Chapter 2

Identity and Access Management

There are many configurations of identity and access management (IAM) systems, and to some extent, each organization’s IAM system will be unique, developed and deployed to suit the organization’s own specific requirements. But all systems do need a provisioning tool, a data store, and one or more access control facilities.

IAM Core Capabilities

There are six core features of an identity and access management (IAM) environment; this chapter will look at how they have developed over the past few years and what should be expected of any new deployments. Figure 2.1 illustrates a simplified IAM environment, showing these core features.

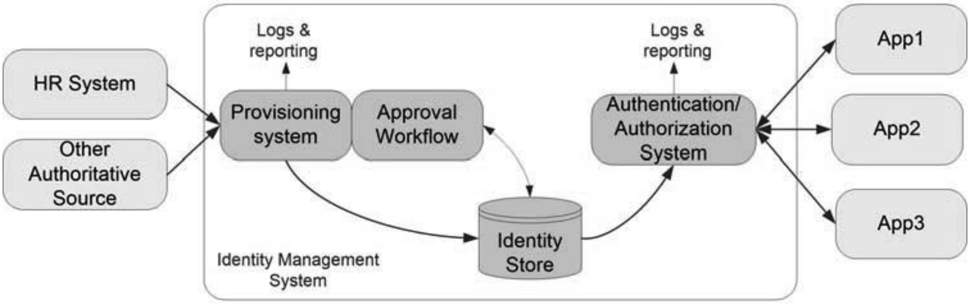

Figure 2.1: Simplified IAM environment

Table 2.1 summarizes the capabilities of the six core features. (We’ll explore these in more detail later in this chapter.)

Table 2.1: Core features of an IAM environment

Core Feature	Description
Provisioning and human resource management system (HRMS) integration	The input to an identity store is via some form of provisioning tool. This can be an application form that is placed on a help desk queue or sent to a system administrator to be actioned. Preferably, this will be an online form, which can 1) ensure that all the relevant data has been collected and 2) apply some “rules” to determine the data repositories to write user data to. In some cases, this will be to an access control list (ACL) within an application. In other cases, it will be an Active Directory (AD) group, and in other cases it will remain in the identity store for access by an authentication system.
Approval workflow and self-service	Identity management (IdM) systems incorporate a workflow system to collect approvals for the requested system access entitlements. This approval might be granted by a user’s manager, a system owner, or a departmental delegate, depending on the access control policy set by the organization. Applications should be entered by the requesting staff member, not their manager or clerical staff, in order to eliminate keying errors and the cost of administration.
Role optimization and entitlement management	The workflow should, ideally, have role optimization built into the rules, which simplifies the granting of entitlements and enables business units to coordinate the combination of access rights to eliminate high-risk entitlements (i.e., those that might break separation-of-duties constraints).
Access management	There is a range of access management models that might need to be supported. The two most common models are AD groups, whereby a user’s ID is written to the AD group membership for the system(s) they are entitled to access, or a Web access management facility that manages access to multiple Web-based applications. Increasingly, applications are relying on authentication/authorization systems that access the identity store to determine a user’s attributes during the login procedure.
Attestation reporting and governance, risk management, and compliance (GRC)	At the very least, IAM systems must provide attestation reports that are sent periodically to managers with a request for them to check the access rights of their staff and attest to the fact that it is correct. Today’s IAM systems will often include recertification facilities to update entries that are wrong or need changing. Reporting functions also support audits and other governance activities.
Analytics and dashboard	Modern IAM systems include out-of-the box analytics and dashboard functionality that management can use to monitor the efficacy of the IAM system and to identify conditions that are causing delays or indicate surreptitious activity.

Significant progress has been made in each of these areas, some more than others, and we need to ensure we leverage this new capability.

Provisioning

The provisioning activity is a core function of an identity management system. It is the act of establishing records in the organization's identity store(s), adding a new user to the identity store facility to give them access to an application, a system, or a protected resource. It could be as simple as adding a user's unique identifier to an AD group or as complex as making an entry in an application's database to grant a new user the required level of access.

Many organizations use Microsoft's AD, NetIQ eDirectory, Oracle Unified Directory, or another directory or database product to store staff data. *Provisioning* is the act of entering this data for new employees and establishing their entitlements. *De-provisioning* is the act of removing data for staff who leave the company, which will remove their access permissions. It is important that provisioning occur as soon as possible, so that the new employee can log in to systems on his or her first day of work (sometimes called *zero-day start*). It is even more important that a staff member's access is disabled as soon as they leave an organization's employ.

> The source of provisioning data is typically the company's human resources (HR) system. Generally, this is the "authoritative source" of most identity data.

Authoritative data sources are typically the systems used to initially collect data; they are production systems that should not be used by other applications for identity management purposes. This is the function of the "source of truth" for the enterprise, which is typically a directory that serves as the identity repository for systems and applications in the enterprise.

A provisioning tool will add, change, or remove identity attributes, with the appropriate approval, in the authoritative data sources. If data is changed in an authoritative source, it is detected and copied to dependent identity stores via a synchronization tool that moves the new data out to the appropriate identity stores. The provisioning tool is interfaced to all the identity stores it provisions into and provides reporting and reconciliation functions to ensure that data integrity is maintained.

In the past, this synchronization would be performed nightly, whereby the identity management tool would interrogate the HR system and write updates to the identity repository. Increasingly this is becoming a real-time operation, where updates occur on event (i.e., when the HR system records a change).

Most organizations have a provisioning tool that automatically onboards new employees, based on their employment record in the HR system. The HR system is the logical source of identity data since it is the place where employee information (e.g., name, address, date of birth) is collected and assigned to a department and position.

The next step is to use this identity information to provision staff accounts in the computer systems that they need to access. In many instances, an employee's position can be used to determine the access they should be granted to applications and systems that are used in performing the role to which they have been assigned. Typically, there is also a mechanism in place to automatically create a home drive for a new employee, and an email address, too, but the provisioning system needs to go further than that. Once a person's job is defined, they should automatically get access to the systems they need for their job responsibilities. This task should not be left up to a manual process, for several reasons:

- It wastes too much time on the part of the system administrator and staff, who do not want to spend time filling in forms to get access to the required systems.
- As soon as someone has to key in information, errors can be introduced; the best person to enter the name correctly is the employee.
- If account details are entered manually, someone must remember to remove the person's access when they leave the company.

Provisioning data should also be collected as early in the onboarding process as possible. If an organization uses a recruiting system into which applicants enter their biographic detail, this record should be used to establish the HR system record, which should then become the authoritative record for the provisioning system.

At a recent assignment at a large airline in Asia, an audit picked up a severe failure in the identity management environment. It was found that there were three times as many accounts in the organization's AD than there were staff members. An investigation revealed that a system modification that had been performed on the 10-year-old, homegrown identity management system had broken the de-provisioning feature, and accounts were no longer removed when staff left the airline's employ. The system had no documentation, and the application was no longer manageable.

Approval Workflow and Self-service

Some managers are fearful when they hear about automating their company's provisioning process. It conjures up visions that anyone will be able to request and gain access to a protected resource. Nothing could be further from the truth; automated provisioning means that the collection of user data, approval of access, and addition of the user's ID in the repository of approved users all happen automatically. All access requests must receive appropriate approval before being acted upon.

Today's identity management systems must support a workflow capability that accommodates the activities shown in Table 2.2.

Table 2.2: Identity management workflow

Activity	Description
User entry of details	The user is the best person to enter his or her details into an identity management system. The user generally knows how to spell his or her name and address correctly, which will avoid rekeying errors. The user can also be required to provide all the necessary information to substantiate his or her account establishment, avoiding a costly administration effort to chase information that should have been provided with the initial request.

Activity	Description *Table 2.2 continued*
Approval workflow	Ensuring that all user requests are appropriately approved is an important component of any provisioning system. A workflow must be configured to properly apply the approval rules established by the business. The requirement could be a manager's approval for the requested access, system manager approval, or in sensitive situations, two approvals might be required. In some cases, approvals can be granted from a "pool" of managers—for example, approval by two of five managers might be required. Knowing whom to escalate approvals to is also important. These rules need to be configured into the workflow engine. The result is that all system access will be properly approved and logged, with access entitlements granted according to corporate policy.
Entitlement catalog	As a user selects his or her required system access, the workflow needs to know the options for the requested service. The workflow can then display the options to be selected by the requestor. This means that the workflow engine must know the entitlements available within the various applications to which a user can request access. The workflow should be able to provision these entitlements either via an AD group or directly into the target application's user account repository. Ideally the IdM directory or authorization service should be used (see chapter 4 for more detail).

> In one state government health department that used a manual provisioning process, a survey of clerical staff identified that it took, on average, four queries of the data provided in a health professional's application for system access, before the application could be acted upon. In the worst case, it took 12 queries to correct errors and collect sufficient data for the provisioning activity.

A good workflow solution presupposes a relatively sophisticated identity management environment:

1. The workflow configuration must be able to identify the approval requirements for the requested access. For instance, if a new user starts as an accounts payable clerk in the finance department, the workflow must be configured to route the approval request to the finance manager. In some cases, there will be multiple approvers. In most cases, the approval request will need to be escalated to the manager's supervisor if it is not approved within a certain timeframe. There are multiple ways in

which this may be achieved: via a reference table of authorized approvers, via the “manager-is” setting in a user’s directory entry, or to a defined delegate authorized to approve system access for a department or workgroup.

2. Access to a user repository of user attributes, typically a directory, is required by the authorization service. This data will normally be collected via the provisioning system and written to a directory. This data can then be leveraged by relying-on applications when a login request is received. For instance, if an application login requires a second factor such as a PIN sent to the user’s mobile phone, the provisioning system should capture the mobile device number at the point of provisioning and write it to the directory as a user attribute.
3. Access to a set of challenge-response questions to validate a user’s identity might be required by the workflow engine. Challenge-response questions are a common method employed to securely reset a user’s password, significantly reducing password management calls to the service desk. Collecting challenge-response questions at the point of provisioning is one way to achieve this.

Role Management

A typical shortcoming in IAM environments is in the area of role management. Roles can be associated with a person’s position in the company and stored as an attribute in the identity store. This can then be used for provisioning into relying applications for access control decisions. For instance, a user might need to be in the appropriate project team before they are allowed access to the team’s document store; a role attribute in the identity store can be used for this purpose.

Difficulty with Roles

Role-based access control (RBAC) is often held up as the optimal approach to user provisioning. Indeed, being able to set a person’s system access entitlements based on their role in the organization is a laudable objective, but there are multiple ways to do this, and they are not all equal in their efficacy or effectiveness.

Ideally roles should be attributes held against a staff member’s record in the corporate directory. This attribute should instruct the provisioning workflow regarding the entitlements associated with a user’s role. For instance, if a person is starting employment as a finance manager, they should get access to the finance management system with manager privileges and should be provisioned into the general ledger, accounts payable, and accounts receivable subsystems. They should also be added to the finance department's file share and the manager's corporate document repository.

Unfortunately, in many cases relying applications do not have the ability for an external workflow to set entitlements. (A *relying application* is one that relies on the corporate identity data store for access control decisions.) A common approach is to use AD groups to manage application access instead. In this instance, when a user attempts to access a specific application, or feature within an application, the system checks the AD group in question. If a user's ID is in the appropriate AD group, they will be granted the requested access; otherwise access will be denied. AD groups can be established for each application, or for each access level within a particular application, and users are provisioned into them either via a workflow or by a manual entry of the user's ID into the group. The former is preferable, but in some cases the workflow will send an email to the system owner and, if approved, the person will be manually added to the group.

The use of AD groups, even manually provisioned, is generally a better option than maintaining an access control list (ACL) within an application, access control based on roles is even better, as shown in Figure 2.2.

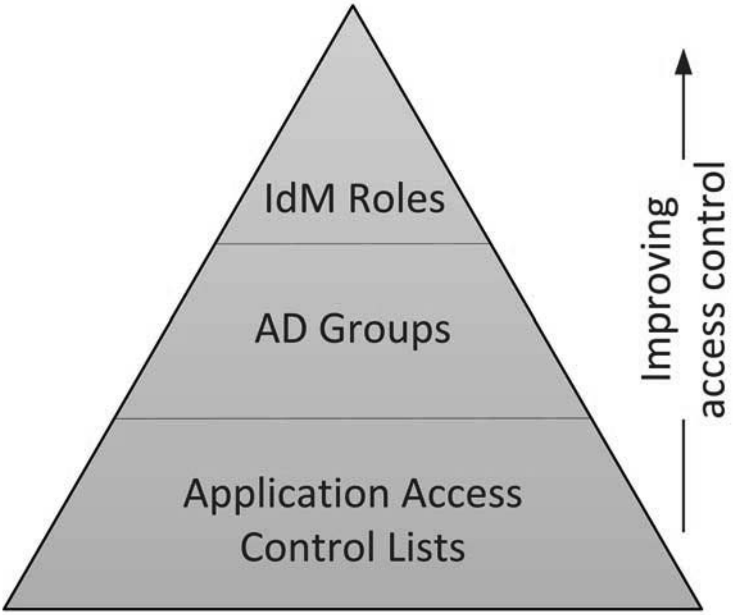

Figure 2.2: Access control hierarchy

If the groups are in AD, when the user leaves the organization their account can be disabled, which will remove their access to systems and applications. If direct entry into the application's database is used, a mechanism is required to ensure that access is removed once the staff member leaves the organization. Experience suggests that this often does not occur.

Access Management: Authentication and Authorization

The core activity of any identity management environment is to support the organization's authentication services. This most commonly refers to the logon procedure that users must complete before being granted access to a company's computing resources. User authentication is the act of verifying a person's credentials. The credentials being validated could be identity information, qualifications, or authorization level. Authentication, as the word implies, is the act of verifying a person's access entitlements as they attempt to access restricted resources.

Authorization is generally more fine-grained; it is the act of granting access to a specific resource such as a computer application or a restricted-access building. By definition, the identity management systems must be able to provide the necessary information about a person to allow the resource (e.g., computer application) to determine the correct level of access to be granted to the user. This is an area of intense change, where applications no longer maintain ACLs of users but instead rely on the identity management authorization system to make the decision regarding a user's access request (see chapter 4 for more detail).

Attestation Reporting and GRC

Anecdotal evidence suggests that fewer than 5 percent of managers receive regular reports on their staff's access entitlement to protected systems and resources. Most companies realize that this is a deficiency that requires correction. Modern IdM systems can generate attestation reports and, in many cases, implement automatic recertification whereby a user's access is confirmed and the approval is logged, and inappropriate access rights are removed.

> It is usually poor practice to rely on audit processes to identify incorrect entitlements settings. Relying on corrective action resulting from audit failures is a poor use of the audit process and will result in sub-optimal IdM management.

A common problem identified by the attestation report is “orphaned accounts,” which occurs when there is a failure to remove access to a system, when a staff member leaves the company. In organizations with a good de-provisioning process, orphaned accounts are generally not a problem.

Note: Access entitlement verification should not be left to periodic audits. This is a “line” function (i.e., the responsibility of the business unit), not a “staff” function (i.e., the responsibility of a governance work group).

Governance, risk, and compliance are processes to ensure companies are not unnecessarily exposing themselves to liabilities. GRC processes are “staff” functions, not “line” functions. Ensuring that adequate governance and risk management are applied to a business process will typically fall to an administrative workgroup that will verify compliance of a line-of-business process with corporate policy. This will typically occur in a periodic audit that will report to management on any business system that is poorly designed or contravenes policy. Paradoxically, common problems are instances in which policy has not been set. Often, in the absence of policy, IT staff effectively set access control rules, and, in many companies, governance policy is not formalized, so there is no methodology to evaluate risk.

With the increasing importance of leveraging identity information, GRC requires improvement.

Analytics and Dashboard

The days of batch processes in identity management are behind us; thus, provisioning end user accounts should be a real-time process. This means that we should be able to monitor processes and report on their health continually. Most modern IdM solutions can display the number of new accounts provisioned within a configurable timeframe and alerts on events, such as denied approvals or the granting of high-risk entitlements. In addition, it is possible with most authentication services to log system access and display in real time the number of authentications to selected services.

At an Australian university, the dashboard feature of the new IdM system allows management to view statistics on the authentication process during enrollment time, when the system access load increased by an order of magnitude. This enables the deployment of sufficient system resources to ensure satisfactory performance for all users.

Migrating to the Cloud

The accelerating use of cloud services is rapidly changing the face of identity management and complicating the provisioning task. Cloud services are often engaged by business groups without the rigorous procurement processes applied to other IT infrastructure or services. This means that integration into the company's processes and procedures is often overlooked. A common practice is to engage a cloud-based application provider and then synchronize user identities and associated attributes to the cloud. Doing so is unwise because as more services are engaged, more instances of a company's identity data store are synchronized to cloud services. The resultant proliferation of identity stores in the cloud represents a security risk and possibly a violation of privacy regulation.

Cloud IdM Service

One of the seminal activities in preparation for migrating to the cloud is to deploy a properly planned IdM service for cloud-based applications. It is not appropriate to expose the corporate directory to the Internet, for several reasons. Similarly, it is not realistic to expect a cloud-based application to send a user lookup request to the corporate network, and wait while the request transits the firewall and the load balancer before it gets serviced. Applications expect millisecond responses, which requires a planned configuration that reduces network latency to the degree possible. Since most identity stores are on-premises, this means establishing a cloud-based identity service. Many cloud service providers offer identity as a service (IDaaS) functionality. When preparing for a cloud migration, the organization will need to select an appropriate supplier to host the identity provider service and then require all software as a service (SaaS) providers to use the service. This will typically mean that all cloud-based application suppliers should support standards such as SAML for exchanging authentication messages. When selecting an identity provider supplier, it is also a good idea to stipulate support for standards such as SCIM for the exchange of identity detail.

An interesting phenomenon that we're seeing increasingly adopted within managed services environments is the trend toward industry-based federations. In such federations, multiple companies, each with its own identity provider services, operate in an industry collaboration to provide access control to information that is of importance to all the members. This is most appropriate for a cloud environment that supports an industry or supply chain in which each participant maintains its own identity provider service and each participant agrees to trust each other's identity repository. The commercial airline industry is a case in point. All major airlines participate in a system that provides access to maintenance data and other operational requirements.

One area of concern when moving to the cloud is single sign-on (SSO). SSO should be addressed; organizations with SSO for on-premises applications should not ignore this requirement when adopting cloud services. Users who have enjoyed SSO on-premises via their Web access management solution will resent the loss of this facility when moving to cloud-based services and will consider it a step backward. The cloud migration planning and selection stage should include the requirement for federation of the identity provider service (IdP) and SSO support across both on-premises and cloud services. For more detail, please refer to chapter 5.

Internet of Things

Increasingly it is necessary to support more than identities within our IAM environments. People want access to monitoring devices that collect data and to control devices that switch something on and off. The sharp reduction in the cost of such devices means that they will continue to become more accessible and economically justifiable.

This means the organization must control access to these devices and make it easy for users to share collected data securely and to delegate access to control devices when necessary. An organization's IAM environment should be the source of access control data for managing these devices.

It's equally important to protect the access from devices to corporate systems. Doing so will often require a signed/encrypted interface, which should use the same technology as interactive access (i.e., use encryption or signing keys). Consider adding devices as entities in the identity data store in order to manage certificates for encryption and digital signing. For more detail, see chapter 7.

Cloud Protocols

The development of protocols to support cloud technology and mobile devices has happened very quickly. We now have the following open standards that support these technologies.

OpenID Connect

OpenID Connect is a widely used technology for user authentication in cloud environments. It started life as OpenID, and in 2014 underwent a major release to provide more extensive functionality, which has made it the technology of choice for third-party authentication to Web applications. This makes it particularly useful for large user populations that require the use of an authentication service.