

Fischer International Identity

BUILT FOR BUSINESS... YOURS™

WHITE PAPER

Fischer Identity as a Service™

Architecture Overview for Client Organizations

This paper is intended for Architects, Business Analysts and others who are interested in the technical aspects of Fischer's Identity as a Service™ Technology and its implementation in end-user organizations.

Table of Contents

1. Introduction	1
2. Fischer Identity as a Service™ Environment	1
2.1 Password Reset and Synchronization Service	1
2.2 Immediate Offboarding Service	2
2.3 Role & Account Management Service	2
2.3.1 Approvals.....	3
2.4 Automated Role & Account Management Service	4
2.4.1 Provisioning Initial Passwords	5
2.5 Privileged Account Access Service	6
2.6 Identity Compliance Service	7
2.6.1 Recertification	8
3. Architecture and Capabilities	8
3.1 Server Architecture	8
3.1.1 Interoperability Engine.....	8
3.1.2 Data Repository	9
3.1.3 Multi-Organization / Multi-Tenant Operation.....	9
3.1.4 Enterprise Business Roles	9
3.1.5 Separation of Duties	9
3.1.6 Role Matrix Management	9
3.2 Connectivity	10
3.3 High Availability	11
3.4 Scalability	11
3.5 Extensibility.....	11
3.6 Security.....	11
4. Implementation and Use	11
4.1 Establishing Connectivity.....	12
4.2 Configuring the Solution	12
4.3 Onboarding Users	12
4.3.1 Account Reconciliation	12
4.3.2 Account Self Claiming	13
4.3.3 User Self Registration.....	13
4.4 Validating User Entitlements	13
4.5 Mergers and Acquisitions	13
4.6 Business Divestitures	13
4.7 Changing Business Directions for Procurement Models	13

1. Introduction

Identity as a Service™ (IaaS™) technology enables IT organizations to concentrate their resources on supporting business requirements and to avoid capital expenditures by affordably hiring identity management expertise and facilities using a software as a service (SaaS) model. This white paper describes the managed services available to end-user organizations, the technology that enables the services as well as implementation and use of the managed services.

2. Fischer Identity as a Service™ Environment

Fischer offers several managed services for identity management. With IaaS™ technology, servers are located at a service provider's location, and the provider can have administrators specialized and dedicated to specific functions and/or to specific organizations that it supports. With the exception of performing a brief installation of the Global Identity Gateway, all implementation and management is conducted remotely from the providers' locations via secure communications across the Internet. Fischer's unique architecture, which is described in section 3 of this paper, enables the IaaS™ model.

The following sections describe the identity management services available through IaaS™ technology.

2.1 Password Reset and Synchronization Service

Many organizations struggle to balance the high cost of resetting forgotten passwords and securing their IT assets by requiring sufficiently-strong password policies. The problem is often compounded by policies that force users to periodically change their passwords, as some users resort to writing their passwords in non-secure locations, sometimes even within plain sight. Yet clamping down too tightly can mean lost productivity and even lost sales as workers impatiently wait for the help desk or administrators to reset their passwords.

The Password Reset and Synchronization Service delivers a comprehensive set of password-management capabilities including password self-service, policy enforcement and password synchronization. Users can reset their forgotten passwords without contacting the helpdesk, either by responding to a configurable set of challenge-and-response questions, or by authenticating with their credentials for a registered user account on another connected system. For example, persons forgetting their mainframe account passwords can authenticate to the service with credentials from their LAN accounts to be allowed to reset the passwords for their mainframe accounts. Organizations control the number and content of challenge-response questions as well as whether end users can create some of their own questions. No workstation changes (e.g., GINA changes) are required for any interface.

Multiple systems with compatible password policies can be grouped together so that when a user changes or resets his password for the group, the passwords of all the user's accounts in the group are changed together. Any new or changed password must conform to the organization's established password policies. When systems have incompatible password policies, they can be placed in separate groups or managed individually. Self-service password capabilities can be configured to be "transparent" through Windows and NetWare pre-login screens and password change screens so that passwords changed through Windows or NetWare are automatically synchronized to the user's accounts on other connected systems. Other available self-service access channels are the Fischer Self-Service Portal and a secure password kiosk. The kiosk allows users to reset their forgotten LAN passwords even when they are locked out of their LAN-authenticated workstations. Configurable password expiration reminders can be sent via email in advance of password expiration so users can think of a new password that can be remembered without being pressured to immediately create a new password when the current one expires. The Secure Password Kiosk is a browser-based facility for resetting forgotten passwords. When the Password Reset and Synchronization Service is combined with the Automated Role & Account

Management Service, the kiosk can be used as a first-time authentication mechanism that allows users to set their own initial passwords as described in section 2.4.1: Provisioning Initial Passwords.

The Password Reset and Synchronization Service can support any password policies to address your organization's requirements. Password policies ensure that only passwords of acceptable strength are allowed, and policies can be defined enterprise-wide, for groups of connected systems with common policies, for individual systems, or for groups of users within a single connected system. Password policy factors include thresholds for the minimum and maximum number of characters, minimum number of uppercase characters, minimum number of numerals, etc. A dictionary check prohibits the use of specified passwords, and your organization can submit additional words in any language for the service provider to add to the dictionary. Password history can also be enforced to prevent reuse of passwords and all password activities and resets are monitored and audited. When the service is combined with the Automated Role & Account Management Service, organizations can automatically disable accounts on connected systems whose passwords have not been changed within a specified period of time.

Getting started: in addition to the implementation procedures described in sections 4.1 through 4.3, the service provider needs to know your organization's password policies and which systems will have their passwords synchronized.

2.2 Immediate Offboarding Service

Most organizations have lost track of some accounts so that when a user leaves the organization, some accounts can remain accessible to the former insiders or to others who have learned the account credentials. There is usually no way to track who is using these "orphan" accounts. This has become an important problem for many organizations as orphan accounts enabled an insider at a large French bank, Société Générale, to create a 4.7 billion euro (more than \$6 billion) loss for the organization.

The Immediate Offboarding Service enables organizations to quickly prevent users from accessing any of their registered accounts. When someone leaves an end-user organization, the organization can quickly block access to the accounts without overlooking any accounts and without delays caused by locating and involving administrators from every system where the user has accounts. The service provides a help desk interface that enables the organization to block the usage of any combination of a user's accounts. By using this procedure, no data is lost and the organization can also have the service provider reassign the account to someone else. When combined with the Password Reset and Synchronization Service, the help desk interface also enables help desk representatives to securely reset the passwords for multiple accounts when end users call to have forgotten passwords reset. Help desk privileges can be limited to password resets, which virtually eliminates the potential for the expensive errors that are possible when help desk representatives have administrative rights on the systems.

Getting started: in addition to the implementation procedures described in sections 4.1 through 4.3, the service provider needs to know which administrators or help-desk representatives require access to this functionality.

2.3 Role & Account Management Service

Processes for fulfilling requirements for IT and business assets can be cumbersome, as users and their managers sometimes need to guess who to ask to provide a specific resource. Without tracking, manual processes sometimes enable users who know the right administrators to access more resources than they should. Further, the practice of providing a new person with the same entitlements as an existing user is notorious for providing too many entitlements. Also, when resources are provided by administrators, they sometimes lose track of which accounts they have provided to each person, and most organizations don't have a central repository for listing all of the resources provided to each person.

This results in compliance violations, excessive permissions, and accounts remaining active after a person leaves the organization or changes positions. The Role & Account Management Service provides the flexibility and control to meet the requirements of all-sized organizations as it enables them to start with a less-automated provisioning process that can later be more fully automated if desired. The Role & Account Management Service automates the provisioning and approval process and tracks the communication of provisioning workflows so that users and their managers don't need to guess who should receive their requests. End users request resources through a self-service portal that can be configured to show each user only the resources he or she is authorized to request. Requesters can add comments to a request. The requests are sent to one or more persons for approval, then routed to the appropriate administrators of IT systems or to resource owners of business assets for fulfillment. These persons indicate to the system when the resources have been fulfilled so that it can track who has been given access to each resource. The service can send email notifications regarding the status of requests to requesters, approvers, administrators, and recipients, and authorized persons can be permitted to review the status of tasks through the web-based interface. Tracking resources as they are provisioned is important for governance and reporting as auditors often demand reports regarding who has access to each resource. Figure 1 depicts the Role & Account Management process.

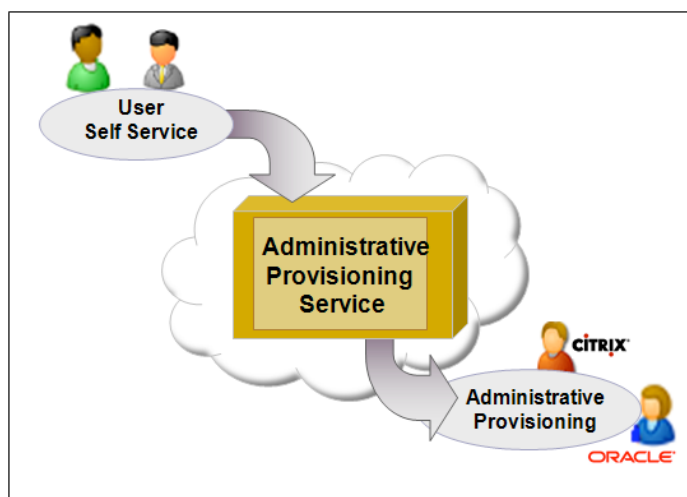


Figure 1: Role & Account Management Process

Getting started: The Role & Account Management Service can be implemented very rapidly since it requires no connectivity to the systems and applications being managed, and the basic configuration doesn't require any business rules. In addition to the implementation procedures described in sections 4.2 and 4.3, service providers need to know information about approvers as well as the IT administrators or resource owners who will receive requests.

2.3.1 Approvals

Requests for resources, whether automated or initiated through user self-service requests, can be configured to require one or more approvals. Provisioning services support a flexible approval model that includes dynamic discovery of approvers, automated escalation, combined serial and parallel approvals, the ability to require a quorum number of approvers (e.g., 3 of 5 approvers), etc. Approval tasks can be distributed to an individual or to multiple persons. Approvers and others are notified about new requests through e-mails that provide links to the appropriate approval requests. If an approval task is not approved within a specified time frame, it is timed out and escalated. New notices or reminders are sent to the escalated approvers.

Dynamic discovery of approvers eliminates the need to reconfigure the IaaS™ platform for personnel changes since the service can automatically send requests to the correct approver based on criteria such as the requester's current manager, the current system owner, etc. Requests can be approved as is, denied, escalated multiple times, or edited to reflect the correct entitlements before approval. Approvers can include optional comments and can also be permitted to delegate or transfer their requests to another approver. Requests can be automatically escalated if an approver does not act (or if not enough approvers act) on a request within a specified period of time.

2.4 Automated Role & Account Management Service

While many organizations lack sufficient provisioning processes, even organizations that have repeatable processes for onboarding and offboarding user access to IT assets and other resources have found that the processes are costly and require time and resources. Also, as with any manual processes, they can introduce errors. Manual processes usually don't assure that records presented to auditors for regulatory compliance are accurate since they rely on the accuracy of manual tasks. Manual processes can also fail to properly offboard users when they leave the organization or no longer need a specific resource.

The Automated Role & Account Management Service removes the requirement for manually requesting resources as well as the requirement for manually fulfilling requests for IT resources. The service can be rapidly configured and implemented so that events on authoritative sources like HR or contractor systems can automatically initiate real-time workflows to drive processes on connected systems. Processes like creating new user accounts on connected systems can be initiated immediately or can be automatically scheduled for the future based on the effective dates of new hire events in the HR system or other authoritative sources so that users don't have to wait for the resources they require. Deprovisioning processes include all accounts and can also automatically revoke temporary access to accounts based on the scheduled departure dates of contractors.

The Automated Role & Account Management Service enforces an organization's policies so that only authorized persons are permitted access to resources, which eliminates orphan accounts. This also enables compliance with regulations such as Sarbanes-Oxley, HIPAA, etc. While authoritative changes on connected systems can initiate identity management processes, non-authorized changes, such as creating a rogue administrative account, can be configured to trigger alerts and/or revert the unauthorized actions to the appropriate status. Service providers can preview the impact of any new provisioning policies and workflows on the environment to ensure that provisioning policies deliver the desired results before committing to changes.

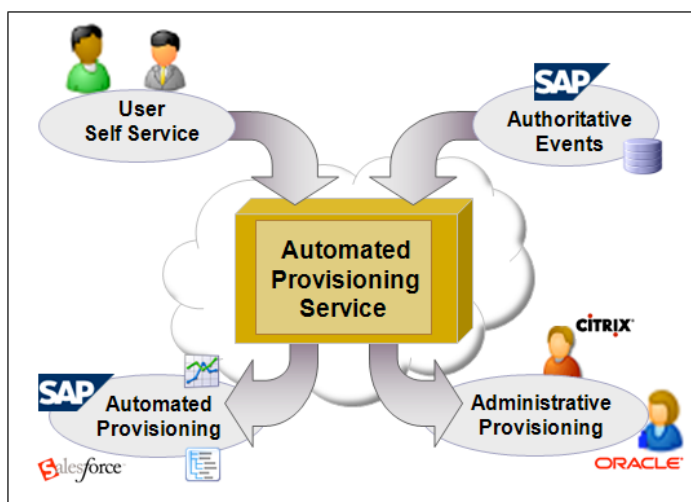


Figure 2: Provisioning Process Initiation and Execution

The Automated Role & Account Management Service includes the full capabilities of the Role & Account Management Service so that resources provisioned or deprovisioned can include IT resources such as accounts and privileges as well as non-IT resources like credit cards, offices, etc. Figure 2 shows that the service handles all combinations of provisioning process initiation and execution. Resource fulfillment is tracked so that when employees or contractors leave the enterprise, they are automatically deprovisioned from IT resources and the service can automatically alert the appropriate people or departments to retrieve physical resources. Organizations can also establish grace periods so that users can be entitled to continued access to specified resources for specific periods of time.

- Provisioning processes can call on helpdesk applications to automatically open, update or close tickets based on the current status
- Accounts and privileges can be provisioned on virtually any connected system including systems like SharePoint Services and Salesforce.com.
- It provisions file shares on Windows and NetWare servers, including the management of ACLs such as READ, WRITE, etc.
- The service can be configured to require approvals and notifications for any provisioning actions.
- All activities are automatically logged to the audit database for proof of regulatory compliance and other requirements.

Getting started: in addition to the implementation procedures described in sections 4.1 through 4.3, the service provider needs your organization's requirements for identity management:

- What triggers the provisioning process for each resource and what approval process is required?
 - For automated requests, which system is the authoritative source, and which actions trigger specific provisioning events? For instance, "hiring" a new employee in the HR system could trigger the provisioning service to provision several accounts for the user based on the organization's business policies, while "terminating" an employee in the HR system could trigger the provisioning process to remove access to all of the user's accounts.
 - For self-service requests, will an end user request the resource for himself, will the person's manager or someone else request the resource on behalf of the user, etc.?
- Which provisioning actions should be automated and which should be performed by administrators? When a system has very few users, it can be cost effective to administratively provision users rather than to automate the process.
- Business and security policy information enables the service provider to automate processes for determining who should have access to each resource. This information also enables service providers to validate existing user entitlements as described in section 4.4.
- Since automated provisioning works with both roles and rules, roles are not required.

2.4.1 Provisioning Initial Passwords

Fischer Secure Password Kiosk provides an alternative to existing mechanisms for distributing initial passwords to new users. This approach is especially useful for securely onboarding new users in remote locations (e.g., extranet, multi-nationals, etc.). Information can be extracted from multiple authoritative sources such as the HR system to populate a configurable number of challenge questions. First-time users would be advised to use the kiosk to answer challenge questions and to create their own initial passwords.

The "kiosk" is actually used from any LAN-attached workstation, including the user's own workstation. For initial use of their accounts, users authenticate themselves to the IaaS™ end-user interface by answering questions about themselves based on the data extracted from the authoritative sources. Once authenticated, end users are directed to securely create their own passwords that adhere to password policies. This method assures that the end user is the only person who knows an account password, and it works without delay, even across enterprises and domains in other parts of the globe.

A related alternative does not require any user data to be extracted from systems. This process enables one or more managers to create Identity accounts for new users and to randomly generate authentication keys for the accounts. The managers must inform the users about the authentication keys so the end users are authenticated before creating their new passwords.

2.5 Privileged Account Access Service

Privileged accounts such as Administrative, Super User, Root, and Fire-Call provide the nearly-unlimited access to system resources that is essential for everyday and emergency IT operations. However, these accounts are typically shared, resulting in multiple persons sharing (or knowing) the credentials for a single account on a system or application. Without privileged account management, it's often impossible to determine which individual actually performed an activity such as creating a new account or changing permissions for an account. Auditors have reported material deficiencies as this violates regulations such as Sarbanes-Oxley and HIPAA.

The Privileged Account Access Service provides the additional control, auditing and compliance needed to manage high-privilege and other shared accounts for any connected systems. Organizations can use small pools of privileged accounts that can be requested as needed, including for emergency or "fire-call" purposes. No one can view or otherwise be privy to the passwords of protected accounts. The end result is a complete audit record of exactly who can perform administrative tasks, when, who approved their access, as well as why the accounts were requested.

In addition to authorizing specific users to log into the service, organizations can pre-authorize users to request specific resources, and can limit the periods of time when the resources can be requested. Users can view which resources can be requested along with the status of each resource; they can request available resources as well as specifying the duration and reason for their requests.

Depending on the configuration, requested resources can be immediately granted, granted upon approval, or granted after a specified period of inaction by approvers (e.g., if an approver does not deny access within 10 minutes, automatically grant access to requested fire-call accounts). This option for configuring fire-call accounts is especially useful for emergency conditions that happen in the middle of the night, since pre-assigned persons can automatically receive the access they need for a limited period of time, even when an approver is not available. Organizations can also choose to force users to change the password when receiving access to an account.

Users may relinquish accounts to make them available for others and system owners can revoke access to accounts at any time. Account access can also be configured to expire so that the user loses access upon expiration and expired accounts can also be locked and/or have their passwords automatically changed to secret values. System owners need to periodically revalidate the use of accounts that will be used indefinitely.

Getting started: in addition to the implementation procedures described in sections 4.1 and 4.2, the service provider will need to know the system owner for each connected system. Typically, the system owner is the person who performs minor "administrative" tasks as needed, such as when an account can be requested and by whom, as well as assigning accounts to account types defined by the service provider that specify criteria such as durations that can be requested, whether the password must be changed upon use, etc.

2.6 Identity Compliance Service

Organizations of all sizes must comply with a variety of regulations and organizations that conduct business in multiple countries must typically comply with regulations from all the countries. Complying with regulations and preparing for audits is costly: auditors demand proof of compliance and audit procedures dictate that organizations that have difficulty answering initial questions are asked even more questions, thereby increasing the time and expense of audits. In addition to requiring historical records for who could access each resource, regulations typically require proof of procedures and proof of controls regarding who can access each resource.

The Identity Compliance Service can be deployed standalone or combined with other managed services to automate preventive, detective and corrective controls, enforce business policies and avoid weaknesses and deficiencies in information technology general controls. Organizations can choose from a variety of capabilities depending on their compliance objectives. When combined with other services, the compliance service can detect compliance violations related to password policies, separation of duties (SoD) and the authorization of resources, including high-privilege and shared accounts. In standalone mode, files are periodically extracted from other systems and applications and are input into the compliance service.

Each managed service logs its own activities to the consolidated audit database, which is used for compliance and other purposes. All audit data is stored in a database with a wide variety of data views for each managed service. End-user organizations can receive formatted data in a report-compatible format, such as CSV or XML, to be used with DB reporting tools like Crystal Reports as well as with security event monitoring systems.

Preventing and correcting compliance violations, as well as providing some types of detective controls, requires the compliance service to be combined with other managed identity services.

- Gap analysis detects orphan accounts and excessive privileges. Provisioning workflows can alert appropriate people about the exceptions and/or correct the problems.
- Compliance analysis can be combined with provisioning to provide a history of each person's entitlements and who approved them.
- Accounts across multiple connected systems can be associated with a single user so that SoD violations can be detected. When combined with provisioning, the compliance service can prevent SoD violations by making authorization to some resources mutually exclusive.
- When combined with provisioning, automatic or manual corrections can be performed on accounts and privileges on connected systems to comply with business policies.
- When inappropriate administrative changes are made to connected systems, provisioning workflows can be triggered to automatically reset attributes to pre-determined values.

Testing compliance with business policies requires the service provider to understand the organization's business policies and to define them to the system. When implementing corrective controls, service providers and end-user organizations can preview the impact of changes before committing to any potential corrections. Once accepted, refined policies and cleansed identity data can be automatically loaded into the data store. Exception workflows can bring all affected users and accounts into compliance with the organization's provisioning-related policies.

Getting Started: implementation tasks depend on which compliance capabilities are required. It can be used standalone by importing files from systems / applications. Testing for orphan accounts needs user information to be matched with accounts. Validating your business and security policies compares information about your users and your policies with accounts and entitlements on systems / applications. Reporting the history of who had access to each resource, who approved the resource, etc., uses the audit trail from the provisioning service.

2.6.1 Recertification

Recertification workflows compare users' actual accounts and privileges on connected systems with the enterprise's access policies, and information can be made available to business-process owners and managers for review. This information can show who has access to each resource, who approved the resources, etc. so that the appropriate line managers and others can validate that each person has access to the right resources. Organizations can automatically send email reminders to managers when it's time for recertification so that reviewers can accept / decline the resource entitlements for each individual. When required reviews have not occurred by specified dates, follow-up reminders can be sent and the process can be escalated. Requested changes can be routed to system owners or others for additional review / approvals, and approved changes can be automatically routed to the service provider for deprovisioning. Any changes are logged in the audit database for use by auditors. Recertification requires the Automated Role & Account Management.

3. Architecture and Capabilities

Fischer's architecture enables all managed services to share common components and capabilities including the interoperability engine, connectivity, secure web services, administration, audit database, etc. This provides synergies for service providers to implement multiple managed services, which reduces costs for end-user organizations. The architecture also incorporates the use of virtual data so that real-time data can be fetched and utilized for event processing.

3.1 Server Architecture

IaaS™ servers securely reside at the service provider's facilities. Servers use a tightly-integrated J2EE framework and a service-oriented architecture (SOA) that enables additional components to be quickly added as your organization's requirements change.

3.1.1 Interoperability Engine

The Interoperability Engine uses ETL technologies for automated schema discovery, data transformation, policy execution, reconciliation functions and other processes. All logic, data mapping and other data-related and user-related functions occur within the engine rather than in being distributed to remote components that would require persons to implement and manage the components in remote locations. Also, the flexibility of the engine supports situations where the data *should not* be consistent across all connected systems and applications, such as for compliance across some legal or national boundaries.

The engine enables event-driven, scheduled and request-driven procedures including dynamic policy and workflow processing so that different policies and workflows can be implemented depending on the situation. Data from multiple sources can be used in combination to make workflow and provisioning decisions, such as modifying users' access privileges as they change roles within the organization. Also, when a role changes, workflow processes can adjust the privileges for each user assigned to the role. Assigning entitlements to user profiles is determined through both rules and roles to facilitate the use of both high-level and fine-grained authorizations. Roles can be configured as either static or dynamic as the engine can retrieve data in real-time as workflows are executed. The engine enables Fischer Provisioning to be integrated with third-party workflow systems through web service calls as well as through database integration. Workflow queuing assures that all workflows are executed in the appropriate sequence.

The interoperability engine maps user accounts on connected systems to user profiles that are managed in the identity database regardless of whether the accounts were created natively on the connected systems or through the provisioning process. The managed services can also be implemented to enable users to claim their own resources by supplying their authentication credentials for each account. Fischer does not require the same unique identifier to be used across connected systems. Automated schema discovery automatically retrieves the current schemas and formats of data from connected systems so that implementers can use a point & click interface to specify requirements.

3.1.2 Data Repository

The IaaS™ platform is a 3rd-party JDBC-compliant database that can be deployed in a high-availability environment and is used for all storage except authentication. This includes audit, policy, configuration, connector, workflow, roles, rules, account mapping, connected-system data, etc. It stores the state and status of each workflow and other processes for use in auditing and in recovery. Since business logic, workflows, roles, rules, configurations, etc. are stored in the database, they persist when the Identity servers are upgraded. A 3rd-party LDAP directory is used for authentication to any IaaS™ service.

3.1.3 Multi-Organization / Multi-Tenant Operation

Fischer's secure multi-tenant architecture uses a master-client organizational concept that supports multiple end-user organizations in a single high-availability (HA) environment. Secure multi-tenancy enables faster implementation and economies of scale that lower costs for end-user organizations. Policy-based security separates and protects data, workflows, configurations, audit information and administrative permissions for each end-user organization. Anything that can be accomplished in a single-tenancy implementation can also be securely accomplished in a multi-tenancy implementation. The duties of configuration, audit and monitoring can be segregated for each organization and more than one server or HA environment can securely share the same database store. Sophisticated controls manage workloads for each organization so that extremely high demand by one organization does not affect other organizations by controlling the percentage of resources that each organization can consume.

3.1.4 Enterprise Business Roles

IaaS™ technology supports dynamic and static roles, including and all four levels of the National Institute of Standards and Technology (NIST) RBAC model, but does not require roles. This is a best-practices solution since most organizations have not fully defined their business roles and would not be able to quickly implement a solution that requires roles. Role access privileges can span multiple connected systems in multiple locations and domains, and the solution facilitates creation, modeling, testing and maintenance of roles. Each role can manage permissions to IT and physical resources based on the values of multiple attributes like title, department, and location. For example, when a person changes departments, his permissions can be automatically modified to reflect the change.

3.1.5 Separation of Duties

Separation of duties (SoD) rules can specify that any number of policies is mutually exclusive, and rules can be prioritized for which policy (if any) is to be executed when policies would produce an SoD conflict. When implemented with the Automated Role & Account Management Service, SoD violations can be prevented and appropriate persons can be notified when an SoD conflict occurs.

3.1.6 Role Matrix Management

Since formal organization structures rarely show the real relationships of how business is conducted, support for multiple operating views is required to enable organizations to flexibly manage entitlements. IaaS™ technology manages membership in roles and other decision criteria through the combination of multiple attributes as well as by considering elements such as each manager's span of control, and for which periods of time end users, managers, and approvers should have their privileges. For example:

- Temporary privileges can be granted to members of a project team for the duration of a project
- Employees can be provisioned to be located in one country but report to line managers in other countries
- Approvers can be dynamically determined, even when the approver is not the person's official manager, such as in cases where a particular request requires approval by a higher-level manager

3.2 Connectivity

The IaaS™ platform uses Fischer's patented Global Identity Architecture™ (GIA) to quickly connect a wide assortment of systems and applications for bi-directional communication. The solution works without regard to the number of domains and firewalls, or the complexity of each organization's IT environment. It also incorporates the entire enterprise, from older legacy systems to "state of the art" web services-enabled applications. IaaS™ has the only technology that can operate without duplicating a significant part of the solution or introducing additional components, such as proxy and VPN servers, as well as adding complexity, cost, and the risk of opening potential security holes

The IaaS™ platform connects to 3rd-party systems and applications via native application programming interfaces (APIs), standards such as SPML, web services, XML, as well as interaction with the underlying repositories of connected systems. Changes on connected systems can be detected in near real time to trigger appropriate workflows based on the organization's business rules. For example, when a person is terminated in the HR application, workflows can immediately revoke access to all accounts on connected systems or can take other actions as appropriate. No software is required on connected systems with the exception of applications that require a local component for any communication.

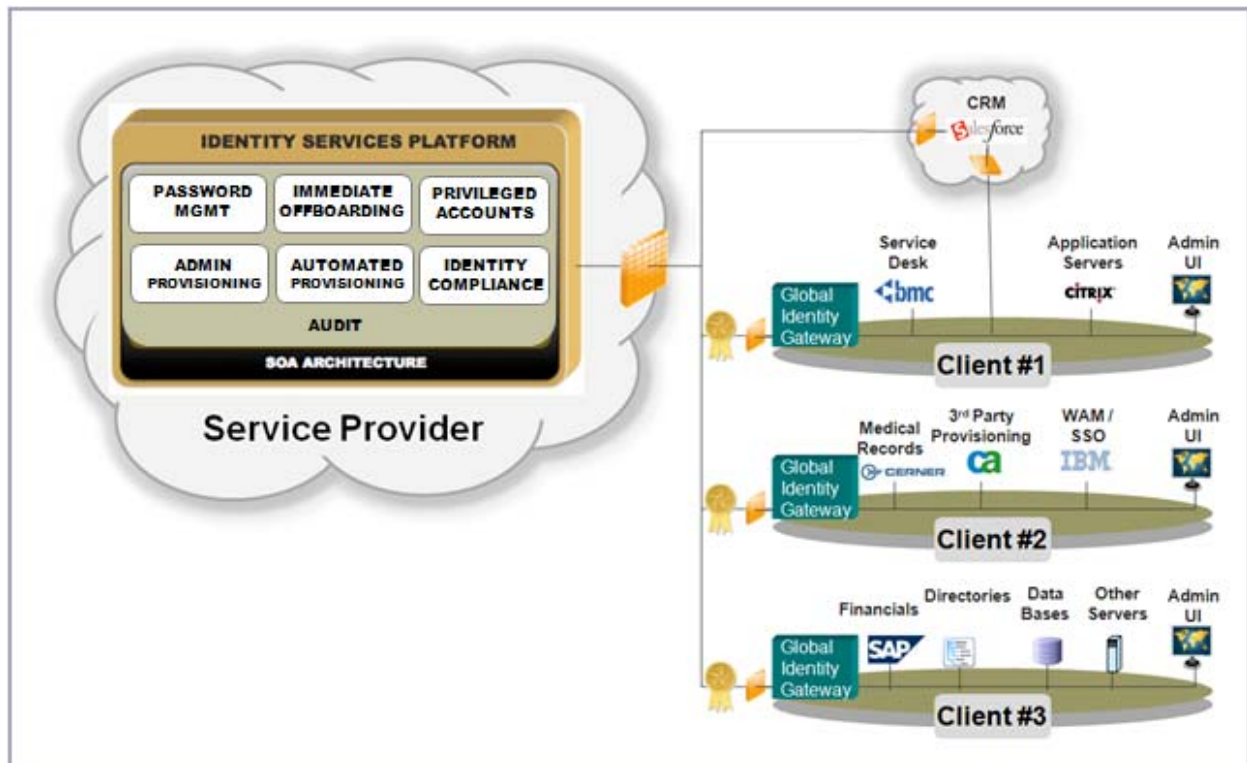


Figure 3: Global Identity Architecture

A Global Identity Gateway (GIG) is installed in each end-user organization domain or other domain where systems or applications are to be connected. The exception is in connecting to applications with web-services connectivity. Each GIG can handle connectivity for any combination of connected systems as it contains the capabilities of all Fischer connectors. It is not required for browser access by end users. The GIG uses web services to avoid firewall-related security problems and web-services security (WSS) to secure the channel, so no new firewall ports need to be opened for the GIG as long as web services are permitted to pass through the firewall. The GIG can be deployed as an appliance or downloaded and installed using standard installation procedures. Configuration and administration are accomplished remotely from the service provider without the need for expensive travel, which reduces costs to end-user organizations.

3.3 High Availability

laaS™ technology has been designed for reliability and availability with no single point of failure. It eliminates or significantly reduces downtime from both unplanned outages and planned outages. The plug & play framework eliminates the need to disrupt production when adding roles, policies, rules, workflows, connected systems, users, additional end-user organizations, etc., and the suite does not require periodic batch jobs to update data or to synchronize servers.

The managed services use J2EE high-availability techniques and coordination for a load-balanced high-availability cluster of servers. The use of load-balanced data repositories avoids having a single point of failure. The suite automatically stores the state and status of each workflow and other process for use in auditing and in the unlikely need for recovery. Also, workflows can be configured to periodically check connected-system availability at end-user organizations and send alerts in the case of a failure.

3.4 Scalability

laaS™ technology delivers unlimited scalability as it scales both vertically and horizontally. Service providers can add servers to the high-availability cluster without disrupting any end-user organizations. Organizations can quickly expand their use of laaS™ technology to encompass mergers, acquisitions and other growth.

3.5 Extensibility

laaS™ technology uses an SOA framework to ensure extensibility. User interfaces can be changed, and .NET and Java web services clients can be used with laaS™ technology or can be embedded within an organization's web portal. Additional logic can be added and connectivity can be extended to more systems and applications, including legacy applications. laaS™ technology can also interoperate with web SSO, enterprise SSO and authentication products.

3.6 Security

laaS™ technology reduces risk and assures data integrity by protecting the use of assets such as systems, data and applications. It enforces business and security policies so that only specifically authorized persons can access resources. Data is protected in transit and when stored.

For a discussion of the security aspects of laaS™, please refer to Fischer Identity as a Service™ Security Overview.

4. Implementation and Use

An laaS™ implementation is much faster than traditional identity management solutions, and entails three major activities: establishing connectivity between the service provider and the end-user organization, configuring the identity management application(s) and onboarding the user population. Service providers can perform all implementation activities without traveling to any of the locations of end-user organizations.

4.1 Establishing Connectivity

Global Identity Gateway can be deployed as an appliance or quickly installed by your organization via standard installation procedures. The service provider remotely configures the gateway and the public key infrastructure used to secure the channel, and needs information to connect to your systems. You can choose to enter your own credentials for the administrative accounts used by the solution so that the service provider never learns the credentials. Some 3rd-party systems also require local components to be installed (e.g., natively connecting to RACF). Specific business cases can require triggers to be implemented on databases or other applications like SAP.

4.2 Configuring the Solution

Fischer has created repeatable services to reduce configuration time, improve quality and decrease costs. These services can be deployed in as little as a few days and up to 30 days. Although the services themselves are fixed in scope with limited choices, each of the services can serve either as a full solution or as a starting point for more complex scenarios with unlimited choices, and multiple service offerings can be combined. All configuration is performed remotely by the service provider using interfaces that don't require any scripting and securely enable reuse, which lowers costs for end-user organizations. Each organization needs to provide information for the configuration, such as data related to approvers, administrators (for administrative provisioning with the Role & Account Management Service), system owners (for Privileged Account Access Service), etc.

4.3 Onboarding Users

Manual account reconciliation is difficult and error-prone due to the vast number of accounts, cryptic or inconsistent account naming convention, workforce size, etc. When deploying IdM (or when extending IdM to new connected systems or new user communities), existing accounts must first be reconciled to specific users. The onboarding process discovers orphan accounts (accounts having no explicit owner). It also provides a starting point for determining provisioning and compliance policies regarding who should have access to each resource. IaaS™ technology provides multiple options for onboarding new and existing users. None of the following onboarding processes requires that the account names on the connected systems match each other.

4.3.1 Account Reconciliation

Manual account reconciliation is difficult and error-prone due to the vast number of accounts, cryptic or inconsistent account naming convention, workforce size, etc. When deploying identity management (or when extending it to new connected systems or new user communities, including those from mergers and acquisitions), existing accounts must first be reconciled to specific users. The account reconciliation process discovers orphan accounts (accounts without an explicit owner), helps cleanse data by identifying accounts that violate the organization's naming conventions, and provides a starting point for determining provisioning and compliance policies regarding who should have access to each resource.

Account reconciliation enables service providers to automatically perform multi-stage correlation to reconcile accounts and users based on any number of user / account naming conventions and attributes. For instance, administrators of legacy systems might have used a variety of algorithms to create user IDs. Account reconciliation can iteratively match users and accounts through any number of algorithms such as:

- ID=employee number or ID=firstName.lastName qualifier OR ID=firstInitial lastInitial lastName qualifier
- AND the phone-extension listed on the system matches the person's actual phone extension, etc.

Accounts can be correlated to users based on a combination of any number of attributes that might be available from a connected system such as user ID, name, email address, phone number, etc. The process automatically matches as many users and accounts as possible and lists other accounts as possible orphan accounts to be investigated.

4.3.2 Account Self Claiming

Account self claiming can be used independently or after account reconciliation or user self registration to enable end users to specify which accounts they use on each connected system. Users can select from a list of systems and applications and present valid credentials to validate their claims. As an option to encourage end-users to claim their accounts, organizations can have their service provider disable any accounts that are not claimed within a specified period of time until the users claim them.

4.3.3 User Self Registration

Self registration allows new (unknown) users to register for identity management accounts so they can use the user self-service provisioning interface to request IT accounts and other resources. This is especially useful for simplifying user onboarding for administrative provisioning (Role and Account Management Service) and for SMB organizations. It is also helpful for larger organizations to enable business partner employees to register for accounts. Of course, approvals by one or more persons can be required for a user's registration to be processed.

4.4 Validating User Entitlements

As part of the onboarding process, actual entitlements can be compared to your organization's policies to validate that each user has appropriate accounts and entitlements on connected systems. At the completion of this assessment, most organizations choose to create a list of potential problems for someone to research and correct as needed, but it's also possible to automatically change entitlements to match the organization's policies. Validating user entitlements can be performed as part of the onboarding process, but it can also be performed periodically for identity compliance.

4.5 Mergers and Acquisitions

If your organization merges or acquires another company, the onboarding processes can be altered and repeated to initiate the identity management process for the people in the acquired organization. For instance, the onboarding process can be used to map people and sub-organizations (departments, workgroups, locations, etc.) from the acquired organization to specific sub-organizations within the parent organization. The onboarding process can then be used to provision appropriate resources to each person and to register accounts for password reset and synchronization.

4.6 Business Divestitures

If your organization sells a subsidiary, division, or other part of the organization, a process can be used to reorganize identity management for that part of the organization so that it can be extracted from the parent organization. Your organization can negotiate with the acquiring organization to determine what types of identity management data to provide to the acquiring organization. The offboarded organization (or its new parent) can decide whether identity management for the new organization is handled through IaaS™ or whether Fischer identity management is deployed in house at the divested organization.

4.7 Changing Business Directions for Procurement Models

If your organization changes strategic direction and decides to run your identity management solution in house, your organization's IaaS™ infrastructure and data can be efficiently redeployed using an in-house model.

Fischer International Identity
3073 Horseshoe Drive South
Naples, Florida 34104
+1 239-643-1500
www.FischerInternational.com



Built for Business... *Yours*™

Document: MCW-09-120A July, 2009

Copyright © 2009 Fischer International Identity, LLC. All rights reserved.

Fischer International, Fischer International Identity, Managed Identity Services, Managed Identity Services Technology, Identity as a Service, IaaS, the Fischer International Logo, Global Identity Architecture, Built for Business...Yours, and all other Fischer product or service names are the trademarks and/or registered trademarks of Fischer International Identity.
