



# Agilysys...

Identity Simplification and  
Consolidation through  
Virtualization and Synchronization

Jeff Smith  
Software and Services Engineer  
Agilysys Sun Technology Solutions

## Summary

Organizations that hope to simplify their identity and access management (IAM) infrastructures can begin at the data layer where identities are stored. Directory virtualization provides new opportunities to move towards this goal by abstracting disparate sources of identity data and presenting them in a unified fashion, all while providing the needed flexibility to consume identity data as it is needed by other parts of the IAM stack. By synchronizing appropriate pieces of data in the identity store layer, an organization can reduce complexity, increase data reliability and reduce risk.

## Challenges in the IAM Data Layer

At the foundation of any IAM infrastructure is the identity store layer. This layer is comprised of the systems where identity information is stored. This “store” is not one monolithic data store, but is composed of multiple applications, databases and directories. As is the case for any application or system, IAM applications such as web access management (WAM) are only as good as the underlying data on which they depend and operate. Also needed is the capability of accessing this data in a unified manner, and in multiple structures.

In the early days of lightweight directory access protocol (LDAP) adoption, one of the benefits sought was a solution to what was called the N+1 directory problem; every application added to the network introduced a new repository of users which had to be separately provisioned and managed. LDAP directory services could help mitigate this problem by providing a centralized directory service which would store all of the user identities. Web and other network applications could use these services to lookup and authenticate users and retrieve other attributes of the user, freeing them of the need to maintain their own user repositories.

While this goal has been partially realized by deploying LDAP directory servers, the vision of one single enterprise directory for all identity data has proven to be overly simplistic. Although directory services have added much value in this area, in real-world implementations identity data still lives across multiple databases, directories and applications throughout the infrastructure. Meta-directory products followed in an attempt to automate the correlation and synchronization of data from multiple sources to form a master directory view, but meta-directories proved overly complex, inflexible and difficult to manage.

Another issue is that it is difficult to design a directory, keep it simple enough to be manageable, while still meeting the requirements of all the applications that need to consume its data. Various applications have the need to see differing data sets and/or to see that data in differing structures. For instance, an applications requirements of identity data for simple authentication will be different than that for evaluating policy based on entitlements, which may be nested somewhere in another application regional database management system (RDBMS) schema. Best practices in directory design usually call for a relatively flat directory tree structure as opposed to deeply nested hierarchies. Often the temptation arises to house more and more information in the directory, to accommodate specific applications, and with that, good directory design can be broken.

**“...[with LDAP] in real-world implementations identity data still lives across multiple databases, directories and applications throughout the infrastructure.”**


As the sophistication and scope of IAM technology has evolved, further needs have arisen to access data on the user that is either not practical or impossible to store in a central directory. For example, access management and role management solutions may need to access role, group and entitlement information that are stored deep in the structure of a complex RDBMS or enterprise resource planning (ERP) system.

## The Virtual Directory Solution

Most of us have a conception of what virtualization means, however it's always helpful to be clear on definitions. About.com's definition is sufficient for the purpose of this discussion: "Virtualization is a technique for hiding the physical characteristics of computing resources to simplify the way in which other systems, applications, or end users interact with those resources."

While it may be impossible or undesirable to physically merge or consolidate the aforementioned multiple identity sources, building a virtualization layer to front them can effectively combine them to appear unified to other components of the identity infrastructure as well as other applications. When applied properly, the virtual directory approach reduces complexity, increases the flexibility of the directory service and — when combined with data correlation and possibly synchronization among the data stores — increases data reliability (therefore reducing risk).

Virtual directories allow for flexibility; identity information from multiple sources can be combined, such as from existing directories, databases and applications, structured in varying ways as needed, and made accessible to applications via common directory access protocols such as LDAP and directory services markup language (DSML) (see *Figure 1*). It still allows us to build directory servers behind the virtualization layer to gather and consolidate identities, but those directories can be designed to their strengths without running the risk of being made overly complex to the point of breaking best practices for directory design (e.g. flatter directory tree structures).



*Figure 1. A virtual directory layer fronting identity sources.*

Multiple views of the identity data through a virtual directory layer offer immediate benefits. Applications that need to see a flat tree structure for fast lookup or authentication can have that view, while other views can easily be built to serve applications that need additional pieces of data that are sourced deeper in relational sources and/or have needs for data presented in a more hierarchical structure. Disparate directories [for example Active Directory (AD) and Sun directory] can be joined to provide an integrated virtual view for applications that need them. This is particularly elegant for joining directories that house separate user communities.

To an extent, the virtual directory solution helps mitigate the ongoing ‘directory vs. RDBMS’ debate. At the end of the day, each has its place and serves differing needs. The virtual directory allows each to exist as they should and allows us to practically and easily combine them when desired.

In addition to providing virtual views of disparate data, simplicity and reliability can be introduced in the identity store through reconciliation and linking of like identity data, as well as synchronization of selected identity attributes and fields. When one user has accounts in multiple sources, each with a different identifier, we need to be able to tie them together through the creation of a global identifier with pointers to the various accounts that belong to the same identity.

**“...when a user’s telephone number is updated once through a writeable virtual view, the view can in turn update three separate sources on the backend with the same piece of data .”**

Account correlation amongst identity stores can be accomplished in multiple ways. Some virtual directory platforms provide identity correlation tools to create a global ID to uniquely identify the user and tie together his/her disparate identities stored across multiple sources. Also a provisioning tool, such as Sun’s Identity Manager, provides account reconciliation features, and creates a “virtual user” entry in it’s repository which contains the unique identifiers to all the resources for which that user has accounts, and which are under the control of Identity Manager.

Through a thorough inventory and understanding of the identity stores, attributes can be identified which should be synchronized across multiple sources to reduce chances for making decisions based on bad data. With these pieces of data identified, this type of field/attribute level synchronization can be accomplished through virtual views. For example, when a user’s telephone number is updated once through a writeable virtual view, the view can in turn update three separate sources on the backend with the same piece of data. Other technologies, such as Sun’s Identity Synchronization for Windows can also be used to synchronize data between AD and Sun directories, or Sun directories and other LDAP compliant directories.

Provisioning can also be simplified through the virtual directory layer — where a provisioning solution may have adapters to multiple backend repositories, the number of adapters may be reduced by provisioning through virtual directory views.

WAM deployments can be made much simpler by the proper application of a virtual directory layer. WAM/single sign-on (SSO) deployments that need to coordinate authentication to many different identity stores can become quite complex.

## A Synchronization Success Story

Many organizations with large UNIX networks are in the process of moving to an LDAP directory for naming services. This migration can be from an environment currently using naming services such as Network Information Service (NIS or NIS+) or one using individual flat file configurations on each server. An Agilysys customer who planned to do such a migration had the additional requirement of having user credentials come from their corporate Active Directory. For over 300 UNIX servers, they wanted to keep the provisioning and ownership of the user identity data at their corporate AD. This would eliminate an additional provisioning point, and help in meeting audit requirements for Payment Card Industry (PCI) Data Security Standard compliance.

For this solution, a Sun directory server infrastructure was created to serve as the UNIX naming service. The Sun Identity Synchronization for Windows product (a component of Sun's Directory Server Enterprise Edition) was used to facilitate one way synchronization of selected users from AD to the Sun directory. Users are selected based on membership in an AD group to denote UNIX users. When a user is created, modified or deleted in the AD, the operation flows to the UNIX directory. The Identity Synch component also has a method for synchronizing AD passwords to the Sun directory without requiring any additional software installation on the Active Directory side (See Figure 2).

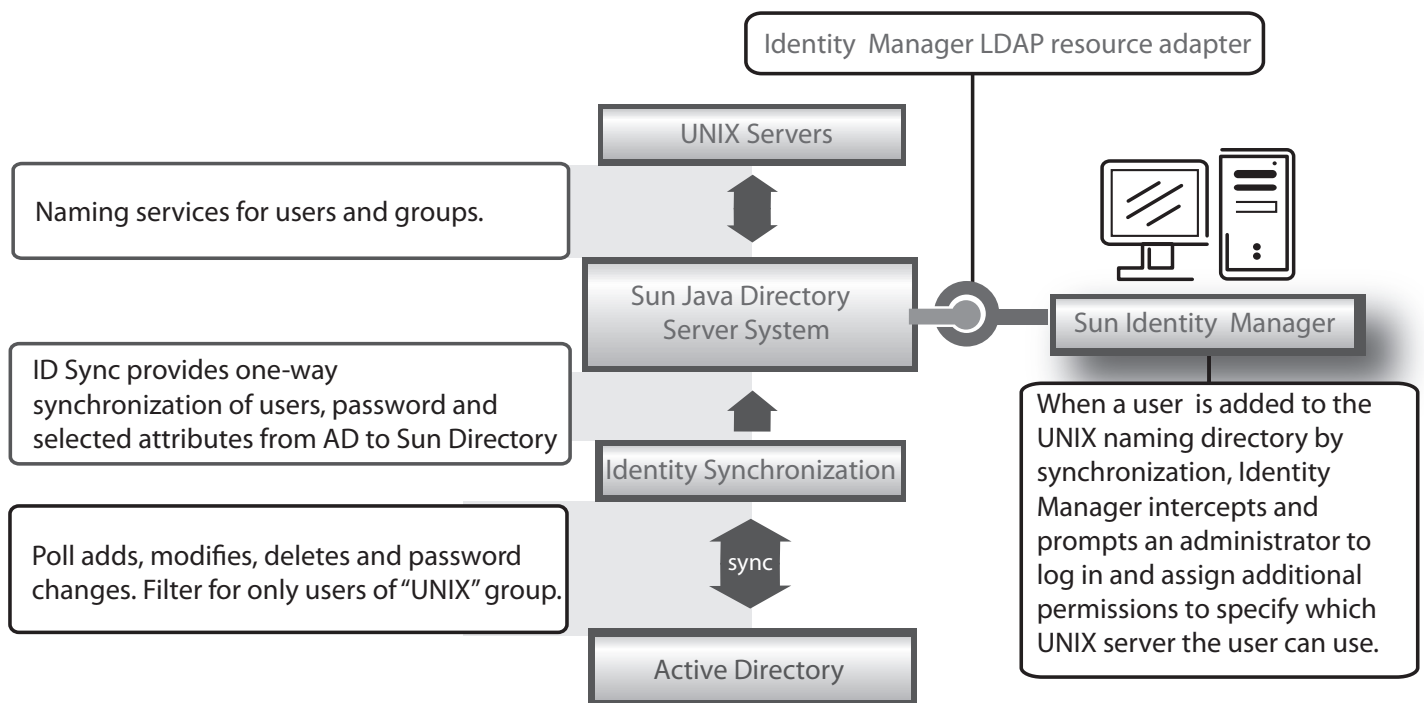


Figure 2. Identity synchronization of AD and Sun LDAP for UNIX user accounts.

This solution keeps the provisioning of users and passwords completely at Active Directory and provides consistency of logins between the two environments.

Additional provisioning in this environment also needs to take place in order to assign the UNIX users to specific UNIX servers or server groups. This stage of user administration is best done by the UNIX administrators, not AD administrators. This was accomplished by using the Sun Identity Manager, which picks up user creations and notifies an admin that the new user has been created in the UNIX environment. The administrator then logs in to an Identity Manager interface, in which he/she can give the user access to the appropriate UNIX servers. Besides facilitating access control for new user creations, the UNIX administrator can use Identity Manager anytime changes are required for server access.

In this solution, much of the provisioning to the UNIX infrastructure was eliminated, as it was leveraged from the AD, which is authoritative for corporate users. This guarantees a correlation between account IDs, passwords and other attributes between AD and the UNIX directory. The solution simplifies provisioning for UNIX administrators, simplifies the user experience (one login and password) and helps with compliance issues by concentrating the authoritative store for user access and credentials in one place.

## Conclusion

If you're planning on simplifying and consolidating an IAM infrastructure, start with the foundational layer of the identity store. Rather than search for the monolithic master directory, directory virtualization and synchronization offer ways to add simplicity, flexibility and unification to the identity data infrastructure while leaving the data where it most belongs.

A virtualization layer allows services in the underlying data store to be designed to their strengths – directories remain directories and databases remain databases – but are still able to deliver the data to applications the way they need to see it through common protocols.

Understanding and accounting for all the identity stores in the infrastructure is of paramount importance. Accounts belonging to the same person need to be identified and reconciled as such. Data that is redundant across multiple sources should be identified and evaluated for synchronization. Virtual directories can synchronize identity sources by updating through a virtual view which joins the sources. Also a synchronization product such as Sun's Identity Synchronization for Windows, can be leveraged to provide some reliable, easy to manage synchronization between different directories.

# Agilysys...

## Who We Are and What We Do

This white paper and the research survey behind it are sponsored by Agilysys, Inc., a leading provider of information technology (IT) solutions serving corporate and public-sector customers with special expertise in select markets, including retail and hospitality. Agilysys provides technology solutions—including hardware, software and services—to help customers resolve their most complicated IT needs. Our expertise includes enterprise architecture and high availability, infrastructure optimization, storage and resource management, virtualization, identity management and business continuity; along with software and services designed specifically for the retail and hospitality markets. We operate from locations throughout North America, and in the United Kingdom and China, with headquarters in Cleveland.

© Agilysys, Inc. 2009

Produced in the USA. All rights reserved.

Directory Server Enterprise Edition, Directory Server, Identity Synchronization for Windows and Identity Manager, and Active Directory logos and/or trademarks are trademarks or registered trademarks of Sun Microsystems and Microsoft, Inc. respectively.

Agilysys reserves the right to change specifications or other product information without notice. References in this publication to Agilysys products or services do not imply that Agilysys intends to make them available in all countries in which Agilysys operates. AGILYSYS PROVIDES THIS PUBLICATION “AS IS” WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. Some jurisdictions do not allow disclaimer of express or implied warranties; this statement may not apply to you.

This publication may contain links to third party websites that are not under the control of or maintained by Agilysys. Access to any such third party site is at your own risk and Agilysys is not responsible for the accuracy or reliability of any information, data, opinions, advice or statements made on these sites. Agilysys provides these links merely as a convenience and the inclusion of such link does not imply an endorsement.

All performance information was determined in a controlled environment. Actual results may vary. Performance information is provided “AS IS” and no warranties or guarantees are expressed or implied by Agilysys. Buyers should consult other sources of information including product benchmarks, if available, to evaluate the performance of a product they are considering buying.



*For more information about how Agilysys Sun Technology Solutions can help with your company's needs, call 732.692.1919 or visit [sun.agilysys.com](http://sun.agilysys.com).*