, discovery of the OpenID Provider, and advanced session management, including logout.
- OAuth provides a method for users to grant third-party access to their resources without sharing their passwords. OAuth has been replacing the traditional proprietary authentication tokens schemes used to different SaaS and PaaS APIs. A determination of implementing either HTTP Basic Authentication or one of the above mechanisms including OAuth should be made depending on the functionality required. While HTTP Basic Authentication is acceptable for desktop or mobile access, a mechanism such as OAuth becomes necessary for web application and services interaction.

The Identity Provider (IdP) needs to provide several capabilities in order to satisfy the requirement of SSO in the cloud.

1. The first capability is federated access via an IdP or an identity federation setup in the cloud.
2. The second capability is the replication of user identities from enterprise user directories into provider directories and applications.
3. For Identity providers providing features for User account management and updating access privileges and other attributes is an important consideration.

4. A security token service will also add value for an Identity provider for providing a variety of token exchange to support multiple standards and support customers for various different technologies for authentication and Single Sign-On.

Identity providers must support multiple versions of federation standards as not all environments use the latest standards and backward compatibility should be provided for in terms of standards.

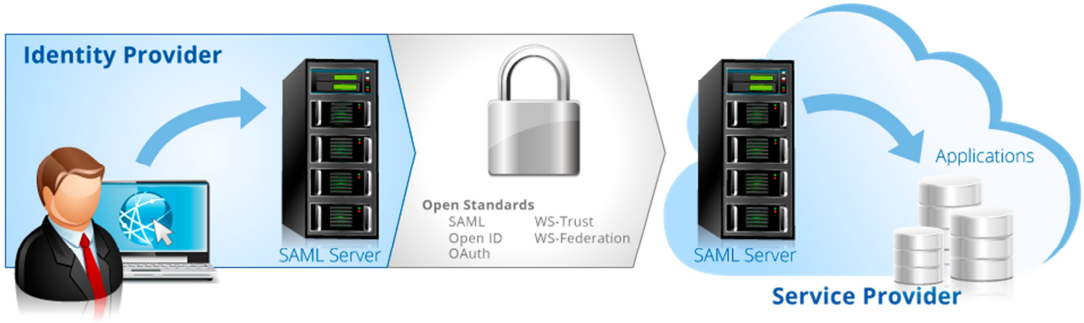

Figure 6: Considerations of Standard for IdP and SP Interaction (SAML is shown as an implementation example)

4.2.7 Web Single Sign-On

In the cloud environment, there are often a number of web applications and services designed to aid users, thus requiring authentication. In some cases, it may be convenient and secure to use a centralized SSO infrastructure bound to the central authentication authority. Web Single Sign-On provides SSO infrastructure for web applications. However, Web SSO has proven to be a very fragile mechanism for access management. Web SSO may be used for authentication of users for generic systems, but for business applications, a Federated SSO is highly recommended.

4.2.8 Authorization (Both User and Application/System)

Authorization is an important mechanism for access control. Both users and applications or systems must be authorized in order to access resources or other services. Authorization requirements include establishing trusted identity profiles and creating (or reusing existing) access control policies.

Authorization and access control are of key importance to a sound cloud security implementation (for PaaS, IaaS and SaaS). Cloud security should include both coarse-grained and fine-grained authorization, and should support authorization both for human users and machine-to-machine interactions. The authorization feature has to work with the other security features, especially authentication, message protection, and incident monitoring.

Authorization policy requirements should be defined centrally in a form that can be understood and managed *by human administrators.* Due to the dynamic nature of access in the cloud, authorization policies should be easy

to govern and manage. Manually managing a purely technical authorization rule set is extremely difficult; there are too many rules to manage and update on a regular basis, which makes a manual approach time-consuming and error-prone. This challenge grows with increasing system size and/or cloud mash-up interconnectedness, because authorization rules will be needed to control access between many nodes.

Alternate processes include:

- Model/metadata-driven security policy automation approaches are used in "model-driven security" products, which infer and generate the matching technical authorization rule set with as little human intervention as possible.
- Intelligent grouping mechanisms for resources which should be made available in a web service directory hierarchy
- Specific visualization tools which can be used to simplify the generation of the technical rules

Technical authorization rules should be stored centrally for each trust domain in a standards-based rule repository. This will enable consistent, unified, manageable administration. An exemplary standard for technical authorization rule representation is XACML.

Depending on performance/robustness requirements, Policy Decision Points (PDPs) can be deployed centralized or decentralized. If a decentralized/local Policy Enforcement Point (PEP) is used, technical authorization rules should be distributed from the policy repository, serving as a the Policy Administration Point (PAP) across the PDPs located on each cloud node on which the policy should be enforced by Policy Enforcement Points (PEPs). The distribution mechanism can either directly push access rules to the PDPs (i.e., messages between the policy repository and the PDPs of the protected resources), or can generated authorization tokens with the embedded policy to users, who can send those tokens along with their cloud service requests ("authorization-based access control," ZBAC). If a centralized PDP deployment option is chosen, the PEP will need to query the PDP whenever an access decision should be made.

Technical authorization rules should be determined and enforced at runtime on the protected cloud node for all incoming requests by PDPs & PEPs. PDPs will frequently require access to identity and authentication related information sources (so-called Policy Information Points, PIPs) to make decisions.

Authorization related incidents must be reported and logged at runtime in a timely fashion.

4.2.8.1 Authorization Token Management and Provisioning

One way to handle authorization is through the use of authorization tokens that enable applications and processes to access protected cloud resources. Many major websites and organizations use OAuth 2.0 protocol for this purpose. The IETF OAuth Working Group currently is developing the OAuth 2.0 specification. OAuth 2.0 provides specific authorization flows for various clients including web applications, desktop applications, and mobile devices. To access a protected resource at a resource server, a client obtains an authorization token (access token) from the authorization server with the resource owner's approval, and then uses the token to access the resource.

For websites and organizations that require higher levels of security protection for their resources hosted in the cloud, the Security Assertion Markup Language 2.0 (SAML 2.0) may be a better choice. SAML 2.0 is an OASIS standard for exchanging authentication and authorization information cross different security domains. In this case, a SAML authorization assertion is the token for accessing protected resources.

The authorization token must be securely managed and provided to the correct recipients to prevent unauthorized access to a resource. The token must be managed throughout its lifecycle, which includes creating, storing, using, refreshing and deprovisioning.

4.2.9 Support for Policy and Regulatory Compliance Monitoring and/or Reporting

Misuse of user access is a major cause of breaches, and leads not only to loss of reputation, but also financial loss in terms of penalties and loss of customers. Various regulations require all user access and activity to be controlled and monitored. This includes implementing a system to monitor user accounts, as well as monitor the activities performed by these accounts.

4.2.9.1 Federated Provisioning of Cloud Applications

- Self-Service request processing, like password reset, setting up challenge questions, request for role/resource, etc.
- Privileged user management/privileged user password management
- To comply with regulatory mandates, all sufficiently sensitive IT operations must implement user and root user control policies, with conforming management and control functions in place to secure systems and mitigate external and insider threats.
- From an enterprise perspective, when deploying any private and/or public cloud-based computing system(s), the same data access policies and regulatory conformance requirements should extend into the private and/or cloud-based operations.
- In publicly deployed cloud systems, privileged IT users may come from both the enterprise and the cloud service provider. To maintain conformance with regulatory requirements, privileged user access and entitlements for cloud services must be managed to conform to established enterprise data access policies. It is critical that the agreed-upon internally or externally established SLAs between the enterprise and the cloud provider meet or exceed the enterprise's general requirements.

It also should be noted that additional privileged accounts can be defined in the service or infrastructure, which are not directly associated with a specific user. These may be application or system accounts, emergency accounts and other types of accounts. These accounts can be used either by a customer privileged user or by a service provider privileged user. It is important to control these accounts and ensure they are properly managed.

4.2.10 Access and Activity/Session Monitoring

Various regulations require privileged access and activity to be controlled and monitored. This includes implementing a system to monitor all use of privileged accounts, as well as monitor the activities performed by these accounts. In some cases, monitoring can also extend to regular user accounts.

4.2.11 Tamper-Proof Audit

All collected audit data, especially data related to privileged account use and activity, must be stored in a tamper-proof fashion. This aspect supports non-repudiation, enabling the completion of an audit trail, collection of forensic evidence and enforcement of SLAs and internal policies.

4.2.12 Policy Management

Policies must be managed in such a way that they are relevant and understandable to the business, but at the same time must be combined with the matching technical enforcement. Policies that express business requirements and understanding are useful as an intermediate step, but are ultimately only useful when concretely enforced. Similarly, technical policies are only useful if they reflect business requirements, and if they can be reliably enforced and monitored.

In a cloud context, policy management deviates from traditional on-premise deployments in that the stakeholder that needs to manage, understand (and sign off) the policy is typically different from the stakeholder that is responsible for the technical implementation of that policy.

A particular policy management SecaaS is "Policy as a Service," where policies are managed and distributed by a Policy SecaaS (e.g. for access control), and enforced within a cloud or on-premise environment. The policy can either be very technical (e.g. fine-grained authorization rules for direct technical enforcement) or closer to the business (e.g. generic security policy models that are turned into the matching technical rules and configurations using model-driven security approaches and then enforced).

It is absolutely critical that policy management is actually manageable. While this may sound as if it is stating the obvious, this is typically highly complex in today's interconnected, dynamically changing IT world. In a cloud environment with potentially many interconnected modules, keeping policy management is consistent and unified often improves manageability.

4.2.13 Role-Based Access Controls (RBAC)

Role-Based Access Control (RBAC) is the process of assigning pre-defined roles to users in place of granular access entitlements or privileges. Usually roles are structured in a hierarchical fashion and comprised of application access and underlying entitlements. This process makes the user access lifecycle management very efficient, and also ensures users have clean access. Currently, businesses using cloud computing environments most likely see a heterogeneous mix of cloud and on-premise applications, resulting in a wide variety of access models. RBAC can very effectively streamline this process by assisting in providing structured grouping and definition of access to various applications, which may include both cloud and non-cloud applications.

Not many cloud applications are built with RBAC in mind, and most grant access in terms of fine grained entitlement and privileges. The idea of RBAC is twofold: 1) to abstract the level of details from the user while assigning access, and, 2) to have a true structure of access predefined. The benefit of RBAC supporting cloud applications is that while assigning access to users, the requestor or approver does not need to know if the application is cloud hosted or not. This not only creates a layer of abstraction for the end user, but also makes it easier for the requestor/approver to request the roles and not worry about the underlying details.

RBAC includes a complete Role Lifecycle, which includes Role Mining, Role Engineering, Role Assignment, Role Certification, Role Maintenance and Role Decommissioning. This forms a very comprehensive process of user access definition and proactively eliminates the issues of redundant or excessive access.

Just like traditional enterprise RBAC for on premise infrastructure, RBAC for cloud-based applications uses application roles and access control lists. This ensures that a cloud application can be integrated well into an existing RBAC implementation for provisioning.

One of the deviations from RBAC from on-premise or traditional computing to that in the cloud is that users not only have the ability to request/provision access, but they also have the power to request/provision infrastructure or complete application environments. RBAC, with respect to cloud infrastructure, would ensure that the cloud management and orchestration software packages would guarantee that these functionalities are encapsulated in the form of roles. For example, an administrator for a cloud-based application would be able to create and manage (edit) virtual application instances/environments, while users would have the ability to instantiate application environments from their existing read only copies available through cloud management software.

Another cloud-based RBAC implementation in the realm of cloud environment provisioning is the ability of application owners from a specific department within an organization to define virtual network boundaries, creating security zones within which applications would run and be accessible by users. This would be a level of delegated administration that the application owners would have based on their roles.

One of the unique requirements and RBAC implementation specific to a multi-tenant architecture would be restrictions based on client guest management requirements. In this implementation scenario, the access to specific cloud environments in a given operations center would be based on the roles of the specific staff. The request processing, approval and fulfillment of the cloud environment requests would be done based on the RBAC-based roles of the operations staff.

Traditional user access provision and management is mostly related to user specific access and granularity; however, this becomes less effective and efficient with cloud as the scale, flexibility and granularity of that access increases manifold, so thorough consideration should be given to RBAC in the security model and architecture in cloud-based environments.

4.2.14 Centralized Directory Service

Directory service is one of the basic blocks of security in enterprise and in the cloud. Directory service is the natural choice for an Identity Provider. It assists other features from the security stack, from provisioning to

fraud detection. Directory services aim to prevent identity proliferation and maintain a secure authoritative source of identity. In the cloud, a directory service offering should be extended to virtual directory services.

A directory service provides an organized repository of information stored and identified by a unique identifier and location. Lightweight Directory Access Protocol (LDAP), based on the X.500 standard, is a primary protocol for directory service. Each entry in an LDAP directory server is identified through a Distinguished Name (DN). In a cloud environment, directory services would continue to be heavily utilized by the Identity and Access Management framework as a security repository of identity and access information. Directory services can be used by identity and access management systems tightly integrated as a unified framework, as well as independently serving multiple IAM service consumers as a trusted source of user attributes for security functions. Cloud-based directory services can abstract the underlying directory platform and provide a uniform interface for user information look up, regardless of whether the underlying directory is LDAP-based, a virtual directory, meta-directory or a database.

4.2.15 Privileged Accounts

When implementing a cloud-based service, it is important to identify the various privileged accounts and the users who will have access to them. Such accounts may include:

- **Customer administrative accounts** – User accounts used by the customer to access and control the cloud service. Typical privileges include configuration changes and management, introduction of users, acquiring and reducing resources and functionality, changing the scope of consumed services etc.
- **Service provider administrative accounts** – Accounts used by provider employees to configure and administer the service provided to customers. These accounts hold the highest level of privileges, controlling the infrastructure of the service.
- **Application accounts** – Accounts used by customers' applications to perform operations in the cloud environment.
- **System accounts** – In an IaaS service, these accounts are defined in the virtual machines themselves.

Unlike user accounts, which operate within the limits and roles granted by the service to perform their business function, the privileged account types described above have the necessary permissions to substantially change the service provided, the infrastructure on which the service is running and the consumption of the service.

4.2.15.1 Privileged Account Control

Privileged accounts require additional controls to prevent abuse and ensure security, including:

- **Control** – control over the privileged access, use of privileged accounts and privileged actions,
- **Monitoring** – including logging and continuous monitoring of the privileged session, and
- **SIEM (Security Information and Event Monitoring systems) Integration**

It is worth noting that different privileged accounts may require different authentication methods such as certificates for application accounts, strong authentication for administrator accounts, etc., and different monitoring and control requirements.

4.2.15.2 Privileged Account Management

A privileged account management system comprises, at least, the following:

- **A secure credentials repository** – customers should look for an encrypted, highly secure solution, as the stored credentials hold the highest level of privileges
- **Automatic policies and workflows** – the system should support business and security workflows. For example, a common regulatory requirement is scheduled credentials replacement. Other policies may require a change after every use (creating one-time passwords), integration with ticketing systems (only allowing the use of the privileged credentials when there is an open ticket), dual-control (requiring additional approval for such use), and so on.
- **An interface to access the stored credentials** – the system authenticates the user and only provides the credentials that user has permissions for.
- **Monitoring, logging and auditing** – the system logs every access and use of the managed credentials, stores the information in a tamper-proof repository and facilitates audit processes.

4.2.15.3 Additional Aspects

Additional aspects include:

- **Automatic privileged account discovery** – this is very important in virtual environments, on which most cloud services are based, as they are highly-dynamic in their nature and require constant discovery of new systems and their relative privileged accounts.
- **Integration with SIEM** – as privileged access is more sensitive in nature, it is highly relevant for integration with organizational SIEM systems to discover and alert on various security threats.

Understandably, it is a challenge for the cloud customer to control the use of privileged accounts by service provider personnel. A common way to control this use is to employ contractual obligations (SLAs), requiring the service provider to implement a privileged account management system. Such requirements are common for insurance and audit purposes, both for cloud service providers and customers.

5.0 References and Useful Links

5.1 References

Federal Financial Institutions Examination Council. (2001, August 8). Authentication in an Electronic Banking Environment. Retrieved from http://www.ffiec.gov/pdf/pr080801.pdf

Geyer, C. (2009, December 23). SAML Specifications. Retrieved from http://saml.xml.org/saml-specifications

Recordon, D. and Hardt, D. (2012, May). The OAuth 2.0 Authorization Framework (draft-ietf-oauth-v2-26). Retrieved from http://tools.ietf.org/html/draft-ietf-oauth-v2-26

5.2 Useful Links

Walking from Cloud to Cloud: The Portability Issue in Cloud Computing
http://digital.law.washington.edu/bitstream/handle/1773.1/447/Carpenter,%206%20Wash.%20J.L.%20Tech.%20%26%20Arts%201.pdf?sequence=3

Federated Identity Management
http://en.wikipedia.org/wiki/Federated_identity_management

Getting to the Problem of the Root
http://www.gartner.com/it/content/668800/668817/ks_sd_may.pdf

Access Policy Management: Authorizing for Access
http://technet.microsoft.com/en-us/library/cc700801

