Using KERBEROS to Harden the Active Directory System (LDAP) in a Domain Used to Support Grid/Clustering Activity

Dennis Guster, Cory Hemminger, Joseph Meunier, Christopher Schroeder
Business Computing Research Laboratory
St. Cloud State University,
St. Cloud, MN 56301
dcguster@stcloudstate.edu

Abstract

The advent of distributed processing and global information systems has driven the need for fast, efficient and secure global authentication systems. Typically distributed processing implies that any given application will require computing resources from multiple computing nodes, and hence for the sake of convenience will require single sign-on capability for the end-user. The user database to support this global authentication process, because it encompasses all hosts in a domain is a prime attack target and requires substantial resources if it is to be adequately protected. To illustrate these concepts a case study was used in which the characteristics of a computing domain were described in detail and the conversion process of this domain from a simple NIS global authentication system to an extremely robust LDAP/Kerberos system was discussed. It was determined that the added complexity and extra work required to implement the LDAP/Kerberos system was well worthwhile due to the vast increase in robustness and scalability observed. Further, this task was carried out at the same time the production hosts were converted from individual physical hosts to virtual machines to provide a "greener" computing environment. Even though this conversion added to the workload the fact that both processes were starting from scratch made it easy to coordinate the needed linkages between the two.

1 Introduction and related Background

The advent of distributed processing and global information systems has driven the need for fast, efficient and secure global authentication systems. Typically distributed processing implies that any given application will require computing resources from multiple computing nodes, and hence for the sake of convenience will require single sign-on capability for the end-user. Further, the computing nodes maybe dissimilar in nature such as windows versus Unix machines. The need of providing single sign-on with cross-platform capabilities is confirmed by Elson, 2003. This work further recognizes the need for a sophisticated security strategy within global authentication and suggests Kerberos as a viable security solution. The complexity of a global authentication system is significantly greater than providing authentication on a host-to-host level. Harbitter and Menasce, 2001 present the importance of scaling in the design of a global authentication system. They state that the overhead to provide secure authentication transactions tends to be greater than in a single host based system. One of the mechanisms that they suggest can provide reliable security for large domains is Kerberos. They further state that it is already a well accepted authentication standard and offers excellent possibilities for large scale global authentication systems.

While clustering and grids offer exciting possibilities in regard to solving computationally intensive problems their capabilities are not without adding overhead to the domain. Often what was a small or medium sized domain becomes large when clustering/grids are added and those devices are shared across the Internet. This distributed clustering environment may result in a plethora of new hosts and users to be integrated into the active directory system. Which means scaling may very well become an issue.

While Birrell et al, 1986 agrees with the concerns presented earlier related to scaling they also emphasize the importance of realizing that no matter what global authentication structure is utilized it will need to function in continuous use for long periods of time so resiliency and reliability are crucial.

The large scale of many global authentication systems makes performance a paramount issue. Aslan, 2004 recognizes the robustness of Kerberos, but points out some of its drawbacks related to performance particularly in configurations when time synchronization across hosts in different domains are required. Besides support for traditional client server models Kerberos based authentication systems have been used successfully in grid applications. Moralis et al, 2004 reports that Kerberos can supply extremely robust security while providing the performance critical to the grid through application of quality of service mechanisms (QOS).

Therefore, this paper will focus on the upgrading process from an instructional computing domain with clustering capability currently running NIS (network information system) to a LDAP domain using Kerberos for authentication. The conversion process used will be delineated and comparisons will be made in regard to scaling, performance, ease of use and robustness of security. Also, in the final section appropriate discussion/conclusions will be presented.

2 Current Environment

The problem with the current environment stems from two different perspectives. First, the hardware is dated and relies on a separate physical host hardware model. This model while offering decent performance is hard to manage, draws significant 110 volt current and requires significant cooling capacity. Preliminary investigation regarding using virtualized servers has proven promising and it was decided to integrate virtualization into the global authentication update process. Second, the current global authentication method, Network Information System (NIS) is wrought with numerous problems both from a security and scalability perspective.

2.1 NIS vs. LDAP

The authors originally selected NIS in an effort to get their cluster/grids up and running as fast as possible. It was during the process of teaching security courses, which had auditing components, that the true degree of vulnerabilities became exposed. Specifically, NIS has a number of vulnerabilities which were identified by Guster, et al, 2008 and are expanded on below.

NIS	
Information Presentation	Clear text
Environment Design & Method	LAN and broadcast method
Authentication	No certificate based authentication
Port Assignment	Dynamic port assignment upon boot up

Table 1: NIS Vulnerabilities

NIS is used with a Unix system and passwords are only stored in encrypted form, the sniffed or ypcat obtained password will be encrypted; still their compromise is a major concern. With all sophisticated hacking tools available on the net and the use of distributed processing to speed up their use, it may be a matter of hours before the encrypted password could be broken (Guster, Safonov, Podkorytov, and Hall 2004). Also, NIS performs dynamic port assignment upon boot up. It requires the port mapper (111), which remains static but randomly assigns ports to the NIS master server; the RPC then allows password modification via the transfer port to the NIS slave. This situation makes it difficult to document port assignments via the services or services.local files and in the development of a firewall strategy. Conversely, LDAP is designed to combat many of these problems and the fine points also identified by Guster, et al, 2008 are explained below.

LDAP	
Information Presentation	Traffic is encrypted, meaning password portion is double encrypted
Environment Design & Method	Network traffic is unicast to the target via a DNS lookup
	Auditing and logging capabilities are enhanced by the software's design
	Provides a much higher degree of scalability
Authentication	Clients/servers are authenticated via a certificate authority
	No direct command line access to the password file
Port Assignment	Ports are fixed to 389 or 636 for the SSL version

Table 2: LDAP strengths

LDAP also provides a much higher degree of scalability, roughly 2 million database entries versus just 10,000 flat file entries for NIS. Further, the object oriented design database structure and the ability to distribute those entries across a network provide a robust environment that should be pertinent for some time. While LDAP in its native form is a vast improvement as compared to NIS it still has the limitation of having the password (though encrypted) travel across the network. Further, keys remain static once configured so each succeeding session uses the same key structure. Shi et al, 2006 recognize that while LDAP in its basic form is a powerful system but that there are still vulnerabilities which are often exploited on an enterprise level.

While these vulnerabilities are troublesome, the fact that it allows global authentication and single sign-on capabilities still makes it well suited to protect individual data. It has been noted that LDAP "accommodates the need of high level…security, single sign-on, and centralized user management" and "offers security services and integrated directory with excellent capability in storage management of user information in an effective directory structure" (Sari and Hidayat 2006 p. 307). This ultimately permits end users to

be more productive as it allows them to “determine application[s], service[s] and server[s] to be accessed, and user privileges.” Moreover, in the article, “Stop Data Thieves Who Get…,” (2005, p.6) it is suggested that LDAP’s authentication capabilities allow servers to “intercept” those trying to access information and “approves their privileges to see data.” This ultimately, as Ferguson (2006) comments on in her article, will allow end users to benefit from more security in a “highly regulated environment” and allow companies to control and manage users (p.12).

Given the endorsements above it makes sense to use LDAP as the cornerstone of a global authentication strategy, but devise a way to help mitigate its vulnerabilities. Moralis et al, 2004 feels that Kerberos can supplement existing security mechanisms resulting in an extremely robust security strategy. Further, this enterprise level security strategy can adapt to a highly networked environment and still maintain the performance critical in network communication through application of quality of service mechanisms (QOS).

2.2 Physical Hosts vs. Virtualization

For years, the traditional model used in enterprise computing typically features many different physical hosts. For the most part the purpose of each physical host is related to the application that it is hosting. Providing this physical separation is typically done for a variety of reasons such as gaining better performance, not disturbing a well functioning existing application or for security purposes. While the reasons cited do have a degree of merit, this physical model has several disadvantages such as it requires a lot of physical space to store multiple physical servers, the power consumption and heat generated can be substantial and the complexity and time required to manage multiple physical hosts can be problematic as well. To attain this desired physical separation it is common to host DNS (domain name service), SMTP (email), WWW (web server) and whatever application servers there might be on completely separate physical hosts. Smith and Nair, 2005 suggest that there is an alternative to this physical related model. Specifically, they state that virtualization can effectively support individual processes on a single physical machine that is logically divided into separate or “virtualized” zones.

Given the communication model used in the vast majority of today’s enterprise systems in which local communication is provided via a LAN(local area network) and communication to the outside world is provided by a WAN(wide area network) Internet security is major concern. This means the concept of separation of applications cannot be abandoned. However, virtualization can be very effective in achieving this separation through the concept of logical rather than physical zones. Specifically, Boyd and Dasgupta, 2002 state virtual operating systems can provide isolation and allow those isolated zones to safely share the physical resources contained on a single physical computer.

Given the potential advantages of virtualization it makes sense to pursue this architecture as the physical infrastructure for the future of any enterprise computing environment and it

has been selected specifically for our domain. However, virtualization, although still offering host isolation, still requires that powerful active directory/global authentication capabilities be put in place. In fact, there are added incentives to adding the LDAP/ Kerberos enhancement at the same time as upgrading to a virtual hardware solution. First, the NIS software does not need to be uninstalled and any limitations its design might present will not affect the LDAP upgrade. Second, it is more than likely the upgrade will take place on new hardware which provides an excellent opportunity to run dual system for a short time while the bugs are identified and re-mediated on the new system.

2.3 Applying LDAP to a Global Environment

Since the days when computing metamorphosed from the flat world of stand-alone computing to the multi-dimensional world of distributed computing there has been a constant quest to provide an effective/robust/efficient means to provide authentication to a group of computers. As previously stated NIS did provide those basic services, but was lacking in security characteristics as well as scalability. LDAP addresses the concerns present in NIS and based on our implementation can be improved even further from a security perspective by hardening the authentication process with Kerberos. To provide some idea of the need for the LDAP/Kerberos combination a look at the computing domain in which it is implemented is warranted. Figure 1 provides a basic mapping of the domain. Within the domain there are different zones, simplistically the Figure can be broken into three parts. First, those virtual hosts that are contained in the physical host Host1. Second the virtual hosts housed in the physical host Host2. Last, at the bottom of the Figure a number of separate independent physical hosts such as Cluster1, Command1, Gateway and so forth are depicted.

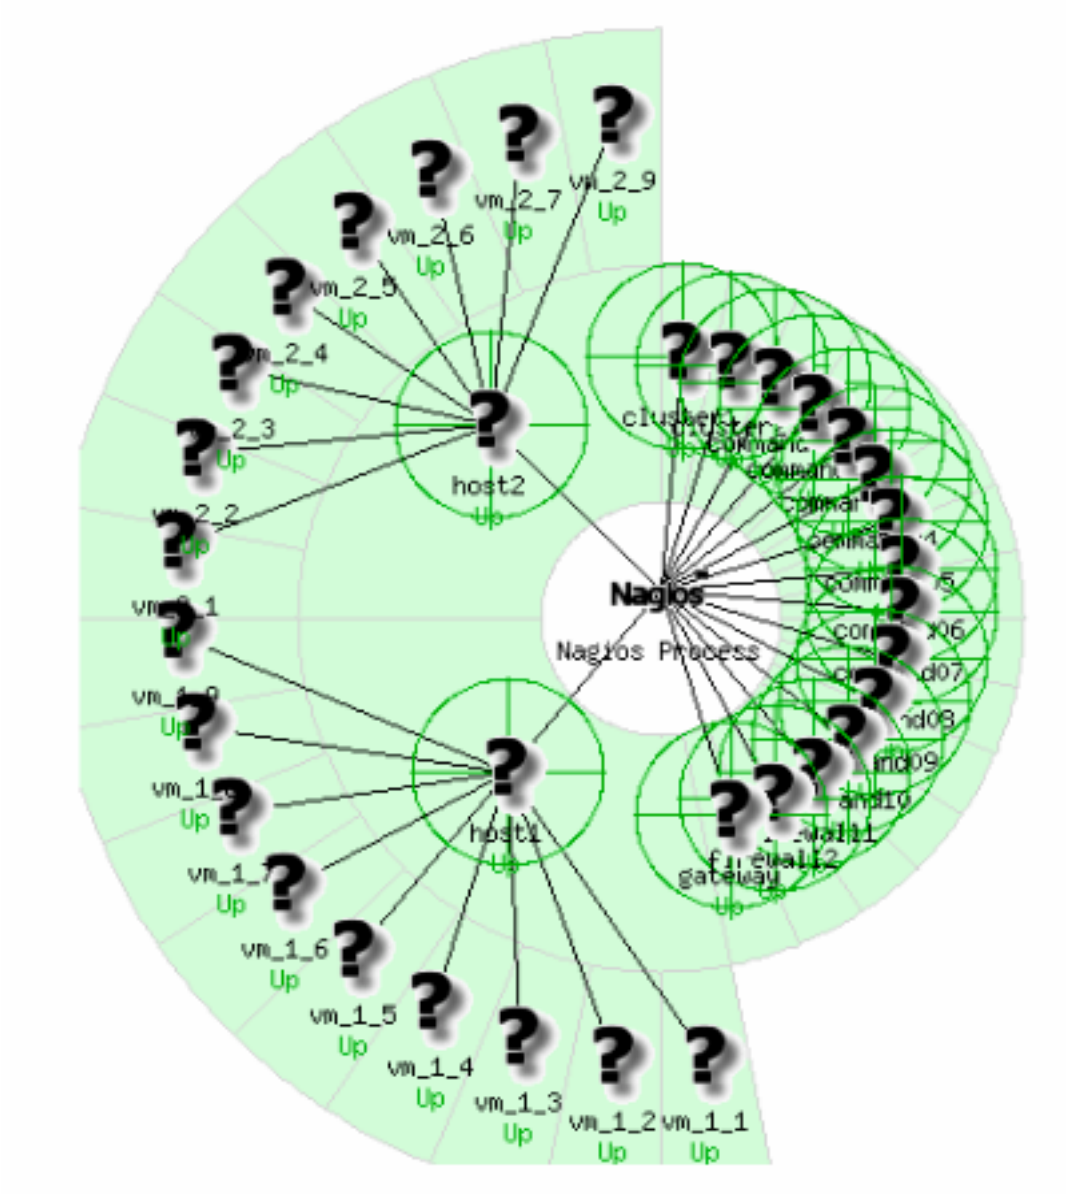

Figure 1: Network Map

It is import that all of these diverse systems communicate with one another and that a user in the domain have access to each and every host if warranted by the security policy. In other words, in the active directory scenario a user doesn't log directly into a host, but rather into a database (the LDAP) that keeps a list of authorized users and which particular hosts they are allowed to access. In fact, a user's account/login information may not reside at all on the host to which they wish to connect. One of the prime advantages to this method is central administration and a single point to protect.

This network mapping is described in more detail by the service overview depicted in Figure 2. This figure provides detailed tables for six different service groupings. First, there are two physical hosts which are logically partitioned, BCRL-in-a-box and BCRL-in-

a-box2. BCRL-in-a-box is designed to support 9 virtual hosts and they are all up and running and each is supporting at least one service. BCRL-in-a-box2 is designed to be the backup or replica to BCRL-in-a-box and therefore looks very similar in function to BCRL-in-a-box. The third service group is the computers used to support the cyber warfare command center, it is currently in a development state and while the ten separate physical hosts are functioning additional services are expected to be added. The firewalls group contains two functioning production firewalls. The HTTP server group provides web hosting and is made up of three functioning production physical hosts. Last the Pingable Servers group is used to allow testing and diagnostics of network connections and consists of three physical hosts. As previously stated, these server groups make up the enterprise computing domain. Imagine how difficult it would be if a substantial user base had individual accounts on each host and each host's authentication services had to be managed on each single (physical or virtual) host. So therefore, in this multi-host enterprise environment based on Figure 2 there are: potentially thirty eight hosts that would need to be managed independently if it were not for the capabilities of a global authentication system.

In addition, to providing single sign-on through global authentication a global file system is also integrated into this domain via a network file system (NFS) mount so that a users default home directory follows them no matter which host they are accessing

While the domain in the case study is not starting from scratch in this endeavor (originally a series of independent hosts) the upgrade from NIS to LDAP is not with without difficulties. The greatest concern is taking the NIS flat file data structure and converting it to a relational data base structure that will support quick access and grow as the domain's user base does thereby fostering long term scalability.

Figure 2: Network Overview

Also, as previously stated, the addition of Kerberos, while greatly adding to the robustness of the authentication process, greatly adds to the complexity of the conversion process. Of particular note is the process of getting these two independent software components to effectively talk to one another. This crucial communication process required much analysis and the development of a complex script that served as the communication interface between the two software components.

The processes undertaken in that script follows:
#!/bin/sh
make working directory and change to it

get uids from LDAP into a file & cleanup file ldapnames
ldapsearch: perform simple search

```
# get only users with a uid
# modify output to only show actual names
# > ldapnames -- write to file

# get principals from Kerberos into a file & cleanup file kerberosnames
# use kadmin as principal ldap/admin

# -list all principals
# delete first line
# delete all hosts and service principals
# remove trailing realm
# remove remaining Kerberos and admin principals
# > ldapnames -- write to file

# compare files to find new users & write to new users file
# print differences
# print only lines with '> ' users that need to be added
# remove '> ' from lines
# > newusers -- write to file

# compare files to find old users & write to file old users
# print differences
# print only lines with '> ' users that need to be added
# remove '> ' from lines
# > oldusers -- write to file

# read newusers and add principals to Kerberos with passwords related to usernames
# read oldusers and removes principals from Kerberos
# remove working directory and exit
```

A quick analysis of this documentation reveals that there are routines for adding/deleting users and maintaining existing users. Further, the documentation reveals numerous calls to both LDAP and Kerberos to extract and reformat the data required so that they can communicate with each other. It is this union that provides a high performance scalable global authentication system that provides robust protection of key user data elements such as the password. Unlike the native version of LDAP in which the passwords travel the network Kerberos is able to evaluate passwords locally and then issue a “ticket” for that single session. This process hardens the whole authentication structure which is especially critical when all user information is centralized into a single database. In the single host model if any given host is compromised only that host is affected. Conversely, in the global authentication model if the database is compromised all the hosts in that domain could potentially be affected. Therefore, adding the extra complexity of Kerberos to the equation is well warranted.

3 Summary, Discussion and Conclusions

The change from stand-alone computing to enterprise level distributed computing has necessitated new thinking in regard to how user authentication data is stored and managed. Instead of managing a flat user file on each host all user authentication data is now stored on a centralized system to provide ease of management and single sign-on to end users across hosts. This system because it encompasses all hosts in a domain is a prime attack target and requires substantial resources if it is to be adequately protected. To illustrate these concepts a case study was used in which the characteristics of a computing domain were described in detail and the conversion process of this domain from a simple NIS global authentication system to an extremely robust LDAP/Kerberos system was discussed. It was determined that the added complexity and extra work required to implement the LDAP/Kerberos system was well worthwhile due to the vast increase in robustness and scalability observed. Further, this task was carried out at the same time the production hosts were converted from individual physical hosts to virtual machines to provide a “greener” computing environment. Even though this conversion added to the workload the fact that both processes were starting from scratch made it easy to coordinate the needed linkages between the two.

There is a lesson to be learned from the data presented herein and that is that there are vulnerabilities in the two most commonly used global authentication software packages, NIS and LDAP. While LDAP is a vast improvement over NIS our analysis still revealed vulnerabilities particularly in how the passwords are handled. So therefore we ascertained that upgrading to LDAP alone would not be sufficient to protect our most valuable asset, which is our user/password database so we opted to harden it with Kerberos. While the literature almost universally agrees with this concept, there is not currently agreement on the best manner to accomplish this. We evaluated a couple of well respected methods such as the one recommended by Sun Micro-systems, but found that none of them precisely met our needs. So therefore, we devised our own custom interface between the software components. While this can be very positive in the sense that the staff that develops such as interface will have an excellent understanding of its intricacies. However, whenever new code is written it is important to thoroughly test it under a variety of condition both from a functional and security perspective. In our case, except for a few minor bugs the system is functioning effectively and based on our security logs attacks are fewer and less effective. So therefore, based on our experience, hardening LDAP with Kerberos proved effective from three perspectives: performance, scalability and security. However, the lack of an accepted standard of integrating the two software component makes the solution complex and difficult to implement. In our case we had system programmers on staff. It is not an undertaking that is easy to pursue without staff support in systems programming.

References

Aslan, H. (2004). Logical Analysis of AUTHMAC_DH: a New Protocol for Authentication and KeyDistribution. Computers and Security,23, 290-299.

Birrell, A., Lampson, B., Needham, R. & Schroeder, M. (1986). A Global Authentication Service Without Global Trust. Proceedings of the IEEE Symposium on Security and Privacy, 223-230.

Boyd, T. and Dasgupta, P. (2002). Process migration: A Generalized Approach Using a Virtual Operating System, *Proceedings of the 22nd International Conference on Distributed Computing Systems*, 385.

Elson, D. (2003). Active Directory and Linux. http://www.securityfocus.com/print/infocus/1563.

Ferguson, R. (2005) "Lawson Settles on WebSphere", *Eweek*, Vol. *22*, No. 32, pp 11-12.

Guster, D. C., Hall, C., Herath, S., Jansen, B., & Mikluch, L. (2008, April 24). *Analysis of vulnerabilities in a global authentication system in a university based distributed processing laboratory*. Presentation at the 3rd International Conference on Information Warfare and Security, Omaha, NE.

Guster, D., Safonov, P., Podkorytov, D., and Hall, C.. (2004) "Security Against Hacking Attacks: Application of Distributed Processing and Software Modifiers in Defense of Password Files", *International Business Trends: Contemporary Readings*, Academy of Business Administration, Las Vegas.

Harbitter, A. & Menasce, D. (2001). Performance of Public-Key-Enabled Kerberos Authentication in Large Networks. Proceedings of the IEEE Symposium on Security and Privacy, IEEE Computer Society Press.

A. Moralis, V. Pouli, M. Grammatikou, S. Papavassiliou, V. Maglaris, "A Security Architecture using Symmetric Cryptography and Kerberos-based approach for Performance Improvement in Grids", 2nd International Workshop on Distributed Cooperative Laboratories: Instrumenting the Grid (INGRID 2007), Santa Margherita Ligure - Portofino, Italy, April 16-18, 2007.

Sari, R. F. and Hidayat, S. "Integrating Web Server Applications with LDAP Authentication: Case Study on Human Resources Information System of UI", *Communication and Information Technologies*, pp. 307-312.

Shi, H., Yanchun, Z., Jingyuan, Z., Beal, E. and Moustakas, N. (2006) *First International*

Multi-Symposiums on Computer and Computational Sciences, Hangzhou, China, pp. 552-559.

Smith, J. and Nair, R. (2005). The Architecture of Virtual Machines, *Computer*, 38(5), 32-38.

"Stop Data Thieves Who Get…" (2005) *ComputerWorld,* Vol. 39*,* No. 47, pp 6.