Overview Guide

Identity Manager 4.0.2

June 22, 2012

Novell

Legal Notices

Novell, Inc. makes no representations or warranties with respect to the contents or use of this documentation, and specifically disclaims any express or implied warranties of merchantability or fitness for any particular purpose. Further, Novell, Inc. reserves the right to revise this publication and to make changes to its content, at any time, without obligation to notify any person or entity of such revisions or changes.

Further, Novell, Inc. makes no representations or warranties with respect to any software, and specifically disclaims any express or implied warranties of merchantability or fitness for any particular purpose. Further, Novell, Inc. reserves the right to make changes to any and all parts of Novell software, at any time, without any obligation to notify any person or entity of such changes.

Any products or technical information provided under this Agreement may be subject to U.S. export controls and the trade laws of other countries. You agree to comply with all export control regulations and to obtain any required licenses or classification to export, re-export, or import deliverables. You agree not to export or re-export to entities on the current U.S. export exclusion lists or to any embargoed or terrorist countries as specified in the U.S. export laws. You agree to not use deliverables for prohibited nuclear, missile, or chemical biological weaponry end uses. Please refer to the International Trade Services (http://www.novell.com/company/policies/trade_services) for more information on exporting Novell software. Novell assumes no responsibility for your failure to obtain any necessary export approvals.

Novell, Inc.
1800 South Novell Place
Provo, UT 84606
U.S.A.
www.novell.com

Online Documentation: To access the latest online documentation for this and other Novell products, see the Novell Documentation Web page (http://www.novell.com/documentation).

Novell Trademarks

For Novell trademarks, see the Novell Trademark and Service Mark list (http://www.novell.com/company/legal/trademarks/tmlist.html).

Third-Party Materials

All third-party trademarks are the property of their respective owners.

Contents

About This Guide

This guide introduces you to Novell Identity Manager, a WorkloadIQ product that manages identity and access across physical, virtual, and cloud environments. This guide explains business issues that Identity Manager can help you solve while reducing costs and ensuring compliance. It also contains a technical overview of the Identity Manager components and tools you can use to create your Identity Manager solution. The guide is organized as follows:

Audience

This guide is intended for administrators, consultants, and network engineers who require a high-level introduction to Identity Manager business solutions, technologies, and tools.

Documentation Updates

For the most recent version of this document, see the Identity Manager Documentation Web site (http://www.novell.com/documentation/idm402/index.html).

Additional Documentation

For documentation on Identity Manager drivers, see the Identity Manager Drivers Web site (http://www.novell.com/documentation/idm402drivers/index.html).

1 Identity Manager and Business Process Automation

The information in this section identifies some of the business processes you can automate through the implementation of a Novell Identity Manager system. If you are already aware of the business automation solutions provided by Identity Manager, you might want to skip to the technical introduction provided in Chapter 4, "Identity Manager Architecture," on page 27.

Managing identity needs is a core function of most businesses. For example, imagine that it's early Monday morning. You scroll down the list of requests in your queue:

- Jim Taylor's cell phone number has changed. You need to update it in the HR database and four other independent systems.
- Karen Hansen, just returning from an extended leave of absence, has forgotten her e-mail password. You need to help her retrieve or reset it.
- Jose Altimira just hired a new employee. You need to give the employee network access and an e-mail account.
- Ida McNamee wants access to the Oracle financial database, which requires you to get approval from three different managers.
- John Harris just moved from the Accounts Payable department to the Legal department. You need to give him access to the same resources as the other members of the Legal team and remove his access to Accounts Payable resources.
- Karl Jones, your own boss, saw a copy of Ida McNamee's request for access to the Oracle financial database and is concerned about the number of people with access. You need to generate a report for him that shows everyone who has access to the database.

You take a deep breath and start in on the first request, knowing that you'll be hard-pressed to keep up with all of the requests, let alone have time to finish the other projects assigned to you.

If this sounds like a common workday for you or someone in your organization, Identity Manager can help. In fact, the core Identity Manager capabilities, introduced in the following illustration, can help you automate all of these tasks and more. These capabilities—workflows, roles, attestation, self-service, auditing, and reporting—use multi-system data synchronization driven by your business policies to automate the processes involved in provisioning users and managing passwords, two of the most difficult and time-consuming duties of an IT organization.

Figure 1-1 *Identity Manager Core Capabilities*

The following sections introduce you to these Identity Manager capabilities and how they can help you successfully meet the identity needs of your organization:

1.1 Data Synchronization

If your organization is like most, you have identity data stored in multiple systems. Or, you have identity data stored in one system that you could really use in another system. Either way, you need to be able to easily share and synchronize data between systems.

Identity Manager lets you synchronize, transform, and distribute information across a wide range of applications, databases, operating systems, and directories such as SAP, PeopleSoft, Salesforce, Microsoft SharePoint, Lotus Notes, Microsoft Exchange, Microsoft Active Directory, Novell eDirectory, Linux and UNIX, LDAP directories.

Figure 1-2 *Identity Manager Connecting Multiple Systems*

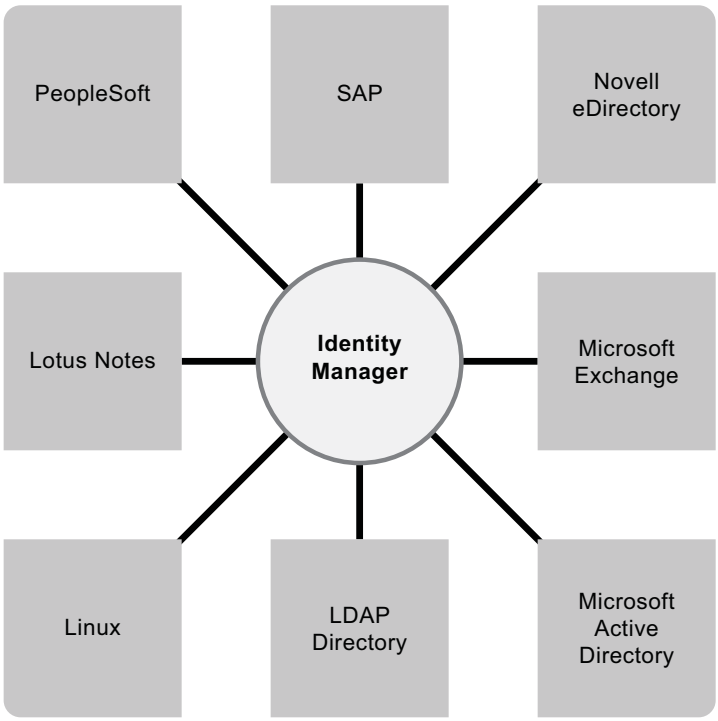

You control the flow of data among the connected systems. Among other things, you determine what data is shared, which system is the authoritative source for a piece of data, and how the data is interpreted and transformed to meet the requirements of other systems.

In the following diagram, the SAP HR database is the authoritative source for a user's telephone number. The Lotus Notes system also uses telephone numbers, so Identity Manager transforms the number into the required format and shares it with the Lotus Notes system. Whenever the telephone number changes in the SAP HR system, it is synchronized to the Lotus Notes system.

Figure 1-3 *Data Synchronization between Connected Systems*

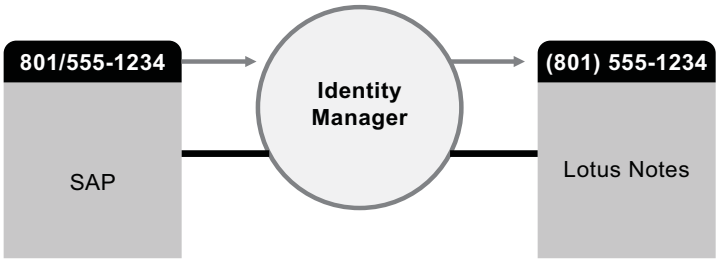

Managing data for existing users is just the beginning of the data synchronization capabilities of Identity Manager. In addition, Identity Manager can create new user accounts and remove existing accounts in directories such as Active Directory, systems such as PeopleSoft and Lotus Notes, and operating systems such as UNIX and Linux. For example, when you add a new employee to your SAP HR system, Identity Manager can automatically create a new user account in Active Directory, a new account in Lotus Notes, and a new account in a Linux NIS account management system.

Figure 1-4 *User Account Creation in Connected Systems*

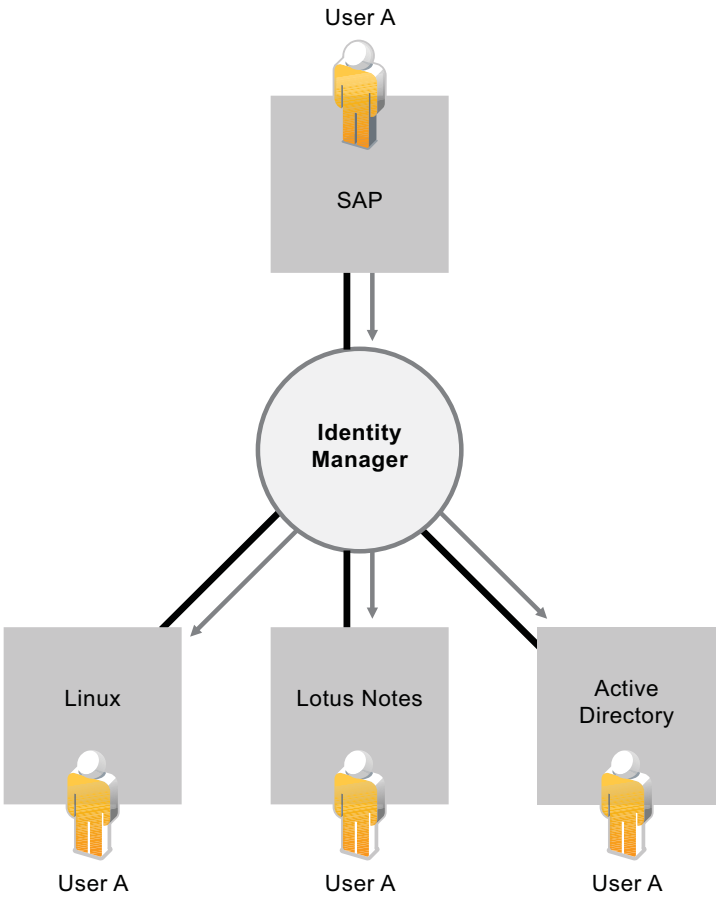

As part of its data synchronization capability, Identity Manager can also help you synchronize passwords between systems. For example, if a user changes his or her password in Active Directory, Identity Manager can synchronize that password to Lotus Notes and Linux.

Figure 1-5 *Password Synchronization among Connected Systems*

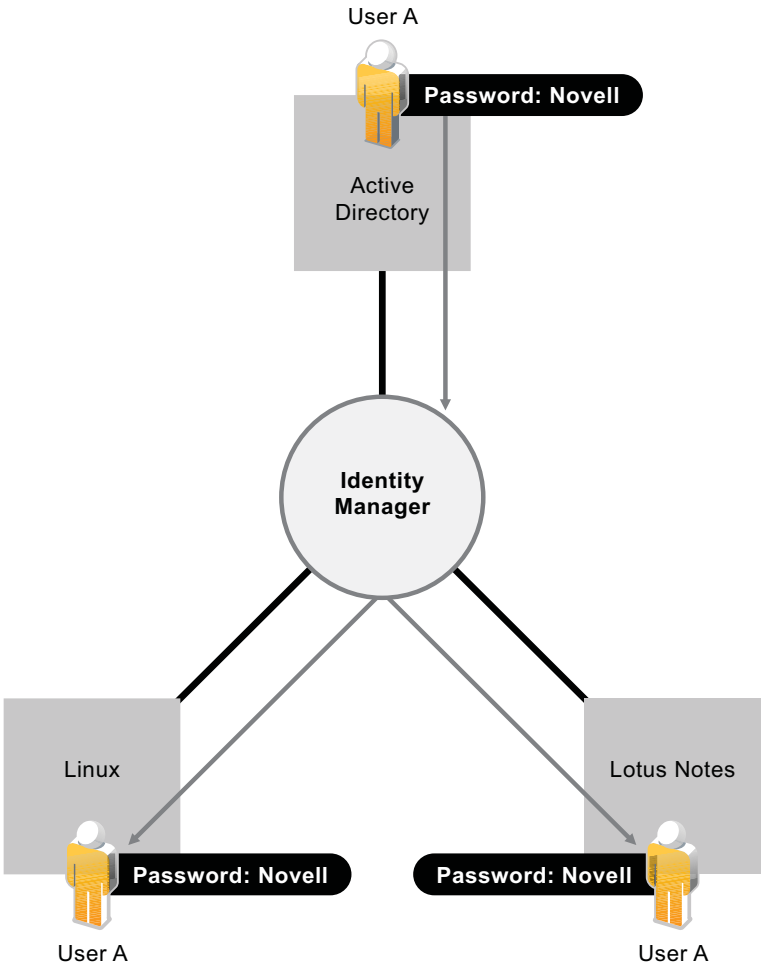

1.2 Workflow

More than likely, user access to many of the resources in your organization doesn't require anyone's approval. However, access to other resources might be restricted and require approval from one or more individuals.

Identity Manager provides workflow capabilities to ensure that your provisioning processes involve the appropriate resource approvers. For example, assume that John, who has already been provisioned with an Active Directory account, needs access to some financial reports through Active Directory. This requires approval from both John's immediate manager and the CFO. Fortunately, you've set up an approval workflow that routes John's request to his manager and, after approval from his manager, to the CFO. Approval by the CFO triggers automatic provisioning of the Active Directory rights needed by John to access and view the financial documents.

Figure 1-6 *Approval Workflow for User Provisioning*

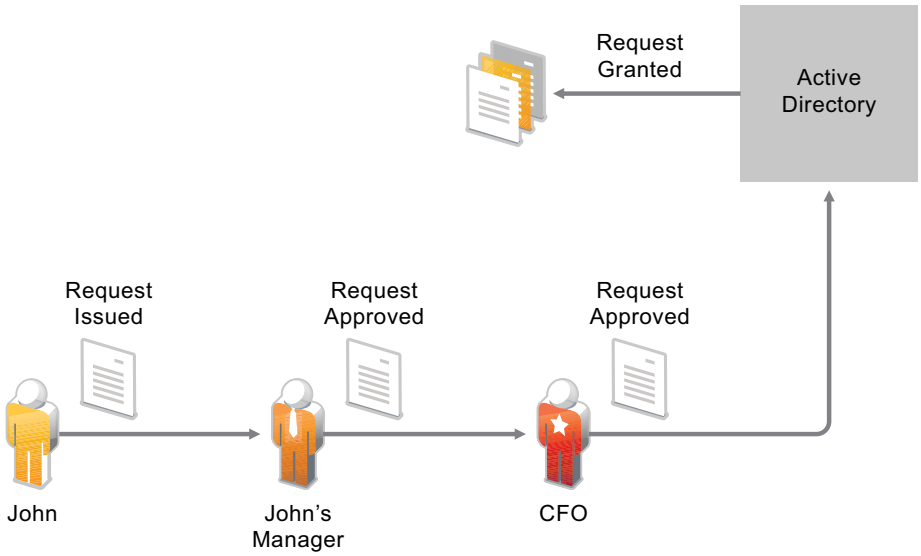

Workflows can be initiated automatically whenever a certain event occurs (for example, a new user is added to your HR system) or initiated manually through a user request. To ensure that approvals take place in a timely manner, you can set up proxy approvers and approval teams.

1.3 Roles and Attestation

Users often require access to resources based upon their roles in the organization. For example, a law firm's attorneys might require access to a different set of resources than the firm's paralegals.

Identity Manager lets you provision users based on their roles in the organization. You define the roles and make the assignments according to your organizational needs. When a user is assigned to a role, Identity Manager provisions the user with access to the resources associated with the role. If a user is assigned multiple roles, he or she receives access to the resources associated with all of the roles, as shown in the following illustration:

Figure 1-7 *Role-Based Provisioning of Resources*

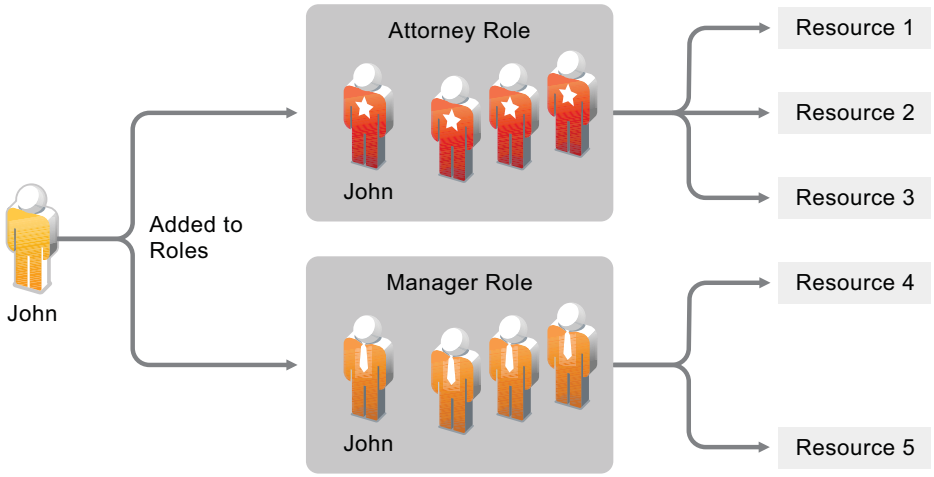

You can have users automatically added to roles as a result of events that occur in your organization (for example, a new user being with the job title of Attorney added to your SAP HR database). If approval is required for a user to be added to a role, you can establish workflows to route role requests to the appropriate approvers. You can also manually assign users to roles.

In some cases, certain roles should not be assigned to the same person because the roles conflict. Identity Manager provides Separation of Duties functionality that lets you prevent users from being assigned to conflicting roles unless someone in your organization makes an exception for the conflict.

Because role assignments determine a user's access to resources within your organization, ensuring correct assignments is critical. Incorrect assignments could jeopardize compliance with both corporate and government regulations. Identity Manager helps you validate the correctness of your role assignments through an attestation process. Using this process, responsible individuals within your organization certify the data associated with roles:

- **User profile attestation:** Selected users attest to their own profile information (first name, last name, title, department, e-mail, and so forth) and correct any incorrect information. Accurate profile information is essential to correct role assignments.
- **Separation of Duties violation attestation:** Responsible individuals review a Separation of Duties violation report and attest to the accuracy of the report. The report lists any exceptions that allow a user to be assigned conflicting roles.
- **Role assignment attestation:** Responsible individuals review a report listing selected roles and the users, groups, and roles assigned to each role. The responsible individuals must then attest to the accuracy of the information.
- **User assignment attestation:** Responsible individuals review a report listing selected users and the roles to which they are assigned. The responsible individuals must then attest to the accuracy of the information.

These attestation reports are designed primarily to help you ensure that role assignments are accurate and that there are valid reasons for allowing exceptions for conflicting roles.

1.4 Self-Service

You probably have business managers and departments clamoring to manage their own users' information and access needs instead of relying on you or your staff. How many times have you heard "Why can't I change my own cell phone number in our corporate directory?" or "I'm in the Marketing department. Why do I have to call the Help Desk to get access to the Marketing Information database?"

With Identity Manager, you can delegate administrative duties to the people who should be responsible for them. For example, you can enable individual users to:

- Manage their own personal data in the corporate directory. Rather than having you change a cell phone number, they can change it in one place and have it changed in all the systems you've synchronized through Identity Manager.
- Change their passwords, set up a hint for forgotten passwords, and set up challenge questions and responses for forgotten passwords. Rather than asking you to reset a password because they've forgotten it, they can do it themselves after receiving a hint or responding to a challenge question.
- Request access to resources such as databases, systems, and directories. Rather than calling you to request access to an application, they can select the application from a list of available resources.

In addition to self-service for individual users, Identity Manager provides self-service administration for functions (management, Help Desk, and so forth) that are responsible for assisting, monitoring, and approving user requests. For example, consider the scenario used in Section 1.2, "Workflow," on page 11 and shown below.

Figure 1-8 *Provisioning Workflow with Self-Service*

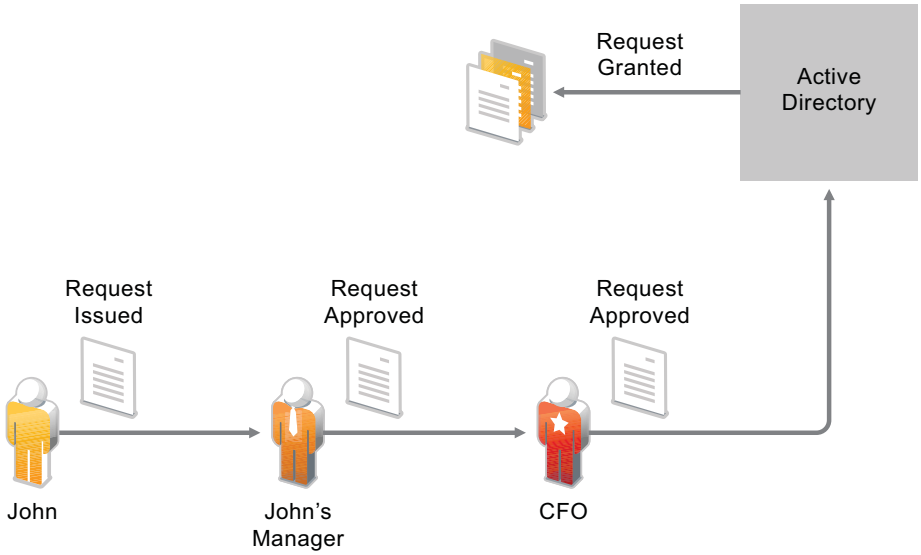

Not only does John use the Identity Manager self-service capability to request access to the documents he needs, but John's manager and the CFO use the self-service capability to approve the request. The established approval workflow allows John to initiate and monitor the progress of his request and allows John's manager and CFO to respond to his request. Approval of the request by John's manager and the CFO triggers the provisioning of the Active Directory rights needed by John to access and view the financial documents.

1.5 Auditing, Reporting, and Compliance

Without Identity Manager, provisioning users can be a tedious, time-consuming, and costly effort. That effort, however, can pale in comparison to verifying that your provisioning activities have complied with your organization's policies, requirements, and regulations. Do the right people have access to the right resources? Are the wrong people shut out of those same resources? Does the employee who started yesterday have access to the network, his e-mail, and the six other systems required for his job? Has the access been canceled for the employee who left last week?

With Identity Manager, you can relax in your knowledge that all of your user provisioning activities, past and present, are being tracked and logged for auditing purposes. Identity Manager contains an intelligent repository of information about the actual state and the desired state of the Identity Vault and the managed systems within your organization. By querying the warehouse, you can retrieve all of the information you need to ensure that your organization is in full compliance with relevant business laws and regulations.

The warehouse gives you a 360-degree view of your business entitlements, providing the knowledge you need to see the past and present state of authorizations and permissions granted to identities in your organization. With this knowledge, you can answer even the most sophisticated Governance Risk and Compliance (GRC) queries.

Identity Manager contains predefined reports that let you perform queries against the Identity Information Warehouse to demonstrate compliance for business, IT, and corporate policies. You can also create custom reports if the predefined reports don't meet your needs.

2 Identity Manager 4.0.2 Features

Novell Identity Manger 4.0.2 provides an intelligent identity framework that leverages your existing IT assets and new computing models like Software as a Service (SaaS) by reducing cost and ensuring compliance across physical, virtual, and cloud environments. With Novell Identity Manager solutions, you can make sure that your business has the most current user identity information. You can retain control at the enterprise level by managing, provisioning, and deprovisioning identities within the firewall and extending to the cloud. Identity Manager can also help you to extend your compliance management to the cloud.

Identity Manager 4.0.2 offers you integrated identity management, roles management, reporting, and package management capabilities for preconfiguring and customizing Identity Manager driver policies. You can also apply security policies across various system domains. Identity Manager allows you to manage the user life cycle in growing regulatory requirements, and applies more granular protection with more strategic user provisioning to satisfy the growing security concerns within the firewall or in the cloud environment. The intelligent identity framework helps you use your existing infrastructure with new computing models like SaaS.

2.1 Identity Manager 4.0.2 New Features

- **New Drivers:**
  - **Bidirectional eDirectory Driver:** The Identity Manager bidirectional eDirectory driver synchronizes data between the Identity Vault and eDirectory. For more information, see *Identity Manager 4.0.2 Driver for Bidirectional eDirectory Implementation Guide*.
  - **Sentinel Identity Tracking Driver Implementation Guide:** The Sentinel Identity Tracking Driver provides integration with Identity Manager and Sentinel to track user account information. Each user account can have multiple account identifiers for each system in the Identity Manager solution. The driver tracks each account identifier and sends that information to Sentinel. Sentinel can run reports to correlate each account identifier with a specific user. For more information, see *Driver for Sentinel Implementation Guide*.
- **Password Management Features:**
  - **Password Policy Enhancements:** Identity Manager now supports three new password policy syntax options:
    - Use Microsoft Complexity Policy
    - Use Microsoft Server 2008 Password Policy
    - Use Novell Syntax

    For more information, see "Supported Password Policy Syntax" in the *Identity Manager 4.0.2 Password Management Guide*.

- **Role Mapping Administrator Features:**
  - **Code Map Synchronization:** Identity Manager 4.0.2 provides facilities for keeping the code map tables synchronized between the Role Mapping Administrator and the Roles Based Provisioning Module. While creating mappings in the Role Mapping Administrator, you can trigger a code map refresh in either the Role Mapping Administrator or the Roles Based Provisioning Module if a mismatch is discovered in the code maps. A code map refresh can run for a long time if executed for all drivers and entitlements. Therefore, the Role Mapping Administrator gives you the ability to trigger a refresh for only the entitlements for which a mismatch was discovered. The Roles Based Provisioning Module also provides new SOAP endpoints for triggering code map refreshes. For more information on the changes to the Role Mapping Administrator, see "Creating Role Resource Mappings" in the *Novell Identity Manager Role Mapping Administrator 4.0.2 User Guide*. For more information on the new SOAP endpoints, see "Resource Web Service" in the *User Application: Administration Guide*.
  - **Resource Creation Settings for the Role Mapping Administrator:** The Role Mapping Administrator now provides several resource creation settings that enable you to prefix the driver name and logical system name when generating resources. The resource creation settings also allow you to select a resource category for automatically generated resources. For more information, see "Customizing the Resource Names" in the *Novell Identity Manager Role Mapping Administrator 4.0.2 User Guide*.
- **Roles Based Provisioning Module Features:**
  - **Approval Processes for Role Revocation:** The Roles Based Provisioning Module now provides support for the use of approval processes with role revocation. The approval process used for role revocation requests, as well as the list of approvers, is the same as for role grant requests. If you have indicated that you want the approval process to execute the standard role assignment approval definition, this process will be used. Alternatively, you can specify a custom approval process for both role grant requests and role revocation requests. For more information, see "Defining the Approval Process for a Role" in the *User Application: User Guide*.
  - **Optimization for Role Delete Operations:** The Roles Based Provisioning Module has optimized the process of deleting roles. When you instruct the User Application to delete a role, it first sets the role status to Pending Delete. The Role and Resource Service driver then notes the change of status and performs these steps:

    1. Removes the resource assignments for the role
    2. Deletes the role itself

    For more information, see "Deleting Roles" in the *User Application: User Guide*.
  - **SAML Support for 64-bit Platforms** SAML support for single sign-on has been added for 64-bit Linux and Windows platforms. For more information, see "Single Sign On (SSO) Configuration" in the *User Application: Administration Guide*.
- **Identity Reporting Module Features:**
  - **Support for EAS on Red Hat Enterprise Linux:** The Event Auditing Service (EAS) now runs on Red Hat Enterprise Linux 5.7 and 6.0 (32-bit and 64-bit).
  - **New Role Hierarchy Report:** The Role Hierarchy Report has been added to the Identity Reporting Module. This report displays the contents of the role hierarchy, as well as the resources associated with each role. For more information, see the *Using Identity Manager 4.0.2 Reports*.
  - **Ability to Purge Historical Data from the Reporting Database:** The Identity Reporting Module now provides the ability to purge historical data from the reporting database. When the reporting module executes a data purge operation, it only purges data from the history tables that is older than the retention value you specify. Any historical data that is

more recent than the retention interval permits will be retained. The purge operation does not remove any of the current state data. For more information, see "Configuring Settings and Data Collection" in the *Identity Reporting Module Guide*.

- **Designer Features:**
  - **REST Activity Support for Workflows** Designer 4.0.2 now includes a new activity in the Provisioning Request Definition editor that enables users to call REST endpoints or resources when processing workflow data. Using the REST activity, workflows can exchange data with REST services both inside and outside of the organization, and users can use data received from a REST service as decision support information on approval forms.

    For more information, see the Rest Activity section in the *User Application: Design Guide*.
  - **Integration Activity Improvements for Workflows** Designer 4.0.2 provides several improvements to the Integration activity in the Provisioning Request Definition editor, including resolving animation issues and reducing the size of deployed PRDs. In addition, the Integration activity now allows users to more easily generate SOAP requests for the activity using the Designer user interface.

    For more information, see the "Adding an Integration Activity" section in the *User Application: Design Guide*.
  - **Performance Improvements in Designer** Designer 4.0.2 provides several performance improvements, including enhanced performance while using the different editors included in the product, improved rendering of configuration pages in the user interface, improved Project Checker speed, and resolved memory issues. For more information, open Designer 4.0.2 and click *Help* > *What's New* in the toolbar.
  - **Designer Optimization for Optional Import of Roles and Resources** Instead of being required to automatically import large numbers of roles and resources from the Identity Vault when they configure a project in Designer, users can now configure Designer to not automatically import the Role Catalog. If a user does not need to import roles or resources, they can select the *Do not import role catalog (excluding system roles)* option in the *Novell* > *Provisioning* > *Import/Deploy* page of the Designer Preferences. Designer will then not automatically import the Role Catalog, saving users time and avoiding the need to manage those roles and resources in Designer.

    For more information, see the "Configuring Roles" section in the *User Application: Design Guide*.
  - **Removal of Unused Packages from the Package Catalog in Designer** If a user has a large number of unnecessary packages imported into the Package Catalog of a project, Designer 4.0.2 provides the option to clean up unused packages from the Catalog, removing any imported packages that are not installed on any driver, driver set, or Identity Vault from the project.

    For more information, see the "Removing Packages from the Package Catalog" section in the *Designer 4.0.2 for Identity Manager 4.0.2 Administration Guide*.
- **Analyzer Features:**
  - **Performance Improvements in Analyzer** Analyzer 4.0.2 offers improved performance by using the batching feature with MySQL database server. Analyzer can now import a million records at one time into the Analyzer database. For more information, see "Database Settings" in the *Analyzer 4.0.2 for Identity Manager Administration Guide*.

2.2 Identity Manager 4.0.1 Features

- **Resource Request Activity:** The Resource Request activity allows you to automate the granting or revoking of resources to users. For example, you might write a provisioning request definition that provisions all of the resources a new employee needs on his or her first day. Using the resource request activity, you can automate the approval of that employee for specified resources. For more details on resource request activity, see "*Resource Request Activity*" in the *User Application: Design Guide*.
- **Telemetry:** Identity Manager Telemetry is a new job introduced with Identity Manager 4.0.1. The job functions as a usage counting tool or a license monitoring tool that provides value to the Identity Manager customers, because they can add more licenses or retire unused licenses. The customers can also leverage benefits such as inactive user pricing.

  The Telemetry job collects details about the installed Identity Manager software and hardware, and usage of Identity Manager drivers in the customer environment. After the customer registers with the Novell Customer Center, the information is sent to Novell. This information allows Novell to support the customer better, develop and test Identity Manager more efficiently and effectively, and make important decisions in the future. For more information, see the *Identity Manager 4.0.2 Jobs Guide*.
- **Reports:** The following reports have been added to the Identity Reporting module:
  - **User Status Change within the Identity Vault:** Displays significant events for the Identity Vault users.
  - **User Password Change within the Identity Vault:** Displays all user password changes within the Identity Vault.
  - **Access Requests by Recipient:** Displays resource assignment workflow processes grouped by recipients.
  - **Access Requests by Requester:** Displays resource assignment workflow processes grouped by requesters.
  - **Access Requests by Resource:** Displays resource assignment workflow processes grouped by resources.

2.3 Identity Manager 4.0 Features

In addition to the newly added features listed earlier in this section, Identity Manager 4.0.1 also includes the following features that were introduced in Identity Manager 4.0.

- **Comprehensive Out-of-the-Box Reporting:** The integrated reporting module of the Novell Identity Manager 4.x product suite enhances visibility into compliance across in-house and cloud deployments. The reporting features let you see a user's identity state and access rights, or report on a user's actions and provisioning history. For more information, see the *Identity Reporting Module Guide*.
- **Enhanced Integration:** For creating a new Identity Manager solution where all of the components reside on the same server, Novell Identity Manager 4.x includes an integrated installer that simplifies the installation process and lets you set up your system more quickly. Instead of installing each Identity Manager component separately, you use the integrated installer to install all of the components in one operation. For more information, see the *Identity Manager 4.0.2 Integrated Installation Guide*.

- **Package Management:** Identity Manager 4.x includes a new concept called Package Management. This is a system for creating, distributing, and consuming high-quality building blocks of Identity Manager policy content. For more information about Identity Manager packages, see *Configuring Packages* in the *Designer 4.0.2 for Identity Manager 4.0.2 Administration Guide*.
- **Cloud-Ready Drivers:** Identity Manager 4.x offers several drivers for out-of-the-box integration with SaaS. The drivers give you seamless integration with SaaS and the hosted solution by providing capabilities such as provisioning, deprovisioning, request/approval processes, password changes, identity profile updates, and reporting. New SharePoint and Salesforce.com drivers help your enterprise identities to integrate with cloud applications. For more information about cloud-ready drivers, see the *Identity Manager 4.0.2 Driver for Salesforce.com Implementation Guide* and the *Identity Manager 4.0.2 Driver for SharePoint Implementation Guide*.
- **Embedded Identity Vault:** The architecture of Novell Identity Manager 4.x products includes an optional built-in Identity Vault so you don't need to create and manage a separate directory structure for identity purposes. In addition, the Novell Identity Manager 4.x family of products includes drivers to easily integrate the Identity Vault with other repositories of identity information in your enterprise, such as Active Directory or various databases. For more information, see the *Identity Manager 4.0.2 Integrated Installation Guide*.
- **Simplified Identity and Roles Management:** The Novell Identity Manager 4.x family of products simplifies the integration of different role repositories into one consolidated location, which means you don't need to manage separate sources of identity information. By using the Role Mapping Administrator with its new intuitive interface, you can even map third-party roles and profiles to Novell Identity Manager 4.x. For more information, see the *Novell Identity Manager Role Mapping Administrator 4.0.2 User Guide*.
- **Enhanced Tools:** Designer is an important tool that includes the business and technical information to create an Identity Manager solution that fits your needs. Several enhancements have been made to Designer 4.x. See the list of Designer enhancements at *What's New* (http://www.novell.com/documentation/designer401/resources/whatsnew/index.html). For information about Designer features and administration, see the *Designer 4.0.2 for Identity Manager 4.0.2 Administration Guide*. In addition, Identity Manager contains a tool to help you simplify the process of analyzing and cleaning your data. For more information, see the *Analyzer 4.0.2 for Identity Manager Administration Guide*.

3 Identity Manager Family

In order to meet different customer needs, the Identity Manager family is divided into three different product groups:

- Identity Manager Advanced Edition
- Identity Manager Standard Edition
- Compliance Management Platform

Identity Manager features that are available in the Identity Manager Standard Edition are also included in the Identity Manager Advanced Edition, along with additional features. Identity Manager Advanced and Standard Edition features are also included in the Compliance Management Platform, along with additional tools.

Figure 3-1 *Identity Manager Product Groups*

For a comparison of the Identity Manager features available in the Advanced and Standard Editions, see Identity Manager Version Comparison.

3.1 Identity Manager Advanced Edition

The Identity Manager Advanced Edition includes the complete set of features available with the product and is mainly targeted for enterprise class user provisioning. It includes the identity self-service features of Standard Edition as well as the full range of workflow-based provisioning features. The Advanced Edition gives you the ability to initiate workflow approval processes, provision roles and resources, and take advantage of the compliance functions. The Advanced Edition also includes the Work Dashboard.

The Identity Manager Advanced Edition is available as a separate ISO.

NOTE: A 90-day evaluation package is available for the Identity Manager Advanced Edition.

3.2 Identity Manager Standard Edition

To meet varying customer requirements, Novell has introduced the Identity Manager Standard Edition. The Standard Edition includes a subset of the features available in the Identity Manager Advanced Edition.

The Standard Edition continues to provide all the features that were present in the previous versions of Identity Manager:

- Identity synchronization
- Rule-based automated provisioning
- Password management and password self-service
- Identity self-service with existing white pages and organization charting functionality

  NOTE: Integration modules continue to remain the same for both Identity Manager Advanced and Standard Editions.

In addition to the preceding list, the Standard Edition also includes the following features that are provided in the Advanced Edition:

- User interface look and feel
- Reporting module
- Content packaging framework
- Support for REST APIs and single sign-on (SSO)
- Analyzer tool for reconciliation

Identity Manager Standard Edition is available in a separate downloadable ISO. To upgrade from the Standard Edition to the Advanced Edition, use the Identity Manager Advanced Edition ISO. You need to apply the correct activation to be able to upgrade to the Advanced Edition. For more information on upgrading from the Standard Edition to the Advanced Edition, see the *Identity Manager 4.0.2 Upgrade and Migration Guide*.

You cannot use an Identity Manager Standard Edition ISO to switch from an existing Identity Manager Advanced Edition. To switch from the Identity Manager Advanced Edition to the Standard Edition, uninstall the Advanced Edition from your server and then install the Standard Edition ISO from the Identity Manager media.

The following functionality is not available in the Identity Manager Standard Edition:

- The Role Mapping Administrator (RMA) is not available.
- The following limitations apply to the User Application:
  - **Identity Self-Service tab is the only tab available to business users:** In the Standard Edition, if you log in to the User Application as a business user, the *Identity Self-Service* tab is the only tab you see. If you log in as a User Application Administrator, you also see the *Administration* tab.
  - **Roles and resources are not supported:** The use of roles and resources requires the Advanced Edition. The *Roles and Resources* tab is not available in Standard Edition.
  - **Compliance tab is not supported:** *The Compliance* tab requires the Identity Manager 4.0.2 Advanced Edition. The *Compliance* tab is not available in Standard Edition.
  - **Work Dashboard is not available:** The *Work Dashboard* tab is not available in the Standard Edition.
  - **Custom roles are not supported:** The ability to define custom roles is not available. The Standard Edition supports only system roles.
  - **Workflows are not supported:** The ability to initiate approval workflows is not supported.
  - **REST APIs:** The REST APIs related to roles, resources, and workflows are not licensed for use with Identity Manager Standard Edition. The Password Self-Service REST APIs are licensed to be used with the Standard Edition.
  - **Security model is simplified:** The Standard Edition offers the security model at a granular level to avoid the unintentional usage of the features provided in the Advanced Edition. You need to assign only the following administrator roles:
    - **User Application Administrator:** A User Application Administrator is authorized to perform all management functions related to the Identity Manager User Application. This includes accessing the *Administration* tab of the Identity Manager user interface to perform any administration actions that it supports.
    - **Report Administrator:** This user has full range of capabilities within the Reporting domain. The Reporting Administrator can perform all actions for all objects within the Reporting domain.
    - **Security Administrator:** This role provides members the full range of capabilities within the security domain. The Security Administrator can perform all possible actions for all objects within the security domain. This role can delegate and grant user access to all Identity Manager Advanced Edition features; therefore, it is separated from User Application administration and report administration roles.

NOTE: For testing purpose, Novell does not lock down the security model in the Standard Edition. Therefore, the Security Administrator can assign all Domain Administrators, delegated administrators, and also other Security Administrators. However, using these advanced features is not supported in production, as indicated in the End User License Agreement. In production environments, all administrator assignments are restricted by licensing. Novell can collect monitoring data in the audit database to ensure that production environments comply. Also, Novell recommends that only one user be given the permission as the Security Administrator.

For more information on User Application features, see the *User Application: Administration Guide*.

- The following limitations apply to the Identity Reporting Module:
  - **Managed System Gateway Driver is disabled:** The Managed System Gateway Driver can pull information from any managed system that has been enabled for data collection in Identity Manager 4.0.2, as long as it supports entitlements.

    The Managed System Gateway Driver is disabled in Identity Manager Standard Edition.
  - **Reports show Identity Vault data only:** The reports generated with the Identity Manager Standard Edition show Identity Vault data only, and do not show data about managed (connected) systems.
  - **Reports do not show historical data:** The Standard Edition does not provide the ability to collect historical state data for reporting. With the Standard Edition, you can only see current state data.
  - **Some reports are not available:** Several new reports have been added in Identity Manager 4.0 and 4.0.1 versions. The Identity Manager 4.0.2 Standard Edition does not include reports that are applicable to connected system and historical data.
  - **Some reports contain no data:** Some of the reports are meaningful only if you have purchased the Identity Manager Advanced Edition because these reports use data that is not available in the Standard Edition, such as roles, resources, and workflow processes.

3.3 Compliance Management Platform

The Novell Compliance Management Platform combines Novell identity, access and security management products with a set of proven tools that simplify the implementation and management of the solution. The platform integrates identity and access information with security information and event management technology to provide a real-time, holistic view of all network events across an enterprise. This tight integration delivers powerful risk management capabilities to ensure that business policy becomes automated IT practice. For more information, see the Compliance Management Platform Website.

3.4 Activating the Identity Manager Standard Edition and Advanced Edition

The Identity Manager Advanced Edition and Standard Edition must be activated within 90 days of installation, or they will shut down. The Identity Manager Advanced Edition and Standard Edition ISOs will work completely for 90 days. At any time during the 90 days, or afterward, you can choose to activate Identity Manager products. For more information, see "Activating Novell Identity Manager Products" in the *Identity Manager 4.0.2 Framework Installation Guide*.

If you apply a Standard Edition activation to an existing non-activated Advanced Edition system, the Metadirectory server and the drivers stop functioning.

NOTE: If you have both the Identity Manager Advanced Edition and the Identity Manager Standard Edition, ensure that you use the right activation on the right server.

---

4 Identity Manager Architecture

The following diagram shows the high-level architecture components that provide the Novell Identity Manager capabilities introduced in Chapter 1, "Identity Manager and Business Process Automation," on page 7: data synchronization, workflow, roles, attestation, self-service, and auditing/reporting.

Figure 4-1 *Identity Manager High-Level Architecture*

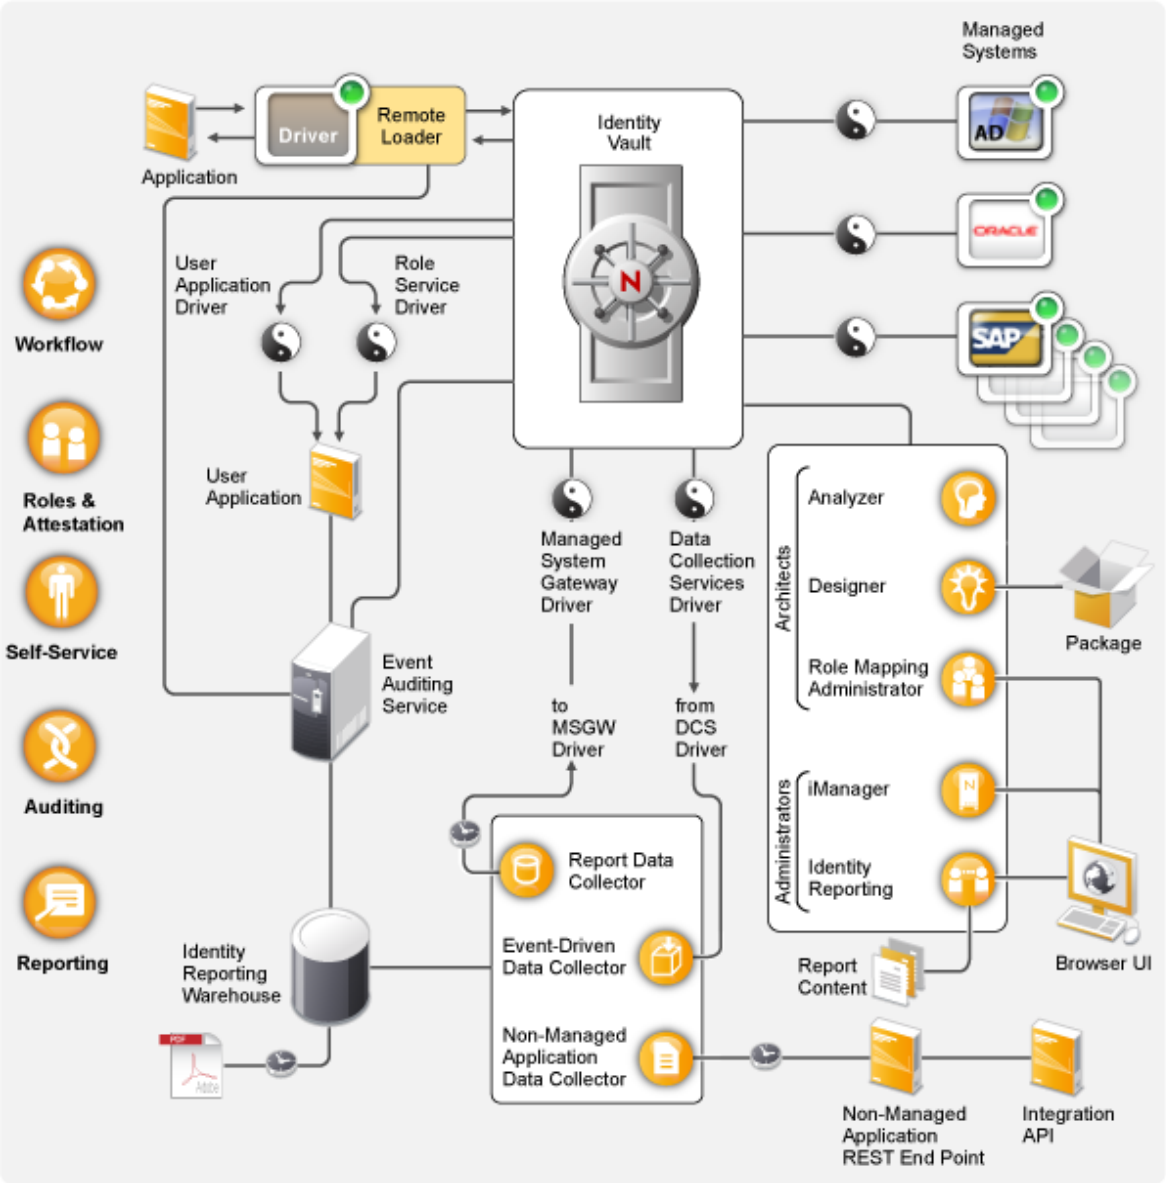

Each of the components is introduced in the following sections:

4.1 Data Synchronization

Data synchronization provides the foundation for automating business processes. In its simplest form, data synchronization is the movement of data from the location where a data item is changed to other locations where the data item is needed. For example, if an employee’s phone number is changed in a company’s Human Resources system, the change would automatically appear in all other systems that store the employee’s phone number.

Identity Manager is not limited to the synchronization of identity data. Identity Manager can synchronize any type of data stored in the connected application or in the Identity Vault.

Data synchronization, including password synchronization, is provided by the five base components of the Identity Manager solution: the Identity Vault, Identity Manager engine, drivers, Remote Loader, and connected applications. These components are shown in the following diagram.

Figure 4-2 *Identity Manager Architecture Components*

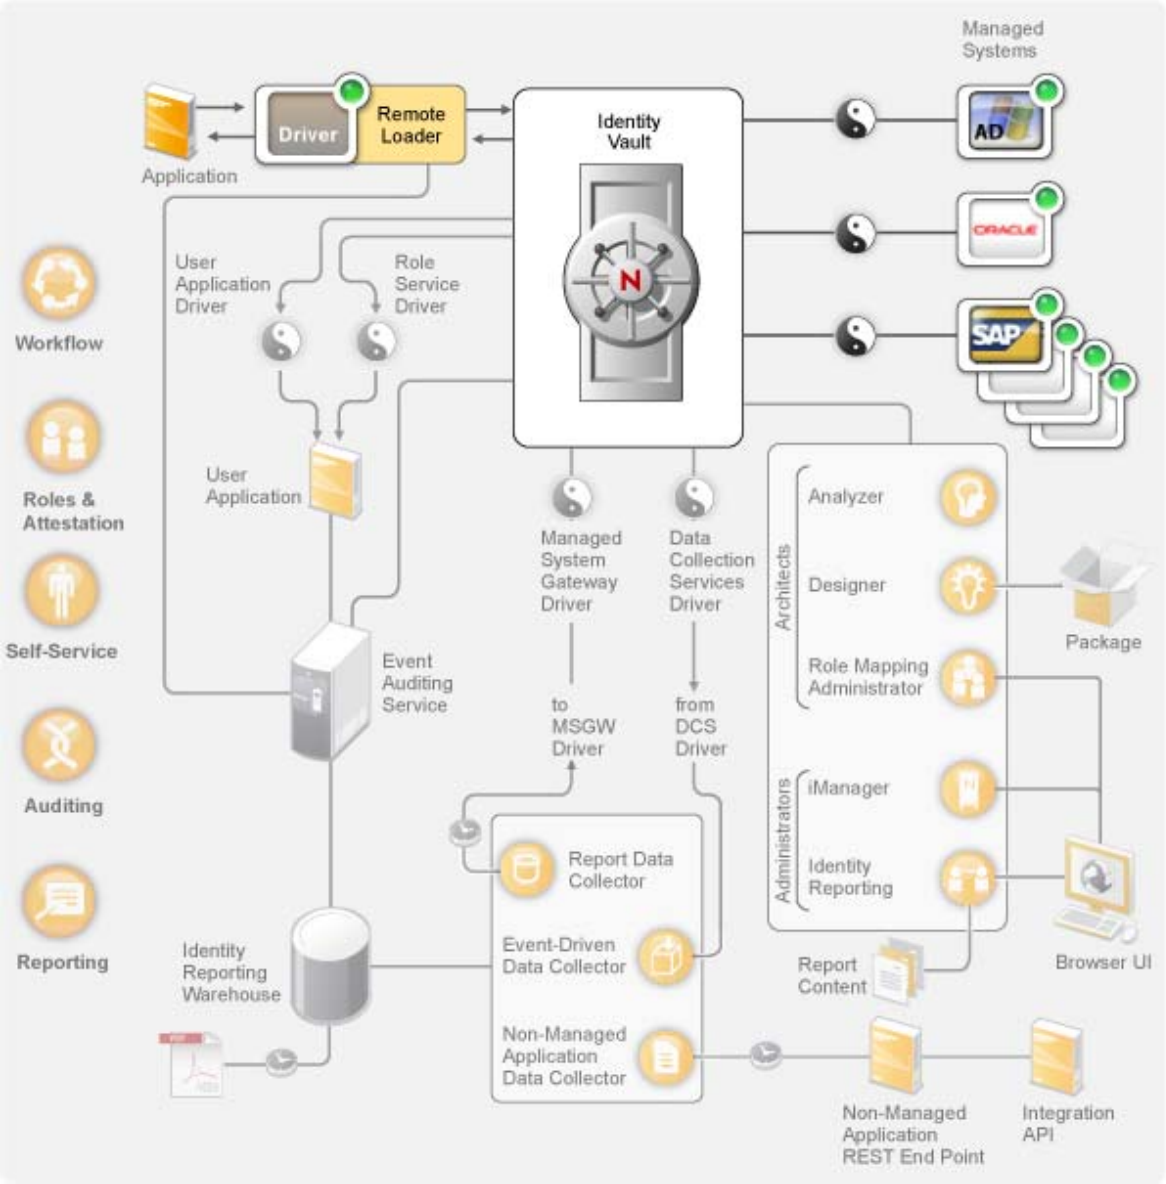

The following sections provide descriptions of each of these components and explain the concepts you should understand to effectively synchronize data among systems in your organization:

- Section 4.1.1, "Components," on page 29
- Section 4.1.2, "Key Concepts," on page 30

4.1.1 Components

Identity Vault: The Identity Vault serves as a metadirectory of the data you want synchronized between applications. For example, data synchronized from a PeopleSoft system to Lotus Notes is first added to the Identity Vault and then sent to the Lotus Notes system. In addition, the Identity Vault stores information specific to Identity Manager, such as driver configurations, parameters, and policies. Novell eDirectory is used for the Identity Vault.

Identity Manager Engine: When data changes in the Identity Vault or a connected application, the Identity Manager engine processes the changes. For events that occur in the Identity Vault, the engine processes the changes and issues commands to the application via the driver. For events that occur in the application, the engine receives the changes from the driver, processes the changes, and issues commands to the Identity Vault. The Identity Manager engine is also referred to as the Metadirectory engine.

Driver: Drivers connect to the applications whose identity information you want to manage. A driver has two basic responsibilities: reporting data changes (events) in the application to the Identity Manager engine, and carrying out data changes (commands) submitted by the Identity Manager engine to the application.

Remote Loader: Drivers must be installed and run on the same server as the application to which they are connecting. If the application is located on the same server as the Identity Manager engine, all you need to do is install the driver to that server. However, if the application is not located on the same server as the Identity Manager engine (in other words, it is remote to the engine's server rather than local), you must install the driver and the Remote Loader to the application's server. The Remote Loader loads the driver and communicates with the Identity Manager engine on behalf of the driver.

Application: A system, directory, database, or operating system that a driver connects to. The application must provide APIs that a driver can use to determine application data changes and effect application data changes. Applications are frequently referred to as *connected systems*.

4.1.2 Key Concepts

Channels: Data flows between the Identity Vault and a connected system along two separate *channels*. The *Subscriber channel* provides data flow from the Identity Vault to a connected system; in other words, the connected system subscribes to data from the Identity Vault. The *Publisher channel* provides data flow from a connected system to the Identity Vault; in other words, the connected system publishes data to the Identity Vault.

Data Representation: Data flows through a channel as *XML documents*. An XML document is created when a change occurs in the Identity Vault or the connected system. The XML document is passed to the Identity Manager engine, which processes the document through the set of filters and policies associated with the driver's channel. When all processing has been applied to the XML document, the Identity Manager engine uses the document to initiate the appropriate changes to the Identity Vault (Publisher channel), or the driver uses the document to initiate the appropriate changes in the connected system (Subscriber channel).

Data Manipulation: As XML documents flow through a driver channel, the document data is affected by the *policies* associated with the channel.

Policies are used for many things, including changing data formats, mapping attributes between the Identity Vault and the connected system, conditionally blocking the flow of data, generating e-mail notifications, and modifying the type of data change.

Data Flow Control: *Filters*, or *filter policies*, control the flow of data. Filters specify which items of data are synchronized between the Identity Vault and a connected system. For example, user data is typically synchronized between systems. Therefore, the user data is listed in the filter for most connected systems. However, printers are generally not of interest to most applications, so printer data does not appear in the filter for most connected systems.

Each relationship between the Identity Vault and a connected system has two filters: a filter on the Subscriber channel that controls data flow from the Identity Vault to the connected system, and a filter on the Publisher channel that controls data flow from the connected system to the Identity Vault.

Authoritative Sources: Most items of data associated with identity have a conceptual owner. The owner of a data item is considered the *authoritative source* for the item. In general, only the authoritative source for a data item is allowed to make changes to the data item.

For example, the corporate e-mail system is generally considered the authoritative source for an employee's e-mail address. If an administrator of the corporate white pages directory changes an employee's e-mail address in that system, the change has no effect on whether the employee actually receives e-mail at the changed address because the change must be made to the e-mail system to be effective.

Identity Manager uses filters to specify authoritative sources for an item. For example, if the filter for the relationship between the PBX system and the Identity Vault allows an employee's telephone number to flow from the PBX system into the Identity Vault but not from the Identity Vault to the PBX system, then the PBX system is the authoritative source for the telephone number. If all other connected system relationships allow the telephone number to flow from the Identity Vault to the connected systems, but not vice versa, the net effect is that the PBX system is the only authoritative source for employee telephone numbers in the enterprise.

Automated Provisioning: Automated provisioning refers to Identity Manager's ability to generate user provisioning actions other than the simple synchronization of data items.

For example, in a typical Identity Manager system where the Human Resource database is the authoritative source for most employee data, the addition of an employee to the HR database triggers the automatic creation of a corresponding account in the Identity Vault. The creation of the Identity Vault account in turn triggers the automatic creation of an e-mail account for the employee in the e-mail system. Data used to provision the e-mail system account is obtained from the Identity Vault and might include employee name, location, telephone number, and so forth.

The automatic provisioning of accounts, access, and data can be controlled in various ways, including:

- **Data item values:** For example, the automatic creation of an account in the access databases for various buildings can be controlled by a value in an employee's location attribute.
- **Approval workflows:** For example, the creation of an employee in the finance department can trigger an automatic e-mail to the finance department head requesting approval for a new employee account in the finance system. The finance department head is directed by the e-mail to a Web page where the department head approves or rejects the request. Approval can then trigger the automated creation of an account for the employee in the finance system.
- **Role assignments:** For example, an employee is given the role of Accountant. Identity Manager provisions the employee with all accounts, access, and data assigned to the Accountant role, either through system workflows (no human intervention), human approval flows, or a combination of both.

Entitlements: An entitlement represents a resource in a connected system, such as an account or a group membership. When a user meets the criteria established for an entitlement in a connected system, Identity Manager processes an event for the user that results in the user being granted access to the resource. This, of course, requires that all of the policies be in place to enable access to the resource. For example, if a user meets the criteria for an Exchange account in Active Directory, the Identity Manager engine processes the user through the set of Active Directory driver policies that provide an Exchange account.

The key benefit of entitlements is that you can define the business logic for access to a resource in one entitlement rather than multiple driver policies. For example, you can define an Account entitlement that gives a user an account in four connected systems. The decision to provide the user with an account is determined by the entitlement, which means that policies for each of the four drivers do

not need to include the business logic. Instead, the policies only need to provide the mechanism for granting the account. If you need to make a business logic change, you change it in the entitlement instead of in each driver.

Jobs: For the most part, Identity Manager acts in response to data changes or user requests. For example, when a piece of data changes in one system, Identity Manager changes the corresponding data in another system. Or, when a user requests access to a system, Identity Manager initiates the appropriate processes (workflows, resource provisioning, and so forth) to provide the access.

Jobs enable Identity Manager to perform actions not initiated by data changes or user requests. A job consists of configuration data stored in the Identity Vault and a corresponding piece of implementation code. Identity Manager includes predefined jobs that perform such actions as starting or stopping drivers, sending e-mail notifications of expiring passwords, and checking the health status of drivers. You can also implement custom jobs to perform other actions; a custom job requires you (or a developer/consultant) to create the code required to perform the desired actions.

Work Orders: Typically, changes to data in the Identity Vault or a connected application are immediately processed. Work orders enable you to schedule tasks to be performed on a specific date and time. For example, a new employee is hired but is not scheduled to start for a month. The employee needs to be added to the HR database, but should not be granted access to any corporate resources (e-mail, servers, and so forth) until the start date. Without a work order, the user would be granted access immediately. With work orders implemented, a work order is created that initiates account provisioning only on the start date.

4.2 Workflow, Roles, Attestation, and Self-Service

Identity Manager provides a specialized application, the User Application, that provides approval workflows, role assignments, attestation, and identity self-service.

The standard User Application is included with Identity Manager. The standard version provides password self-service to help users remember or reset forgotten passwords, organization charts to manage user directory information, user management functionality that enables creation of users in the Identity Vault, and basic identity self-service such as management of user profile information.

The User Application Roles Based Provisioning Module is a part of Identity Manager 4.0.2 Advanced Edition. A standard User Application with advanced self-service, approval workflow, roles-based provisioning, Separation of Duties constraints, and attestation capabilities is included. The Identity Manager 4.0.2 Advanced Edition contains both the standard and the roles based provisioning module capabilities.

Figure 4-3 *Identity Manager User Application*

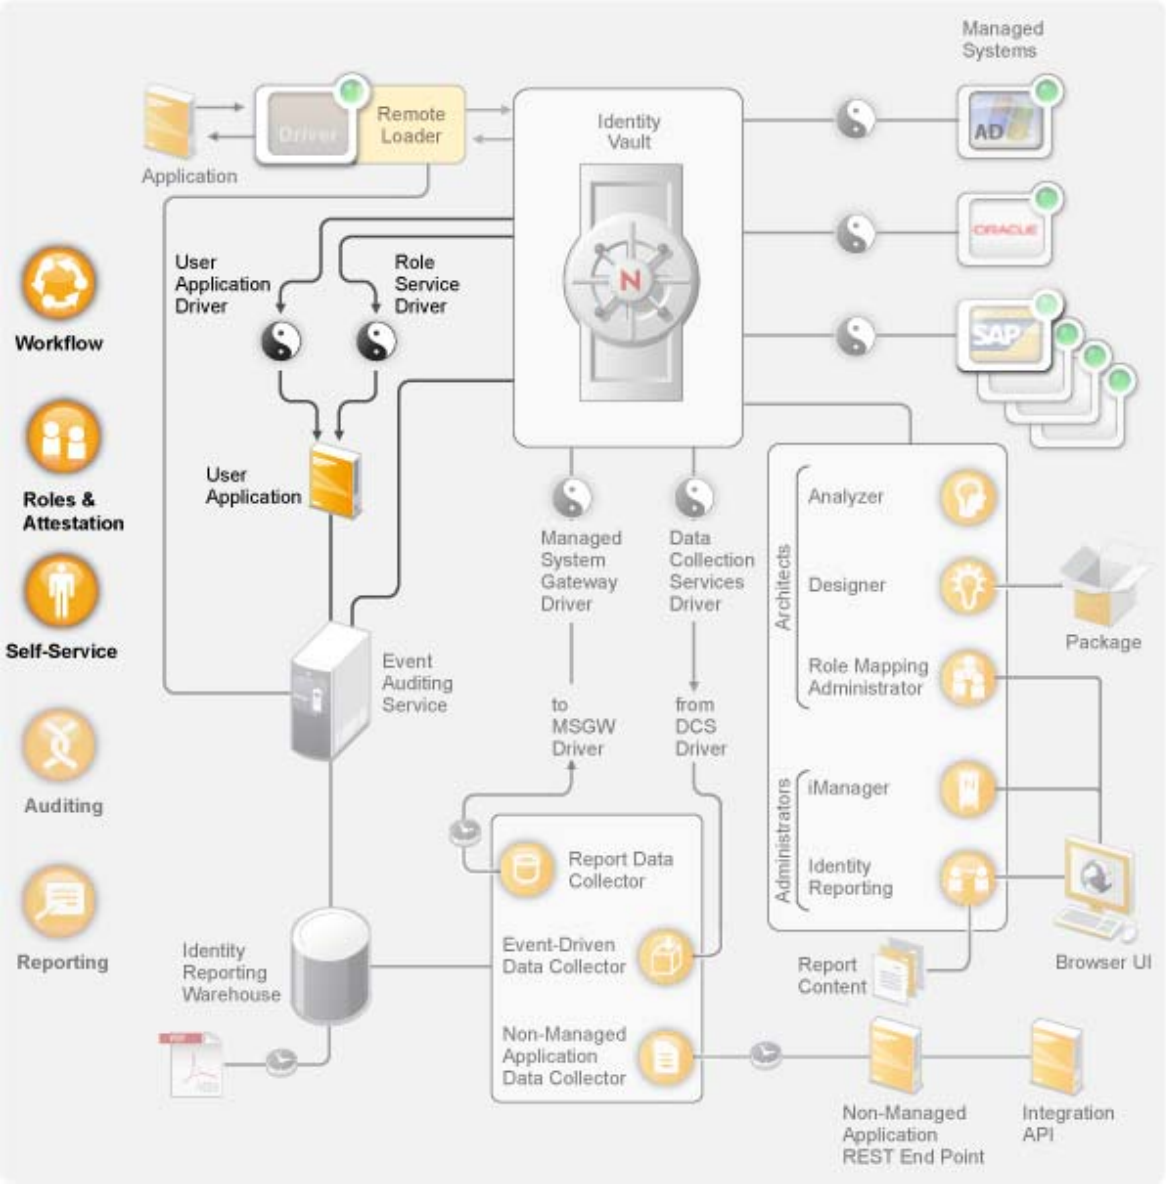

The following sections provide descriptions of each of these components and explain the concepts you should understand to effectively implement and manage the components:

- Section 4.2.1, “Components,” on page 33
- Section 4.2.2, “Key Concepts,” on page 34

4.2.1 Components

User Application: The User Application is a browser-based Web application that gives users and business administrators the ability to perform a variety of identity self-service and roles provisioning tasks, including managing passwords and identity data, initiating and monitoring provisioning and role assignment requests, managing the approval process for provisioning requests, and verifying attestation reports. It includes the workflow engine that controls the routing of requests through the appropriate approval process.

User Application Driver: The User Application driver stores configuration information and notifies the User Application whenever changes occur in the Identity Vault. It can also be configured to allow events in the Identity Vault to trigger workflows and to report success or failure of a workflow's provisioning activity to the User Application so that users can view the final status of their requests.

Role and Resource Service Driver: The Role and Resource Service driver manages all role and resource assignments, starts workflows for role and resource assignment requests that require approval, and maintains indirect role assignments according to group and container memberships. The driver also grants and revokes entitlements for users based on their role memberships, and performs cleanup procedures for requests that have been completed.

4.2.2 Key Concepts

Workflow-based Provisioning: Workflow-based provisioning provides a way for users to request access to resources. A provisioning request is routed through a predefined workflow that might include approval from one or more individuals. If all approvals are granted, the user receives access to the resource. Provisioning requests can also be initiated indirectly in response to events occurring in the Identity Vault. For example, adding a user to a group might initiate a request to have the user granted access to a specific resource.

Roles Based Provisioning: Roles based provisioning provides a way for users to receive access to specific resources based upon the roles assigned to them. Users can be assigned one or more roles. If a role assignment requires approval, the assignment request starts a workflow.

Separation of Duties: To prevent users from being assigned to conflicting roles, the User Application Roles Based Provisioning Module provides a Separation of Duties feature. You can establish Separation of Duties constraints that define which roles are considered to be in conflict. When roles conflict, Separation of Duties approvers can approve or deny any exceptions to the constraints. Approved exceptions are recorded as Separation of Duties violations and can be reviewed through the attestation process described below.

Roles Management: Management of roles must be done by individuals assigned to the Roles Module Administrator and Roles Manager system roles.

The Roles Module Administrator creates new roles, modifies existing roles, and removes roles; modifies relationships between roles; grants or revokes role assignments for users; and creates, modifies, and removes Separation of Duties constraints.

The Roles Manager can do the same things as the Roles Module Administrator with the exception of managing Separation of Duties constraints, configuring the Roles system, and running all reports. The Roles Module Administrator has unlimited scope within the Roles system, but the Roles Manager scope is limited to specifically-designated users, groups, and roles.

Attestation: Role assignments determine a user's access to resources within your organization, and incorrect assignments could jeopardize compliance with both corporate and government regulations. Identity Manager helps you validate the correctness of role assignments through an attestation process. Using this process, individual users can validate their own profile information and Roles Managers can validate role assignments and Separation of Duties violations.

4.3 Auditing and Reporting

Auditing and reporting is provided by the Identity Reporting Module, a new feature for Identity Manager 4.0.2, as shown in the following diagram:

Figure 4-4 *Identity Manager Auditing and Reporting*

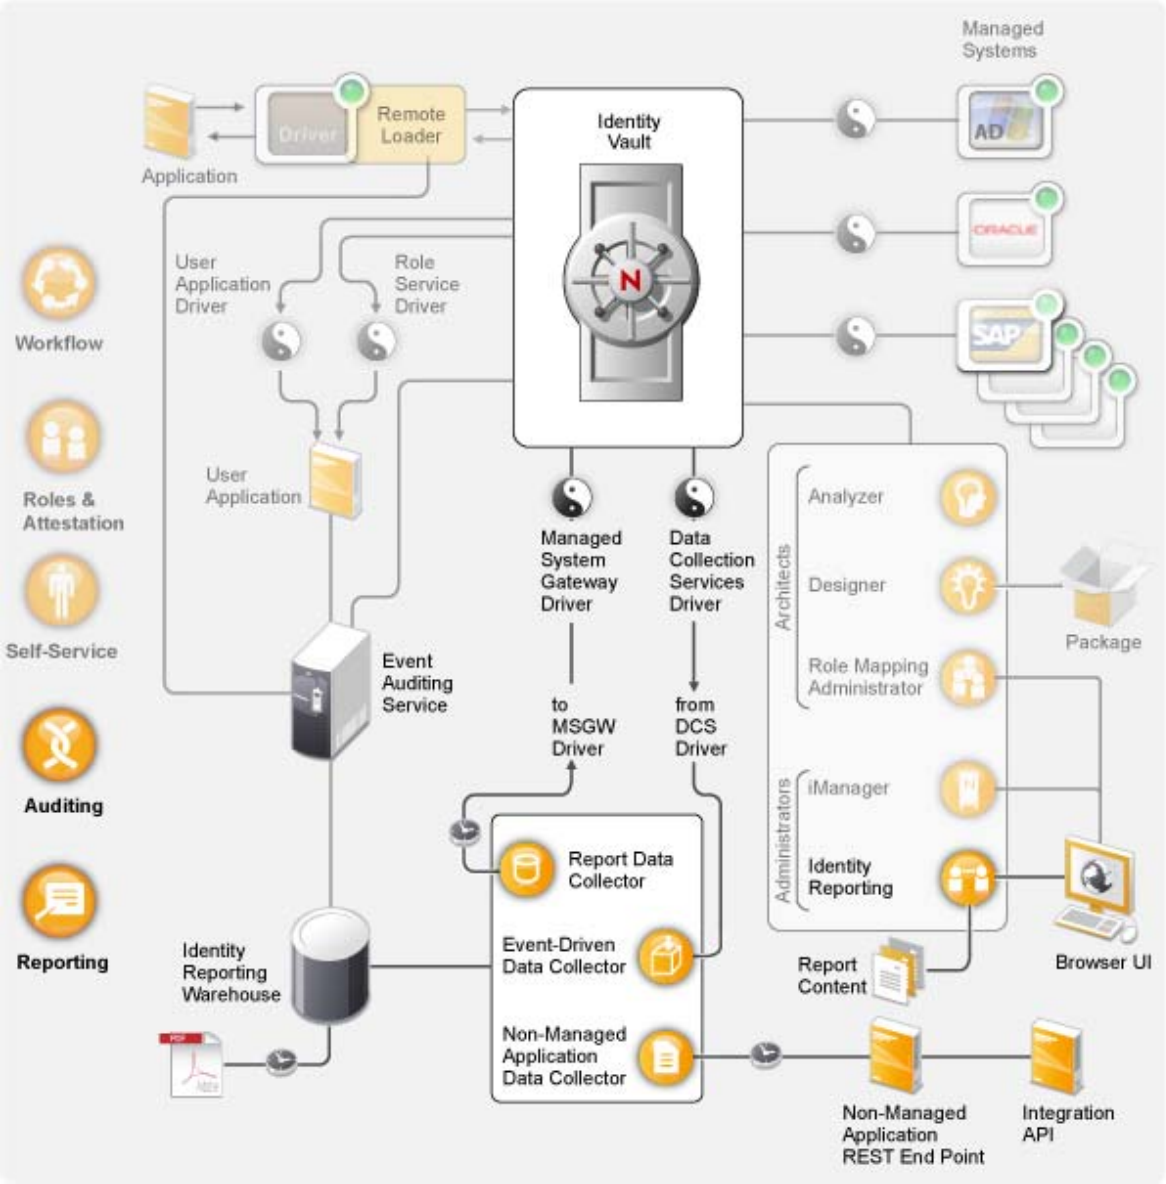

The Identity Reporting Module generates reports that show critical business information about various aspects of your Identity Manager configuration, including information collected from Identity Vaults and managed systems such as Active Directory or SAP. The Identity Reporting Module uses the following components to manage the data:.

Event Auditing Service: A service that captures log events associated with actions performed in the reporting module, such as the import, modification, deletion, or scheduling of a report. The Event Auditing Service (EAS) captures log events associated with actions performed within the Roles Based Provisioning Module (RBPM) and the Role Mapping Administrator (RMA).

Identity Information Warehouse: Repository for the following type of information:

- Report management information (such as report definitions, report schedules, and completed reports), database views used for reporting, and configuration information.

- Identity data collected by the Report Data Collector, Event-Driven Data Collector, and the Non-Managed Application Data Collector.
- Auditing data, which includes events collected by the Event Auditing Service.

The Identity Information Warehouse stores its data in the Security Information and Event Management (SIEM) database.

Data Collection Service: A service that collects information from various sources within an organization. The Data Collection Service includes three subservices:

- **Report Data Collector:** Uses a pull design model to retrieve data from one or more Identity Vault data sources. The collection runs on a periodic basis, as determined by a set of configuration parameters. To retrieve the data, the collector calls the Managed System Gateway driver.
- **Event-Driven Data Collector:** Uses a push design model to gather event data captured by the Data Collection Service driver.
- **Non-Managed Application Data Collector:** Retrieves data from one or more non-managed applications by calling a REST end point written specifically for each application. Non-managed applications are applications within your enterprise that are not connected to the Identity Vault. For more information, see "REST Services for Reporting" in the *Identity Reporting Module Guide*.

Data Collection Service Driver: A driver that captures changes to objects store in an Identity Vault, such as accounts, roles, resources, groups, and team memberships. The Data Collection Service driver registers itself with the Data Collection Service and pushes change events (such as data synchronization, add, modify, and delete events) to the Data Collection Service.

The information captured records changes to these objects:

- User accounts and identities
- Roles and role levels
- Groups

  NOTE: The reporting module does not support dynamic groups and only generates reports on static group data.

- Group memberships
- Provisioning Request Definitions
- Separation of Duties definitions and violations
- User entitlement associations
- Resource definitions and resource parameters
- Role and resource assignments
- Identity Vault entitlements, entitlement types, and drivers

Managed System Gateway Driver: A driver that collects information from managed systems. To retrieve the managed system data, the driver queries the Identity Vault. The data retrieved includes the following:

- List of all managed systems
- List of all accounts for the managed systems
- Entitlement types, values, and assignments, and user account profiles for the managed systems

Identity Reporting: The user interface for the reporting module makes it easy to schedule reports to run at off-peak times to optimize performance. For more information about the Identity Reporting Module, see the *Identity Reporting Module Guide*.

Reports: Identity Manager contains predefined reports to display the information in the Identity Information Warehouse in useful and consumable ways. You can also create custom reports. For more information about the reports, see *Using Identity Manager 4.0.2 Reports*. For information about custom reports, see "Creating Custom Report Definitions" in the *Identity Reporting Module Guide*.

Non-Managed Application REST End Point: A non-managed application is an application that is not connected to an Identity Vault, but nonetheless includes data that you want to report on. By defining a REST end point for an application, you make it possible for the reporting module to collect data from this application.

Integration API: The Identity Reporting Module provides a set of REST APIs that allow to implement a REST end point for a non-managed application, as well as write a custom reporting application.

5 Identity Manager Tools

Identity Manager provides tools that help you create and maintain your Identity Manager solution. Each of the tools has a specific function.

Figure 5-1 *Identity Manager Tools*

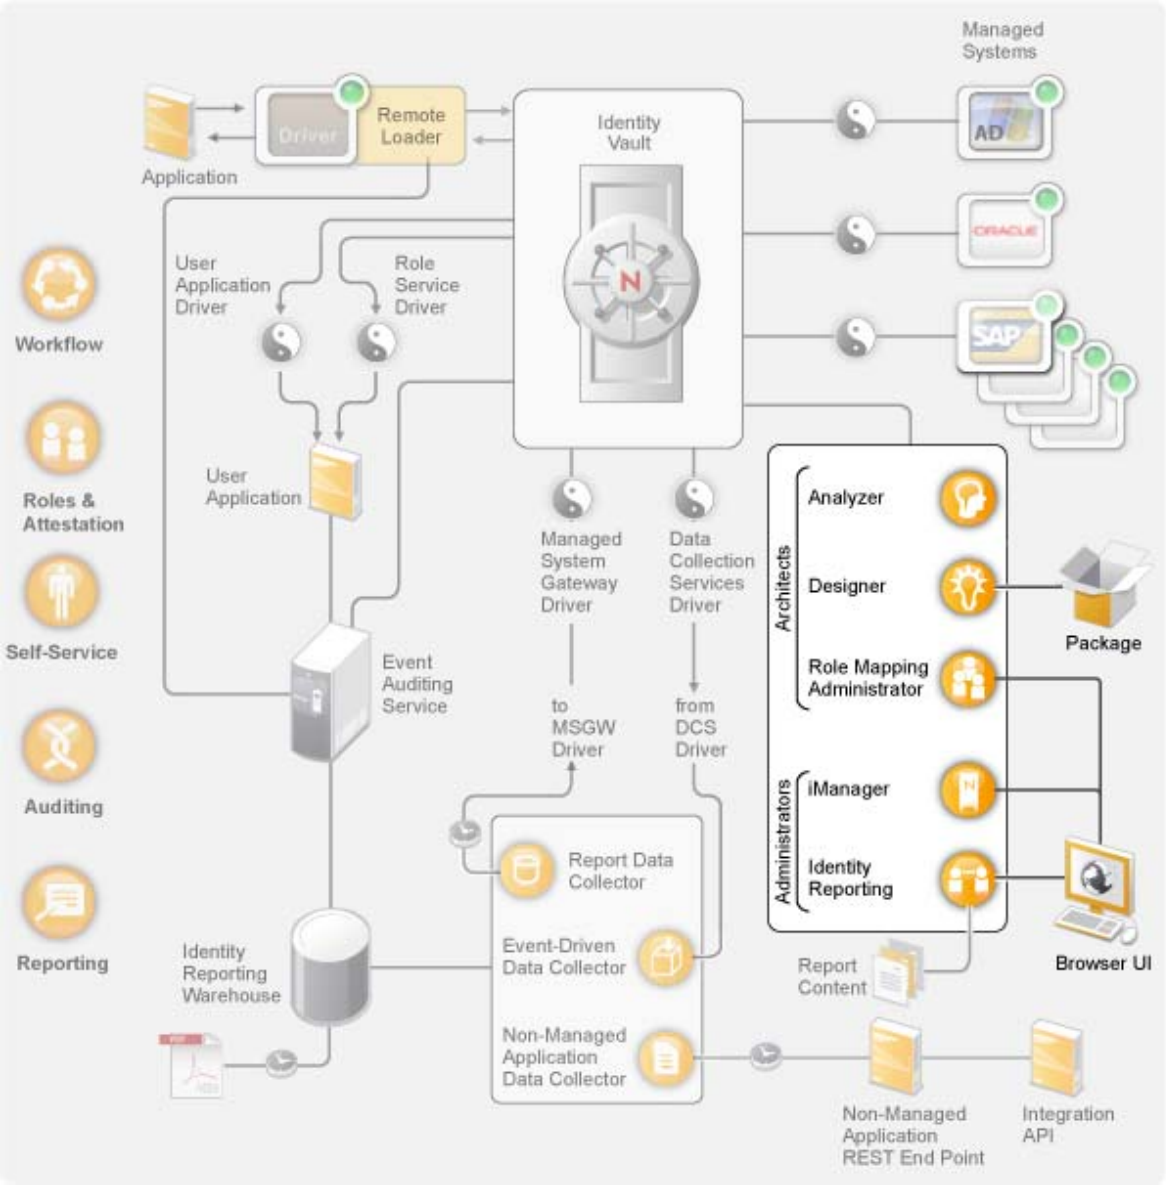

You use Designer to design, create, and configure your Identity Manager system in an off-line environment and then deploy your changes to your live system. Designer also provides you the package management capabilities for pre-configuring and customizing Identity Manager driver policies. Analyzer is used when you create your Identity Manager solution to analyze, clean, and prepare you data for synchronization.

The Role Mapping Administrator is used to create and manage roles throughout your Identity Manager solution.

You can use iManager to perform the similar tasks as Designer and also monitor the health of your system; however, package management is not supported in iManager. We recommend that you use iManager for administration tasks and Designer for configuration tasks that require changes to packages, modeling, and testing prior to deployment.

More information about each of these tools is provided in the following sections:

5.1 Analyzer

Analyzer is an Eclipse-based identity management toolset that helps you ensure that internal data quality policies are adhered to by providing data analysis, data cleansing, data reconciliation, and data monitoring and reporting. Analyzer lets you analyze, enhance, and control all data stores throughout the enterprise.

Figure 5-2 *Analyzer for Identity Manager*

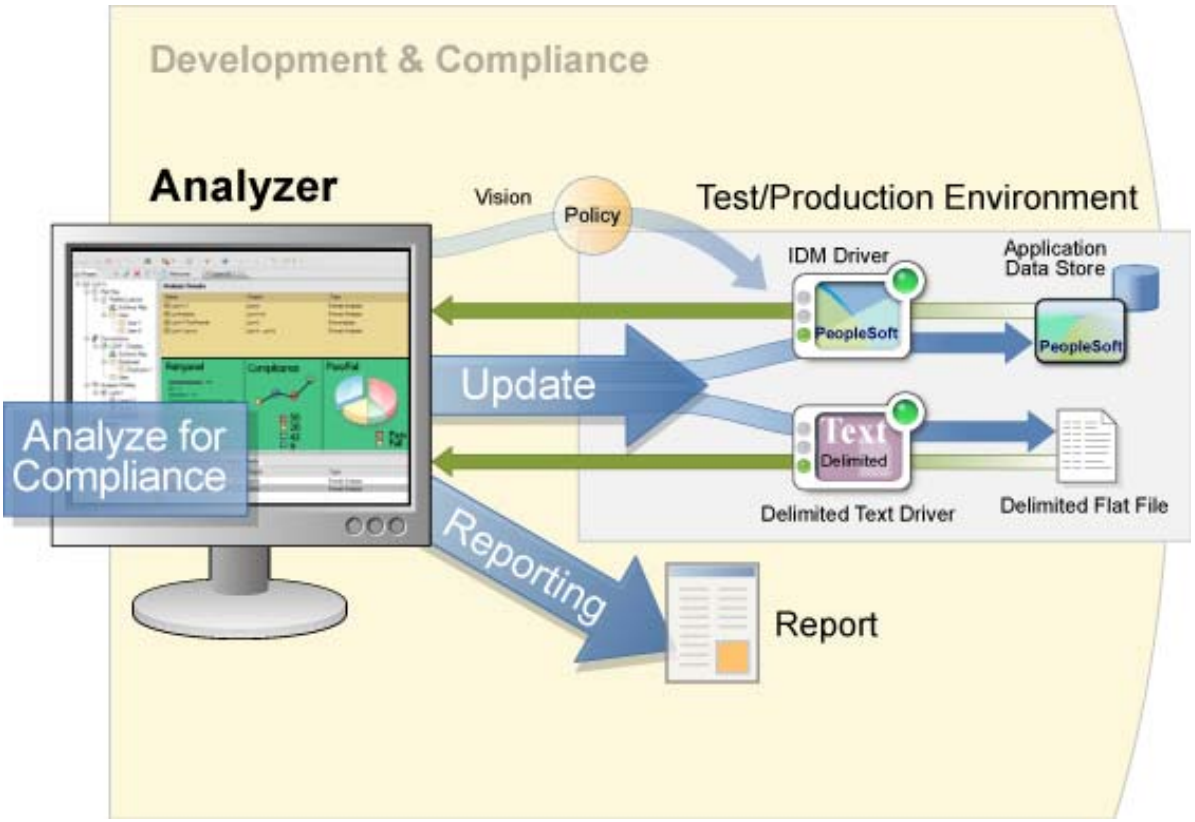

5.2 Designer

Designer is an Eclipse-based tool that helps you design, deploy, and document your Identity Manager system. Using Designer's graphical interface, you can design and test your system in an offline environment, deploy the system into your production environment, and document all details of your deployed system.

Figure 5-3 *Designer for Identity Manager*

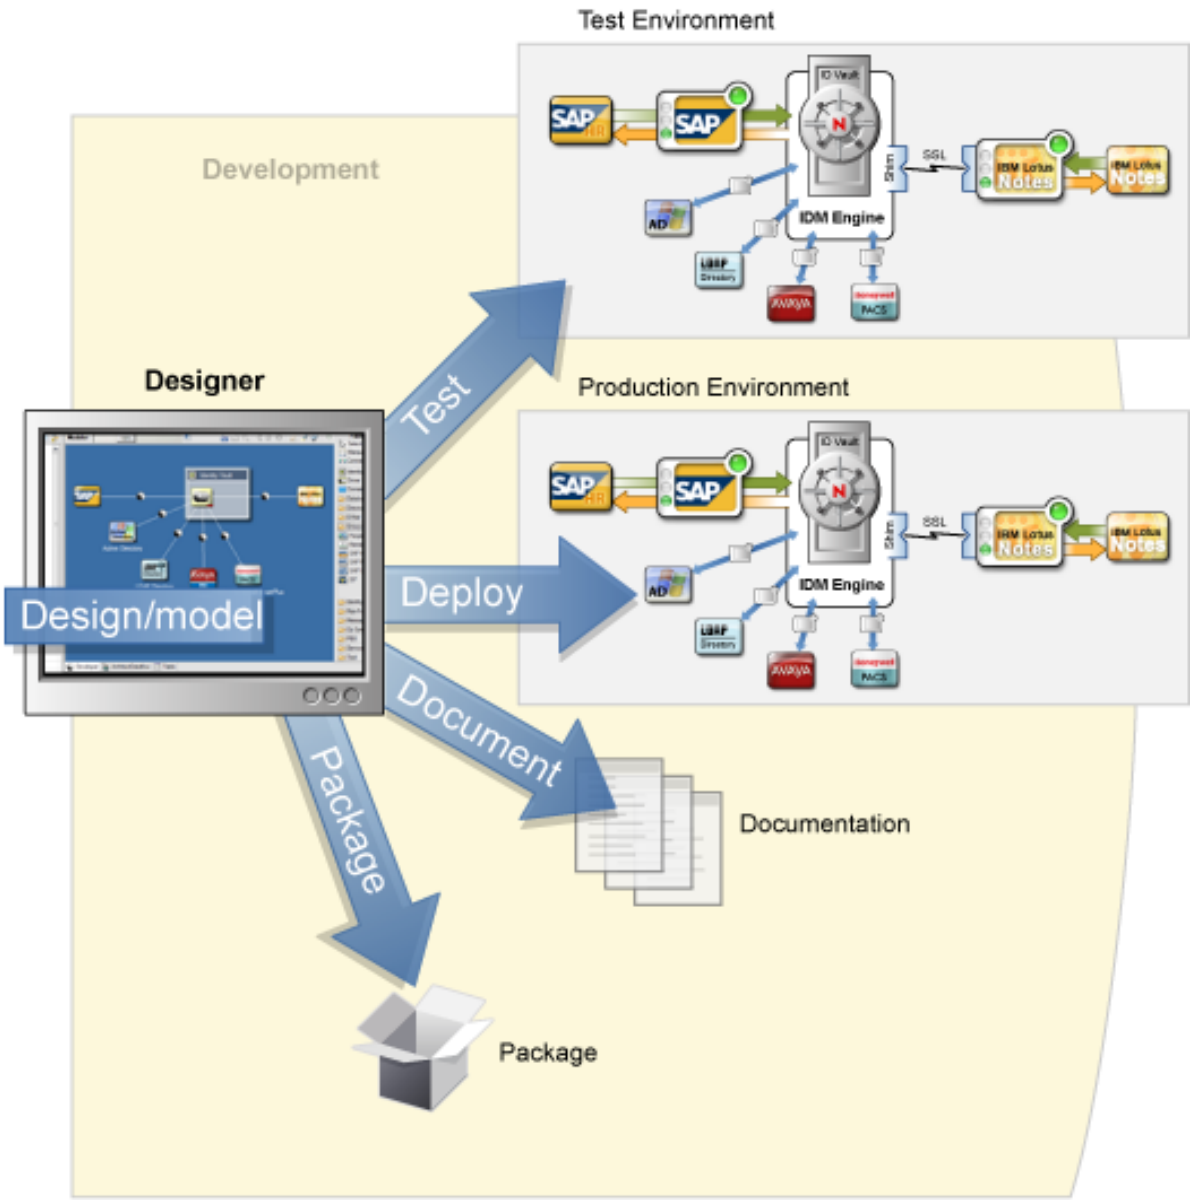

Design: Designer provides a graphical interface through which you can model your system. This includes views that allow you to create and control the connections between Identity Manager and applications, configure policies, and manipulate how data flows between connected applications.

Deploy: The work you do in Designer is deployed to your production environment only when you initiate the deployment. This gives you the freedom to experiment, test the results, and resolve any issues before going live in your production environment.

Document: You can generate extensive documentation that shows your systems hierarchy, driver configurations, policy configurations, and much more. Basically, you have all the information needed to understand the technical aspects of your system while helping you verify compliance with your business rules and policies.

5.3 iManager

Novell iManager is a browser-based tool that provides a single point of administration for many Novell products, including Identity Manager. By using the Identity Manager plug-ins for iManager, you can manage Identity Manager and receive real-time health and status information about your Identity Manager system.

Figure 5-4 *Novell iManager*

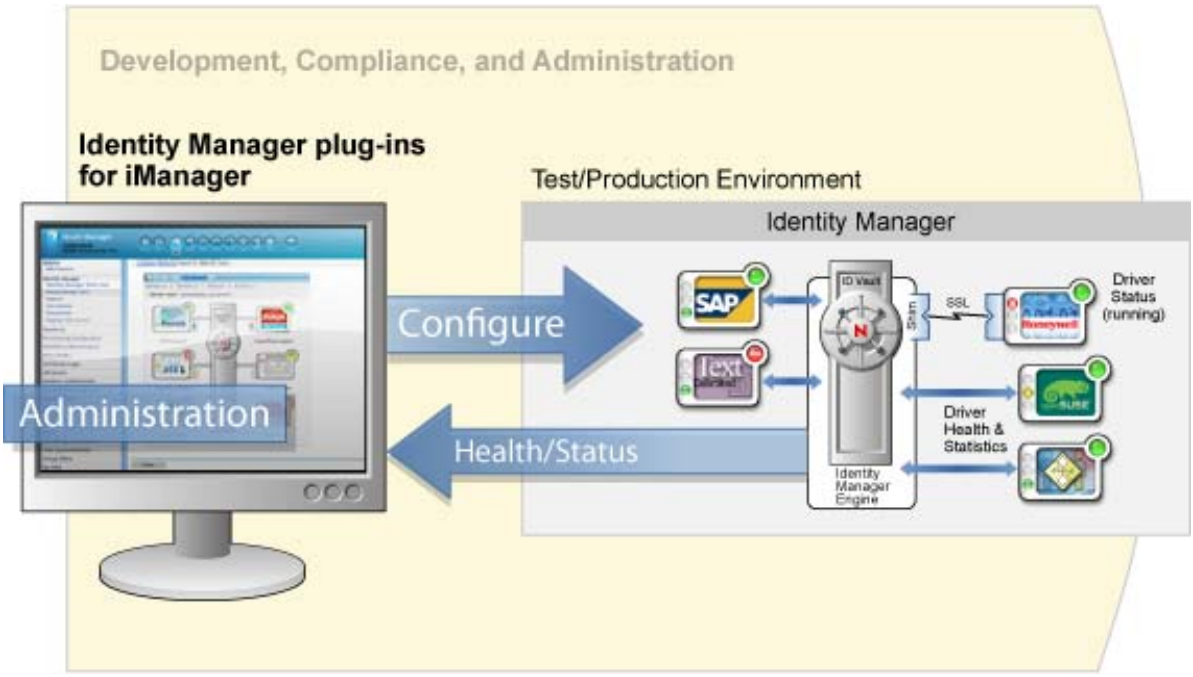

5.4 Role Mapping Administrator

The Role Mapping Administrator is a Web service that discovers authorizations and permissions that can be granted within your major IT systems. It allows business analysts, not just IT administrators, to define and maintain which authorizations are associated with which business roles.

Figure 5-5 *Role Mapping Administrator*

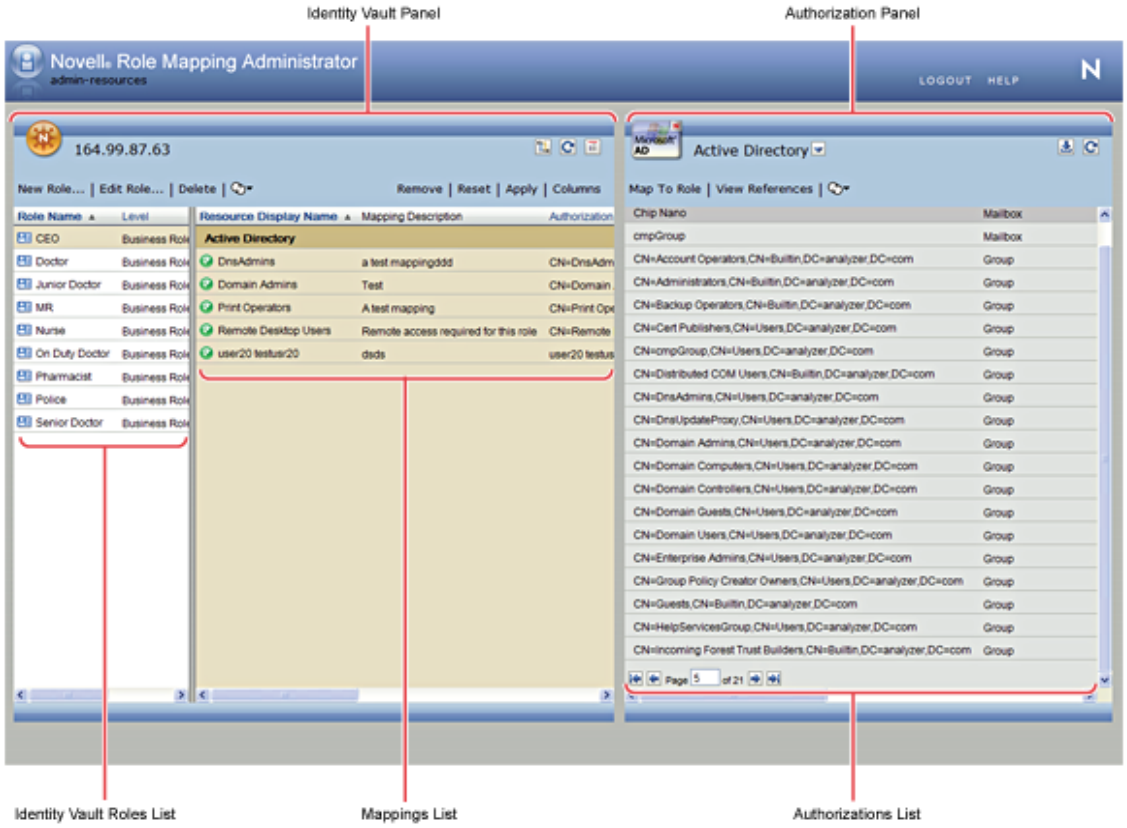

5.5 Identity Reporting

The Identity Reporting Module generates reports that show critical business information about various aspects of your Identity Manager configuration, including information collected from Identity Vaults and managed systems such as Active Directory or SAP. The reporting module provides a set of predefined report definitions you can use to generate reports. In addition, it gives you the option to import custom reports defined in a third-party tool. The user interface for the reporting module makes it easy to schedule reports to run at off-peak times to optimize performance.

Figure 5-6 *Identity Reporting Module*

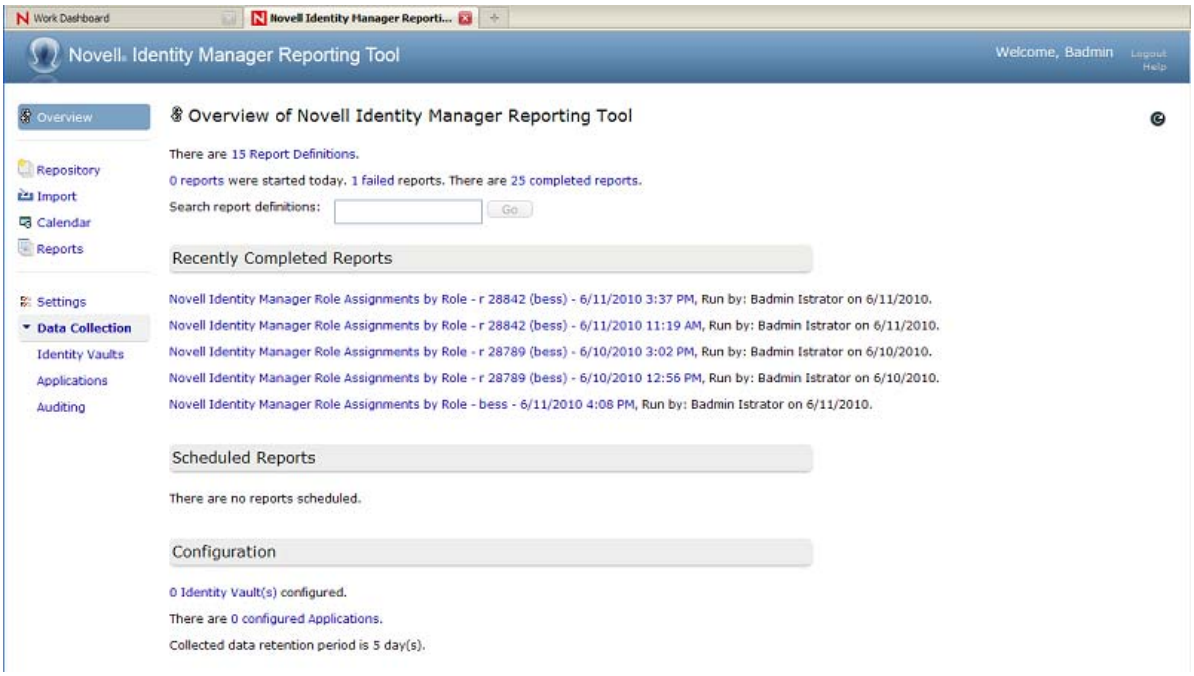

The reporting module provides several open integration points. For example, if you want to collect data about third-party applications that are not connected to Identity Manager, you can implement a custom REST endpoint to collect data from these applications. In addition, you can customize the data that is pushed to the Identity Vault. After this data is available, you can write custom reports to see this information.

6 Identity Manager Operations

After learning Identity Manager components and how they work together in Identity Manager, it is important to understand the basics of how data is synchronized between the Identity Vault and the connected system (external application). To recapitulate, Identity Manager has four main components:

- The Identity Manager engine that provides the framework.
- The Identity Manager policies that control the mapping of attributes and classes and the matching and creation of entries.
- Event filters that control the direction of data synchronization.
- The Identity Manager driver shim that serves as an interface between the application and the Identity Manager engine.

This section contains a brief explanation of the Identity Manager processes applied to data that flows between Identity Manager and the connected system.

6.1 The Identity Vault

The Identity Vault is a repository of identity information. It is also called the Novell eDirectory tree. The Identity Vault stores information specific to Identity Manager, such as driver configurations, parameters, and policies.

The Identity Vault has an extensive schema which may be customized. The Identity Vault can be viewed narrowly as a private data store for Identity Manager or more broadly as a metadirectory that holds enterprise-wide data, that you want to synchronize amongst applications including various directories, databases, phone systems, operating systems, and Human Resource systems. The data in the vault is available to any protocol supported by eDirectory, including NCP (NetWare Core

Protocol), LDAP, and DSML. Identity Manager eases the administrative efforts of large enterprises by preventing administrative effort duplication. For example, data synchronized from a PeopleSoft system to Lotus Notes is first added to the Identity Vault and then sent to the Lotus Notes system.

A typical Identity Manager environment has an Identity Vault at the center with other applications connected to it. The Identity Manager architecture can be thought of as multiple one-to-one relationships or a hub-and-spoke relationship. Each individual relationship is between the Identity Manager, Identity Vault, and a specific connected application.

Figure 6-1 *Fishbone Diagram*

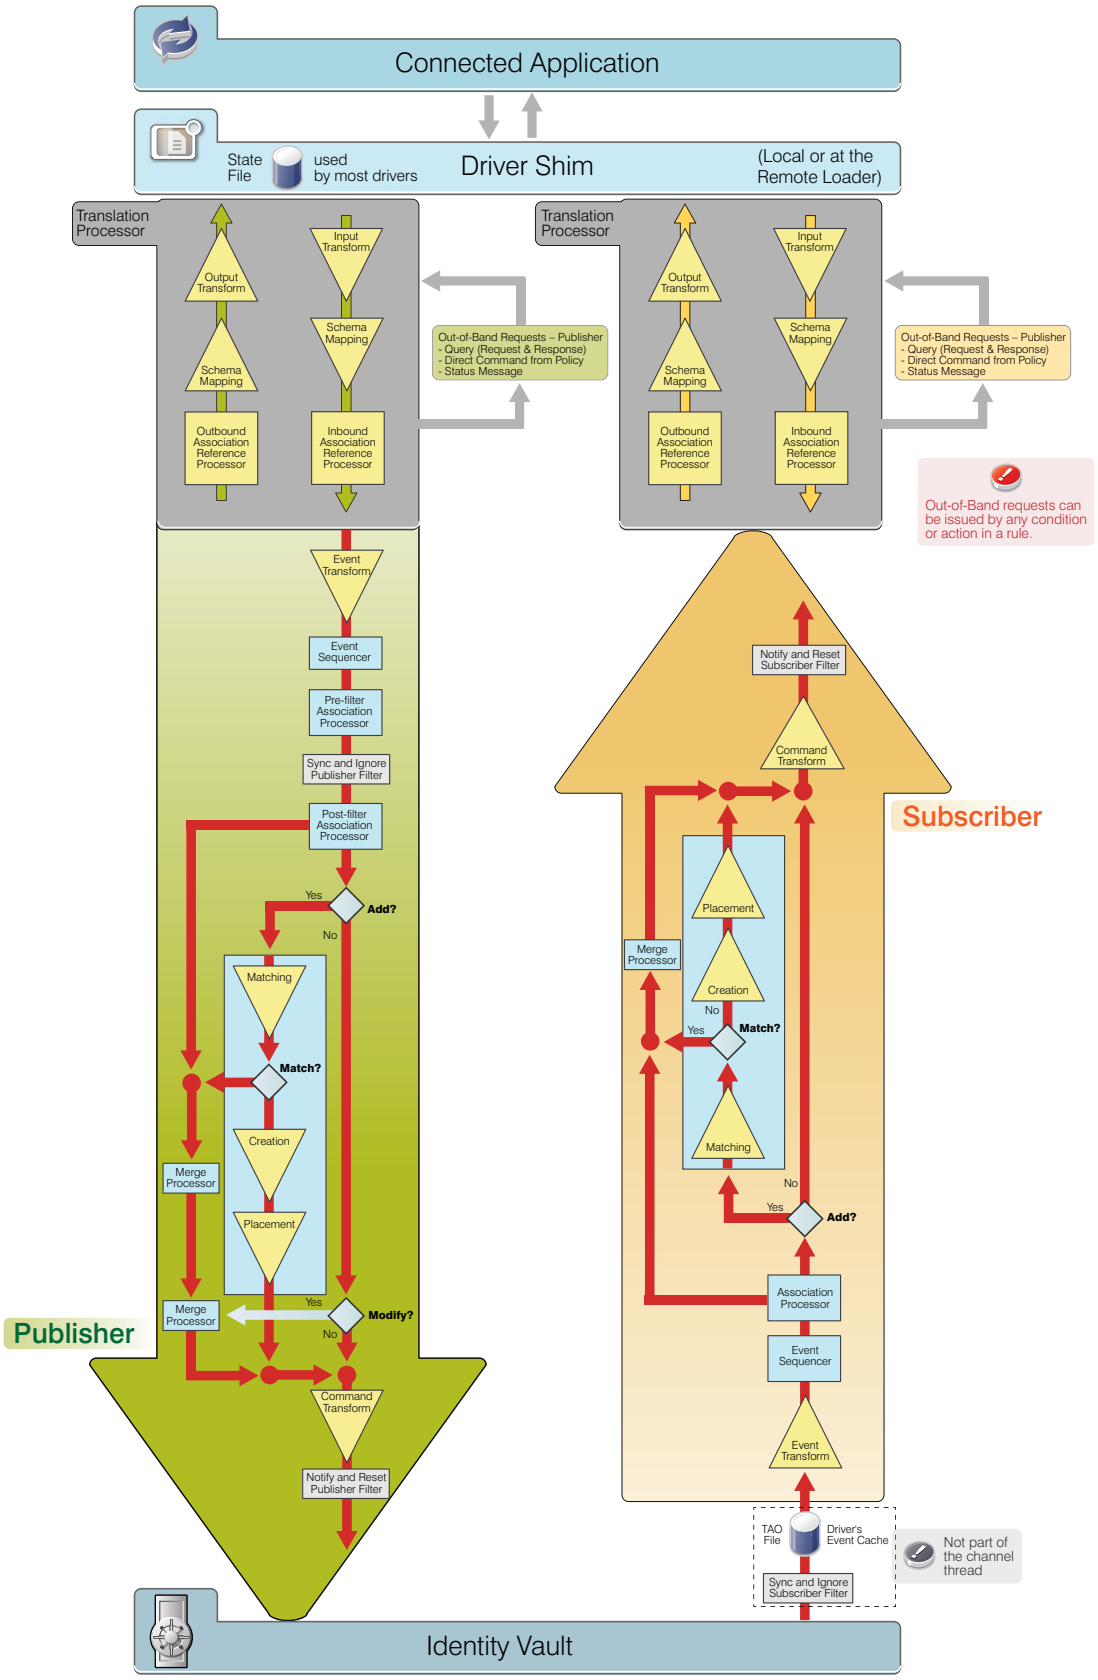

A driver is an application shim combined with policies that allows Identity Manager to communicate with an external application in order to synchronize data between the application and the Identity Vault. Note that the term driver and shim are interchangeable. In Figure 6-1, the shim is located at the top, linked to an external application and the Identity Vault. Between the driver shim and the Identity Vault are the rules which manage the data.

Data flows through an Identity Manager system in the form of XML documents. Identity Manager has a vocabulary of XML named XDS which is used to represent the state of objects and data operations with the corresponding attribute values.

The Identity Manager engine uses the shim to deliver and consume information with a connected system. It uses the driver configuration rules to decide how and what to do.

Drivers connect to the applications in order to manage objects and entities. A driver has two basic responsibilities:

- Report data changes (events) in the application to the Identity Manager engine.
- Carry out data changes (commands) submitted by the Identity Manager engine to the application.

The combination of a connected system driver, application connection information, and a set of policies is referred to as a driver configuration. Driver configurations are stored in a set of directory objects in the Identity Vault. The DirXML-Driver object contains other objects that define the policies and parameters associated with the configuration.

The driver configuration defines a data pipeline between a connected system and the Identity Vault. The driver configuration defines what might be synchronized and how to map eDirectory schema to a connected system schema or metadata. For example, in an HR application, a user's first name might be referred to as First Name and Given Name in the Identity Vault. In the namespace of the connected system, you refer First Name, but in the name space of the Identity Vault you refer Given Name. In Identity Manager, most of the time you work with the attribute names in the Identity Vault namespace.

A relationship is established between an Identity Vault object and an connected system object when the two objects represent the same entity. This relationship is called an association and is stored in the Identity Vault on the associated Identity Vault object. The association establishes a relationship between the Identity Vault object and the object in the connected system. Key values that uniquely identify objects in connected systems include Global Unique Identifiers (GUIDs), DNs, primary keys in databases, and so on. Each driver is coded to use a specific key.

6.2 The Shim

A shim is compiled code that handles translating commands and data between the connected system and Identity Manager.

The driver shim is often written in Java, which uses native application programming interface (API) calls that the system makes available to developers. APIs can include LDAP standard calls, native Windows Active Directory calls, and JDBC connections for SQL databases. The shim has the following responsibilities:

- Translating what the application understands to a standard XML document
- Creating and maintaining the connection to the connected application
- Managing the commands sent from Identity Manager and the connected application
- Monitoring the connected application for changes

For example, if the connected system is an HR system, and a new person is hired, the shim needs to build an XML document that describes this information. In Identity Manager terminology, this is an Add event and an XML document is built to describe this event to the Identity Manager engine. The event is submitted to the engine and a new user is created in the specified location.

After the new user object is created in the Identity Vault, an event is generated for other drivers that monitor changes to user objects. For example, if you have the GroupWise driver deployed, an Add event is generated for the GroupWise driver to create an e-mail box for the new user.

6.3 Channels

The flow of data between the Identity Vault and a connected system has two directions named Publisher and Subscriber. These directions are named from the point of view of the connected system:

- The Subscriber channel is the channel in which data flows from the Identity Vault to the application via the shim. The applications subscribe to data from the Identity Vault.
- The Publisher channel is the channel in which data flows from the application to the Identity Vault. The applications publish data to the Identity Vault.

There are cases in which policy might cause data to conceptually flow backward in a channel. This is referred to as channel write-back.

6.4 Events and Commands

The distinction between Events and Commands is subtle but important. The report of a change in data at the channel input is an event. Events occur both in the Identity Vault and in the connected system. Examples of events include:

- Creation of an object
- Modification of object attribute values
- Changing of an object's name
- Movement of an object within the object hierarchy
- Deletion of an object

An event coming from the Identity Vault sent over the Subscriber channel is eventually turned into a command to be submitted to the driver shim to cause some change in the connected system. An event coming from the application sent over the Publisher channel is eventually turned into a command to be submitted to the Identity Vault to synchronize the change that occurred in the application.

Commands are the output of a driver channel. When the shim sends an event notification to Identity Manager, the shim is informing Identity Manager of a change in data that occurred in the connected system. Identity Manager then determines, based on configurable policies, which commands, must be sent to the Identity Vault. When Identity Manager sends a command to the shim, Identity Manager has already taken an Identity Vault event as input, applied the appropriate policies, and determined that the change in the connected system represented by the command is necessary.

From the point of view of the overall system, if a command from one driver on its Publisher channel is creating or updating an object in the Identity Vault, it might cause events to be submitted on the Subscriber channels of other drivers in the system. This allows changes to cascade, flowing to all connected systems.

6.5 Schema Mapping Policy

The Schema Mapping policy applies to both the Subscriber channel and to the Publisher channel. The purpose of the Schema Mapping policy is to map schema names, particularly attribute names and class names, between the Identity Vault namespace and the connected system namespace. The Schema Mapping policy is applied before the Output Transformation policy when Identity Manager submits or returns a document to the shim and after the Input Transformation policy when the driver submits or returns a document to the Identity Manager.

Referring to the example of a new hire used earlier, the HR system uses First Name and the Identity Vault uses Given Name, when both refer to the same attribute. The Schema Mapping policy handles the change in names between the connected system's namespace and the Identity Vault's namespace.

The Schema Mapping policy is bidirectional. It overlaps both channels. On the Publisher channel, the connected system's names are mapped to the Identity Vault. On the Subscriber channel, the Identity Vault names are mapped to the connected system.

6.6 Event Transformation Rule

The Event Transformation rule operates on events reported on a channel input. The Subscriber and Publisher channels usually have different Event Transformation rules. The purpose of the Event Transformation rules is to modify the report of the events before the events are processed further by Identity Manager. Note that Merge operations do not transit the Event Transformation rule.

There are many common applications for the Event Transformation rules, including:

- Scope filtering (for example, only allow events on objects in a particular subtree, or with a particular attribute value)
- Custom event filtering (for example, disallow moves or deletes)
- Transforming the event directly into a custom command to be passed to the connected system
- Generating additional events

6.6.1 Publisher

The input to the Publisher channel is a description of an event coming from the connected system. The purpose of the Event Transformation rule is to modify that event description. This is applied after the Input Transformation policy and Schema Mapping policy, but before any other policy-based event processing. The policies implemented in the Event Transformation rule act on the event, such as Add, Delete, or Modify, and not on the data in the event. This is the place where policies are applied to events. For example, you can apply a policy that blocks add events.

If an Add operation is converted into a Merge operation, the current document is discarded, and the filter is used to query to both the connected system and to Identity Vault for all values. The setting for each attribute in the filter is used to decide what to do with the data. The options include overwriting the source information with the information from the destination, overwriting the destination with the source, combining the two and updating both with the results, or doing nothing.

If an Add event contains an association value, the Identity Manager engine turns it into a Modify event.

6.6.2 Subscriber

The input to the Subscriber channel is a description of an event coming from the Identity Vault. In many cases, the filter might be used to determine the types of objects you want, and the attributes of those objects, but the Event Transformation policy can be used to further customize the events. This can be referred to as scope filtering, and it allows for much finer control of what gets through.

For example, you can use filter to specify user objects. It assumes that you want all users synchronized. If a connected system is limited to a subset of all users, then the Event Transformation policy is used to decide if an event for an object is in scope or not. For example, if your connected system should have only users with a department attribute of Sales in it, then a rule on the Event Transformation policy to block any event that is for a user that does not have Sales as its department can accomplish this goal.

6.7 Filter

The filter controls the flow of data between the Identity Vault and the connected system. The filter plays several roles in an Identity Manager driver configuration. Figure 6-1 shows filter in four places representing most of its roles, but there is really only one filter for the driver.

The driver filter specifies the classes of objects and the attributes of those objects for which Identity Vault processes events and commands for both channels. The filter instructs the Metadirectory engine about events and information the driver's configuration is interested in. From the Identity Vault side, events are queued for the driver if they match an object class in the filter, and if they match an attribute that is set to Sync, Notify, or Reset. Events that occur in the Identity Vault that do not match the data types specified in the filter are ignored by this driver. Similarly, for the application, events that occur that do not match the data types specified in the filter are ignored, though the shim might still have to examine them to see if they need to be handled. For example, if the Identity Manager driver configuration should synchronize only user information, the filter specifies User objects and modification to other Identity Vault objects is ignored. From the possible User class attributes, the filter specifies the selected attributes, such as CN, Given Name, Surname, and Telephone Number. Modifications to other user class attributes is ignored.The user object class and set of related data attributes are listed in the filter for most connected systems.

While the channels allow for data flow, policies and filters are placed in the channel to regulate what gets through and how it looks when it reaches the destination. For example, by configuring the driver filter you can block an attribute value, such as a telephone number from reaching the Identity Vault from the connected system or vice versa. This helps to regulate whether the Identity Vault or the connected system is the authoritative source to meet specific business requirements. For example, if the filter for the relationship between the PBX system and the Identity Vault allows an employee's telephone number to flow from the PBX system into the Identity Vault but not from the Identity Vault to the PBX system, then the PBX system is the authoritative source for the telephone number. If all other connected system relationships allow the telephone number to flow from the Identity Vault to the connected systems, but not vice versa, the net effect is that the PBX system is the only authoritative source for employee telephone numbers in the enterprise.

- Section 6.7.1, “The Sync Attribute,” on page 53
- Section 6.7.2, “The Notify Attribute,” on page 54

6.7.1 The Sync Attribute

On the Publisher channel, when an event has been queued for the channel to process and it has passed through the Input Transformation rule, the Schema Map, and the Event Transform, the Sync attributes are selected from the input document, and any attributes not set to Sync or Notify are

removed. Attributes that are set to Reset are also handled by querying Identity Vault for the correct value, and having the correct value sent back to the connected system to undo the change that has just been made.

On the Subscriber channel, the Sync filter works the same way it works for the Publishes channel. The only difference is that events are coming from the Identity Vault instead of the connected system.

6.7.2 The Notify Attribute

Notify is a way for attribute data to be used in the event document, without it actually being synchronized to the Identity Vault. For example, you need a person's first name, middle name, and last name from your HR system in order to create an account, but you do not actually want to store the middle name in the Identity Vault. By setting the middle name attribute to Notify, you can access the attributes value without having to store it in the Identity Vault. Any attributes set to Notify are stripped out of the document prior to being submitted to the destination.

6.8 Add Processor

6.8.1 Publisher

The Add Processor is used to decide if an event is an add document. This is a branching point in the driver's processing of the event. An add document is redirected to the Matching rule. The shim supplies the association value, allowing the Identity Manager engine to quickly and easily find the correct object in the Identity Vault. Associations are created as a match between two objects or when an object is newly created in either the Identity Vault or the connected system. After an association is formed between objects, this association remains in effect until the objects are deleted or the association is deleted by the administrator. Well-designed Matching rules automate the creation of associations between existing objects in the Identity Vault and the connected system. For more information, see Section 6.13.3, “Associations,” on page 60.

If it is not an add document, it moves on to the Command Transformation as the next step.

6.8.2 Subscriber

On the Subscriber channel, the Add Processor is used to decide if an event is an add document. This is a branching point in the driver's processing of the event. An add document is redirected to the Matching rule. The Identity Manager engine uses the association value of an Identity Vault object to allow the shim to modify the correct object in the connected system. For more information, see Section 6.13.3, “Associations,” on page 60.

If it is not an add document, it moves on to the Command Transformation as the next step.

When a Modify event does not contain an association that resolves to an actual object when it arrives at the Add Processor, the Identity Manager engine attempts to create it. This is the Synthetic Add process, and it can happen on either the Publisher or the Subscriber channel.

The Identity Manager engine uses the Modify event to figure out which object to work with. It then uses the filter to query back to get all attributes that are available for that object that are set to Sync or Notify on the current channel. It discards the Modify event and builds an Add event to replace it. The Add event is forwarded through the Matching, Create, Placement, and Command Transform rules

(on the Subscriber it also goes through the Schema Map and Output Transform). The Event Transformation rule is not applied for Synthetic Adds. This is due to the rule location before the Add Processor.

6.9 Matching Rule

Matching rules establish links between an existing object in the Identity Vault and an existing object in the connected system. The matching rules specify which class and attribute values must match for an object in the Identity Vault and an object in the connected system to be marked as corresponding entries.

A good matching rule requires you to investigate both systems involved, and find the data that guarantees a 1:1 mapping between them. Attributes such as employee ID number, email address, and badge number are some of the more common pieces of data used for matching criteria. If there is no single attribute available, the combinations of attributes might be used. Matching on Surname only is not a good criteria. For example, in larger organizations, there might be a possibility that two employees have the same last name. Matching on Surname + Given Name would produce higher quality matches and matching on Surname + Given Name + Department would further increase the probability of correct matching. If a match is successful, an association between the two objects is created. If a match is not successful, the Create rules are used.

- Section 6.9.1, “Publisher,” on page 55
- Section 6.9.2, “Subscriber,” on page 55

6.9.1 Publisher

The Matching rule is used to link an object in the Identity Vault with the corresponding object in the connected system. For example, if you are connecting an existing HR system to an existing eDirectory system, there are people in the HR system, and users in the Identity Vault, and they both represent the same user. The Matching rule contains rules which allow Identity Manager to determine that "Joe Doe" in HR system is "jdoe13" in the Identity Vault.

The Matching rule uses matching criteria and queries Identity Vault looking for a matching object. The Matching rule returns zero when no object is matched, so that the Add event continues to be processed. It returns one when one matching object is found, which means that the object in the input document matches an object in the Identity Vault. After the objects are matched, the data between the two objects is merged based on filter settings. If the Matching rule finds more than one matching object, the Identity Manager engine treats this as an error and quits the transaction. You should either modify the Matching rule or manually handle this conflict.

6.9.2 Subscriber

On the Subscriber channel, the Matching rule works on the Add events and uses the Identity Vault data to query the connected system looking for matching objects.

6.10 Create Rule

The Create rules are applied to the Add events when the Matching rules fail to find a match. The Create rules specify the minimum set of data that an event must have before an object can be created in the Identity Vault or the connected system.

6.10.1 Publisher

If the Matching rule does not find a matching object in the Identity Vault, the Create rule is applied to the document to ensure that the document contains sufficient information. It is also used to supply default values for attributes, and it might specify a template to be used in the creation of the new object. From the Identity Vault side, a user object must have a name (CN or UID) and it must have a Surname. While this might be enough for Identity Vault objects to be created, most organizations might need additional information before creating an account. The driver can reject documents that do not contain sufficient information to continue processing.

The Create rule can also veto an Add event if the Add event fails to meet the conditions imposed by the Create rule. For example, if the Create rule requires an object to have a telephone number and it doesn't have one, the Add event is vetoed.

The discarded events are reprocessed when the additional attribute information is added in the connected system. This results in a Modify event without an associated object, which the Add Processor converts to a Synthetic Add.

6.10.2 Subscriber

The Create rule in Subscriber works the same as the Publisher channel in determining if an event has sufficient information to create an object in the connected system. This requires knowledge of the connected system and its technical or business requirements.

The Create rule is often used to examine the attributes available for the new object (from the source event) and vetoes the creation of the new object if one or more required attributes is missing. The most common example on the Subscriber channel is to require a password. Normally a user is created in the Identity Vault in two stages, first as a user object and then as a second operation a password is set. It is very common to see that the object is created, but the Create rule fails due to the lack of password which is a required attribute for creating a new user object. A moment later when the password event comes through, the new object is successfully added.

6.11 Placement Rule

6.11.1 Publisher

If the Matching rule determines that there are no matching objects, and the Create rule verifies that the event meets the minimum requirements, the Placement rule specifies which object to create, and where to place. For the Publisher channel, the object created is placed in the Identity Vault using the naming rules contained in policy. For example, the Identity Vault places all objects in the `Data\Users\` container.

6.11.2 Subscriber

On the Subscriber channel, the Placement rule works the same as the Publisher channel. Placements in connected systems can be simple or complex. A simple Placement rule places all objects in the same location.

Placements in the connected systems require a detailed knowledge of how objects are represented in the connected system.

A more complex example might use an HR location code attribute to determine where to place an object. You can use a Mapping Table to help establish the relationship between a placement location and an attribute value.

6.12 Command Transformation Rule

The Command Transformation rule operates on commands that are about to be issued to a channel output. The Subscriber and Publisher channels usually have different Command Transformation rules. The purpose of the Command Transformation rule is to provide final processing on commands before the commands are sent to the Identity Vault or to the connected system.

Some possible applications for the Command Transformation rule include:

- Changing the command type (for example, an object delete command might be transformed into a modification that will cause the object to be archived)
- Blocking commands
- Adding additional commands
- Controlling the output of the Identity Manager engine's Merge process

- Section 6.12.1, “Publisher,” on page 57
- Section 6.12.2, “Subscriber,” on page 58

6.12.1 Publisher

All events pass through the Command Transformation rule. This is where the earlier branch at the Add Processor rejoins the flow. Most of the driver policies reside in the Command Transformation, because conceptually this is where the conversion from event to command happens. Up to this point, the document has been describing an event that has happened in the connected system. Now that event is converted into a command and applied to the Identity Vault. It is the last chance to modify a command before it is applied to the Identity Vault.

6.12.2 Subscriber

On the Subscriber channel, most of the driver policies reside in the Command Transformation, because conceptually this is where the conversion from event to command happens and before the Schema Mapping policy is applied. Both the Schema Mapping policy and the Output Transformation policy are executed after the Command Transformation policy on the Subscriber channel. Up to this point, the document has been describing an event that occurred in the Identity Vault. Now this event is converted into a command and applied to the connected system.

6.13 Rules, Policies, and Style Sheets

Within any one of the rules covered above (Input Transform, Command Transform, etc.) are zero or more policies. Some of these may come from the preconfigured driver import used as a starting point. Others could be customizations of the driver configuration.

A rule in Identity Manager is a collection of policies. Each rule has conditions that have to be met, and actions to be carried out when the conditions are true. The grammar of the conditions is meant to be human readable and generally to make sense. For example, a condition of "if object class equal user" would be True if the object being described in the current document is a User object, and would be False if the object is a Group. Conditions are made up of Condition Groups. Within a Condition Group, all of the conditions can be combined with And or Or, and the result must be True for the Condition Group to evaluate to True, otherwise it evaluates to False. Multiple Condition Groups can also be combined, using And and Or. Once the Condition Group(s) have evaluated to True, the Rule's Actions are performed.

Figure 6-2 *Policy Builder Conditions*

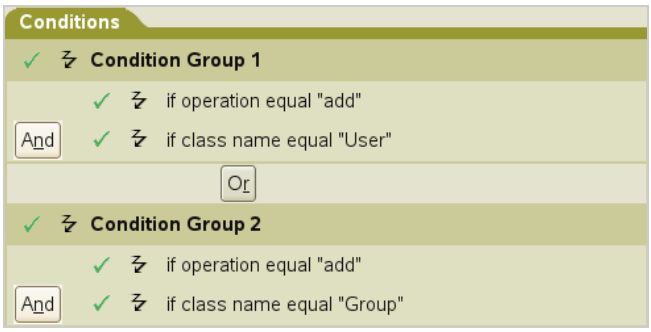

Actions can act on the current document, and in many cases this is sufficient. But Actions can also query the source or destination, which could be Identity Vault or the connected system, depending on which channel you are on, for additional information. Actions can change the current document in to a modified version of itself, or can block it entirely. Actions can be used to make different documents. A document describing a delete event in a connected system could be tested by a Condition (if operation equal delete), and acted on by a set of Actions to prevent the associated object from being deleted in Identity Vault (veto), but to modify that object to remove its association value for this application (remove-association), and to disable it (set destination attribute value Login Disabled = True).

Figure 6-3 *Policy Builder Actions*

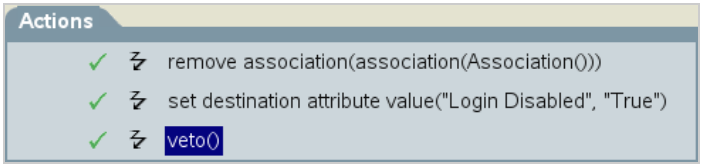

Policies define what data is transferred and how the data is synchronized between the connected system and the Identity Vault. A default set of policies is available with the driver configuration. Other policies might be local customizations of the driver. You can write policies using the DirXML Script, XSLT, or ECMA Script.

The purpose of a policy is to make changes to the input document and produce an output document. Most policies are evaluated either on the Subscriber channel or on the Publisher channel. The Schema Mapping policy, the Input Transformation policy, and the Output Transformation policy are evaluated on both channels. For example, one organization might use the `inetOrgPerson` as the main user class, while another organization might use User. A policy can be implemented to add the phone number change to an `inetOrgPerson` for the first organization, and a separate rule can be implemented to make it work for the User class. Policies make schema transformations, specify matching criteria to determine if an object already exists in the connected system or the Identity Vault, and many other things. Because of this, an Add event reported by your connected system might end out as a Modify operation in the Identity Vault, if a matching policy determines that the object you added already exists in the Identity Vault.

On the Subscription channel, when a new user is created in the Identity Vault and you want it created in the connected system, before sending this command to the driver, the Identity Manager engine calls a series of policies. These policies define the way objects are created and determine if a corresponding user already exists in the connected system, make decisions about placement, provide default values for required attributes that are not specified, and so on. This Add event might be transformed into a Modify event if the object exists in the connected system. Attributes that were not contained in the original event could be added to conform with the object creation model of the connected system.

Style sheets define XSLT transformation rules. Style sheets transform input or output commands into a different command, change an event from one type to another, or perform other arbitrary XML transformations. For more information, see Identity Manager Style Sheets.

6.13.1 Input Transform Rule

The Input Transformation rule applies to both the Subscriber channel and to the Publisher channel. The purpose of the Input Transformation rule is to perform a preliminary transformation on all XML documents sent to Identity Manager by the connected system and returned to Identity Manager from the connected system. The Input Transformation rule is applied to the XML documents sent to `XmlCommandProcessor.execute` and `XmlQueryProcessor.query` when called by the connected

system and to the XML documents returned from `SubscriptionShim.execute` and `XmlQueryProcessor.query` when called by Identity Manager engine. The Input Transformation policy is applied before the Schema Mapping policy.

The Input Transform rule is often used to transform data from the application format into the Identity Vault format. When the Input Transformation is used for data format transformations the Output Transformation policy usually performs the data transformation in the opposite direction (transforms data from the Identity Vault format to the connected system format). This rule operates in the connected system's (application's) namespace. For example, it might be used to reformat data, such as changing a phone number that is formatted as 1(815)555-1212 to 1-815-555-1212.

The Input Transformation rule is also used to perform actions in response to the results of commands sent to the shim. Note that the schema names are always in the application namespace in the XML processed by the Input Transformation policy.

It is also possible to use the Input Transformation rule to transform an arbitrary XML format native to the connected application to the format expected by Identity Manager. Such transformations must be written in XSLT because DirXML-Script operates only on the Identity Manager specific XML vocabulary that is specific to Identity Manager. A few examples are the Delimited Text driver and the SOAP driver.

6.13.2 Output Transform Rule

The Output Transformation policy applies to both the Subscriber channel and to the Publisher channel. The purpose of the Output Transformation policy is to perform a final transformation on all XML documents sent to the shim by the Identity Manager engine and returned to the shim by Identity Manager engine. The Output Transformation policy is applied to the XML documents sent to `SubscriptionShim.execute` and `XmlQueryProcessor.query` when called by the Identity Manager engine and to the XML documents returned from `XmlCommandProcessor.execute` and `XmlQueryProcessor.query` when called by the shim. The Output Transformation rule is applied after the Schema Mapping rule.

The Output Transform rule is the converse of the Input Transform rule. It modifies the command that is about to be submitted to the shim as required. This usually involves undoing what has been done in the Input Transform rule. If you have an Input Transform rule that converts phone numbers formatted as 1(815)555-1212 to 1-815-555-1212, you need to have an Output Transform rule that converts 1-815-555-1212 to 1(815)555-1212.

You can also use the Output Transformation policy to transform the format used by Identity Manager to an arbitrary XML format native to the connected application. These transformations must be written in XSLT because DirXML-Script operates only on the XML vocabulary that is specific to Identity Manager.

6.13.3 Associations

The Association value is Identity Manager's way of keeping track of which object in the connected system matches an object in the Identity Vault. Each driver handles this slightly differently. In almost all cases, this should be a 1:1 match, so that it is possible to say that "john doe", employee number 1234567 in the HR system matches exactly with the user object "jdoe13" in the Identity Vault, with "doe john" in Active Directory, and jdoe13@example.com" in the e-mail system. Most connected systems have some sort of internal unique identifier, even if it is not the one that you usually see in the system's management tools. eDirectory and Active Directory have a globally unique identifier or GUID. Many HR systems have an employee number. E-mail systems usually have a unique email address value for each person. Identity Manager uses these identifiers to build its association.

Associations are stored in the Identity Vault only. On the Subscriber channel, the Identity Manager engine uses this value to allow the shim to modify the correct object in the connected system. On the Publisher channel, the shim supplies the association value, allowing the Identity Manager engine to quickly and easily find the correct object in the Identity Vault to work with. The following association states are stored in the Identity Vault:

- **0 Disabled**: Changes in the driver objects are not synchronized with the Identity Vault.
- **1 Processed**: Successful association has been created between the driver objects and Identity Vault.
- **2 Pending**: The Identity Manager engine identified a modification to an object, and attempted to match it or create it in the connected system, but was unable to do so.
- **3-Manual**: A manual association was created by the user.
- **4-Migrate**: The account was synchronized or migrated.
- **blank No association**: No association has been created.

6.13.4 Synthetic Adds

When a Modify document without an association encounters the Add Processor, the Identity Manager engine converts the Modify event into a Synthetic Add process. This process happens on the Publisher or the Subscriber channel.

The Identity Manager engine uses the Modify event to figure out which object to work with. It then uses the filter to query back to get all attributes that are available for that object that are set to Sync or Notify on the current channel. It then throws away the Modify document and builds an Add document to replace it. This Add document is then forwarded through the Matching, Create, Placement, and Command Transformation (on the Subscriber it also goes through the Schema Map and Output Transformation).

Figure 6-4 *Subscriber Channel Association Processor*

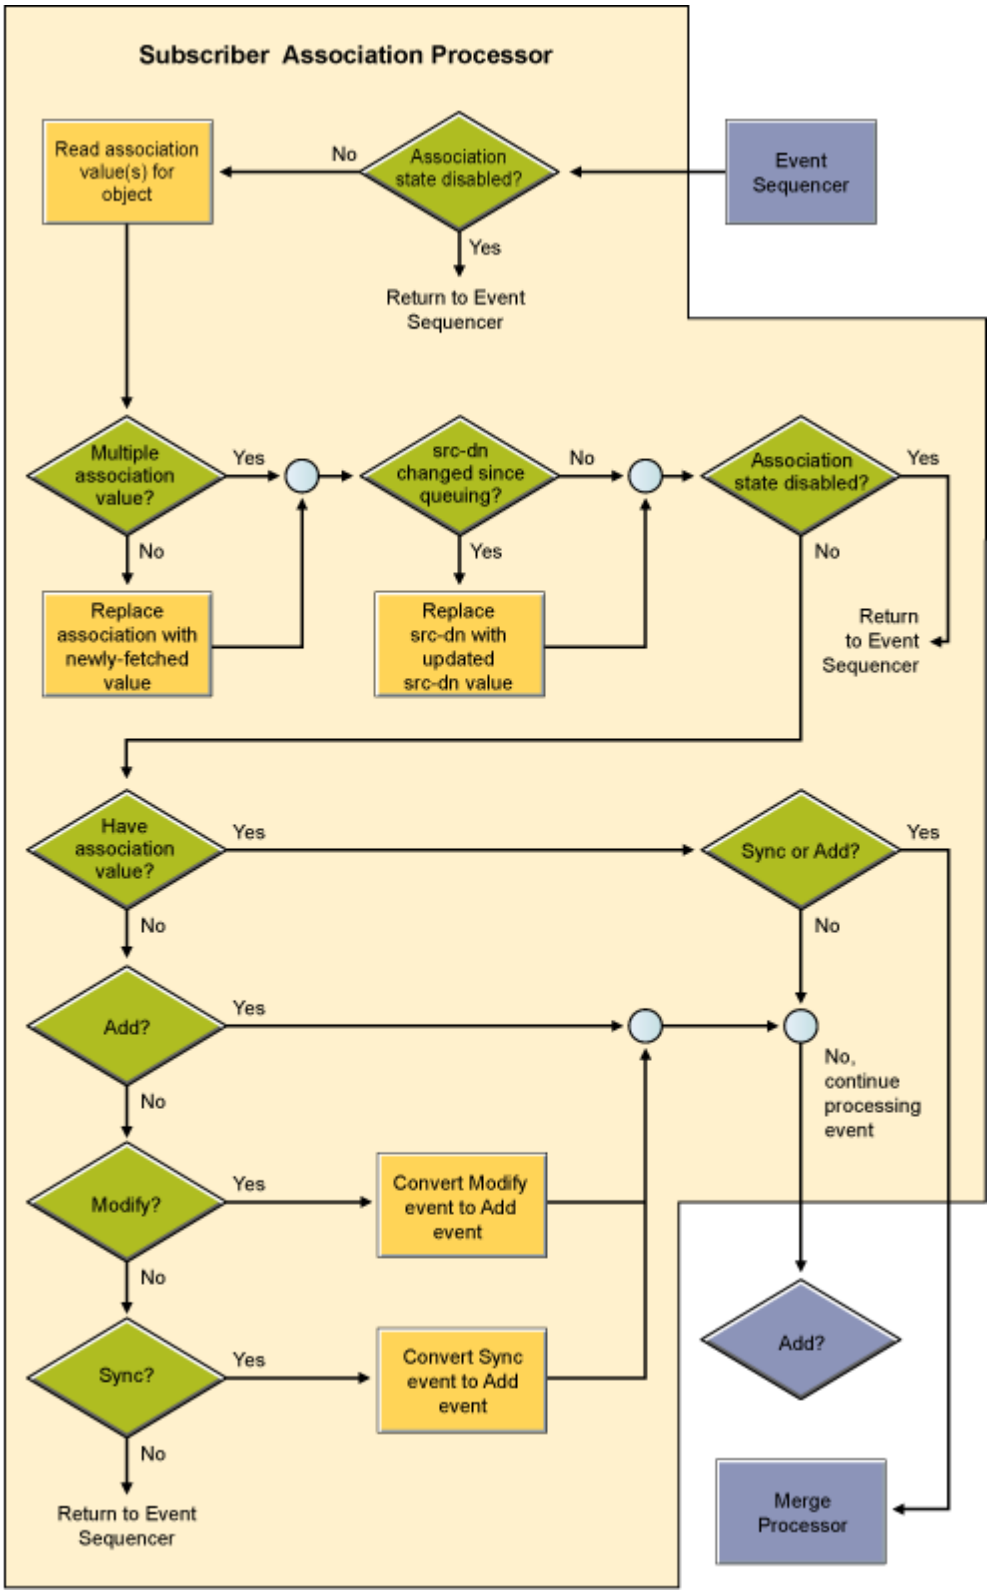

6.13.5 Merge Processing

A Merge operation occurs when the Identity Manager engine converts an Add operation into a Modify operation. This happens most commonly during an initial migration, as the migration sends objects down a channel, and the Matching rule finds an object that it can use to associate with the object being migrated.

In a Merge operation, the current document is discarded again (like the Synthetic Add), and the filter is used to query both the connected system and the Identity Vault for all values. The setting for each attribute in the filter is used to decide what to do with the data. The options include overwriting the

source information with the information from the destination, overwriting the destination with the source, combining the two and updating both with the results, or doing nothing. The following flow charts illustrate the Publisher Merge Processor and the Subscriber Merge Processor.

Figure 6-5 *Publisher Merge Processor*

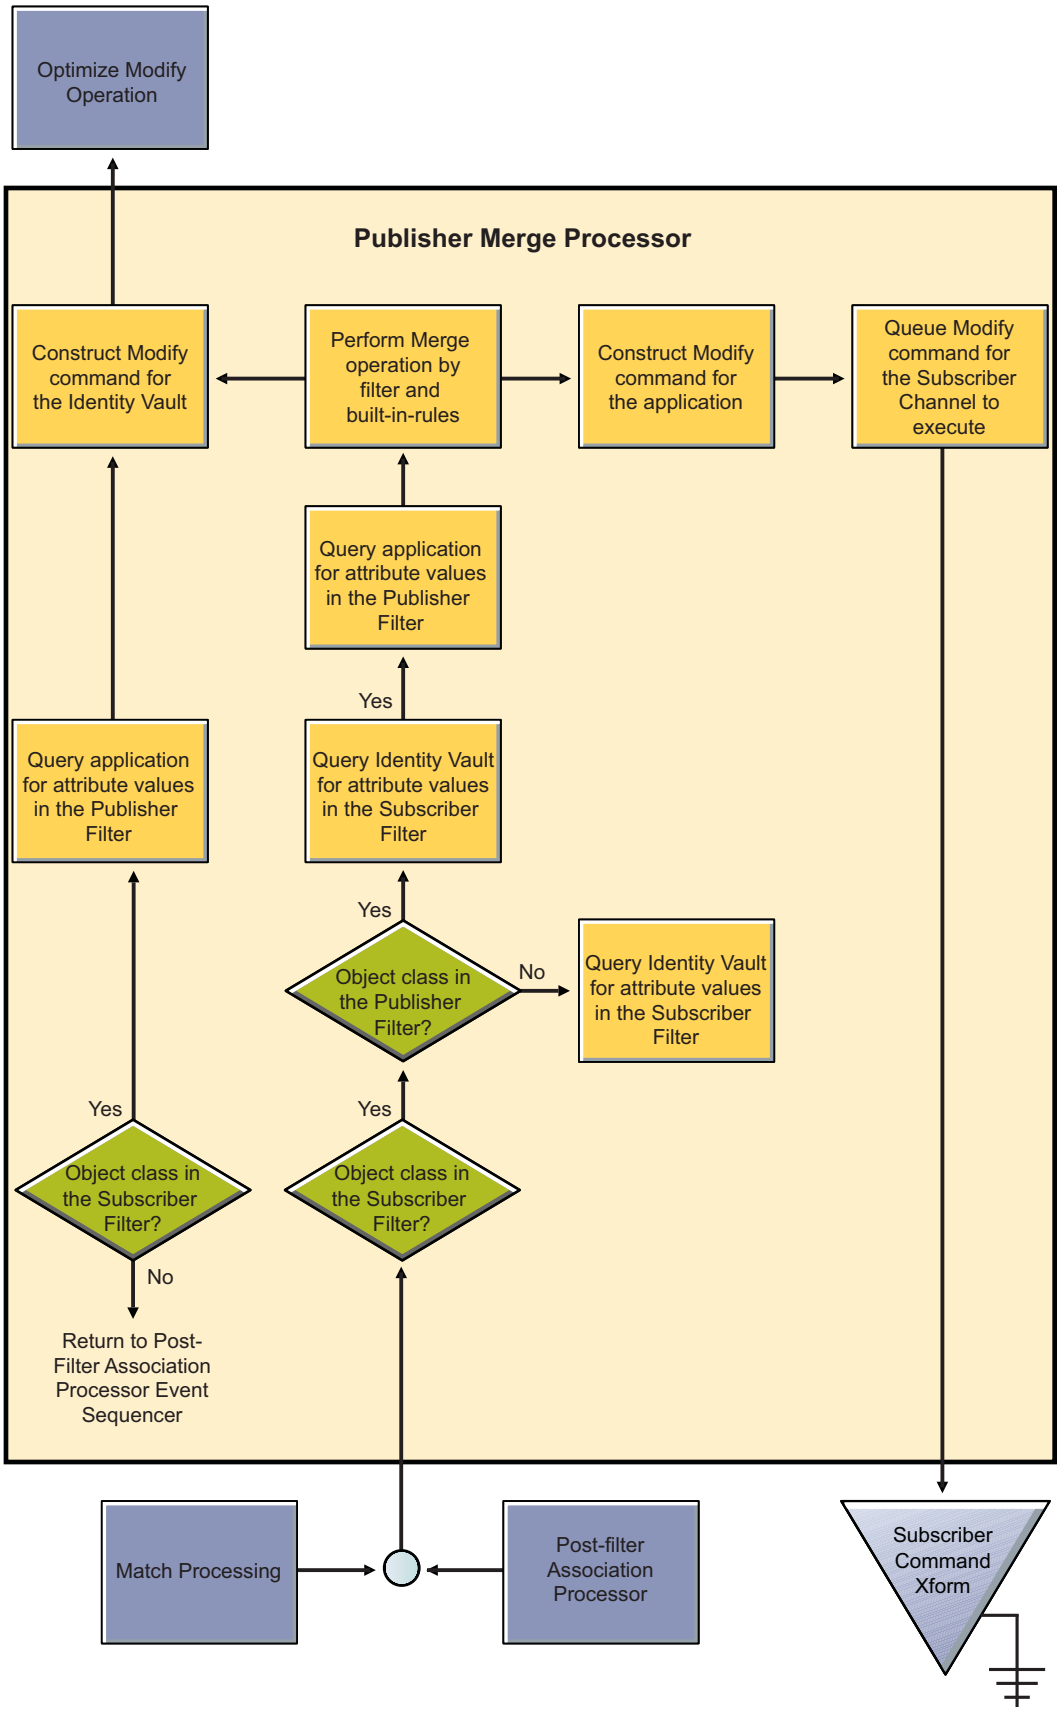
Optimize Modify Operation
Publisher Merge Processor
Construct Modify command for the Identity Vault
Perform Merge operation by filter and built-in-rules
Construct Modify command for the application
Queue Modify command for the Subscriber Channel to execute
Query application for attribute values in the Publisher Filter
Yes
Query application for attribute values in the Publisher Filter
Query Identity Vault for attribute values in the Subscriber Filter
Yes
Object class in the Publisher Filter?
No
Query Identity Vault for attribute values in the Subscriber Filter
Yes
Object class in the Subscriber Filter?
Yes
Object class in the Subscriber Filter?
No
Return to Post-Filter Association Processor Event Sequencer
Match Processing
Post-filter Association Processor
Subscriber Command Xform

Figure 6-6 *Subscriber Merge Processor*

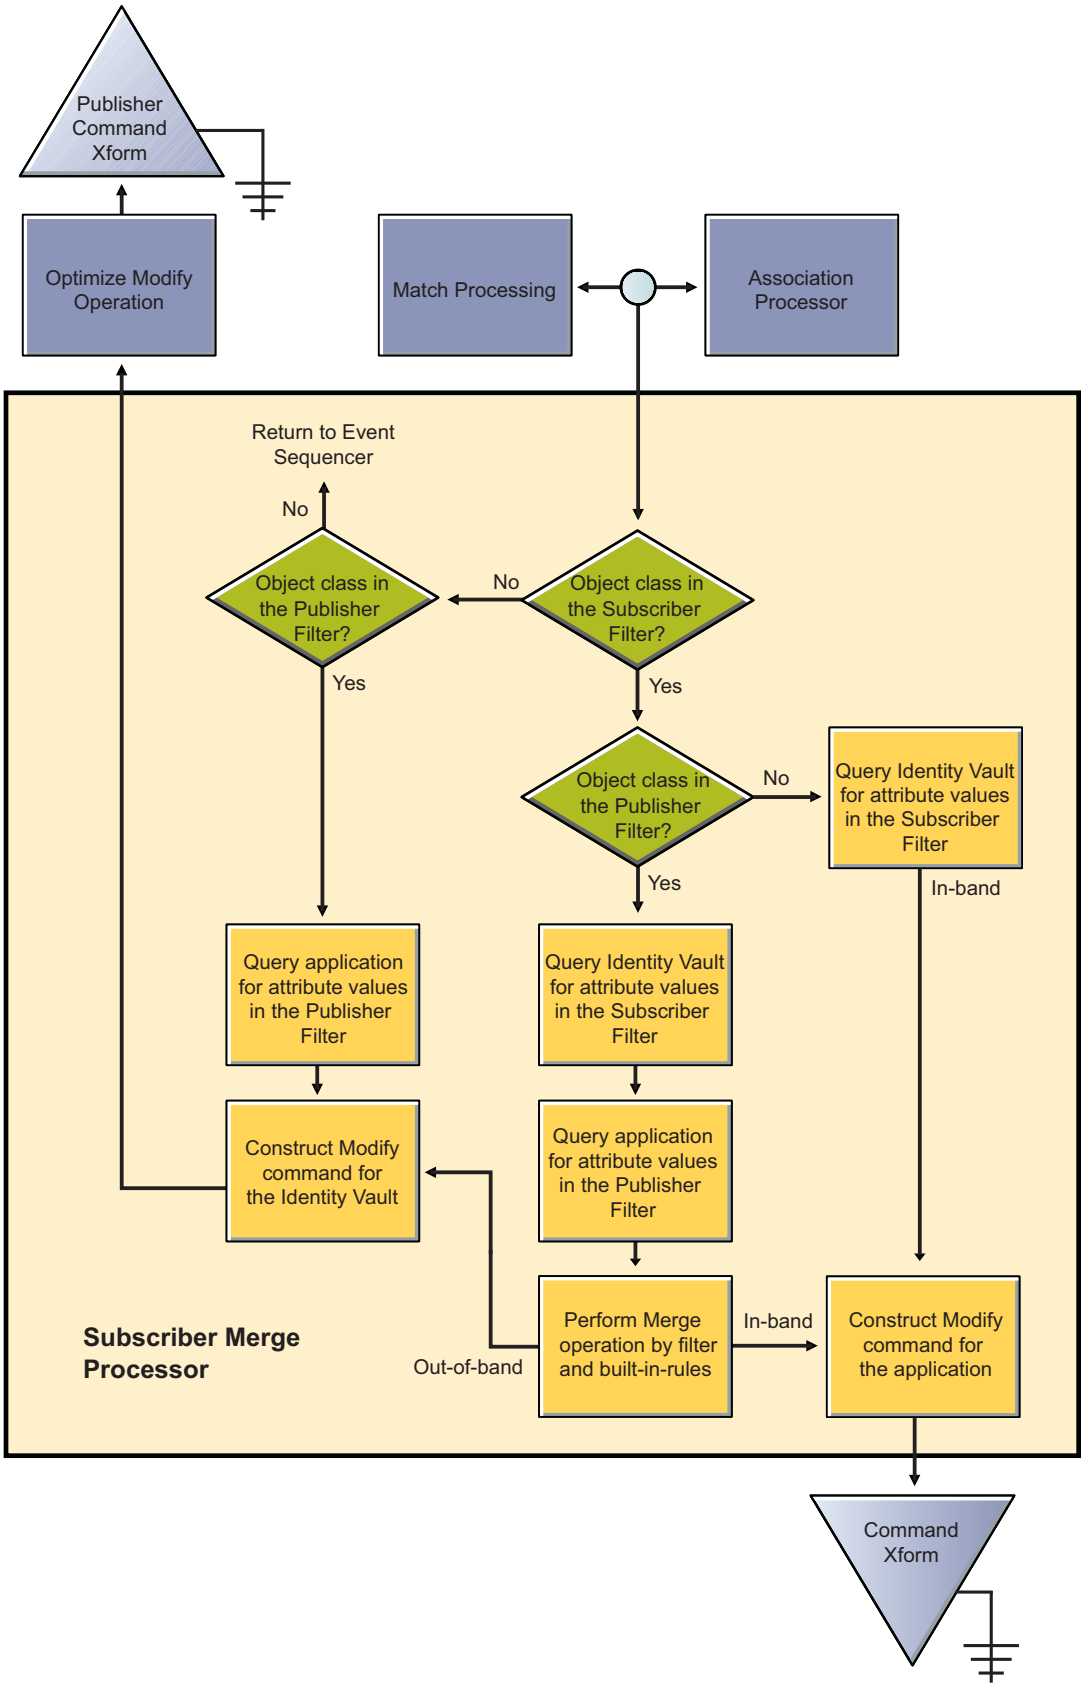

7 What's Next

After you understand the components that make up Identity Manager 4.0.2, the next step is to use the documentation to create your Identity Manager solution. The following sections explain where to find the documentation for the tasks listed:

7.1 Planning an Identity Manager Solution

The first step in designing an Identity Manager solution is to decide exactly what you want your solution to do in your business. Use the "Planning" section in the *Identity Manager 4.0.2 Framework Installation Guide* to create a plan for your Identity Manager solution by using Designer. You can also design your User Application solution by using the *User Application: Design Guide*.

Designer allows you to capture information into a project and share the information with other people. You can also model the solution in Designer before you start making changes. For more information about Designer see *Understanding Designer for Identity Manager*.

7.2 Preparing Your Data for Synchronization

After you create your plan, you need to prepare the data in your environment for synchronization. Analyzer is the tool you use to analyze, clean, and prepare the data for synchronization. For more information, see the *Analyzer 4.0.2 for Identity Manager Administration Guide*.

7.3 Installing or Upgrading Identity Manager

When you have created your plan and prepared the data, you can install Identity Manager. If you have a small to medium IT environment and you haven't used Identity Manager before, it is best to use the integrated installer. The integrated installer installs and configures all components that come with Identity Manager. For more information, see the *Identity Manager 4.0.2 Integrated Installation Guide*.

If you have an existing Identity Manager system or you have a large IT environment, use the *Identity Manager 4.0.2 Framework Installation Guide* to install or upgrade the different Identity Manager components. Each Identity Manager manager component is installed and configured separately, so you can customize your Identity Manager solution.

- For installation instructions, see "Installation" in the *Identity Manager 4.0.2 Framework Installation Guide*.
- For upgrade instructions, see "Performing an Upgrade" in the *Identity Manager 4.0.2 Upgrade and Migration Guide*.
- If you are migrating an existing system to new hardware, see "Performing an Upgrade" in the *Identity Manager 4.0.2 Upgrade and Migration Guide*.
- If you need to migrate the Roles Based Provisioning Module, see the *Identity Manager 4.0.2: RBPM and Reporting Migration Guide*.

7.4 Configuring Identity Manager

After Identity Manager is installed, you must configure different components to have a fully functioning solution.

- Section 7.4.1, "Synchronizing Data," on page 68
- Section 7.4.2, "Mapping Roles," on page 68
- Section 7.4.3, "Configuring the User Application," on page 69
- Section 7.4.4, "Configuring Auditing, Reporting, and Compliance," on page 69

7.4.1 Synchronizing Data

Identity Manager uses drivers to synchronize data between different applications, databases, operating systems, and directories. After Identity Manager is installed, you need to create and configure one or more drivers for each system you want to synchronize data with.

Each driver has a documentation guide that explains the requirements and configuration steps required to synchronize data. The driver guides are located at the Identity Manager 4.0.2 Drivers documentation Web site.

Use the specific driver guide for each managed system to create a driver to synchronize identity data.

7.4.2 Mapping Roles

When you have information synchronizing between the different systems, use the Role Mapping Administrator (RMA) to manage the roles in the different systems. For more information, see the *Novell Identity Manager Role Mapping Administrator 4.0.2 User Guide*.

7.4.3 Configuring the User Application

The next step is to add a business perspective to the Identity Manager solution with the User Application. The User Application enables you to address the following business needs:

- Providing a convenient way to perform roles-based provisioning actions.
- Ensuring that your organization has a method for verifying that personnel are fully aware of organizational policies and are taking steps to comply with these policies.
- Providing user self-service, allowing a new user to self-register, and providing access to anonymous or guest users.
- Ensuring that access to corporate resources complies with organizational policies and that provisioning occurs within the context of the corporate security policy.
- Reducing the administrative burden of entering, updating, and deleting user information across all systems in the enterprise.
- Managing manual and automated provisioning of identities, services, resources, and assets.
- Supporting complex workflows.

The *User Application: Administration Guide* contains the information on how to configure these features of the User Application.

7.4.4 Configuring Auditing, Reporting, and Compliance

The last and most important step in creating your Identity Manager solution is to configure the auditing, reporting, and compliance features so you can verify that your solution is in compliance with your business policies. Use the following guides to set up and configure these features:

- **Auditing:** See the *Identity Manager 4.0.2 Reporting Guide for Novell Sentinel*.
- **Reporting:** See the *Identity Reporting Module Guide* and *Using Identity Manager 4.0.2 Reports*.
- **Compliance:** See "Using the Compliance Tab" in the *User Application: User Guide*.

7.5 Administering Identity Manager

After your Identity Manager solution is complete, there are many different guides that help you administer, maintain, and change your Identity Manager solution as your business changes and grows. The different administration guides are located on the Identity Manager 4.0.2 documentation Web site under the Administration heading.