Identity and Privacy Strategies

Reference Architecture Technical Position

User Provisioning

Version: 2.0, Oct 05, 2005

AUTHOR(S):

Nick Nikols

(nnikols@burtongroup.com)

Statement of Problem

How should enterprises automate the mapping of identities to accounts, credentials, and access rights?

BURTON GROUP 7050 Union Park Center · Suite 510 · Midvale, Utah · 84047

P 801.566.2880

F 801.566.3611

www.burtongroup.com

Publishing Information

Burton Group is a research and consulting firm specializing in network and applications infrastructure technologies. Burton works to catalyze change and progress in the network computing industry through interaction with leading vendors and users. Publication headquarters, marketing, and sales offices are located at:

Burton Group
7050 Union Park Center, Suite 510
Midvale, Utah USA 84047-4169
Phone: +1.801.566.2880
Fax: +1.801.566.3611
Toll free in the USA: 800.824.9924
Internet: info@burtongroup.com; www.burtongroup.com

Burton Group's *Identity and Privacy Strategies* service provides objective analysis of networking technology, market trends, vendor strategies, and related products. The information in Burton Group's *Identity and Privacy Strategies* service is gathered from reliable sources and is prepared by experienced analysts, but it cannot be considered infallible. The opinions expressed are based on judgments made at the time, and are subject to change. Burton offers no warranty, either expressed or implied, on the information in Burton Group's *Identity and Privacy Strategies* service, and accepts no responsibility for errors resulting from its use.

If you do not have a license to Burton Group's *Identity and Privacy Strategies* service and are interested in receiving information about becoming a subscriber, please contact Burton Group.

Table Of Contents

Statement of Problem

How should enterprises automate the mapping of identities to accounts, credentials, and access rights?

Typical Requirements

The process of managing user accounts and access to corporate resources is a daunting task that must overcome the complexity of typical enterprise infrastructure environments. Regulatory, security, and business policies overlay additional controls that must be enforced by the solutions that provide administration of user accounts and entitlements. A product category called "user provisioning," which Burton Group defines as "an integrated set of tools used to manage the lifecycle of users and IT entitlements," addresses these issues. Detailed coverage of provisioning technology and market conditions can be found in the *Identity and Privacy Strategies* report, " Provisioning: Many Product Choices for Enterprises."

This Technical Position helps enterprises design an optimal user provisioning solution that automatically controls access to resources according to policy constraints. Discussion is limited to human users (employees, contractors, partners) and information technology (IT) entitlements (not web services or physical assets). For definitions of "provisioning services," "resource," and other terms, see the *Identity and Privacy Strategies* overview, "Concepts and Definitions."

Make New Employees Productive Faster

In an ideal situation, new employees would have access to all necessary resources the moment they walked in the door on their first day of work. In reality, new employees are delayed for several days, or even weeks, before all resources have been properly allocated. Several factors may conspire against new employees and their attempts to be productive immediately.

- Typically, enterprises do not have streamlined processes for creating all the accounts and granting access to all the applications, databases, web servers, and other resources.
- Many islands of administration may exist for each application or platform, resulting in coordination challenges among several administrators.
- Paper-based forms—which take longer to process, must rely on interoffice or postal mail services, and are prone to error—are still used in many environments.
- Multiple levels of approval may be required before granting access to sensitive systems, resulting in prolonged delays if the process is not automated.
- Poor or incomplete documentation of the necessary resources for each position results in multiple iterations of requests before access settings are completed correctly.
- Backlogs of requests that must be manually completed by administrators may result in a queue for new requests and cause additional delays.

Provisioning products are designed to automate the steps to create new user accounts and establish initial authorization or entitlement settings. But products must be flexible enough to accommodate the intricate nuances of how each enterprise operates. Enterprises require provisioning tools that are readily customized to interface with existing infrastructure systems as well as business processes.

Reduce Administration Costs Through Automation

Manual administration of user accounts is an expensive task for many enterprises that are aggressively seeking to reduce costs and improve operational efficiency. Automation can lower administrative costs by reducing the manual input required for traditional methods of account management. Enterprises require automation in the form of workflow processing to route approval requests to the appropriate manager. Workflow systems must also be able to handle ordering of events, escalate priorities if requests are not completed in a timely fashion, and interface with other systems during processing.

Provisioning products must also provide automation based on rule- or role-driven policies. For example, if a new employee has the “research engineer” role, then the provisioning system should be able to automatically create accounts and grant access on systems that are restricted to users in the research department. Enterprises also need to trigger provisioning actions based on other user attributes in authoritative repositories. Therefore, provisioning products must support rule-driven policies that allocate resources based on specific attributes. For example, all employees that have the “eastern region” attribute could be assigned to the Pittsburgh e-mail server.

Automation is also important for account maintenance and termination. As employees or other users change their relationship with the organization, the provisioning systems must be able to act on attribute changes in key systems. Enterprises normally have a small number of key repositories or systems where most user status changes are made. The provisioning system must be able to recognize these changes and adjust the user's privileges in accordance with the new status. Changes in user status can include promotion to a new position, transfer from one department or location to another, or assignment to different projects. At the end of the user's relationship with the enterprise, the provisioning system must be able to quickly cancel all accounts and entitlements in order to prevent unauthorized access in the future.

Support for Request-Driven Provisioning

In addition to fully automated provisioning actions, enterprises also require tools that support request-driven processing. In many situations, organizations or business units prefer that users be issued basic privileges such as network operating system (NOS) and e-mail accounts. Supervisors or managers then submit requests for additional privileges or accounts, depending on how the user's tasks are assigned.

Some enterprises may decide to permit users to submit requests for the additional resources or privileges required to carry out their job functions. In other cases, users may be able to request optional resources or ask to be included on distribution lists dedicated to topics they're interested in.

Password Management

Efficiency and cost savings can be enhanced when the provisioning solution includes password management capabilities. Calls to the help desk are reduced and workers are more productive when they can reset forgotten or compromised passwords through a self-service web interface. Password synchronization can also play an important role in simplifying the user experience by capturing password changes and forwarding them to other systems on the network.

Password management functions generally fit into the two categories mentioned previously: reset and synchronization. Enterprise requirements for password management may consist of one or both capabilities to satisfy user communities. Password reset, as illustrated in Figure 1, occurs when a user satisfies challenge questions on a self-service webpage and enters a new password. This new password is then forwarded to other systems within the domain based on the user's selection or security policy.

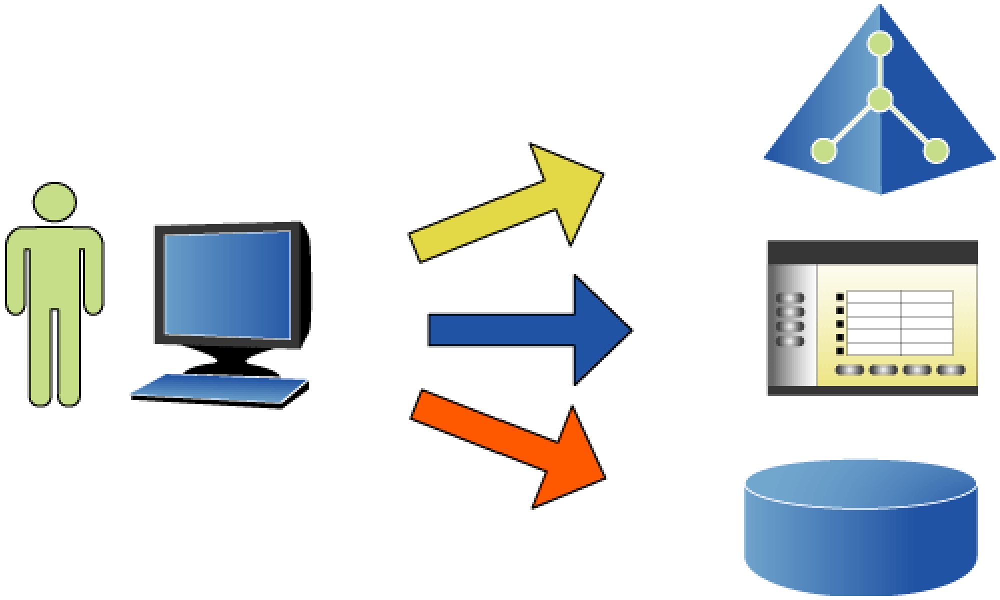

Figure 1: *Password Reset*

Password synchronization, on the other hand, captures changed passwords from one or more systems and then propagates the new password to other systems in the domain. In this scenario, enterprises require password capture capabilities on their key systems, such as NOS, mainframe, or other significant application systems.

Both password synchronization and password reset systems create a risk aggregation problem because the same "weak" credential may be applied across multiple business applications. Single sign-on (SSO) solutions are better able to address this problem by leveraging strong authentication as the primary login and generating unique strong passwords (long and random strings) for each supported application.

Whether an enterprise chooses synchronization or reset, there are additional requirements a password management system must address. Acceptable password syntax is typically dissimilar for different platforms and applications in medium to large enterprises and is driven by a combination of security policy and platform capabilities. Password syntax policies regulate password length; the mixture of letters, numbers, and special characters; how often a password can be reused; whether dictionary words can be used; and several other options. Therefore, the product may have to accommodate different password policies across systems in the infrastructure, if the enterprise does not intend to create a single password policy. Password systems must also securely handle passwords during capture and transport to guard against clear text passwords being revealed. Finally, password systems must have broad support for all the platforms and systems that a given enterprise may have deployed.

Centralized Policy Management

User provisioning functions play a key role in an enterprise's identity management (IdM) infrastructure by automating basic account setup and entitlements administration. These basic functions normally consist of creating the user account, setting the initial password, and adding the new user to group lists on target platforms. Some enterprises desire to manage more policy controls from the provisioning solution by integrating tightly with managed platforms to administer group settings, access control lists (ACLs), and local security rules.

To move beyond basic user administration, provisioning products must provide more functionality in their connectors to support all the necessary security controls on managed platforms. If these functions are not inherently available, enterprises require full-feature application programming interfaces (APIs) and the ability to readily customize the provisioning product. With these enhancements, a provisioning system can broaden and deepen its control of policy settings by managing group and ACL creation, setting access rules, creating role definitions, and providing consistent policy settings on all target systems.

Provisioning systems must be capable of supporting multiple provisioning domains in order to properly align with business requirements. Provisioning domains are defined as a subset of policies within a provisioning system.

Delegation of Provisioning Policy

Although many enterprises prefer centralized control of policy settings, they also require delegation of certain policy administration functions to business units, branch offices, or partner organizations. Provisioning systems must meet this requirement through flexible user interfaces that can be customized for local preferences, and by strictly partitioning policy domains for delegated administrators.

Delegation functions are required for approving create or change requests, changing policy settings, or modifying rule definitions that control automated processes. Delegation partitions can be created by implementing separate physical servers or using security settings to carve out subdomains.

Improve Enterprise Security

Many enterprises suffer from security vulnerabilities introduced by orphaned accounts, which are accounts that are not immediately deleted or deactivated when an employee leaves the company. Most administrative staff members are overwhelmed with account creation and change requests. Consequently, account deletion requests often are relegated to last priority. To address this area, provisioning systems must support a rich set of functions, including delete, suspend, and disable—depending on what the target environment is capable of.

Enterprise security can also be enhanced by consistently enforcing policies across the environment and ensuring that users have all the privileges they require to perform job functions but no extra privileges that would represent a security risk. Consistency is enforced by automating account creation, but provisioning systems must also detect and reconcile changes made on target systems by local administrators. When changes are detected that occurred outside formal processes, a provisioning system should be able to submit these changes through its policy engine and determine whether the changes should be allowed, logged, or rescinded.

Demonstrate Compliance with Legal, Regulatory, and Policy Mandates

Reporting tools are required to show compliance with prevailing laws and regulations, depending on the jurisdiction the enterprise operates in and the type of business it conducts. Enterprise requirements include the ability to determine when access was granted to sensitive systems, which manager approved the changes, and which policy was in effect that permitted the access. Further, many enterprises are required to answer such inquiries as: Who has access to the trading system? What systems do users have access to? How many dormant accounts exist on the payroll system?

Provisioning systems may not address all reporting requirements with out-of-the-box features, but they must be configurable or customizable. Every enterprise has unique reporting requirements that will require changes to the provisioning product or integration with other reporting tools the enterprise may have deployed.

Auditing and Accountability

Auditing functions provide an important feedback mechanism that helps administrators determine whether policy objectives are being carried out. Enterprises require complete audit records to verify privilege settings, check whether users have accounts on proper systems, and generate the other reports mentioned in the "Demonstrate Compliance with Legal, Regulatory, and Policy Mandates" section of this Technical Position.

Audit trails also form a basis for accountability within the organization by tracking who requested access, why the request was granted or denied, and who approved the request. During investigations, audit records are required to conduct thorough analysis of incidents. Due to their reliance on audit records for critical analysis, enterprises require that audit systems and the records they produce be protected from tampering. This implies using cryptographic technologies to sign and/or encrypt data. Enterprises may also require clear separation of duties for administrators where authority to modify audit settings and access audit records is restricted to a subset of administrators who do not perform user management functions. More information on auditing can be found in the *Security and Risk Management Strategies* report, “Auditing and Audit Trails,” and the *Identity and Privacy Strategies* report, “Achieving Organizational Compliance: The Emerging Role of Identity Audit Software.”

Leverage and Improve Enterprise Data

Enterprises that are considering the deployment of provisioning solutions have previously invested in infrastructure systems that contain identity and enterprise data. When planning and designing provisioning deployments, enterprises are seeking to leverage existing investments where possible. Therefore, provisioning systems may have to accommodate installed directories, databases, and other repositories that currently hold valuable identity data.

Approaches will vary between enterprises, but possible scenarios for the provisioning repository architecture include using a centralized or virtualized repository. Enterprises that have established a centralized directory for all users will likely wish to leverage this repository for the provisioning solution. Conversely, enterprises that have identity data distributed in several systems (and prefer this configuration) may be more inclined to use a virtual repository method.

Another issue that plagues most enterprises is corrupted or incomplete identity data. Data cleansing is a significant step in the provisioning deployment process. Therefore, enterprises require that the provisioning system ensure data will be properly maintained going forward.

Alternatives

The following categories of alternative technologies are considered in this Technical Position:

- Homegrown and commercial solutions
- Workflow options
- Internals of the provisioning system
- Policy models
- Password management

Homegrown and Commercial Solutions

Provisioning is not something new for large enterprises; many have built homegrown systems to perform provisioning functions. Enterprises that have built homegrown systems can continue with these investments if the resulting systems meet current requirements. Other enterprises that are just starting out with provisioning may still decide to build a system using scripting tools, messaging systems, or workflow packages that are already deployed in the infrastructure.

Commercial solutions are available from a wide range of vendors, as described in the *Identity and Privacy Strategies* report, "Provisioning: Many Product Choices for Enterprises." In addition to the large number of providers, there are also many choices for how provisioning products themselves are designed and packaged.

Large Array of Provisioning Products

The provisioning market is a point of convergence for several related technologies and markets, including password management, security management, and meta-directory services. As many as 30 vendors *claim* to offer some form of provisioning product, thus providing further testament to increasing consumer interest in this market segment. When moving into the provisioning space, vendors bring with them product approaches that are closely aligned with these starting points and that provide diverse capabilities. In general, vendors in today's market fit into one of the following market categories: pure play, security management, directory tools, password management plus, IdM suites, web access management (WAM)-related, and NOS-oriented. Table 1 illustrates these market segments and the vendors that fit into them. A vendor's relative position in this table does not indicate a superior or inferior status. Vendors are grouped in order to more easily discuss this large contingent.

Security management	**Password management plus**	**Directory tools**	**Pure play**	**IdM suites**	**NOS-oriented**	**WAM-related**
ASG-Entact ID	Blockade Systems	Fischer Intl.	BHOLD	Computer Associates	Abridean	Entrust
Beta Systems	Courion	MaXware	Thor Technologies	HP	BindView	Jamcracker
BMC Software	M-Tech	Microsoft		IBM Tivoli	NetIQ	RSA Security
Bull/Evidian		Radiant Logic		Novell	Quest Software	
OSM				Oracle	Voelcker	
				Sun		

Table 1: *Vendor Segmentation Matrix*

The blurring of product lines and the overlapping strategies of a variety of vendors have led to substantial volatility in the provisioning market. Partnerships, mergers, and acquisitions form, and rumors of acquisitions circulate. Consolidation in provisioning and other IdM market segments will continue as vendors jockey for position and seek to gain a larger market share.

Workflow Options

Enterprises have several alternatives to consider when implementing workflow as part of a provisioning solution. Choices must be made in the following general categories, but it's likely that enterprises will use a combination:

- Static vs. dynamic workflow
- Data- or request-driven workflow
- Approval or rule-based processing
- Graphical or text interfaces
- Integration with existing workflow systems

Static vs. Dynamic Workflow

Provisioning products have varying degrees of flexibility in how workflows are defined to meet the business process activity within the organization. In some cases, all possible event scenarios must be understood by the implementer and hard-coded into the workflow. Using hard-coded workflow events is also known as static workflow.

With dynamic workflow, the provisioning system is able to calculate the actions to take based on environmental variables. Instead of coding all actions for every possible workflow outcome, dynamic systems evaluate conditions and determine the proper actions to take on the fly. The Computer Associates *e*Trust IdentityMinder eProvision workflow is an example of a dynamic system in which, for example, if a user changes location from the New York to the London office, the workflow system calculates all the actions necessary to move the e-mail account to the London-based server.

Data- or Request-Driven Workflow

Workflow processing can be triggered by human requests or by changes to identity data in the authoritative repository. If request-driven workflow is used, managers or users can initiate the action by applying for privilege changes or by asking for accounts on systems—typically through a web-based interface.

Data-driven workflows are started when authoritative identity sources change. Many enterprises will use the human resources (HR) system or enterprise directory as the main input source for provisioning activity. In such cases, changes to the HR system, like adding a new employee, changing a job title, or removing an employee entry, cause the provisioning system workflow to take specific actions.

Approval or Rule-Based Processing

Once workflows have been initiated, enterprises have options for how actions are granted or denied. Enterprises have the option of sending approval requests to the resource owner(s) or determining approval based on a rule or policy definition.

If the enterprise uses approval routing, then each request for access that enters the provisioning systems is sent to the appropriate manager or resource owner for review. If the approver does not grant the request, then a message is sent to the original requestor and the action is not taken. On the other hand, if the request is approved, then further processing continues and the requested action is completed.

More sophisticated workflow systems can route requests to multiple users and complete the request when the first approver reviews the entry. In other cases, if the first approver does not respond in a timely fashion, workflow systems can react by escalating the issue and sending the request to an alternate approver.

In rule-based systems, workflow requests are evaluated using a rule or policy definition that determines the outcome without human intervention. For example, if a user requests access to the research and development system, and the user is a member of the engineering department, then access is automatically granted and the user account is created. Similar rules can be established for data-driven systems—for example: If a new employee is added to HR and the department attribute equals cashier, then add the user to the branch terminal system.

Graphical or Text Interfaces

Most workflow systems use a graphical interface to define and view workflow events. Many administrators prefer this method because it can aid them in visualizing the process flows when developing or troubleshooting events. Some systems, such as M-Tech ID-Synch, use a more text-based approach to workflow administration to give implementers control through scripting languages.

Integration with Existing Workflow Systems

There are cases in which an enterprise has previously invested in workflow systems associated with other applications in the infrastructure. In these situations, the enterprise may prefer to use the skills its staff has gained with existing workflow products for the provisioning deployment. Some vendors, such as BMC and IBM Tivoli, are more adept at integrating external workflows into their products. More information on workflow and the identity lifecycle can be found in the *Identity and Privacy Strategies* overview, "Identity Lifecycle and Workflow: Building an Identity Program."

Internals of the Provisioning System

The architectural construction for provisioning system internals can vary for the data and policy repository, provisioning server, connectors to managed systems, auditing functionality, and reporting capability.

Provisioning Repository

Significant options are available for how enterprises can construct the repository that contains identity and policy data. Options range from a centralized repository that holds all identity and policy data to a virtualized repository that contains only metadata locally in the provisioning server. BMC offers a combined approach in which all user data is maintained in a central database and includes Lightweight Directory Access Protocol (LDAP) access to the repository through a bundling agreement with Radiant Logic.

Sun Microsystems and Evidian use a virtual repository for their solutions, and vendors such as Computer Associates and Courion permit flexibility on how much data is centralized versus distributed in other authoritative systems. Enterprises must also be concerned about how data is synchronized with target systems, in addition to determining the best choice for the provisioning repository. For each scenario, tools and utilities must be able to keep provisioning server settings coordinated with data in the managed systems.

Provisioning Server

Designing the structure for deploying provisioning servers will largely depend on the scale of the enterprise and the geographical distribution of the workforce. Enterprises may choose to implement a centralized server or to distribute multiple servers to provide high availability and scalability for larger environments. Thor Technologies and other vendors permit the implementation of distributed provisioning servers that cooperate to provide redundancy, backup, and share workloads.

Connecting to Managed Systems

Provisioning servers interact with managed platforms through connectors, or agents, that carry out account management or policy setting instructions. There are two basic alternatives for where connectors are installed in the provisioning architecture: The connectors are installed on the provisioning server and operate in a remote fashion, or they are installed on the managed system and interact locally with the target system's APIs. A slight modification to this approach occurs when connectors are installed on a gateway server that acts as an intermediary.

Remote connectors rely on published APIs that can be called securely over network connections. These APIs must also provide all the functions necessary to carry out provisioning actions that the enterprise requires. If remote capabilities are not available or the API cannot be called over a secure channel, then the enterprise can decide to implement local connectors for managed systems.

Provisioning vendors have historically built their connectors in a proprietary fashion, but they have a standards-based alternative with Service Provisioning Markup Language (SPML). SPML was formally ratified in November 2003 and is supported in several commercial products. Using SPML, vendors can build connectors with this standard interface, instead of relying solely on proprietary means. Enterprises can make SPML the preferred choice for connecting to managed systems and only use proprietary connectors as an exception.

Deciding between remote and local connectors involves an evaluation of tradeoffs. First, the enterprise must define required functionality, political constraints, and preference for local/remote connectors. Then the enterprise can evaluate products that most closely match the desired approach, but the enterprise may come to a point where it must decide whether to compromise on functionality or on “remoteness” (for example, one HR system wants all the functionality a local connector can give whereas another HR system does not want anything installed in its system).

Auditing Provisioning Actions

To ensure that provisioning processes are operating according to policy, enterprises must decide how to plan the capture and storage of audit records. Products offer alternative mechanisms for collecting data by providing a centralized audit database or the capability to divert audit records to external systems. Courion AccountCourier permits enterprises to send audit records to external auditing systems, send alerts to event monitors, or take other actions based on audit records.

Policy Models

As discussed previously, policy models will impact how other parts of the provisioning system are architected, particularly the workflow operation. There are three different models for implementing policy in a provisioning system: request-, role-, or rule-driven.

Request-Driven

In this model, access to resources is granted on a per-request basis rather than in a fully automated fashion based on rules or roles. Therefore, users or their managers must submit requests to create accounts, access applications, or change privileges. These requests are then processed according to workflow settings that automatically approve requests or forward them to resource owners for review.

Role-Driven or Rule-Driven

Role- and rule-driven policies are similar, in that they both operate based on attribute settings in identity repositories. Role attributes represent a collection of privileges that are necessary to carry out a particular job function, such as “fulfillment processor.” This collection of privileges could represent accounts on certain systems, entries in group lists, or inclusion in ACL settings.

Rule-driven policies take actions based on the value of other attributes within the identity repository. For example, rules can be established that provision resources based on location, department, cost center, job code, business unit, or title.

Password Management

There are two basic alternatives for password management. The simpler of the two is password reset, in which the user satisfies challenge questions at a self-service portal and the new password is pushed out to other systems. The more complex is password synchronization, in which a password change is captured on one or more systems and sent to other systems under the product's control.

Future Developments

The following developments are important to track in this space, although they are still ongoing or not yet sufficiently developed to include in the current architecture.

Convergence in the Marketplace

The trend for convergence between password management, provisioning, and meta-directory products continues as vendors introduce additional features to level the playing field among the different subcategories. Early signs of the next wave of convergence can be found in the WAM area. OpenNetwork (recently acquired by BMC) introduced basic provisioning features to its product by integrating with Microsoft Identity Integration Server (MIIS) and supporting SPML. Jamcracker has a combined WAM and provisioning product called Pivot Path that is available as part of its service provider offering. Secured Services also has a combined offering that is commercially available.

IdM suite vendors are working to tightly integrate products and technologies into seamless, bundled packages. Over the next few years, IBM, Computer Associates, Novell, and Sun will work to integrate audit services, use of repositories, user interfaces, and other back-end services across product lines. This integration will further blur the lines between product categories and impact integration with best-of-breed products. For more information on audit services, see the *Identity and Privacy Strategies* report, "Achieving Organizational Compliance: The Emerging Role of Identity Audit Software."

Provisioning of Non-IT Assets

Some enterprises are managing non-IT assets such as personal computers, mobile phones, or personal digital assistants (PDAs) through their provisioning systems. However, most organizations are focused on corralling and improving the provisioning of user accounts and privileges.

As enterprises gain control of the IT environment, attention will shift to other assets that can be managed through provisioning. In addition to IT accounts, users need several other resources to effectively perform their job functions. Enterprises can gain additional efficiencies if these non-IT assets are inventoried and controlled via provisioning so that new employees have *all* the accoutrements for the tasks to be carried out. Provisioning products are evolving to handle these assets more effectively; in many cases this may just involve an e-mail to an administrator who manually processes the request, but products that interface with asset management systems more directly to automate the entire process will also soon be available.

With tighter management of these resources, enterprises can also expect additional cost savings. Expenses can be significantly reduced if mobile phones, subscriptions to research services, telephone calling cards, and company credit cards are discontinued immediately upon termination of the user's relationship with the organization.

Provisioning to Support Federated Environments

Identity federation is gaining rapid acceptance in the market, based on Security Assertion Markup Language (SAML) and Liberty Alliance technologies. The early focus of federated scenarios concentrates on SSO across independent security domains. SSO is typically not reliant on persistent identity data on both sides of the federation transaction, but this will change as the market matures and adopts account linking and other advanced federation options. Participants in advanced federation scenarios will need to exchange persistent identity information, and provisioning tools can be used for this function.

Burton Group has published additional information on SAML, Liberty Alliance, and federation in the following *Identity and Privacy Strategies* reports: "SAML 2.0: Convergence Point for Browser-based Federation," "Liberty Alliance: Meeting Early Adopter Requirements," and "Federating a Distributed World: Asserting Next-Generation Identity Standards."

Increased Adoption of SPML

SPML was formally introduced in November 2003 and is not yet universally supported by provisioning vendors. However, there are a small number of vendors who implemented SPML 1.0 in their production environments. As more provisioning vendors adopt the standard, SPML implementations will increase in three ways.

First, vendors can begin to build connectors for managed systems using the SPML framework instead of proprietary methods. Second, enterprises can integrate identity infrastructure components, such as HR systems, to the provisioning server using the SPML Requesting Authority (RA) interface. Finally, enterprises can deploy multiple provisioning servers (using the same or different products) and connect the systems through SPML messaging.

Although vendor support for SPML is an important development, enterprises stand to gain further when major application system providers create SPML Provisioning Service Target (PST) interfaces for their products. As major application vendors integrate PST interfaces into their products, enterprises will have standards-based integration points, which will permit easier implementations and less-complicated configurations.

Standards development continues with the Organization for the Advancement of Structured Information Standards (OASIS) Provisioning Services Technical Committee (PSTC) working on version 2.0 of SPML. The next release will address several issues around searching of large data spaces, handling of complex Extensible Markup Language (XML) objects, and reliance on Directory Services Markup Language (DSML), among other enhancements. Completion of SPML 2.0 is expected in the second half of 2005.

Provisioning of Web Services

Provisioning systems can have an important function in the management of web services applications, as the industry continues to adopt this form of application development. Individual web services components can have identities associated with them, requiring provisioning of accounts and setting of privileges to resources and other web services. Conversely, provisioning systems will control user access to web services-based applications, requiring that provisioning systems have knowledge of and the capability to manipulate their privilege settings. There is a role for provisioning users with access to web services, but there is also the potential for components of typical provisioning systems to themselves appear as web services and consume request/response (possibly SPML transactions) to provision access to an existing platform or application.

WS-Provisioning

WS-Provisioning describes the APIs and schemas necessary to facilitate interoperability between provisioning systems and to allow software vendors to provide provisioning facilities in a consistent way. The specification addresses many of the problems faced by provisioning vendors in their use of existing protocols, commonly based on directory concepts, and it confronts the challenges involved in provisioning the web services described using Web Services Description Language (WSDL) and XML Schema. WS-Provisioning defines a model for the primary entities and operations common to provisioning systems, including the provisioning and de-provisioning of resources and the retrieval of target data and target schema information. It also provides a mechanism to describe and control the lifecycle of the provisioned state.

A draft version of WS-Provisioning was submitted by IBM/Tivoli to the OASIS PSTC as input to SPML 2.0. Tivoli Federated Identity Manager also supports WS-Provisioning to exchange persistent identity information among federated environments.

Evaluation Criteria

Some key evaluation criteria or considerations that architects need to be aware of include the following:

- **Compliance:** What laws, regulations, and/or policies must the enterprise demonstrate compliance with? Privacy regulations could influence domain choices, connector choices, data flows, and approval requirements.
- **Existing infrastructure:** What identity repositories, workflow systems, credentialing systems, and other IdM components exist in the enterprise (or are planned) that must be accommodated by the provisioning system?
- **Existing applications:** What applications or other systems must be provisioned?
- **Business practices and security policies:** How are approvals for access requests handled within the enterprise? What are the barriers to changing existing practices? What parts of the approval process can be automated?
- **Reconciliation:** How often must data be synchronized between target systems and the provisioning repository?

Statement & Basis for Position

There are four groups of position statements for user provisioning.

- Scope of Provisioning Environment

  Should enterprises have one or more provisioning systems?

  Within a given provisioning system, should there be multiple domains?

- Policy Model

  Should provisioning policy be represented by roles or rules?

  Should provisioning decisions be automated or request-driven?

- Password Management

  Should password management be provided on a self-service basis, through password synchronization, or both?

- Internals of the Provisioning System

  Should a single provisioning server or multiple provisioning servers be used?

  How should provisioning servers connect to managed systems?

  What repositories should the provisioning system use as its authoritative source?

  What provisioning events should be audited?

Scope of Provisioning Environment Position

There are two position statements for determining the boundaries of the provisioning environment.

- Number of Provisioning Systems

  Should enterprises have one or more provisioning systems?

- Provisioning Domains

  Within a given provisioning system, should there be multiple domains?

Number of Provisioning Systems Position

The logic for choosing the provisioning systems position is as follows:

IF business units and applications share common policies and are willing to tightly couple provisioning systems
THENimplement a single provisioning system

OTHERWISEimplement multiple provisioning systems

Alternative Number of Provisioning Systems position statements (important: choose only *one*):

Implement a single provisioning system.

Enterprises should attempt to satisfy provisioning requirements with a single provisioning system to reduce the cost of implementation and the complexity of the operating environment. When common security and business policies are shared across the organization, and the organizational domains are willing to tightly couple provisioning systems, a single instance should accommodate most situations. However, it may still be difficult to accommodate different architecture tiers (e.g., NOS, enterprise, or e-business) with a single system.

Note that this position refers specifically to the use of provisioning for account management for workforce members. It does not apply to the provisioning of network elements, firewalls, or other components that may be provisioned by different types of products.

Moreover, this position sets a strategic objective. Tactically, there may be portions of the IT environment that are not well supported by mainstream enterprise provisioning systems, or portions of the IT environment provisioned through legacy provisioning systems to be eliminated over time. A reasonable "80/20" approach may be to create one provisioning that aggregates most account management into one system but allows "niche systems" for specialized requirements.

Or…

Implement multiple provisioning systems.

When the organization has rigid boundaries between business units that have incompatible policies and processes, the enterprise may be better served by implementing multiple provisioning systems. Enterprises may also have higher-risk operations in which it is more cost effective to partition a more robust environment rather than forcing the entire enterprise to run at the highest security level or forcing the high-risk applications to operate in a reduced security setting.

Deploying provisioning systems across a large number of applications or business units involves making some compromises. If political tensions are high between organizational units, then compromise may not be feasible. Further, multiple provisioning domains may be required to meet scalability or performance requirements in large-scale environments.

When multiple provisioning systems from a single vendor are deployed, enterprises can choose to use proprietary mechanisms for linking the systems. Vendor-specific protocols and interfaces can be used to transmit workflow requests, consolidate audit records, or administer product settings.

With the introduction of SPML, disparate provisioning systems and managed platforms can be linked using an industry standard. Using SPML, requests for provisioning actions can be chained from one provisioning system to another using an XML-based request/response protocol. Note that this will become increasingly viable over time as more provisioning vendors adopt SPML.

In addition to SPML, enterprises can use other technologies to create a hierarchy among provisioning systems or components. For example, changes in the HR system can be synchronized with an enterprise directory. Changes in the enterprise directory could then trigger other actions using a provisioning system.

Provisioning Domains Position

The logic for choosing the provisioning domains position is as follows:

```
IF business units or IT organizations don't require direct, day-to-day control over provisioning

-AND-

IF there is no regulatory requirement for separate regional or other domains
  THENimplement a single provisioning domain

OTHERWISEimplement multiple domains
```

Alternative Provisioning Domains position statements (important: choose only *one*):

Implement a single provisioning domain.

A single provisioning domain offers the simplest environment to implement and operate. Deploying multiple domains in a provisioning environment is complex. It may be difficult to determine which resources, accounts, and policies each domain controls. Increasing the number of high-level administrators with control over provisioning policy may weaken the security of the system. Enterprises should resist the temptation to over-engineer the environment to avoid introducing unnecessary risk and complexity.

Or…

Implement multiple domains.

When business units, site administrators, or other organizations controlling IT resources require direct, day-to-day control over provisioning policy, it may not be possible to deploy a single provisioning system without delegating control to these organizations. Also, regulatory considerations in some countries or lines of business may require provisioning to be handled differently than in others. In such cases, the provisioning system might resolve the political issues or regulatory discrepancies by using multiple domains for policy control, each with the ability to perform tasks such as setting account parameters, specifying workflows, creating roles, and appointing delegated administrators.

Policy Model Position

There are two position statements for the policy model used by provisioning systems.

- Implementing Provisioning Policy Using Roles, Rules, or a Combination

  Should provisioning policy be implemented using roles or rules?

- Automated or Request-Driven Policy

  Should provisioning decisions be automated or request-driven?

Provisioning Policy Position

The logic for choosing the provisioning policy position is as follows:

IF the enterprise has job functions that map well to IT access rights
  THENdefine role structure and use roles in combination with rules to grant access

OTHERWISEimplement rule-based policies

Alternative Provisioning Policy position statements (important: *more than one may apply,* depending on enterprise requirements):

Define role structure and use roles in combination with rules to grant access.

Some businesses have stable job functions that can be represented in roles that map well to IT access rights. Role definitions are powerful administration constructs that should be used when possible to drive automated decisions (such as granting all persons with the role of "auditor" an account to use the audit database) or determine a person's ability to initiate provisioning requests (such as enabling audit managers to assign a person an account for the audit database). With role-based provisioning, accounts or privileges can be changed simply by giving a person a new role or by changing a role's definition.

Role definitions can be hierarchical to reflect the structure of the organization, and roles can be defined to mirror business requirements such as separation of duties. Group definitions or attribute values can also be used to represent roles. However, defining roles can be a difficult project that should be approached carefully. Enterprises with job functions that map well to their IT resources and a relatively static organization will have an easier time creating roles. On the other hand, enterprises with project teams that report to multiple managers and applications are very dynamic, so role definitions are more difficult to develop. For more information on roles, see the Reference Architecture Technical Position, "Roles."

At a minimum, basic roles (like "employee" or "partner") should be used to grant access to resources that everyone is assigned by default. Roles are useful because they help minimize what the provisioning system must evaluate with rules.

Or…

Implement rule-based policies.

Often, however, not all provisioning decisions can be based entirely on roles. For example, the procedure for creating an e-mail account for a bank teller in a Florida branch may be different from creating a bank teller's account in Zurich. Usually, role-based implementations for provisioning must be combined with rules to achieve the necessary granularity.

When job functions do not map well to access rights, enterprises should base provisioning actions on various attributes or group memberships in the provisioning repository or in authoritative directory systems.

Enterprises can choose to use rule-based systems without defining roles or to use the two approaches in combination. Rule systems can also be used as an enterprise transitions into a role-based environment.

Automated or Request-Driven Policy Position

The logic for choosing the automated or request-driven policy position is as follows:

```
IF automated provisioning steps can be accurately mapped

-AND-

IF security policy permits automatic assignment of resources
  THENimplement an automated provisioning policy

OTHERWISEimplement a request-driven policy model
```

Note: The answer may be a combination, but some enterprises might have a policy that leans one way or another; typically, basic access may be granted automatically by default, but deeper provisioning must often be done by request.

Alternative Automated or Request-Driven Policy position statements (important: choose *as many as required* for each type of account or access privilege):

Implement an automated provisioning policy.

Automation is a key factor in improving security, increasing efficiency, and reducing manual intervention in user administration activities. Enterprises should seek to leverage automation when they have the authoritative data to trigger actions and company policy permits automatic granting of access to resources.

When provisioning actions are automated, auditing and reporting functions are critical to ensure that policies are being properly carried out. However, even with automated policies, certain actions can be routed to a resource owner for review before granting access to the resource.

Or…

Implement a request-driven policy model.

When full automation is not permitted by company policy, regulations, or the business rules surrounding a particular IT system, provisioning should be implemented using request-based procedures where users or managers submit requests for resource access. The review of access requests can be automated by checking identity attributes or sending the request to a manager for approval. By initiating provisioning actions through requests, enterprises could also begin a transition to a more fully automated system over time.

Password Management Position

The logic for choosing the password management position is as follows:

```
IF policy permits a common password to be set on IT systems by a provisioning system
  THEN IF users require the ability to manage passwords natively in one or more of their managed systems
      THENimplement password synchronization
```

OTHERWISEimplement centralized password self-service

OTHERWISEdo not implement centralized password management

Alternative Password Management position statements (important: *more than one may apply,* depending on enterprise requirements):

Implement password synchronization.

Password synchronization should be implemented when users require the ability to manage passwords natively in managed systems but desire passwords to be the same on all systems. Provisioning products must overcome password syntax differences between platforms during deployment, and some systems may not support the desired password length or complexity, thus rendering synchronization not feasible.

Often, enterprises support a combination of password self-service and password synchronization on systems for which policy allows automated password management.

Or…

Implement centralized password self-service.

Self-service password management capability is required by most medium to very large organizations to reduce calls to the help desk and increase user productivity. Self-service password management typically involves a centralized system hosting a webpage that presents challenge questions which the user must successfully answer before the new password can be issued and modified on all target systems. Nearly all provisioning products offer this level of password self-service.

Some enterprises have deployed Interactive Voice Response (IVR) systems to reset passwords, which the provisioning product may have to interface with. Provisioning systems may also have to accommodate password or personal identification number (PIN) resets on token or smart card systems.

Or…

Do not implement centralized password management.

There may be systems within an enterprise whose security policy does not allow automated, centralized password management.

Internals of the Provisioning System Position

There are four position statements for provisioning system internals.

- Centralized or Distributed

  Should a single provisioning server or multiple provisioning servers be used?

- Connecting to Managed Systems

  How should provisioning servers connect to managed systems?

- Provisioning Repository

  What repositories should the provisioning system use as its authoritative source?

- Provisioning Version Control and Audit Trails

  What provisioning events should be audited?

Centralized or Distributed Servers Position

The logic for choosing the centralized or distributed servers position is as follows:

```
IF there are provisioning domains whose owners require possession of a server
-OR-
IF there are regional IT centers with many provisioned resources
-OR-
IF high volume indicates need for multiple servers
-OR-
IF separation of duties requirements are best implemented through distributed servers
  THENconsider using multiple servers
OTHERWISEuse a single server
```

Alternative Centralized or Distributed Servers position statements (important: *more than one may apply,* depending on enterprise requirements):

Consider using multiple servers.

Deploying multiple servers is a more complex implementation. Multiple servers might be avoided even in regional cases by using remote agents, or in domain cases by convincing business units that policy domains give them enough control. But if you must have multiple servers, there are products that make this possible, and SPML is increasingly creating the opportunity for more distributed operation.

Or…

Use a single server.

Utilizing a single provisioning server is the simpler case that allows faster propagation and troubleshooting. Enterprises should use this configuration whenever possible.

Note, however, that a provisioning system must always have a business resumption capability. Even a single server may need a warm standby (or other means of assuring availability) for emergencies, and if there are multiple servers for domains or regional centers, each must have its own warm standby or other means of assuring availability.

Connecting to Managed Systems Position

The logic for choosing the connecting to managed systems position is as follows:

```
IF required functions can be accomplished securely and efficiently using remote connectors and functionality is available
  THENimplement remote connectors
OTHERWISEinstall local connectors
```

Alternative Connecting to Managed Systems position statements (important: *more than one may apply,* depending on enterprise requirements):

Implement remote connectors.

Remote connectors (also known as agents) are the desirable configuration for most situations due to ease of deployment and ongoing maintenance. Most target systems expose interfaces (such as LDAP, Secure Sockets Layer [SSL], or others) that the connector can access from a remote system to perform basic account creation and modification functions. Remote connectors are also desirable when the target platform permits secure communications between components.

Enterprises may also be forced to use remote connectors when system owners do not permit local connector installation. Even in cases where local connectors may provide more functionality, political circumstances can force the use of remote connectors.

Or…

Install local connectors.

There are several reasons for implementing connectors locally on target systems. Local connectors are required when the target system does not support an interface that can be accessed over the network or if the interface does not provide a secure channel to exchange sensitive data like passwords. Certain functions, such as group creation, may require a local connector that interfaces with platform-specific API calls that aren't available remotely.

For password synchronization, most platforms require that connectors be installed on the target system in order to capture passwords in plain text format. The captured passwords are then encrypted and forwarded to other systems in the network.

Other situations may call for bi-directional capabilities to synchronize data between the target platform and the provisioning repository. Some products may require a local agent to effectively perform this function.

Provisioning Repository Position

Note: Architects may find similarities in the Reference Architecture Technical Position, "Directory Tiers, Instances, and Roles." To ensure architectural harmony, architects should seek to align directory planning and provisioning system planning, where possible.

The logic for choosing the provisioning repository position is as follows:

```
IF the enterprise consolidates identity data in a directory or database that is authoritative for all users
  THENuse the enterprise identity repository as the authoritative source for provisioning

OTHERWISE IF regulatory or other policy allows aggregation of multiple identity repositories

-AND-

IF a comprehensive aggregation of users can be efficiently created and maintained
  THENphysically aggregate user identity from multiple sources into a comprehensive authoritative source for
  provisioning use

OTHERWISEuse virtual directory services to obtain authoritative identity information for provisioning from
multiple sources
```

Alternative Provisioning Repository position statements (important: *more than one may apply,* depending on enterprise requirements):

Use the enterprise identity repository as the authoritative source for provisioning.

An enterprise identity repository, such as a directory service, contains most or all the user attributes that are required to manage users and their entries on target systems. During the implementation of a provisioning solution, an enterprise can use a pre-existing identity repository, or it can create one by using the provisioning system to consolidate identity data from multiple data sources. The provisioning system then uses the enterprise identity repository as its authoritative source for user account creation, termination, and change.

With all users and attributes in a single directory or database, the enterprise has additional flexibility for report generation and tracking of what users have access to. In most cases where a comprehensive enterprise identity service exists, it should not be necessary to duplicate all the information in an internal database or directory associated with the provisioning system.

Or…

Physically aggregate user identity from multiple sources into a comprehensive authoritative source for provisioning use.

With all users and attributes in a single directory or database, the enterprise has additional flexibility for report generation and tracking of what users have access to. Therefore, it is desirable to create a comprehensive identity repository.

Or…

Use virtual directory services to obtain authoritative identity information for provisioning from multiple sources.

In some cases, country regulations, security policies, or other policies prohibit aggregation of all user identities into a common service. In other cases, the number of identities and sources may be so high, the rate of change so volatile, and/or the data definitions so different across domains that aggregations cannot be efficiently created and maintained.

Virtual directory services built into the provisioning system or used by the provisioning system can aggregate sufficiently complete and authoritative identity records at runtime from data in multiple authoritative systems. Certainly, many enterprises do not have a single authoritative repository of comprehensive user identity information, and a virtual approach allows them to operate their provisioning system over multiple repository systems, if required. Typically, the provisioning system needs to maintain mapping information in a directory or database for all users so that it can accomplish runtime identity information joins and/or filter out duplicate records present in more than one source.

The virtual model is also useful when provisioning activities that require calls to external systems for capturing real-time data. Virtual repositories can also be easier to implement while awaiting the start or completion of the large projects required to consolidate identity data into the single store and implement the consistent processes required to maintain that information.

Provisioning Version Control and Audit Trails Position

The logic for choosing the provisioning version control and audit trails position is as follows:

Use version controls and audit trails as appropriate.

Security and regulatory policies will guide organizations in determining what records must be captured. Important audit data typically includes provisioning requests, completed requests, request approvers, and information on any accounts that have been established.

Enterprises should take additional precautions to protect the auditing subsystem and the data contained within it. Data protection includes the use of encryption or signing of audit records to protect them from unauthorized access or alteration. Auditing systems should also be protected such that systems administrators cannot alter captured data and cover up inappropriate or unauthorized activity.

Because provisioning systems are essentially policy management authorities (PMAs), architects are referred to the "Policy Versioning Control and Audit Trail" position statement in the Reference Architecture Technical Position, "User Authorization."

Relationship to Other Components

Directory integration challenges are similar to those found in provisioning projects. For more information, see the Reference Architecture Technical Position, "Directory Integration."

Enterprises can choose to construct multiple provisioning systems to meet specific requirements. For a similar discussion as it relates to directory architecture, see the Reference Architecture Technical Position, "Directory Tiers, Instances, and Roles."

The Reference Architecture Technical Position, "User Management," discusses additional issues and techniques for managing the lifecycle relationship of users to the enterprise.

The "Policy Versioning Control and Audit Trail" position statement from the Reference Architecture Technical Position, "User Authorization," serves as the basis for auditing in this Technical Position, and provides guidance for managing policies across the enterprise.

The Reference Architecture Technical Position, "Roles," discusses how and to what extent roles should be used within IdM, applications, and other systems.

Revision History

October 2005

- Cleaned up discussion of provisioning vendors.
- Added references to the Reference Architecture Technical Position, “Roles.”
- Added references to federation reports.
- Added a section on WS-Provisioning.
- Updated company and product references.
- Added the “Revision History” section.

May 2004

- This is the first iteration of this Technical Position.

Author Bio

NickNikols

Senior Analyst

Emphasis: Directory services, identity management, meta-directory services, virtual directory services, provisioning
Background: Over 13 years of experience developing innovative products and architectures that provide directory services, identity management, provisioning, and directory/application integration.

Primary Distinctions: Principal inventor of Novell's DirXML technology. Led and managed the engineering team that productized Novell's DirXML meta-directory/provisioning offering (now renamed to Nsure Identity Manager). Initiated Novell's involvement with the XSLT and DSML working groups. Invented Novell's Filtered Replication technology, a distinguishing feature of Novell's eDirectory product offering. Recipient of 2000 Novell President's Award. Member of Novell's eDirectory/NDS engineering team, specializing in the cross platform development of Novell's directory offering. Member of Novell's UnixWare engineering team.