SAND2008-5098C

Protecting Sensitive Information in Directory Services

William Claycomb
Sandia National Laboratories
PO Box 5800, MS 0823
Albuquerque, New Mexico, 87185-0823
USA
Phone: 505-284-9949
Fax: 505-284-5619
Email: wrclayc@sandia.gov

Dongwan Shin
Computer Science Department
New Mexico Institute of
Mining and Technology
801 Leroy Place
Socorro, New Mexico, 87801
USA
Phone: 505-835-6458
Fax: 505-835-5587
Email: doshin@nmt.edu

***Abstract*— Directory services are commonly used to store information related to individuals within a corporation. Sometimes, they contain sensitive information, and need to be protected. Existing solutions offer minimal protection against insider attacks, a growing threat to both government and industry data services. We present a solution for data protection that leverages virtual directories, metadirectories, and data encryption to provide a user-centric approach to data protection, delegation, and collaboration. We explore the benefits and vulnerabilities of this solution, and present performance testing results to support our proposal.**

I. Introduction

Directory Services are used to store information about objects within an organization, such as users, computers, etc., and are organized in a hierarchical structure. Often, directory services (or simply, *directories*) are used as authoritative data sources for many applications that require user information, such as web portals, email, and instant messaging. In some cases, the information contained in a directory is considered confidential, such as employee id number, clearance level, or other *personally identifiable information (PII)*. While techniques exist to protect this information, they do not adequately prevent a user with administrative privileges from unauthorized access. Additionally, many companies have several directories, some containing duplicate information. This can arise from inadequate information planning, applications requiring proprietary data sources, or the need to protect specific information at different levels. Determining authoritative data sources and synchronizing data across directories is a challenging and ongoing task for many corporations.

Each object in a directory is described by a set of *attributes*. Examples include name, address, email, or manager name. In some cases, an object's attributes should be publicly available for others to use, such as name and/or email address. However, in other cases, attributes are used for a single application, or closely related groups of applications, and should not be available to others outside a set of authorized users. These attributes are necessary to perform important functions, such as payroll, or for access control decisions, such as clearance level, but should not be used for any unauthorized purpose. In these cases, options exist to protect the information. Access control lists (ACLs) or marking attributes *confidential* [1] are two ways to protect data from ordinary users. In general, however, directories are used to share information, and rarely contain access controls beyond simple user authentication (only users with accounts on the system may access the data). An insider threat, someone with authorized access, could potentially retrieve personal information about every object in the directory. The malicious activities possible with such information could include selling the information to competing companies, foreign governments, or spammers, or even worse - the targeted attack of specific individuals within the company, such as domain-level administrators, known as *context aware attacks*, or *spear phishing* [2].

We present an architecture for protecting individual attributes in directory services from unauthorized access. In standard configurations, clients communicate directly with directory services using the Lightweight Directory Access Protocol (LDAP). Clients connect to a specific port on a specific server, and may authenticate using various methods, including providing a username and password, if necessary. Our architecture places a third component in this standard configuration between the client and server, to handle LDAP communication between them. This third component relies on information provided by the client to encrypt sensitive information. While other solutions have proposed encrypting attribute information, our architecture provides this capability without requiring additional software or hardware on either the client or destination server.

The remainder of this paper is organized as follows. Section II presents an analysis of background material and related solutions. Section III outlines the motivation for our solution, including a description of the threats currently posed by insider attacks. Section IV describes our architecture in detail. In Section V, we analyze the results of our implementation

testing. Section VI discusses the architecture, including various advantages, as well as attack models. Section VI concludes this paper with a glimpse of future work.

II. Background and Related Work

A. Directory Services

Directories are collections of information related to objects in an organization. These objects often include users, computers, or contacts. *Directory Services* are the services which make this data available for use by others. Frequently, the intention is to provide a single point of access for various applications and individuals to find in-formation about users and other objects within an organization [1]. The information contained within the directory may come from direct input, and can be manually maintained, but also may be referenced and managed indirectly from other corporate data repositories, such as databases and other information stores. Commonly used directory services are Microsoft Active Directory [3], IBM Tivoli [4], Apple Open Directory [5], Novell eDirectory (formerly called Novell Directory Services) [6], OpenLDAP [7], Fedora Directory Server [8], and Sun Java System Directory Server [9].

B. Protecting Information in Directory Services

A few techniques exist for protecting the information stored within a directory itself. In general, access control lists (ACLs) can be used to implement some form of protection in most directories. For instance, in OpenLDAP, ACL protection can be applied to individual objects, groups of objects, specific LDAP filters, or a list of attributes [10]. Other techniques are almost exclusively implementation-specific.

Microsoft Active Directory [3] provides additional access control features through the use of *confidential attributes* [1]. This is a setting applied to the *searchFlags* component of individual attributes, and is only supported on Microsoft Windows Server 2003 SP1 and later. When processing confidential attributes, the directory server checks for additional access control rights associated with the requesting user. This particular type of access, called "CONTROL_ACCESS," is granted to administrative accounts by default, but can be delegated to other accounts individually.

Another approach to protecting attributes is encrypting them. Fedora Directory Server [8] has the capability to encrypt all instances of specified attributes. This means that for every object containing such attributes, the data in that attribute is encrypted using a symmetric key known to the directory server. Various encryption methods can be configured, and different attributes can be encrypted using different ciphers. Encryption and decryption are handled by the directory server itself, so access to attributes is not controlled by this method. However, data would be protected from unauthorized access if the directory data was stolen or otherwise compromised.

A third approach to protecting directory attributes is described in [11]. This method is not dependent on a particular directory implementation. Rather, it uses public key infrastructure (PKI) to allow users to control the encryption of attributes related to their own directory information. This solution describes different methods for using PKI to ensure either data authenticity alone, or data authenticity combined with confidentiality. Specific solutions are proposed for scalability and usability purposes.

Additionally, [12] proposes encrypting directory information for users based on a *unique-id* chosen for each user. This method applies primarily to public directory servers, and does not address the issue of preventing unauthorized access so much as it addresses the issue of preventing access to the entire directory. For instance, a company could share contact information publicly, and provide selected clients with appropriate unique-ids, without worrying that the entire directory would be scanned for email addresses. One important aspect of the work is to choose a unique-id well, so that it cannot be easily guessed, but can still be easily shared with authorized users.

C. Metadirectories and Virtual Directories

One way to protect personal information is to reduce the number of different data stores containing personal information. This can be a complicated task, particularly for businesses with many disparate data sources. The International Data Corporation (IDC) and Gartner Groups have found that large corporations may have in excess of 100 data stores containing user information [13]. Additionally, proprietary systems often do not interact with other data sources. Consolidating data into a single data source is often not possible, due to constraints on who may have access to specific information. Technology has emerged to address these issues, specifically by referencing the underlying data sources and presenting end-users with customized views of the data they require, and by synchronizing data between different data sources. Two similar but distinct methods of handling these tasks are *metadirectories* and *virtual directories*.

1) Metadirectories: Analyzing the origin of the word "metadirectory," we see the Greek phrase "meta-" which means "after," or "beyond." In modern English, this term often describes abstraction. Thus, a metadirectory is an abstraction of an actual directory. In this sense, it acts as a directory in some instances, by providing user interaction via LDAP, but does not act as a directory in other instances, because it is not the actual authoritative source of directory information. A metadirectory is used to abstract data from other directories into a single source, which can be used for two purposes.

The first use is for end user reference. Users may access the data collected by a metadirectory via LDAP. In particular, this not only reduces the number of data sources an end user connects to, but enables customization of directory data for individual uses. Therefore, in one sense a metadirectory is a real directory - information is actually stored locally, and is queried directly by end users.

However, this repository is not the authoritative source of the data. Rather, data synchronization must occur between the metadirectory and source directories to ensure consistency and accuracy of the data. It is the synchronization of data

which is its second purpose. When different data sources must store the same information, it is desirable to have a single source of authoritative data, which is synchronized with other data sources. For example, if the HR department is the authoritative source for a user's telephone number, but the company directory application, which uses its own data source, also requires a telephone number to be stored, a metadirectory could be used to automatically synchronize the data from the original source (HR). A more advanced use of metadirectories is for *user provisioning*, which is a modified version of synchronization, where new user accounts are created and prepared for use, based on data in an authoritative source, such as an account database. Examples of metadirectory implementations include Microsoft Identity Lifecycle Manager 2007 [14], Sun ONE Meta-Directory [15], and Critical Path Meta-Directory Server [16].

2) Virtual Directories: "A virtual directory functions as an abstraction layer between applications and data repositories." [17] In contrast to metadirectories, *virtual directories* do not maintain the data in a standalone data source (though some offer *data caching*, which does store data locally for improved performance). Rather, virtual directories reference various data sources and present a consolidated view to the end user. This has the advantage of not requiring data synchronization - the data presented is always real-time, directly from the source. Most virtual directory implementations have the additional capability of acquiring data from sources other than directories, such as databases, and presenting this information to end users via LDAP.

Virtual directory instances can be highly customized to modify, or *transform*, data prior to client use. Additionally, some products offer data synchronization as well, which when coupled with a virtual directory instance, could be used for user provisioning in much the same way as a metadirectory. Virtual directory products currently in use include Radiant Logic's RadiantOne [18], Symlabs Virtual Directory Server [19], and Oracle Virtual Directory [20].

III. Motivation

The threat of unauthorized access of sensitive data by employees or other authorized users, known as "dedicated insiders", is well documented [21], [22], [23]. While the psychology and behavioral factors of these individuals is beyond the scope of this paper, their motivation and level of access should be considered. Additionally, it should be noted that the number of offenses committed by insiders is rising each year [23]. In January 2008, the U.S. Secret Service and CERT issued a report titled "Insider Threat Study: Illicit Cyber Activity in the Government Sector" [14]. This study out-lines a multi-year project, started in 2002, that explores the activity and threats posed by insiders, defined as "employees who have perpetrated acts of harm against an organization via computer, system, or network to include theft of intellectual property, fraud, and acts of sabotage within critical infrastructure sectors." Among the key findings of the study that are relevant to this paper are the following:

- Most of the insiders had authorized access at the time of their malicious activities
- Access control gaps facilitated most of the insider incidents, including:
  - The ability of an insider to use technical methods to override access controls without detection
  - System vulnerabilities that allowed technical insiders to use their specialized skills to override access controls without detection

In addition to outlining the methods and characteristics of the unauthorized access, the study also details findings about the motivation of the insiders, as well as the scope of the problem. In particular, the study notes that "in many cases insiders used authorized access to alter or obtain an individual's personal data in some manner." Theft of personal data was noted as a particularly likely target of insider threats, most often to sell to others for financial gain. The study notes that this is useful in "understanding how access to identity-related data might contribute to insider activity in [the government sector]." Additionally, it was noted that "agencies at all levels of government were targets of insider threat," and that the attacks were successful because of "similar vulnerabilities in business practices and access controls." [21] To address these concerns, the study also presented considerations for government agencies with respect to the protection of data, including the following:

- Electronic storage of citizens' confidential information necessitates accurate, reliable, and confidential record keeping within government databases and computer systems. Policies and technical controls are implemented to provide a safety net for critical data.
- Government agencies at all levels need to remain vigilant against the potential impacts of insider incidents on public trust and the citizens' confidence in government services
- Government agencies should have proactive strategies to protect information entrusted to them
- Federal agencies are required to comply with Title III of the E-Government Act of 2002 known as the Federal Information Security Management Act.

IV. Our Approach

Our approach to protecting sensitive information in directory services is to encrypt that information using user-controlled keys and to provide access to that data using user-controlled delegation. This user-centric approach follows current trends in computer security and privacy, but should not interfere with more traditional approaches to access control. Our approach also maintains usability with existing client applications and source directories. To better understand the overall picture of our solution, it is first important to understand various key components.

A. Encryption

Encrypting sensitive information to protect it from misuse is hardly a new concept. In the simplest application towards protecting information in directory services, an end user

Fig. 1. Basic Client Encryption to LDAP Directory

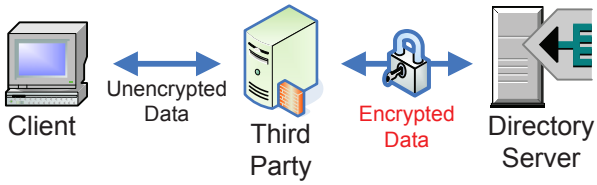

Fig. 2. Basic Third Party Encryption to LDAP Directory

would simply encrypt sensitive information and then store the encrypted data in a directory. To share information, the user would share the encryption/decryption key with another user, who would obtain the encrypted form from the directory and decrypt it locally. This is shown in Figure 1.

However, this approach presents several usability and security problems. First of all, the integrity of the data relies entirely on the shared key. If a malicious user were to obtain this key, or if an authorized user were to share it with an unauthorized party, the information could be compromised. Data integrity could be provided by using an asymmetric encryption algorithm, such as RSA, but this still does not protect the data from unauthorized access.

Secondly, this approach requires the user to perform encryption and decryption before and after retrieving the information from the directory. At best, this could be accomplished by a custom application, which interfaces directly with the client LDAP application. At worst, existing client LDAP applications would need to be rewritten to incorporate encryption and decryption. This is an undesirable situation for which a simple solution exists: add a third party, between the client and server, to handle encrypting and decrypting the data. This is shown in Figure 2.

The third party component could be a custom component, a *proxy*, written specifically for the purpose of handling encryption and decryption of information between the client and directory. However, we find it much more useful to leverage the existing technologies of virtual directories and metadirectories to provide the third party component to the model. The benefits of doing so are numerous, and will be discussed in detail later.

B. Using Virtual Directories

If we consider a virtual directory as the third party - the component responsible for encrypting and decrypting data - we must consider several key aspects, namely: how does the virtual directory obtain key information from the client, how does the virtual directory perform pass-through authentication to destination directories, and how does the virtual directory manipulate the data in the destination directory?

1) Obtaining client key information: When LDAP communications occur between a client and server, several standard pieces of information are transmitted. These components are generally configured by the client application, and can be changed by the end user. They are: username, password, and destination server name and port. We leverage these components to pass encryption information to the server as follows. The destination server and port, D_s, are replaced with the destination server name and port of the virtual directory. This configures the client to communicate with the virtual directory, instead of the original destination directory. The password remains the same as the original password used to authenticate to the original destination directory. We replace the username component with a string which is the concatenation of the following: the original destination directory name, D_s, the client username, U_c, the hash of the original user password, $\{P_c\}_H$, and a client key, S_c. The last two components are encrypted using a secret key known only to the virtual directory server, S_V. The addition of these last components requires additional setup, performed by another application with access to the virtual directory key, S_V, and is also discussed in detail later. The resulting string is called an *authentication string*:

$$U_c|D_s|\{S_c|[\{P_c\}_H]\}_{S_V}$$

2) Performing pass-through authentication: We do not ignore traditional authentication and access control methods with this solution. Unless configured for anonymous authentication (also called *anonymous bind*), the destination LDAP server will expect a client to authenticate prior to data retrieval. Some Virtual Directory implementations allow a static username and password to be used for every transaction, but this defeats the purpose of fine-grained access control. Rather, we will pass the original client username and password, obtained from the authentication string and password provided by the client, to perform an initial bind prior to data retrieval. If this bind is not successful, then no data transmission occurs between the virtual directory and the client.

3) Storing the data: Once the user has successfully authenticated to the destination directory, we use the transformation capabilities of the virtual directory instance to extract the user's secret key, S_c, and password hash, $\{P_c\}_H$. The password hash is used as an additional measure of security against an attack where a malicious administrator may change the user's password and, using the original authentication string, masquerade as the user. While this step may seem redundant, it is necessary because of the nature of LDAP clients. Many LDAP clients allow users to cache login information, including the username. An attacker would need to have no knowledge of the client secret key, S_c, if he used a cached authentication string and a newly-reset password. However, if the client were configured to prompt for a password every time, while still

retaining a cached authentication string, the user's password hash could be checked against the password hash encrypted by the virtual directory's secret key, $\{\{P_c\}_H\}_{S_{pa}}$. In this instance, a changed user password would cause a failure, because its hash would not match the original hash in the authentication string.

Once verified, the user's secret key, S_c, is used to perform encryption or decryption of data stored in the directory. The protocol for reading an encrypted attribute is shown in Figure 3, and the protocol for writing an encrypted attribute is shown in Figure 4.

C. Collaboration and Delegation

One of the key components to our approach, as shown in Figures 3 and 4, is the capability of the user to delegate access to attributes, enabling collaboration with other users. We modify a traditional Access Control List (ACL) model, by identifying the access control entry principal by password hash. If another user is delegated permission to access a particular attribute, the corresponding password hash must exist in the ACL attached to the attribute when stored in the destination directory. This ACL is managed by the virtual directory server, and again would require additional interaction by the attribute owner to manage.

V. Performance Testing

The solution presented here has been implemented for testing and usability purposes. Directory servers were represented using Microsoft Active Directory Administration Mode (ADAM) [24]. The Virtual directory component was modeled via a custom application on a separate system, and clients were simulated using directory services functions in Microsoft Visual Studio .NET 2008.

To accurately compare results of testing different configurations of using virtual directories, three separate ADAM instances were created, to represent the following situations:

- No virtual directory - communication directly between the client and destination directory server
- A virtual directory handing communication between client and destination directory server, but processing no encrypted attributes
- A virtual directory handing communication between client and destination directory server, and processing some encrypted attributes

Data was simulated using real-world directory objects from a corporate Active Directory instance. For each test, 10,000 user objects were created, with 25 attributes populated for each user. Testing both with and without the virtual directory server, as well as with and without encrypted attributes was performed. When using encryption to protect attributes, three attributes per user were stored encrypted. The time to perform each operation was recorded, as well as the overall size on disk of each directory instance.

TABLE I
AVERAGE NEW ACCOUNT CREATION TIME (MS)

Configuration	Time (ms)
No virtual directory - no encryption	92
Virtual directory - not encrypting	99
Virtual directory - encrypting	205

TABLE II
AVERAGE TIME TO MODIFY A NON-ENCRYPTED ATTRIBUTE (MS)

Configuration	Time (ms)
No virtual directory - no encryption	6
Virtual directory - not encrypting	12
Virtual directory - encrypting	12

1) Creating New Objects: Creating new objects in a directory service, known as *account provisioning*, involves two distinct steps: creating the new object, and populating the attributes of that object. For testing purposes, this was measured as one atomic operation. Table I shows the average time necessary for each testing configuration to create new accounts.

2) Modifying an Non-encrypted Attribute: When modifying an attribute, the virtual directory server is able to detect whether or not the attribute is encrypted. If the attribute is not encrypted, the virtual directory simply passes through the modification request from the client to the destination directory server. The time to modify an non-encrypted attribute is shown in Table II

3) Modifying an Encrypted Attribute: To modify an encrypted attribute, the virtual directory is required to decrypt the authentication string, extract the shared secret key of the client, S_c, and check to see if the requesting client is either the data owner, or listed as an authorized user of that particular object attribute. In some cases, the performance is dependent on whether or not the attribute is blank or has been previously populated. The time to complete these tasks is shown in Table III

4) Retrieving an Attribute Value: Retrieving an attribute also depends on the particular configuration and whether or not the attribute is encrypted. The time to retrieve an object attribute is Table IV

5) Directory Size on Disk: The disk space necessary to store a directory services instance can be easily measured when using ADAM. Table V shows the beginning and ending

TABLE III
AVERAGE TIME TO MODIFY ENCRYPTED ATTRIBUTES (MS)

Configuration	Time (blank attribute (ms))	Time (populated attribute (ms))
No virtual directory - no encryption	5	6
Virtual directory - not encrypting	12	12
Virtual directory - encrypting	106	100

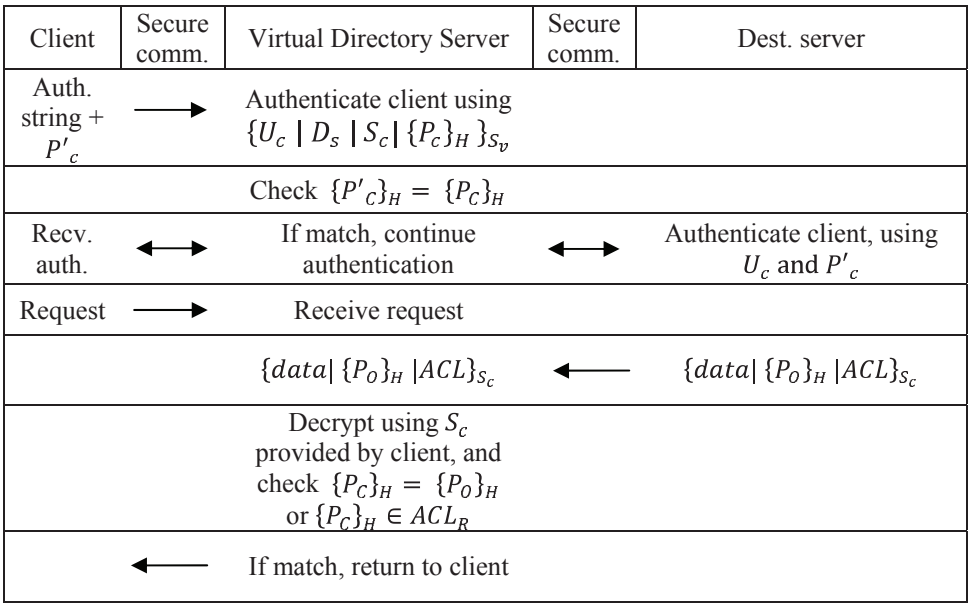

Fig. 3. Reading an encrypted attribute

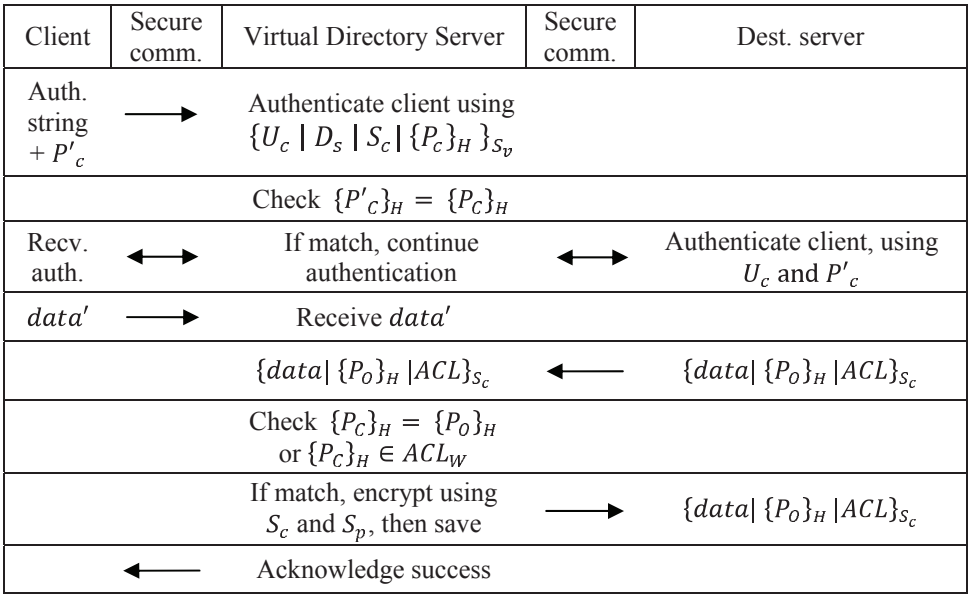

Fig. 4. Writing an encrypted attribute

size of the file used to store the directory for each test configuration. The final file size was recorded after all test accounts had been created and all test attributes modified.

VI. Discussion

Analyzing the solution we present should be approached from several angles. First, what are the advantages to using virtual directories and metadirectories as the encryption provider? Next, what are the benefits and limitations of using encryption to protect the data in directory services? No analysis of data protection would be complete without discussing vulnerabilities and attack models. Finally, how well does the solution perform, particularly in real-world situations?

A. Advantages of using virtual directories

By using virtual directories, we leverage existing technology to overcome barriers such as application reliability and security. Additionally, many virtual directory implementations compliment existing access control methods, by specifying yet another level at which users may be granted permission to access specific objects. Another distinct advantage is that virtual directories are client and destination independent. That

TABLE IV
AVERAGE TIME TO RETRIEVE AN ATTRIBUTE (MS)

Configuration	Time (non-encrypted attribute (ms))	Time (encrypted attribute (ms))
No virtual directory - no encryption	3	3
Virtual directory - not encrypting	6	6
Virtual directory - encrypting	6	98

TABLE V
DIRECTORY SIZE ON DISK (MB)

Configuration	Beginning size (MB)	Final size (MB)
No virtual directory - no encryption	6.3	56.6
Virtual directory - not encrypting	6.3	56.6
Virtual directory - encrypting	6.3	77.6

is, any LDAP client can be configured to use a virtual directory, and virtual directories can be connected to almost any type of directory service, as well as other types of data sources, such as databases.

One additional advantage could be gained by incorporating a metadirectory service into the solution as well. By using the data synchronization component of metadirectories, we could ensure that all instances of a particular attribute related to a certain user were encrypted. This takes data protection one step further, by eliminating the need to individually protect data in each separate data source.

B. Advantages of using encryption

Allowing the user to maintain the encryption/decryption key used in this solution is a user-centric approach [25], [26], [27] to data protection and identity management. In general, user-centric identity management is a method of managing user identities where the users themselves control what information is stored, the actual content of that information, and who is allowed to view the information. One motivation for this concept is privacy, accomplished by giving users the choice about what is shared, and with whom it is shared. [11]. Allowing the users to control the key provides them with complete control over the content of the data, and by including a user-specified ACL in the model, we allow users to specify who is allowed to access that data.

This is a particular advantage when considering one possible threat to conventional ACL-based access control. Administrative users may have permissions to modify ACLs on directory objects, and could grant themselves permission to read attributes intended to be confidential. By encrypting this data, we mitigate this particular threat.

C. Vulnerabilities

Of course, it's still possible for a dedicated attacker to compromise this system by gaining administrative access to the virtual directory server. This type of attack is difficult to prevent in any architecture. However, the type of attack which would compromise the data stored using our solution would be a more sophisticated attack, require more technical knowledge, and be more risky in terms of detection. No longer is a simple ACL modification necessary, now an attacker must either compromise the virtual directory application and intercept unencrypted data in transit, or he must compromise the data during transit or storage, by attacking the SSL connection. This is a much harder attack to undertake, and the risk of detection by network monitoring tools is greater.

A much simpler attack on this solution would be to compromise the user's secret key. However, this would be useless without also compromising the user's original password. Tools such as keystroke logging and administrative access to the user's computer could be used to mount such an attack, but again, this requires more technical skill, and comes with a higher risk of detection.

D. Performance analysis

Examining the performance tables shown in Section IV seems to reveal a large difference between the time it takes to manage encrypted attributes versus the time to manage unencrypted attributes. This is hard to avoid - encryption is not computationally easy - but we believe this large difference is not functionally detrimental to the overall performance of the directory. Often, attributes which need to be protected anyway are rarely accessed. A difference of 100ms is hardly noticeable when the attribute is only accessed a few times per day.

More significant to the performance of the solution in the real world is the user interaction required. An initial interaction is necessary to establish the authentication string. This could be done via secure web services, for example, but still require user configuration of the local LDAP client. Additionally, any authorized password change would require a new authentication string to be issued.

Collaboration and delegation would also be an application management issue. To add a user to the object ACL, the owner would need to use a third-party interface, and would need to have access to the hash of the delegatee's password. Again, password changes would require a modification of the ACL stored on the directory object. For large-scale delegation, this could become unwieldy. However, for sharing information with a few select sources, the benefits of this solution appear to outweigh the administrative overhead.

VII. FUTURE WORK AND CONCLUSION

We have presented an architecture for protecting sensitive information in directory services. This architecture leverages the existing technology of virtual directories and metadirectories as a layer between client and directory service applications. This third layer is responsible for handling communication between client and server, and manages encryption and decryption routines with information provided by the client. The client provides the information using standard LDAP client fields, requiring only a reconfiguration - not a recode - of client applications. By allowing users to control and protect encryption keys, we enable a user-centric model of data

protection, and reduce the threats posed by dedicated insider attacks.

Future work will include additional real-world implementation and testing. Integration with existing PKI infrastructure may also pose an interesting approach, and could help to eliminate some of the overhead associated with user key and password management. Finally, by examining existing data stores and the applications that utilize them, we may come to a better understanding of how sensitive information is distributed over an enterprise-level environment, and may discover new approaches to information protection based on such knowledge.

REFERENCES

[1] "How to mark an attribute as confidential in windows server 2003 service pack 1."

[2] M. Jakobsson, "Modeling and preventing phishing attacks," in *Financial Cryptography*, 2005.

[3] "Windows server 2003 active directory." [Online]. Available: www.microsoft.com/windowsserver2003/technologies/directory /activedirectory/default.mspx

[4] "Ibm tivoli directory server." [Online]. Available: www-306.ibm.com/software/tivoli/products/directory-server/

[5] "Mac os x server open directory." [Online]. Available: www.apple.com/server/macosx/opendirectory.html

[6] "Novell edirectory." [Online]. Available: www.novell.com/products/edirectory/

[7] "Open ldap." [Online]. Available: www.openldap.org/

[8] "Fedora directory server." [Online]. Available: directory.fedoraproject.org/

[9] "Sun java system directory server." [Online]. Available: www.sun.com/software/products/directory_srvr/home_directory.xml

[10] G. Carter, *LDAP System Administration*. O'Reilly, 2003.

[11] W. Claycomb, D. Shin, and D. Hareland, "Towards privacy in enterprise directory services: A user-centric approach to attribute management," in *Proceedings of the 41th IEEE International Carnahan Conference on Security Technology*, Ottawa, Canada, 2007.

[12] A. Berger, "Privacy protection for public directory services," *Computer Networks and ISDN Systems*, vol. 30, pp. 1521–1529, 1998.

[13] M. Chacon, "Unifying diverse directories," *Network Magazine*, vol. 16, pp. 70–75, 2001.

[14] "Microsoft identity lifecycle manager 2007." [Online]. Available: www.microsoft.com/windowsserver/ilm2007/default.mspx

[15] "Sun one meta-directory." [Online]. Available: www.sun.com/software/products/meta_directory/home_meta_dir.xml

[16] "Critical path meta-directory server." [Online]. Available: www.criticalpath.net/pdf/MetaDirectory.pdf

[17] I. Radiant Logic, "Using virtualization to leverage your investment in active directory," Radiant Logic, Inc., Tech. Rep.

[18] "Radiant logic, inc." [Online]. Available: http://www.radiantlogic.com/main/

[19] "Symlabs virtual directory server." [Online]. Available: http://symlabs.com/products/virtual-directory-server

[20] "Oracle virtual directory." [Online]. Available: http://www.oracle.com/technology/products/id_mgmt/ovds/index.html

[21] E. Kowalski, D. Cappelli, T. Conway, B. Willke, S. Keverline, A. Moore, and M. Williams, "Insider threat study: Illicit cyber activity in the government sector," U.S. Secret Service and CERT, Tech. Rep., January 2008.

[22] M. Keeney, D. Capelli, E. Kowalski, A. Moore, T. Shimeall, and S. Rogers, "Insider threat study: Computer system sabotage in critical infrastructure sectors," U.S. Secret Service and CERT/SEI, Tech. Rep., May 2005.

[23] E. Shaw, K. Ruby, and J. Post, "The insider threat to information systems," *Security Awareness Bulletin*, no. 2-98, September 1998.

[24] "Windows server 2003 active directory application mode."

[25] M. Koch and W. Worndl, "Community support and identity management," in *Seventh European Conference on Computer-Supported Cooperative Work - ECSCW 2001*, Bonn, Germany, September 2001.

[26] M. Koch, "Global identity management to boost personalization," in *Ninth Research Symposium on Emerging Electronic Markets*, Basel, Switzerland, September 2002.

[27] A. Josang and S. Pope, "User centric identity management," in *Asia Pacific Information Technology Security Conference, AusCERT2005*, Austrailia, 2005.