SAP NetWeaver® Identity Management Overview

White Paper

September 2007

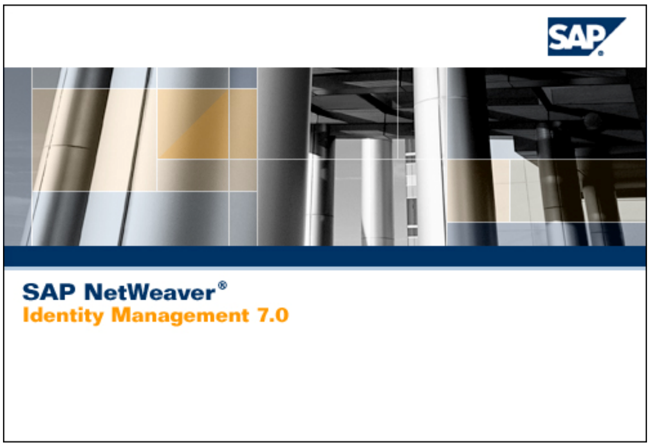

SAP NetWeaver®
Identity Management 7.0

Table of contents

Executive summary

For any organization to achieve a more efficient and secure management of internal and external accounts, an identity management solution is a prerequisite. This includes everything, from the initial steps of coordinating and joining all existing accounts using directory services, to setting up a complete workflow and provisioning system for distributed management of accounts, as well as password handling across multiple applications and federation across security domains.

This document describes the SAP NetWeaver Identity Management products, which cover most of the core functionality for implementing an identity management solution, including a virtual directory.

The Identity Center is used for the initial joining of accounts within the organization, ensuring that the quality of this information is kept high when the organization changes and employees join and leave. Since the security infrastructure depends on information about the accounts being correct when authorizing actions within the organization, it is of utmost importance that the identity information is correct.

The organization's security is only as good as the quality of its identity information.

To implement an advanced workflow solution, using distributed user management and approvals, the Identity Center provisioning and workflow modules are used. The operations to be performed and the applications to be provisioned are described in a task hierarchy, and a web-based application is used for action approvals.

The Virtual Directory Server ensures that correct and relevant information is made available to the users and applications needing the identity data, as well as ensuring that access to sensitive information is restricted.

Introduction

Identity management is a challenge for most organizations today: the larger the organization, the greater the challenge. The user must present his or her identity to get access to the many ICT (Information and Communication Technology) applications within the organization. Examples include the various operating systems, the HR (Human Resources) system, CRM (Customer Relationship Management) systems, databases, directories, physical access control systems, e-mail systems and support systems.

The term *identity management* covers various components. Implementing different aspects of identity management within an organization will yield different advantages.

The term *identity* can also cover many different kinds of identities. Although the initial (and often primary) target for an identity management solution is to handle the organization's employees, handling identities for customers and partners is equally important. In many cases resources such as meeting rooms, PCs and mobile devices, which all may have their own identity in some context, can be included in an identity management solution.

Aspects of identity management

The figure below provides an overview of the various aspects of identity management. Different vendors operate in different areas of the identity management space and very few vendors cover all functionality.

The distinction between the different areas is more blurred than the figure indicates.

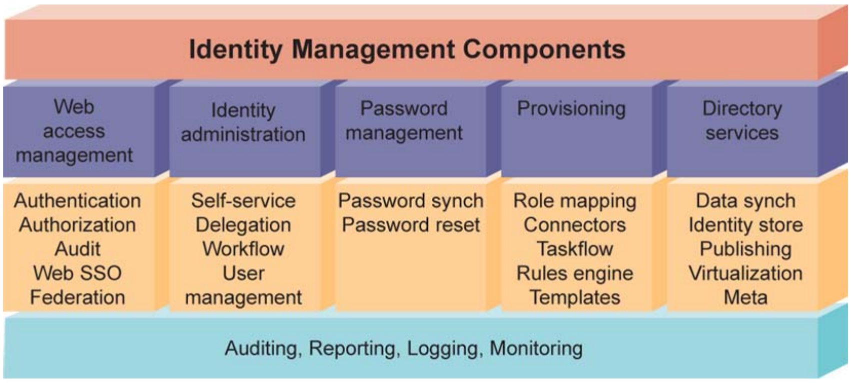

Figure 1: Aspects of identity management

Below is a brief explanation of each of the areas:

Directory services. In most cases, directory services are a prerequisite for the other parts of an identity management solution. The purpose is to ensure that there is one common view of all the identities, the identity store. Building an identity store includes the often complex task of joining the identity information across numerous repositories within the organization, which do not have one common identifier, and ensuring that the information stored about these identities is kept synchronized over time. As part of this process, the authoritative repositories (i.e. which repository owns which attribute) must be defined within the organization.

The identity data is made available by publishing the relevant parts of it to an access point. Virtualization adds the possibility to present different views of the data for different purposes or to different user groups.

The virtual directory can also be part of implementing the directory services, to achieve real-time access to the master data, as well as avoiding duplication.

Provisioning. Provisioning is a more advanced form of updating repositories than a pure directory services solution. Although basic provisioning has been done by the directory services products, a provisioning solution often requires more advanced functionality, such as handling of rules and/or roles for deciding how to update the various repositories. The provisioning connectors are normally more advanced and have better repository awareness than pure directory services, which simply read and write attributes.

Password Management. Password management includes functionality for ensuring that the user's identifier and password remain the same across a number of repositories. It is also responsible for updating the password in all applications when a user changes the password. This aspect includes password policy checking, the legal semantics of the password (i.e. number of characters and requirements for special characters) and the update frequency and password history.

In addition, when users forget their password, there is a need for password resets, either by an administrator or by the user in question.

Identity Administration. This includes functionality for managing the identity information. There are two main approaches to managing the identities. One is to use the existing infrastructure for this administration, using a combination of directory services/provisioning for distributing the accounts and access information to the various repositories. The other is to use a separate application for this purpose. The needs of the organization determine which of these approaches should be used.

In addition, the identity administration may include self-service functionality where users can register themselves in the identity system (which is common in the B2C scenario, where customers register themselves). Users can also manage (parts of) their own identity information. In an organization, employees may be allowed to manage their own telephone number and other personal information.

Web Access Management. This part of identity management includes functionality for authentication and authorization of identities, as well as identity federation across domains. This can in turn be used in a scenario where the user does not have to log in for each application, if a trust relationship is defined. An important part of Web Access Management is the auditing function - the ability to check who did what and when.

Auditing, Reporting, Logging, Monitoring. These are the low-level components of most identity management solutions. Although this aspect is shown as one layer in the figure, this is *not* standardized. Different identity management applications take different approaches to this. However, the functional requirements are the same: to be able to verify operation, generate reports from the identity data, produce and view logging information, and monitor the current status.

A layered identity management architecture

A different way of viewing the identity components is to look at this layered architecture (which is described in more detail in the white paper "Implementing an Identity Management Solution").

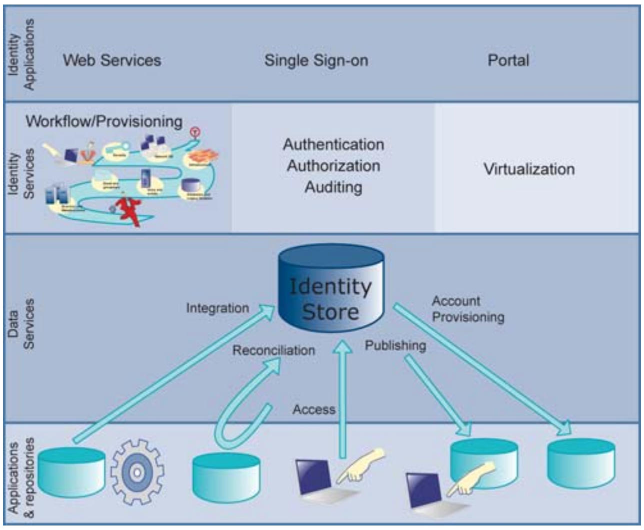

Figure 2: Identity management architecture

The lowest layer contains the existing applications and the next layer, the Data Services layer, contains the identity store. The proper functioning of all the other identity management applications depends on the identity store.

Applications & Repositories. All existing ICT infrastructure in an organization, including the data repositories and the applications/interfaces that are used to access them. These may be business applications of various kinds containing customer and product information; specific applications maintaining identity data, such as human resources applications; or applications used to maintain other types of information, such as document management systems.

Data Services. The Data Services layer builds a uniform, normalized, integrated view of the Applications & Repositories layer. This is achieved through services/functions such as synchronizing, joining and publishing data and providing access to data. The Identity Store is a core component of the Data Services layer. It is used to gather information about all identities throughout all applications in the organization. In many cases, the identity store may also be implemented using Virtual Directory Server.

Identity Services. This layer consists of the services that are offered on the bases of the Data Services layer. These include provisioning, authentication, authorization and virtualization.

Identity Applications. This layer consists of all applications using the Identity Services. This may be the identity-management components of existing applications, as well as functions such as workflow, federation, single sign-on and self-services.

SAP NetWeaver Identity Management

The SAP NetWeaver Identity Management products have their roots in MaXware AS. SAP NetWeaver Identity Management has a strong presence in the identity management space, as shown in the below.

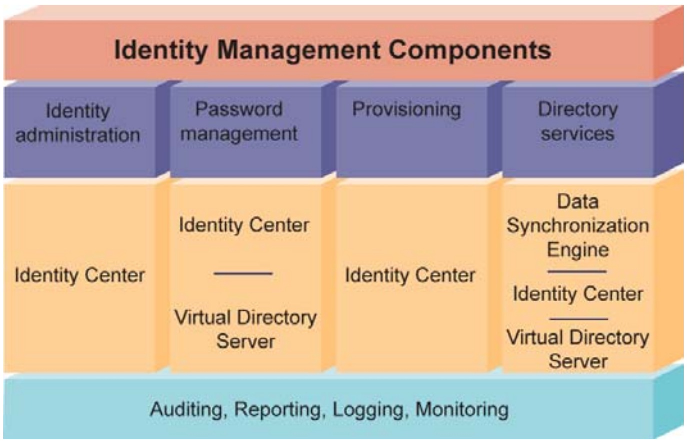

Figure 3: The SAP NetWeaver identity management products

The first MaXware directory services product, the MaXware Data Synchronization Engine, was released in 1996 and has since grown into the high-end identity management product MaXware Identity Center. The MaXware Identity Center includes provisioning, workflow and password management functionality. The MaXware Virtual Directory was introduced in 1998, and has functions to join heterogeneous data across multiple data sources, as well as performing advanced updated functions to multiple repositories.

SAP NetWeaver Identity Management consists of the following components:

Identity Center. The Identity Center is a high-end identity management solution, providing low latency and high availability. It uses a relational database for the configuration data and the logging and status information, as well as for the identity store and all provisioning and workflow states.

Data Synchronization Engine. The Data Synchronization Engine is responsible for any low-level operation on the applications and repositories. It runs as part of the Identity Center.

Virtual Directory Server. A virtual directory provides the organization with real-time access to the identity information, as well as to other critical information, by providing a single access point to all information. The Virtual Directory Server can also be used to control access to the identity data. It is able to present the same data in different ways to different groups of users. It can also be used to write-protect or hide certain attributes, for example when making information available externally.

Data Synchronization Engine

The Data Synchronization Engine is based on creating jobs that consist of passes. Each pass in a job performs one specific task, such as reading from a data source (known as From-passes), or writing to a data source (known as To-passes).

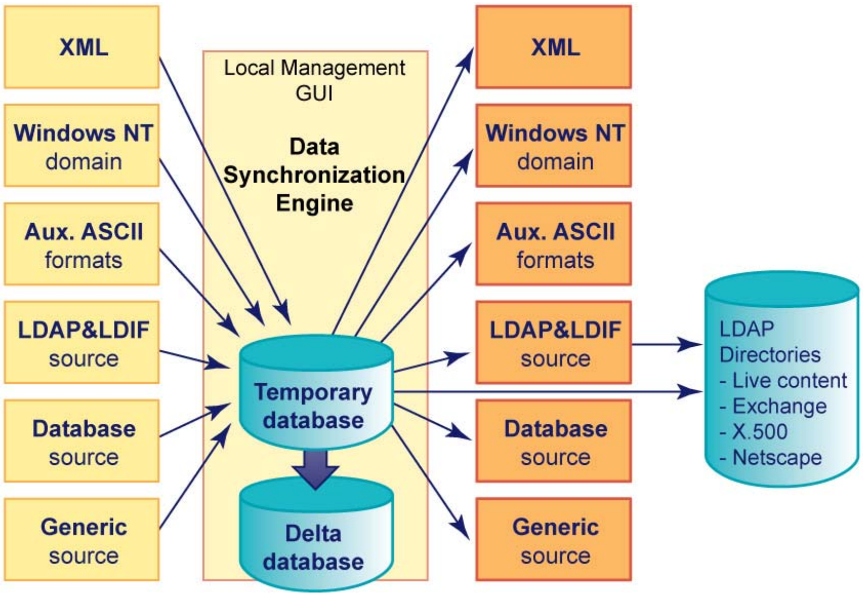

Figure 4: Data Synchronization Engine

The Data Synchronization Engine reads from a repository using one of the From-passes. The data from all from-passes is stored in the internal database and joined when writing to the target repositories. A change-log can be used when reading data, to ensure that only updates are read. The delta mechanism can be used both when reading and writing, to reduce the load on the system, and speed up the process.

The Data Synchronization Engine can also manipulate attributes when processing data or combine data in all conceivable ways, using the SQL engine of the database.

Note that the Data Synchronization Engine is non-intrusive, and is not limited to writing data to a directory. It may just as well be used to synchronize from a text file into an Oracle database, as synchronizing from an XML source into a directory server.

Also note that the Data Synchronization Engine does not require any modifications in the existing repositories.

Identity Center

The architecture of the Identity Center is designed to provide maximum flexibility, scalability and security in a single software solution. This allows identity management across multiple applications and databases both within the organization and in an extranet environment. The Identity Center manages all of its activities from a core database and supports both Microsoft SQL Server and Oracle. All components in the solution interact with the database to ensure that all identity management activities are properly executed.

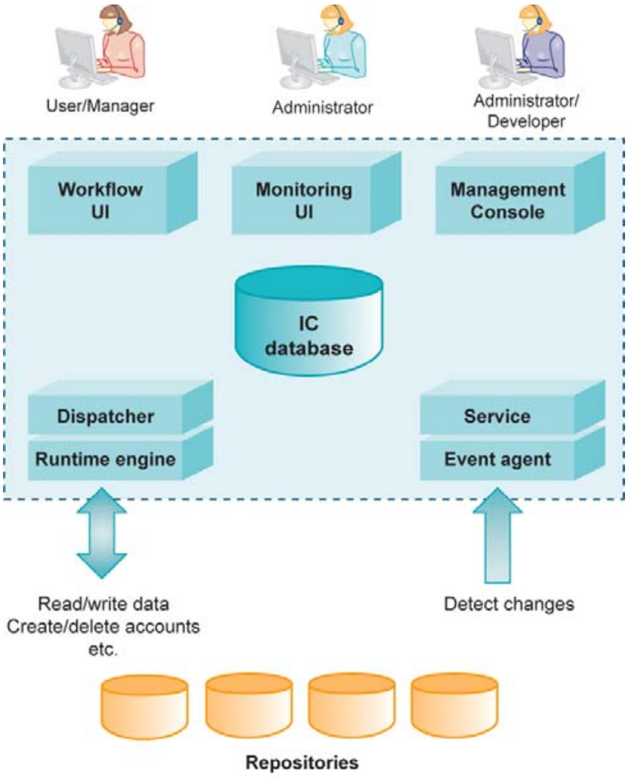

Figure 5: Identity Center

The Identity Center consists of the following components:

- **IC database:** All information about provisioning/workflow tasks and jobs, the identity store, scheduling information, state information and audit/logs is kept in this passive store.
- **Dispatcher/Runtime engine:** These components act as local or remote agents for the Identity Center and are responsible for processing both provisioning and synchronization tasks. They are also responsible for performing reconciliation and bootstrapping. Multiple pairs of dispatcher/runtime engines can be used and dedicated to running specific types of jobs.
- **Service/Event agent:** An event agent can be configured to take action based on changes in different types of repositories such as directory servers, message queues or others. The event agent will detect changes and submit information to the Identity Center. The dispatcher will then initiate execution of a given job. This mechanism is optional and its only purpose is to initiate synchronization based on changes in repositories in addition to the scheduled operations.
- **Workflow UI:** The Workflow web interface is used for all end-user registration/self service, password resets and approval of tasks.
- **Monitoring UI:** The Monitoring web interface is used to provide an overview of the system status, audit and logs during daily operations.
- **Management console:** The Design user interface is used for configuring the Identity Center, including provisioning/workflow tasks and jobs.

The identity store

Any organization depends on the identity information for authorizing entities (for example employees, partners, customers and applications) as well as for authorizing access to applications. The actual identity information, in addition to the unique identifier for the entity, consists of a number of attributes, including the user password or pass phrase used for authentication, the white pages information (telephone number, address, e-mail address, etc.), access control information, and possibly a great deal more.

The main problem is that in many organizations the identity information is stored and/or maintained in each application within the organization. This leads to superfluous maintenance work as well as the risk of storing erroneous information, which in turn may represent a security risk. If the identity information is incorrect, people may be given access they should not have. If information about former employees has not been removed, they may still have access to applications. A strong authentication mechanism (for example using biometrics) is of no use if the identity information is not updated.

The organization's security is only as good as the quality of its identity information.

One of the core functions of an identity management solution is to build a central repository containing the identity information: the identity store. A directory server is not necessarily the best storage for the identity information. The Identity Center has been using a relational database for this purpose for several years. A relational database has several features which are useful for holding the identity information. These features, for example foreign keys and column constraints as well as triggers and stored procedures, are normally lacking from a directory server. Any report generator can also be run on a relational database. The directory services will ensure that the identity store is synchronized with the various repositories and applications within the organization, while the various attributes for the identities may be owned (managed) by different applications.

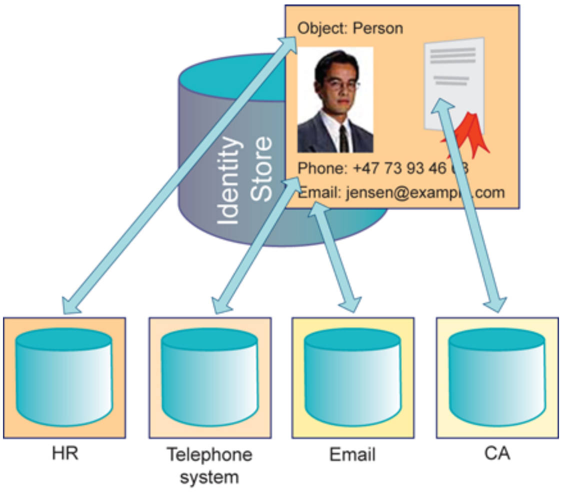

Figure 6: The identity store

The identity store must be correct and complete, as this information will be used for authentication and authorization. This is taken care of by the identity management application, for example the Identity Center. The identity information must be available 24/7, as it will (at least over time) become a critical part of the organizational infrastructure. Most relational databases can be configured to run in a high-availability mode by running on top of duplicated hardware, thus eliminating a single point of failure. There is no preference as to which database to use. In many cases the organization already has extensive knowledge of a relational database, which also may be utilized for the identity store.

We have also recognized the need for making the identity information available to a number of different applications in a number of different ways. A directory server is a good choice for publishing the information, as it offers a standardized way of accessing the information, using the LDAP or the DSMLv2 protocols, which both are on-the-wire standards. In addition, the Virtual Directory Server can be used to give a different view of the same data to different applications or users, as well as retrieving real-time identity information from any type of repository.

Provisioning

The term provisioning is often used to denote user provisioning or account provisioning. The functionality includes creation of accounts, setting initial passwords, setting and modifying access rights, disabling (revoking) an account and deleting an account. The overall purpose is to make sure an identity (for example a user) has the correct access to the applications. However, a provisioning solution should be able to provision any type of information.

The Gartner Group has the following definition of user provisioning:

> *User provisioning products generally leverage a central repository of user attributes to control data in outlying directories or repositories. User provisioning products also include workflow capabilities to apply business rules to the account provisioning process and typically provide user self-service capabilities (e.g., password reset)*
>
> *Gartner 2003*

There are several motivations for an organization to embark on a provisioning project.

- The most common motivation is to reduce the cost of internal maintenance. This can be lowered by simplifying the account management and automating the process of managing the accounts across repositories.
- Provisioning can also shorten the time needed for account creation, which is especially important in a B2C scenario, where the customers will use a web application to create their accounts. The provisioning system will receive information from the web application and create the necessary accounts in the various repositories.
- Another issue is security. In many cases, "old" accounts still exist after they are no longer in use, posing a security risk. In addition, people changing roles within organizations will traditionally never *lose* any access rights from the old role, but only get additional access rights according to the new role. Because of this the access rights accumulate over time, which is in most cases not according to company security policy.
- Another motivation is regulatory requirements. Initiatives like the Sarbanes-Oxley Act[1] and others place high demands on organizations to provide reports on authorization information. Organizations are required to report who has access to which applications, and also who *had* access at a given point in time, as well as who authorized this access. In addition, segregation of duties is becoming more important. An example of this is that a corporate manager with rights to approve a new supplier should not at the same time be allowed to issue purchase orders for this supplier. This is to avoid the temptation for the manager to pass business to a company with whom he/she has personal or family ties.

The boundary between directory services and provisioning is not very clear. Directory services synchronize between different applications and repositories, while a provisioning solution creates and manages accounts. We have supplied provisioning solutions for a number of years, using our identity management solution. The provisioning module contains tasks and rules, and it is specially designed for account provisioning.

In many cases, determining the applications in which to create an account and the access rights to provide depends on some other attributes of the entry. For example, this could be the location of a user, or the user's position in the organization, both of which can trigger the creation of an account and setting of specific access rights within an application.

[1] More information can be found in the white paper MaXware White Paper: Achieving Sarbanes-Oxley compliance

In the provisioning module of the Identity Center it is possible to define a hierarchy of tasks, starting with very basic tasks such as creating an account and setting the account attributes.

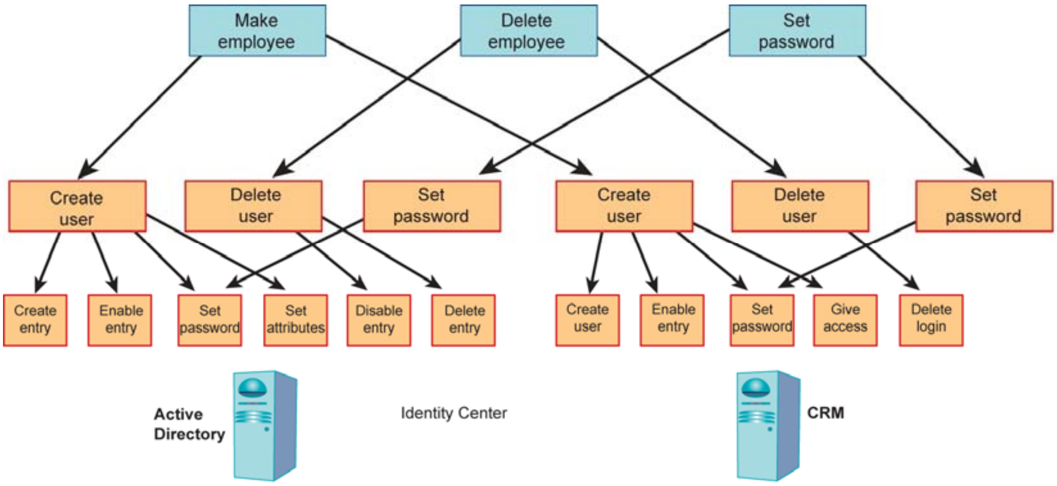

Figure 7: Provisioning task structure

More complex tasks can be built, including tasks that operate on several repositories. Tasks can be defined to run in sequence or in parallel. For each task, error-recovery tasks can be defined, as well as confirmation tasks, which for example can send an e-mail or SMS message when a task is completed.

Persistence plays a key role in the provisioning module. There will always be situations where temporary failures occur, such as network failures or power glitches. Since the SAP NetWeaver Identity Management products uses a reliable relational database for storing the identity information, as well as all provisioning tasks and the state of each of these, it will always be possible to recover from such failures.

Scalability is another important issue. The provisioning module is limited only by the performance and size of the database and the processing power. Most relational databases today can be scaled by adding more hardware. The SAP NetWeaver Identity Management components can be configured to run on multiple computers, making it possible to add more computers as required.

The provisioning module stores all logging and audit information within the database and this information is made available through a web application. In addition, any report generator can be used with the database, to produce any report required.

Workflow

While provisioning performs the task of creating, modifying, disabling and deleting accounts, such operations will require approvals in many cases. The approvals may also require that the approver adds information to the entry, for example the title.

The workflow module of the Identity Center can be used for this purpose.

One typical scenario is where users register themselves in a given system using a web application. A new user will not automatically be added to the internal application; an administrator must approve the user, as shown in the figure below.

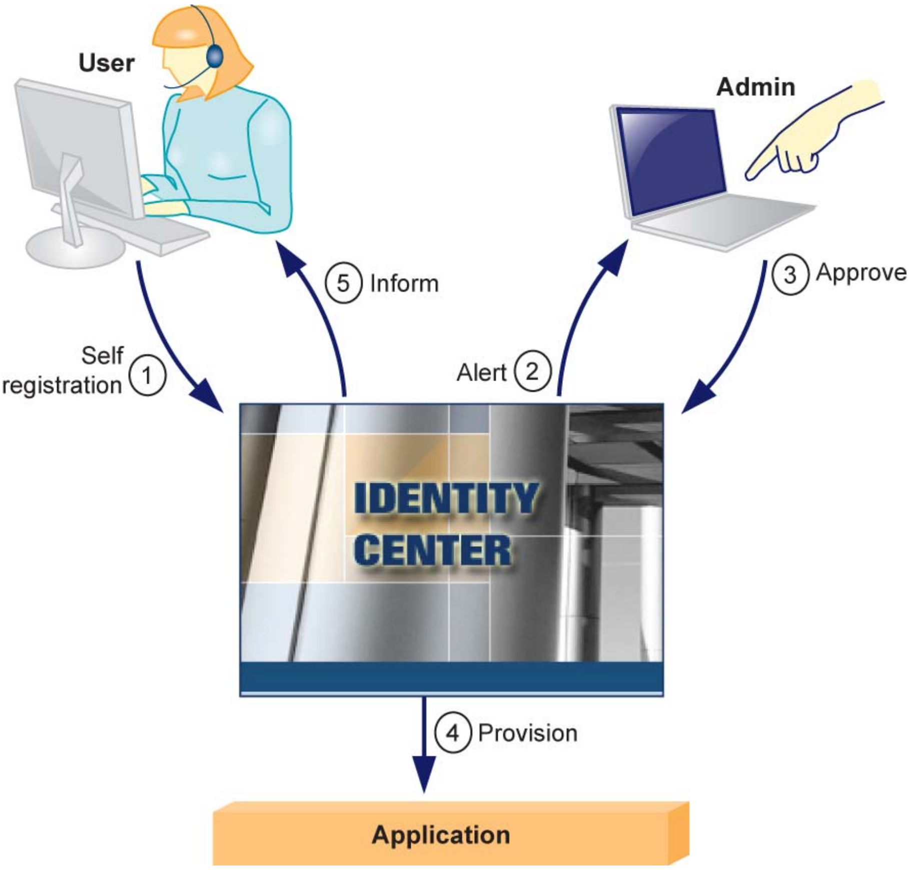

Figure 8: Workflow

The user registers using the web application (1). The workflow module will alert the administrator (2), using e-mail, an SMS to the administrator's mobile phone or any other means which is suitable within the organization. The administrator will approve (or reject) the user registration (3) and possibly supply additional attributes. The provisioning module will create the user within the given applications (4), based on attributes the user has entered, such as location and department. Finally, the user is informed by e-mail or SMS that the account has been created (5).

The Identity Center Workflow module has a web-based front-end for registrations and approvals and makes use of the Identity Center provisioning module for performing the given operations, including the approval process. Which attributes to prompt for, which attributes are mandatory as well as what other information to show can be configured from the Identity Center user interface.

In a more advanced scenario, any member of a given group may be allowed to do approvals. There may also be several approvals, which must all be performed before the provisioning takes place. For example, the department manager may approve the title of the user, while the mail administrator will define the mail server and define the storage quota for e-mail.

Password management

Password management is an important part of identity administration. This is an expensive task for most organizations since it usually requires a help desk staff to respond to requests about forgotten passwords in different systems. While maintaining security is always a priority, a balance must be struck between secure password handling and simplicity for end users.

Password recovery

Many helpdesk calls concern forgotten passwords. The Identity Center includes a kiosk solution for resetting lost passwords.

A user who forgets his/her password can log on with a given user name and reach the Workflow's password recovery task without gaining access to any other resources. This provides a secure way for recovering passwords without assistance from a helpdesk or another internal service desk.

Virtual Directory Server

The Virtual Directory Server provides the organization with real-time access to the identity information as well as to other critical information, by providing a single directory access point to all information. This simplifies client configuration, and also gives the administrator more control of the information.

Using standardized protocols (such as LDAP and DSMLv2) makes access to the information much simpler than using database protocols, which in many cases requires a great deal of client configuration.

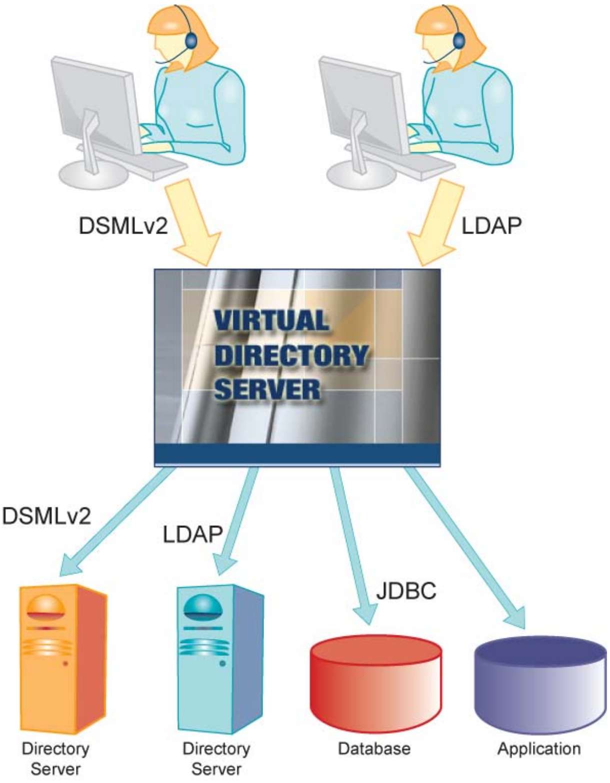

Figure 9: Virtual Directory Server

The Virtual Directory Server connects to any number of repositories, and the information is presented to the user as one joined view. In addition, different users and applications may be given different views of the information. This can be configured to make the information more relevant, depending on the information consumer. For example, applications in the IT department may want to view employee details organized by physical location, while applications in the HR department want to view the information presented by organizational unit. In addition, the IT department may not need employees' private postal addresses, but may be interested in their e-mail server and disk quota details.

The Virtual Directory Server will present the information in different ways, depending on the credentials used to connect. In addition to modifying the way information is organized (the directory tree), the directory may also modify each entry by filtering out attributes or even changing the values of attributes. The latter may be useful when publishing information about employees for external use. There may be a great deal of information (for example organizational affiliation) which the organization does not want to publish externally. In some cases information is substituted (for example, the direct telephone extension is replaced with the switchboard number).

A typical firewall scenario using the Virtual Directory Server is shown in the figure below.

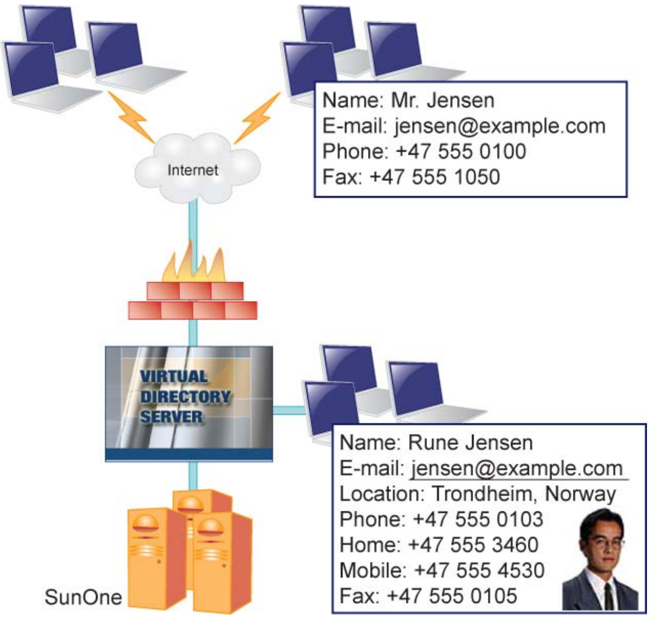

Figure 10: Firewall scenario

The Virtual Directory Server is not limited to reading information, but can also update the repositories, which has proven useful in PKI scenarios (see the white paper "Using a virtual directory in a PKI infrastructure"). When publishing certificates (or other information) there may be a need for analyzing the information, and taking some action based on the information before it is stored. The same information may be stored in several locations for simpler access from the applications.