고성능 Grid 환경에서의 LDAP 시스템의 성능분석

권성호, 김희철, 이용두

Performance Analysis of LDAP System in High Performance Grid Environments

Quan Chenghao, Hiecheol Kim, Yongdoo Lee

Abstract For high performance Grid environments, an efficient GIS(Grid Information Service is required. In the Metacomputing Directory Service(MDS) of the Globus middleware, the Lightweight Directory Access Protocol(LDAP), which is a distributed directory service protocol, is used. The LDAP GIS differs from general purpose LDAP directories in that most of the LDAP operations are write in Grid environments. To get an efficient design of the GIS, it is thus required to analyze the performance of the LDAP system in the context of Grid environments. This paper presents the result of a performance analysis of LDAP systems. The main objective of the evaluation is to see the performance scalability of the LDAP system in the Grid environment where the write operations prevails. Based on these results, we suggest directions of an efficient LDAP-based GIS for a high performance Grid.

1. Introduction

One of key components of a Grid middleware is an information service, which is called a GIS(Grid Information Service). It provides users and applications with the information on the resources and services available in the Grid. Usually, implementations of GISs use LDAP(Lightweight Directory Access Protocol) directories. As an example, MDS(Metacomputing Directory Service) which is the GIS of the Globus Grid middlewre is based on the a LDAP directory. For higher performance of the Grid, it is necessary that the informations serviced by a GIS system should be provided in a timely and reliable manner. In the system with an LDAP-based implementation of the GIS, the performance of the GIS is thus dependent on the performance of LDAP operations over the directories, which further affects the performance of the Grid. The analysis of performance of LDAP system is thus needed in the context of Grid environments.

LDAP is a distributed directory service protocol used primarily with read-most environments to implement Internet directories. Hence, researches are focused mostly to the optimization of read operations. Scientific researches and studies are not still made in the environment where write operations are more than read operations. However, Grid environments entail that write operations occupy more than 90 percents among all the LDAP operations.

The goal of this research is to suggest the direction of the design for the high performance LDAP system in Grid environments. To this end, we simulated a distributed LDAP system to evaluated the performance with respect to the various factors such as the ratio of read operations and the ratio of local operations, etc.

The rest of this paper is organized as follows. In section 2,

* Ph.D. Course, School of Computer and Communication Engineering, Daegu University.
** Professor, School of Computer and Communication Engineering, Daegu University.

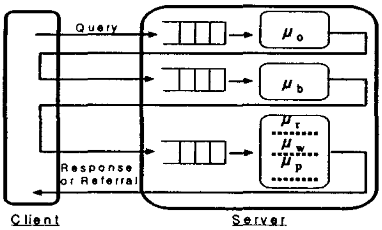

Figure 1. Simulation model

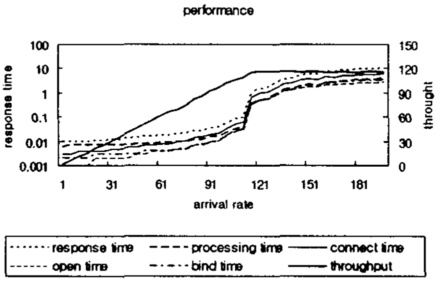

Figure 2. Model verification

a brief outline of the LDAP and the MDS is presented as a background and subsequently some previous researches about performance analysis made to LDAP systems are introduced. In section 3, we describe our simulation methodology applied in our experiment. In section 4, we present the result of our simulation. Finally, concluding remarks are offered with future researches.

2. Related works

LDAP is a distributed directory service protocol[7]. It is based on a client-server model and runs over TCP/IP. Information is modeled in a DIT (Directory Information Tree). By adopting referrals, a DIT is partitioned and distributed across multiple servers. A referral acts like a pointer that can be followed to locate the server on which the desired information is stored. Commonly, LDAP operations are modify, add, delete, compare and search. In the general purpose, search(read) operations dominate the server load, but in Grid environments, modify(write) operation occupies 95% of server load.

Currently, the Globus software uses the OpenLDAP client software to access the MDS[3]. The MDS provides the Grid information service of the Globus project. As the number of entities participating in the Globus computational Grid increases, the response time of the GIS becomes inadequate[1][2]. According to a previous research, the majority of the operations are modifies in Grids[4].

Because directories are meant to store relatively static information and are optimized for that purpose, they are not adequate for storing information that changes rapidly. In other words, an LDAP system in the general purpose is not suitable for high performance Grid environments. In order to improve the performance of the GIS, it is thus required to analyze the performance of LDAP system in the context of Grid environments.

So far, not so many prior researches are made to enhance the performance of LDAP systems. Some research evaluates the performance of LDAP systems[5]. But the experiment assumes the general purpose LDAP setting where search operation dominates the server load. The analysis of workload of LDAP systems in the Grid environments is made[4]. These results do not provide any direct way to enhance the performance of LDAP systems in Grid environments. A replication mechanism is investigated to enhance the performance of OSI directory systems[6]. In Grid environments, replication does not presents much room for performance enhancement, because when data is changed, these changes must be propagated to the replicas of the data and this adds fairly severe overhead when write operations prevail.

3. Simulation Methodology

The system to simulate is a distributed LDAP system in which N nodes are configured in a full mesh. A query that arrives to a node will be serviced by the local server, if the server has the data for the query; otherwise, the query is passed to another node directed by the referral.

A simulation model of an LDAP operation is shown in Figure 1. Each connection in normal LDAP systems may consist of a connect, a bind to an identity, one or more adds, deletes, modifies, or searches, an unbind, and a close. In this model, we suppose only one operation per connection. This supposition can be considered as the worst-case in the execution behavior of LDAP systems. In the model, the query arrival rate has a Poisson distribution represented by parameter λ. The service rate of the LDAP server has parameters with μr, μw, and μp respectively for read, write and pass due to referral.

In our simulation, we use CSIM that is a process-oriented discrete-event simulation package for use with C programs[8].

The experiments consists of two steps. The first step of the experiment is aimed at verifying our simulation model. The simulation is made to a single node environment, and then its result is compared with results from other researches. As shown in Figure 2, it is found that out result conforms to the one reported by other literature[5].

In the second step, we evaluate the performance of the system with a range of values for parameters such as the ratio of read over write, the ratio of local over remote and the

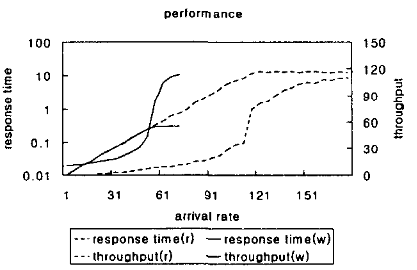

Figure 3. Comparison of read and write

number of node. Especially, we compare the performance of general purpose case, in which the read operation occupies 95% of all the operations, with that of Grid environments in which write operations occupy 95%. Futhermore, we explore the effect of indexing, data localization, and data distribution.

4. Simulation Results

4.1. The ratio of read over write

The experiments are made respectively for two configurations. One is a read-intensive configuration in which 95% of operations are read. The other is a write-intensive configuration in 95% of operations are write. Figure 3 shows average response time of a query. The results show that the response time of the read-intensive configuration is two times better than the write-intensive cases. As expected, the read-intensive configuration is saturated for almost two times bigger arrival rate.

4.2. Indexing

An index is a store of search results to provide quick look-ups for searches. The disadvantage of an index is that they have to be updated whenever an attribute they are indexing is changed.

The effect of indexing is shown in Figure 4 and 5. In the case that read operations occupies 95%, indexing benefits the response time with the average speedup by 1.3 times. On the contrary, in the case that the write operation occupies 95%,

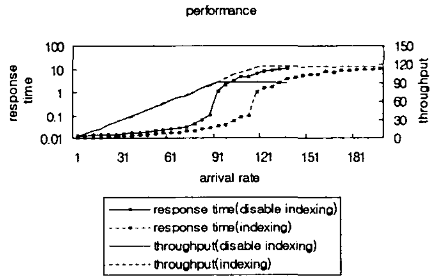

Figure 4. Indexing effect for read

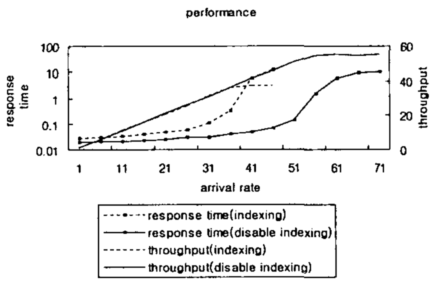

Figure 5. Indexing effect for write

indexing makes the average response time down to 1.5 times. As for the average throughput, no change is shown in both cases, whereas the maximum throughput is 1.3 times and 0.67 times respectively for each case.

In the rest of the simulation, indexing is enabled in the read-intensive configuration, while it is disabled in the write-intensive configuration.

4.3. Data Localization

When a query is made to a server, it will be served by the server if the data is available in it. On the other hand, it will be passed to another server pointed by the referral. In this case the query is called a remote query.

Both response time and throughput become are enhanced as the rate of local operations increases, as shown in Figure 6.

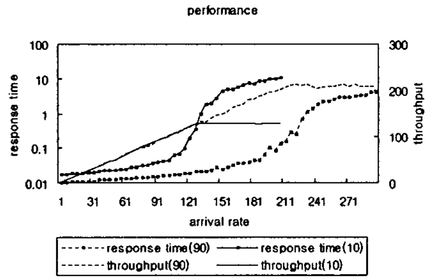

Figure 6. Data Localization

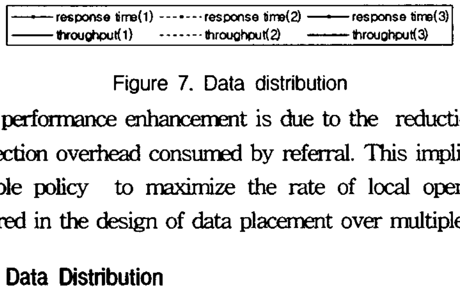

Figure 7. Data distribution

This performance enhancement is due to the reduction of the connection overhead consumed by referral. This implies that a suitable policy to maximize the rate of local operations is required in the design of data placement over multiple servers.

4.4. Data Distribution

To implement a very large directory, it is desired to use data distribution over multiple servers. By distributing data across multiple servers, we can expect better performance as well.

Figure 7 shows the result of our simulation for data distribution. When the rate of local operations is 90%, the maximum throughput due to data distribution is improved by 98% over the configuration without data distribution. In this case, the configuration without data distribution is a system with only one server. When the rate of local operations is 10%, the maximum throughput due to data distribution is improved by 49% over a non-distribution case. Here, the maximum throughput means the throughput when the system is saturated. The result indicates that the performance scalability can be obtained with appropriate data distribution over multiple servers even for Grid environments where write operations are dominant.

5. Conclusions

This paper presents the performance evaluation of a distributed LDAP system. The evaluation is made by using simulation with synthetic load. The main objective of the evaluation is to see the performance scalability of the LDAP system in the Grid environment where the write operations prevails. The result shows that indexing makes average response time worse by 1.5 times in normal Grid environments. The data distribution across multiple nodes presents modest improvement of the performance. Our evaluation result indicates that performance degradation due to the high rate of write operations in Grid environments can be compensated by adopting some advanced methods such as disable indexing, local policy and data distribution. In our future work, it is required to study more advanced topics such as the tradeoff between data distribution and data replication. We believe that our research results can be used as valuable materials for the efficient LDAP-based design of GISs in high performance Grid environments.

References

[1] I. Foster, C. Kesselman, S. Tuecke. The Anatomy of the Grid: Enabling scalable virtual organizations. *Intl. Journal of Supercomputing Applications.* (to appear) 2001.

[2] K. Czajkowski, S. Fitzgerald, I. Foster, C. Kesselman. Gird Information Service for Distributed Resource Sharing. Proc. 10th IEEE International Symposium on High Performance Distributed Computer(HPDC-10), IEEE Press. 2001.

[3] I. Foster and C. Kesselman. Globus: A Meatacomputing Infrastructure Toolkit. International Journal of Supercomputing

applications. 11(2):115-128, 1997.

[4] W. Smith, A. Waheed, D. meyers, J. Yan. An Evaluation of Alternative Designs for a Grid Information Service. *HPDC 2000, IEEE Press*: 185-192, 2000.

[5] X. Wang, H. Schulzrine, D. Kandlur, D. Verma. Measurement and Analysis of LDAP Performance. *International Conference on Measurement and Modeling of Computer Systems (SIGMETRICS'2000)*, Santa Clara, CA, pp. 156--165, Jun. 2000.

[6] YC. Jang, KS Kim, JS. Woo, SS. An. The Replication Mechanism Analysis of OSI(Open System Interconnection) Directory System. *J. ENG. SCI. & TECH.* Vol. 33, pp. 49~61.

[7] H. Johner, L. Brown, FS. Hinner, W. Reis, J. Westman. Understanding LDAP *IBM.* June, 1998.

[8] User's Guide: CSIM18 Simulation Engine(C Version). *Mesquite Software, Inc.*