

Highmark Unifies Identity Data With Oracle Virtual Directory

*An Oracle White Paper
January 2009*

Highmark Unifies Identity Data With Oracle Virtual Directory

Executive Summary	3
The Challenge: A Single Directory Service.....	4
Directory Servers to Directory Service.....	4
Deployment Overview	5
Deployment Scale	5
Hardware.....	5
High Availability.....	5
Architecture Diagram.....	6
Identity Data Access Configuration.....	7
Unifying Identity Data.....	7
Removing Duplicate Identities with UniqueEntry Plug-in.....	7
In The Future: Providing a Single Profile with Join Adapter.....	8
Implementation Process	8
Benefits and Return On Investment	8
Technical Benefits.....	8
Return On Investment.....	9
Conclusion.....	9

Highmark Unifies Identity Data with Oracle Virtual Directory

EXECUTIVE SUMMARY

Oracle Virtual Directory (OVD) has enabled Highmark to reduce the time it took to deploy its new access management solution and portal applications. It also enabled Highmark to clean up its identity data without needing to change the existing source systems.

Highmark Inc. is an insurance company based in Harrisburg, Pennsylvania. They provide services to over 2 million customers including many Web-based applications.

As part of deploying new Web-based applications they adopted a new access management solution (Oracle Access Manager) and the Web applications were deployed upon a new Web-application infrastructure (IBM Websphere application server and portal).

One of the challenges of deploying OAM and Websphere was that identity information was split between two RedHat Directory servers. One RedHat server contained employee attributes and the other contained non-employee credentials for customers and vendors. However, WebSphere and OAM, like other commercial off the shelf (COTS) applications, can only support one user directory. There are also future plans to make use of attributes that exist within a Microsoft Active Directory containing employee information.

Highmark wanted a solution that would allow them to access all of the identity information in these directories as a single source without consolidating the identities into a single directory.

Highmark deployed OVD to provide a single unified interface to all of the directories to avoid consolidation. It allowed them to quickly go into production by providing a unified view of data from both of their Red Hat directories. It also has the flexibility to add in the Active Directory data in the future when they deploy applications that need to make use of data stored in AD without needing to change any of the other existing applications. OVD also eliminated duplicate user identities without requiring any data changes in the back-end servers.

OVD has enabled Highmark to deploy their access management solution to over 100 applications servicing 2 million user accounts in a short time, provided a quick and high return on investment (ROI).

THE CHALLENGE: A SINGLE DIRECTORY SERVICE

There were three primary challenges facing Highmark:

- Accounts in multiple directories (including internal and extranet directory)
- Select number of duplicate accounts between internal and extranet directory
- Applications required a single directory service access point

The identity information is stored in the following repositories:

- RedHat Directory (current production)
- Microsoft Active Directory (planned for future)

There are two RedHat Directories. One RedHat Directory is for staff accounts that contain attributes that are used for applications but not stored in AD. The second directory (Extranet Directory that contains over 2 million user accounts) primarily contains customer accounts but does contain a few staff accounts. The staff accounts cannot be removed because there are still legacy applications that connect only to this data. Active Directory is used for Windows logins but is not yet used with OVD because they do not yet have applications that need the data in Active Directory.

Directory Servers to Directory Service

Highmark needed to simplify their application deployment configuration by providing a single point of contact for all of their identities.

However, consolidating all of the data into a single “super-directory” would be too costly and time consuming. Due to continued support requirements for legacy applications, they could not remove existing directories, and adding another “super-directory” for consolidation would add more infrastructure and more data synchronization to maintain and support.

More importantly, Highmark had existing processes for provisioning accounts into these various systems that were certified for required regulations. If they were to consolidate and synchronize data, besides all of the technical challenges, the regulatory/security rules would have to be re-implemented & certified.

Further complicating the situation was the fact that some users had legitimate accounts in all 3 directories and for various reasons the triple-duplicated accounts could not be removed.

After evaluating OVD, the Highmark IT staff was convinced that OVD provided an easy to use, flexible solution that could be quickly implemented without the need for any new synchronization requirements.

Highmark chose Oracle to solve these problems because OVD:

- Allowed them to unify identities without needing to copy data into a new repository
- Leveraged all of their existing data repositories and certified processes
- Provided the ability to join data and remove duplicate accounts without modifying the source data
- Could be used by new applications such as access management as well as the 100+ legacy LDAP enabled applications

DEPLOYMENT OVERVIEW

Deployment Scale

- Supports over 2 million accounts
- Over 100 applications including IBM WebSphere and OAM

Hardware

Production: 4 x 2 x86 CPU running Linux with hardware load-balancer

Sandbox: 2 x 2 x86 CPU running Linux

Test: 2 x 2 x86 CPU running Linux

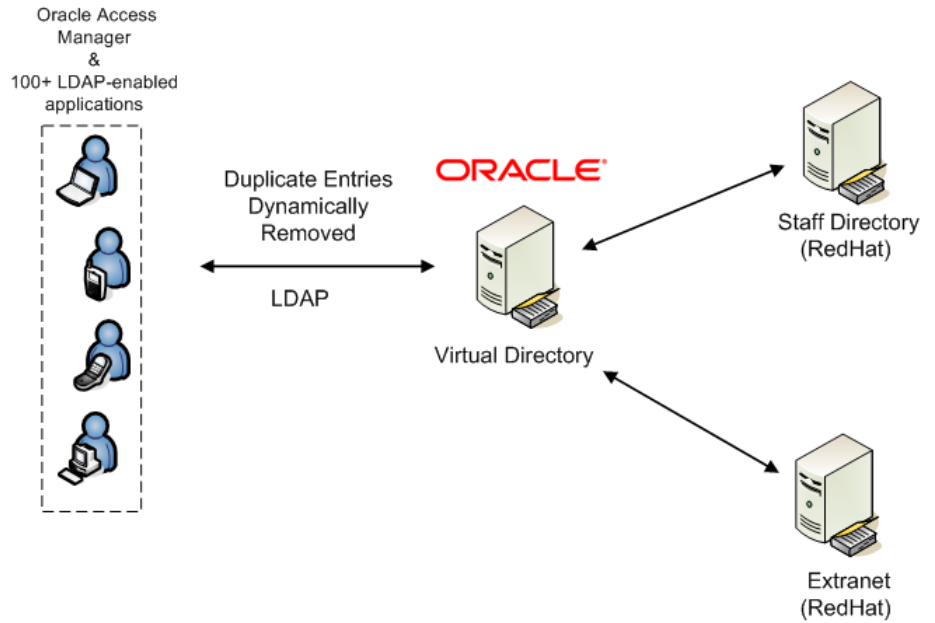
High Availability

The production systems are configured to be highly available by having duplicate configuration and using a hardware load-balancer to balance load as well redirect traffic if an OVD server were unavailable for any reason.

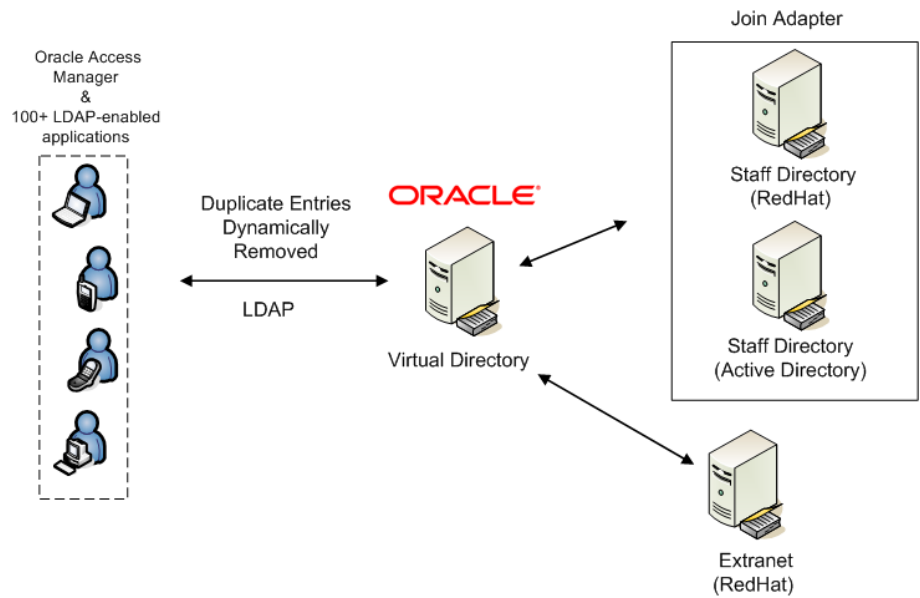
Additionally, OVD LDAP adapters are configured to point to redundant LDAP servers. This includes the ability to send requests to sources located in the primary data-center, but will fail-over to remote facility automatically if LDAP servers in primary facility are unavailable.

Architecture Diagram

The following is the high-level architecture diagram showing Highmark's OVD current deployment.



The following diagram shows Highmark's planned future deployment



Identity Data Access Configuration

Unifying Identity Data

As explained in the earlier sections, there are 2 LDAP sources for identity used today with a 3rd planned for the future.

Highmark configured OVD in the following configuration:

Local Store Adapter to hold a root entry (e.g. dc=highmark,dc=com) – this entry is static and is used as the search-base by applications.

There are two visible adapters to applications – one for staff and one for external. Both appear as virtual branches (e.g. ou=staff,dc=highmark,dc=com and ou=external,dc=highmark,dc=com). This way the data can be searched in parallel by any search that specifies a scope of subtree and a search base of “dc=highmark,dc=com” which is how most applications search an LDAP server.

By configuring OVD this way – applications are able to see data in all of the directories as a single source.

Removing Duplicate Identities with UniqueEntry Plug-in

The Staff and Extranet directories have overlapped user data for a subset of staff. OVD supports using a Join adapter to link identity data when it is split, however, in this case, Highmark only wanted to expose the entries in the Staff directory, not any data for the staff member contained in the Extranet.

Thus Highmark used the UniqueEntry plug-in, which removes entries based on duplicate attribute values. At Highmark the value they chose was uid.

Highmark configured OVD so that whenever one of their applications searched OVD and returned multiple entries – if there were any entries had the same uid value as a previous entry, that entry was not returned to the client application.

This way Highmark was able to avoid having to change back-end data-stores, which was crucial because there were still applications that depended upon that data.

In The Future: Providing a Single Profile with Join Adapter

Within Highmark, staff entries can have split-profiles. A split-profile is where identity attributes are split between two or more sources. In Highmark's case they need to reconcile user data in the Active Directory and a RedHat directory.

OVD provides the Join adapter, which allows the identity attributes to be unified on the fly using a join rule. There are several pre-defined Join rules and Highmark used the pre-defined SimpleJoinRule. The SimpleJoinRule allows entries between the primary and secondary source to be linked together when the value between listed attributes are the same.

For example, if the mail attribute in the secondary is the same as uid in the primary then the join rule would be "mail=uid". Or it can be the same attribute like this "uid=uid".

The benefit of the Join adapter is that when an application retrieves a staff entry, the application sees all of the attributes (assuming it has proper rights to the data) regardless which source they came from.

Implementation Process

The implementation of OVD was completed by Highmark. They did not hire anyone for third-party assistance.

BENEFITS AND RETURN ON INVESTMENT

There were several benefits and a significant return on investment for choosing OVD over a meta-directory approach.

Technical Benefits

- Shorter deployment time – OVD was up and running in a month.
- Eliminated duplicate accounts in real-time without changing data at sources.
- Increased flexibility through decoupling applications from target data sources and through data abstraction and transformation.

Return On Investment

OVD delivered high and quick ROI as a result of short OVD deployment time, accelerated application deployments, and re-use of existing data-stores.

For example, a meta-directory approach would have cost much more. Besides the obvious software costs more hardware would have been needed for storage, backups and monitoring of the synchronization. Additionally, the data would need to be cleaned and normalized before synchronization.

This process would have taken at least 18 months to deploy.

In comparison, OVD was installed and ready within 1 month.

This is because OVD did not require additional storage or data backups. Avoiding additional data storage and backups also provided additional cost savings.

Quick OVD deployment also accelerated application deployments including Websphere, OAM, etc.

CONCLUSION

Highmark is an insurance company that faced a directory services challenge.

They had multiple directories but applications required a single directory service. Additionally, they needed to link attributes from two of the directories to create individual virtual entries, while remove duplicate entries without changing the original data sources.

OVD provided the solution. It allowed them to unify their identity data without needing to consolidate and it allowed them to clean their data without changing original data source. OVD accelerated Highmark application deployment, eliminated additional regulatory certification requirements, and thus provided quick and high ROI.

For more information about Oracle Virtual Directory please see

<http://www.oracle.com/identity>.

ORACLE FUSION MIDDLEWARE

Oracle Virtual Directory Case Study: Highmark, Inc.

January 2009

Author: Mark Wilcox

Oracle Corporation
World Headquarters
500 Oracle Parkway
Redwood Shores, CA 94065
U.S.A.

Worldwide Inquiries:

Phone: +1.650.506.7000

Fax: +1.650.506.7200

oracle.com

Copyright © 2005, Oracle. All rights reserved.

This document is provided for information purposes only and the contents hereof are subject to change without notice.

This document is not warranted to be error-free, nor subject to any other warranties or conditions, whether expressed orally or implied in law, including implied warranties and conditions of merchantability or fitness for a particular purpose. We specifically disclaim any liability with respect to this document and no contractual obligations are formed either directly or indirectly by this document. This document may not be reproduced or transmitted in any form or by any means, electronic or mechanical, for any purpose, without our prior written permission.

Oracle, JD Edwards, PeopleSoft, and Retek are registered trademarks of Oracle Corporation and/or its affiliates. Other names may be trademarks of their respective owners.