

A load cluster management system using SNMP and web

By Myung-Sup Kim,^{*} Mi-Joung Choi and James W. Hong

Clustered servers for Internet service is a popular solution to cope with the explosive increase in client requests. The high probability of service failure in cluster servers make the cluster management system necessary to provide high availability and convenient administrator control. In this paper, we present the design and implementation of a load cluster management system (LCMS) based on SNMP and Web technology. Our LCMS implementation has been deployed on a commercial ultra-dense server. Copyright © 2002 John Wiley & Sons, Ltd.

Introduction

An Internet server farm^{1,2,10} using load balancing is one solution to the enormous growth of requests to Internet server systems. This solution will have wide-ranging ramifications for the set-up of high-performance Internet service sites because of several attractive features. It is a very cost-effective solution. We can compose a load cluster system using off-the-shelf and low-price computers. It also provides good scalability and extendibility, so we can easily add a single server into a load cluster group to increase service throughput, and remove a single point of failure during its operation. From a client viewpoint, a load cluster group is considered to be a single powerful system with a single hostname and IP address. It is not necessary to

change in the client side to obtain service from a load cluster group.

There are several implementation methods to construct an Internet load cluster system, which includes a dispatcher approach,^{10–12} a parallel filtering approach,⁷ and a round-robin Domain Name Server (RR-DNS) approach.^{2–5} But RR-DNS and parallel filtering cannot consider the current exact workload of real servers in the cluster farm for a real server selection algorithm. Other problems arise, such as the difficulty in detecting a single server fault and recovery. Further, in the RR-DNS approach, the catching in local DNS has a locality problem²⁸ where all requests from hosts at a single domain are forwarded to a single server, so it can be easily overloaded. In the dispatcher approach, all requests from the clients are caught by a single host and are forwarded to

Myung-Sup Kim is a PhD candidate in the Department of Computer Science and Engineering, POSTECH. His research interests include distributed processing and network management.

Mi-Joung Choi is currently a graduate student in the Department of Computer Science and Engineering, POSTECH. Her research interests include Web-based network management and policy-based network management.

James W. Hong is an associate professor in the Department of Computer Science and Engineering, POSTECH, Pohang, Korea. He has been with POSTECH since May 1995. He has been very active as a participant, program committee member and organizing committee member for IEEE CNOM sponsored symposiums such as NOMS, IM, DSOM and APNOMS. For the last several years, he has been working on various research projects on network and systems management, which utilize Web, Java and CORBA technologies. His research interests include network and systems management, distributed computing and traffic engineering and planning. He is a member of IEEE, KICS, KNOM and KISS.

^{*}Correspondence to: Dr Myung-Sup Kim, Department of Computer Science and Engineering, Pohang University of Science and Technology, San 31, Hyojadong, Namgu, Pohang, Korea.

[†]E-mail: mount@postech.ac.kr

the appropriate host. The best examples of this type are the L4 switch^{6,7} and the Linux Virtual Server (LVS).⁸⁻¹² This type of load cluster system is widely used because of its good performance and throughput.

To operate this kind of load cluster system, supporting a fault tolerant and stable service is very efficient. The efficient use of resources is also a key point in the design of a load cluster management system. A cluster server is a set of low-performance machines such as PCs but this provides high performance as an Internet server. The possibility of a single server failure is greater than one single server system. In the case of failure of a real server or a load balancer, a highly available cluster management system reconfigures the cluster farm by removing this failed host in the service group or replacing the activity with another host, so the service to the client continues. Much research on the load cluster management system focuses on automatic cluster configuration only, and resource management is not efficient, and the high availability (HA) functionality is separated from the cluster management system.

In this paper, we first discuss the requirements of the load cluster management system and design a load cluster management system which provides efficient resource management and good HA features. We have implemented our LCMS on a Linux platform using SNMP, Java and Web. Our LCMS implementation has been used tested on a commercial ultra-dense server called Netstech EnterFLEX 2100.²⁹

The organization of this paper is as follows. In the next section we present an overview of the load cluster system and discuss some related work on cluster management systems. In the third section we discuss the requirements of the load cluster management system. In the fourth section we present the design of a load cluster management system supporting HA and good reliability in detail. In the fifth section we present implementation issues. In the final section we summarize our work and discuss possible future work.

Related Work

In this section, we discuss some related work on cluster management. First, we present an

example of a cluster system, LVS, and then some existing load cluster management systems and super cluster management systems. We also compare weak and strong points of the systems.

—Linux Virtual Server—

One example of a load cluster system using the dispatcher approach is the Linux Virtual Server (LVS).⁸⁻¹² The LVS consists of one load balancer and several real servers. The load balancer receives all requests from clients and distributes them to a real server using an appropriate scheduling algorithm, such as round-robin (RR), weighted round-robin (WRR), least-connection (LC) or weighted least-connection (WLC). Three types of IP load balancing techniques (packet forwarding methods) exist together in the LinuxDirector (load balancer). They are virtual server via network address translation (NAT), virtual server via IP tunneling, and virtual server via direct routing. The parallel services of real servers can be made to appear as a virtual server on a single IP address, so that the end users see a virtual server, not a cluster of servers. Scheduling granularity is per connection, which can create a sound load balance among the servers. This type of IP-level load balancing is better than application-level load balancing, such as reverse-proxy. Figure 1 illustrates an example of LVS service architecture.

—Load Cluster Management System—

In this section, we discuss two of the most popular load cluster management systems that are currently available in the market: Turbo Cluster Server 6²⁰ from TurboLinux and RedHat HA Server (Piranha)²¹ from RedHat.

Turbo Cluster Server 6 provides a good configuration tool and has a powerful and well-categorized menu system. It also provides a useful tool: a data synchronization module between real servers. Further, it provides a simple web-based configuration and status display tools. However, the configuration menu is consol-based, so a new user cannot easily learn how to operate it. The alarm notification functionality is not provided in this cluster management system.

Using Piranha, we can construct two types of cluster systems: Fail-Over Service (FOS) and

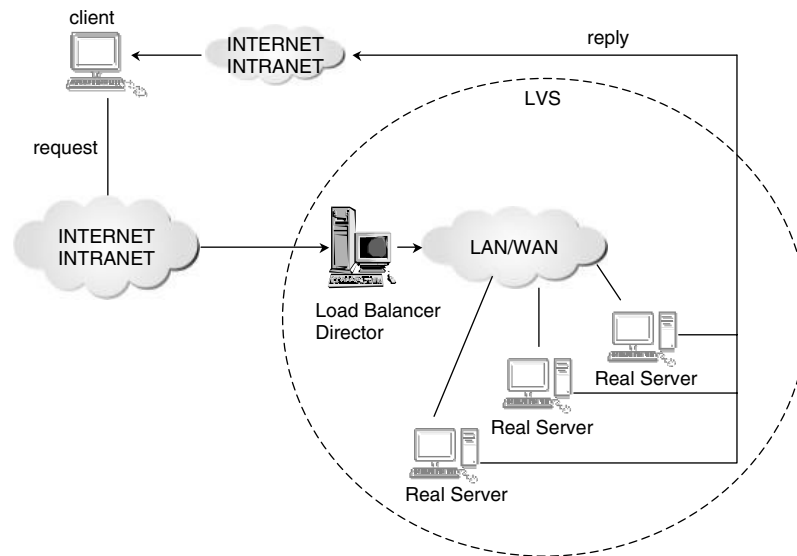


Figure 1. LVS service architecture

Linux Virtual Server (LVS). It provides Web- and Windows-based configuration tools. But these tools are dedicated only to configuration functionality and provide a primitive monitoring and status-reporting menus. The current version of Piranha does not support alarm notification and log files.

Most CMS architectures are based on the manager/agent paradigm and use their own communication protocols between a manager and an agent.

Most CMS architectures are based on the manager/agent paradigm and use their own communication protocols between a manager and an agent.

Turbo Cluster Server 6 and Piranha are a distributed configuration system. Both are very weak in status monitoring and lack a reporting mechanism.

—Super Cluster Management System—

In this section, we discuss a super cluster management system (CMS). Generally, CMS is an

interface to a cluster of computers, which is used for executing parallel programs in a high-performance cluster system using a parallel processing library such as MPI.¹⁷ Until now, we could not locate any integrated management system for the load cluster system, such as CMS on a high-performance cluster. Our LCMS is based on CMS in the basic idea and modified the functionality and management model.

Recently, many CMSs have been proposed and developed. Most CMS architectures are based on the manager/agent paradigm^{14,15} and use their own communication protocols between a manager and an agent. A specific CMS^{15,16} is developed to manage a specialized target, such as an NT cluster, but this specialized system causes a portability problem. The main functions of CMS are status monitoring and job assignments, which do not consider fail-over in the case of a single system failure. GUI-based user interfaces are typically implemented on a Windows application, so the administrator can access and control a cluster system using only a specified computer on which CMS is running. The existing CMS for load balancing mainly focuses on HA functionality to provide a stable service. Therefore, cluster configuration and status monitoring was not considered. Table 1 illustrates the comparison of several super and load cluster management systems. The load cluster management system

	CEO	JobCoNTrol	NCSA Symera	Turbo Cluster	Piranha
Management type	Distributed	Distributed	Distributed	Centralized	Centralized
Communication Method	Proprietary	Proprietary	DCOM	Proprietary	Proprietary
Manager interface	GUI	GUI	GUI	Consol/Web	GUI/Web
Managed target	UNIX/ Windows	Windows NT	Windows	Linux Windows	Linux Windows
Cluster type	HPC	HPC	HPC	LBC	LBC
Cluster Configuration	X	X	X	O	O
Fault tolerant	X	X	O	O	O
Resource management	O	O	O	X	X
Portability	O	X	X	X	X

Table 1. Comparison of cluster management system

(LCMS) presented in this paper is an integrated system with all functionalities to operate a load cluster system. LCMS should provide three major functionalities: load cluster configuration, high availability, current and past cluster system status monitoring. Next, we consider these functional requirements for LCMS in more detail.

Requirements

Before considering LCMS design, we should discuss the functional requirements of LCMS. We offer this discussion as a guideline for further LCMS designs. These requirements should be applicable to all types of load cluster systems and should not be biased towards a certain implementation method. We categorize LCMS requirements into six types:

- **Efficient Resource Management:** This requirement concerns host management and cluster group management. In host management, LCMS should be able to store the information about hosts that are members of a load cluster group. The information can include the host name, the IP address, CPU type, memory size, etc. When a new cluster group or service group is created, this information should be stored. The LCMS should manage all this information and provide a method to lookup or update host and cluster information to administrators.
- **Load Cluster Configuration:** A cluster group consists of two or more hosts where one host acts as a load balancer and the others as real servers. Further, there can be a dedicated backup server in a cluster group. LCMS should provide a way to assign a role to each host when a new load cluster group is created or an already existing cluster group is modified or deleted. And an administrator should be able to check the current status of each host and cluster group.
- **High Availability:** This can be the main functionality of LCMS for a fault-tolerant service. LCMS should check the current state of each host in a cluster group. If a host in a cluster group fails it should recover by load cluster reconfiguration. Reconfiguration can be executed by removing or replacing a failed host with a new one.
- **Effective management Interface:** LCMS should provide GUI-based cluster configuration and

- a visualized status monitoring method for management convenience.
- Security: All administrator, host and cluster information should be stored safely and a user authentication method should be provided. Only an authorized administrator should access the management system. Another security issue on LCMS is that the communication between manager component and agent component should be secured.
 - Minimization of Management Overhead: Management functionality is important but is not the main function of a load cluster system. Therefore, LCMS should not create high network overhead. Management overheads should be distributed among hosts in the cluster group so as not to overburden a specific host.

System Design of Lcms

Figure 2 is an overall architecture of our Load Cluster Management System (LCMS), which is designed to satisfy the requirements we discussed in the previous section. Our LCMS consists of three kinds of manager: Load Cluster (LC) Manager, Load Balancer (LB) Manager and Real Server (RS)

Manager. These three managers have different functionalities and are distributed among hosts in the cluster. This LCMS is designed for the dispatcher types of load cluster.

The LC manager is responsible for configuration. To store host and cluster information, an information repository is used. Administrator information is also stored in this repository. An administrator can add and remove a host into the information repository through the web interface provided by the LC manager and can make a new load cluster group and modify or delete an existing load cluster group. The LB manager is responsible for status monitoring of each cluster group. It periodically checks the status of the load balancer in each cluster group. When a load balancer is down or does not work well for some reason, the LB manager chooses one among all real servers in the cluster group and set this real server as a load balancer. This real server works as the load balancer until the original load balancer wakes up and plays its role again.

When an administrator creates a new cluster group, he must select one host as a load balancer, and the other hosts run as real servers. The RS manager is running at the load balancer in each cluster group. It is responsible for the HA functionality. The RS manager periodically checks whether each real server plays his role well or

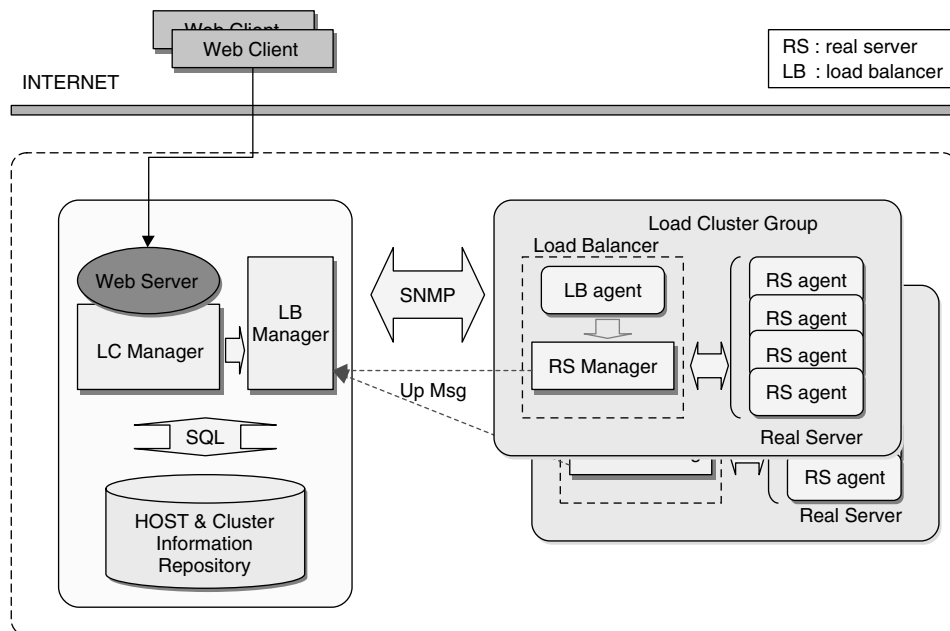


Figure 2. Load Cluster Management System

not. If a real server fails, the RS manager removes this host from the request-forwarding list in a load balancer and includes this host again when this host is running again.

All hosts that belong to a load cluster group have an SNMP agent.¹⁹ There are two types of SNMP agents: an RS agent and an LB agent. Between a manager and an agent, the communication is performed via SNMP over UDP, so LCMS reduces network traffic for management functionality.

The manager interface is composed based on the Web, so the manager can easily access the LCMS from any computer connected to the Internet and perform all management functions, such as displaying historical data, as well as the current status of each host and cluster.

—LC/LB Manager Architecture—

Cluster configuration and resource management are the main roles of the LC manager, which is illustrated in Figure 3. The administrator and host are resources to be managed by the LC manager. The information of these resources is stored in the Information Repository (IR). The LC manager

provides a web-based UI for user authentication, resource management, and cluster configuration.

As an Information Repository the LC manager uses a database. In the following subsection, we discuss in detail how host and cluster information is stored in the IR. The LC manager and DB server can be located in different hosts, which gives workload distribution.

There are three main components in the LC manager: a status monitor, an information manager, and a configuration manager. The configuration manager is responsible for the creation and deletion of a load cluster group, and uses SNMP to communicate with each host and SQL to store cluster information. The cluster configuration sequence is as follows. First, the administrator views the host list stored in the IR with the assistance of the information manager. Next, the administrator selects a set of hosts and selects one host as a load balancer to create a new cluster group. Next, the administrator can create one or more service groups in that cluster group. A host cannot belong to more than two cluster groups. Figure 4 illustrates the cluster and service assignment to hosts in the IR.

Role assignment to each host is made by SNMP. The agent receives SNMP packets and sets the

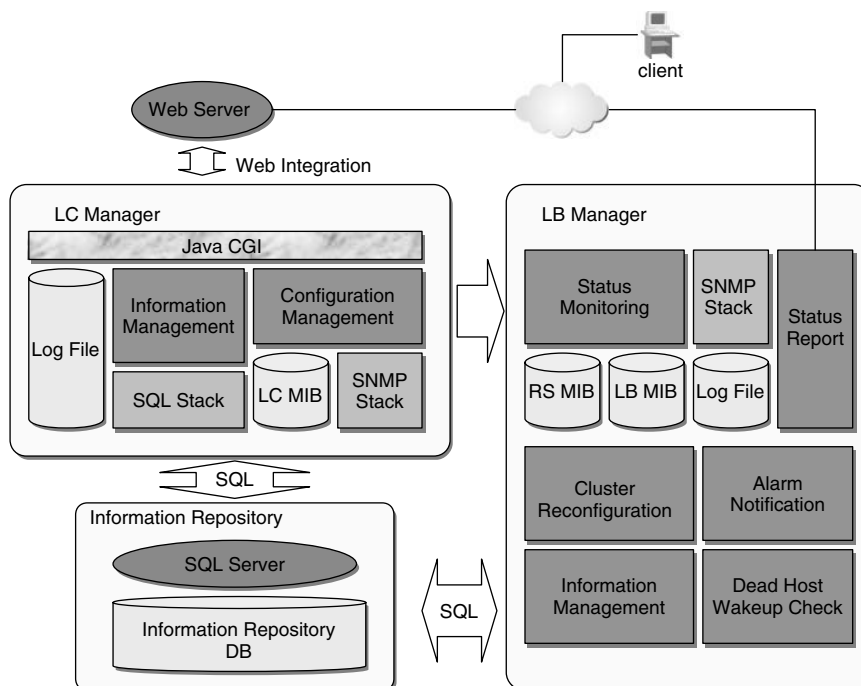


Figure 3. LC and LB manager architecture

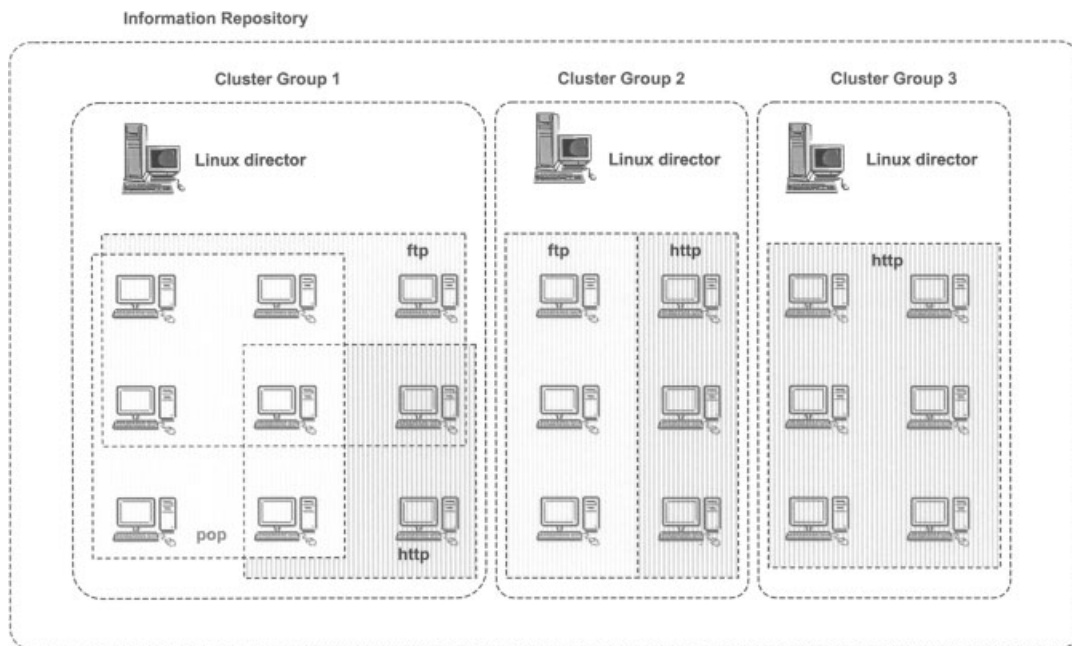


Figure 4. Cluster and service group construction

system to begin its service. All actions by the administrator are recorded into a log file and also are viewed in a web client.

The role of the status monitor in the LB manager is to check the status of a load balancer in each cluster group and take appropriate action when one is down or malfunctions. All significant actions are recorded in a log file and load balancer malfunctions are reported to the administrator by e-mail.

—Information Repository Structure—

To store and maintain host and cluster information we used the DB as an Information Repository. The data stored in the LCMS information repository can be categorized into three groups. One concerns administrator information, the second, host information and the last, cluster group information.

The tables in Figure 5 indicate the attributes of each information group. Host information has several attributes, such as host name, CPU type, memory size, etc. When the administrator adds a new host to an LCMS, the host information is gathered from the newly added host by SNMP, and stored in the DB table. Because one load cluster group can have several service groups and one

service group can have several real servers, there are three tables to indicate cluster information.

—RS Manager Architecture—

The RS manager illustrated in Figure 6 is running at a load balancer in a load cluster group, which is executed when a cluster group is defined and an administrator decides a host as a load balancer. To accomplish this, the LC manager sends a SNMP request message to set the host as a load balancer.

The functions of the RS manager is to monitor the up-to-date status of all real servers in which the service daemon, such as a web server or a streaming server, is running to provide its services to the client. If a real server fails, it removes the real server in its load cluster group and reconfigures the load balancer. This periodic check of the real server state is performed using the SNMP protocol. The SNMP agent in a real server has a module to check the service state of the real server. By this SNMP communication for monitoring the real server, LCMS can reduce network bandwidth overhead and distribute the management workload to hosts in the load cluster group. Another advantage of this manager/agent architecture is that the manager can obtain the exact up-to-date workload of all

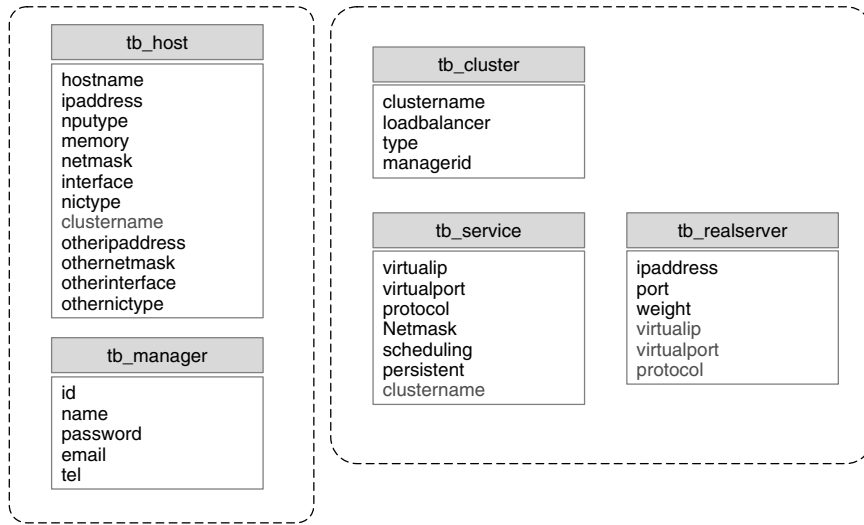


Figure 5. Tables in information repository

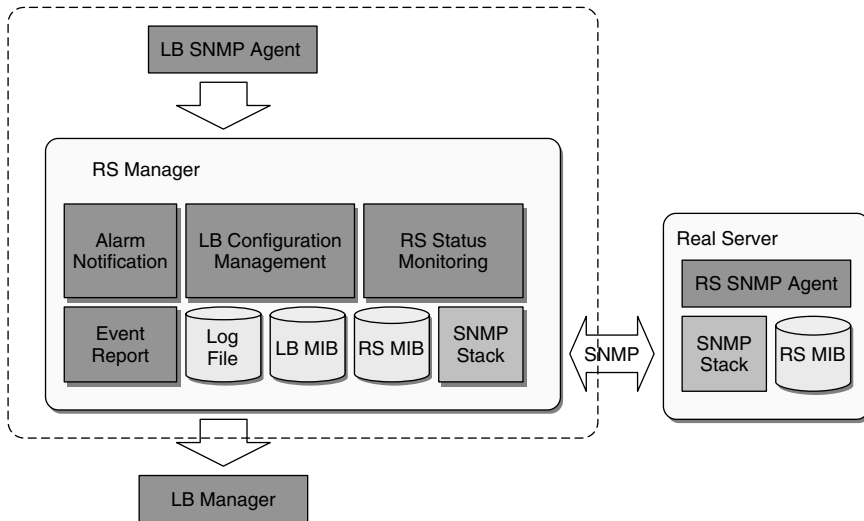


Figure 6. RS Manager architecture

real servers, such as CPU usage, memory usage, the number of current service connections, and outgoing network bandwidth. This information can be used for the request packet-scheduling job. To monitor all real servers, the RS manager uses the RS MIB.

From the time the RS manager is executed, the status of all real servers is periodically monitored and the status data is stored in log files. When a real server fails, the alarm module executes by sending e-mail to the administrator and a message to the LB manager.

—LCMS Agent—

In the LCMS architecture, two types of agents perform different tasks. This corresponds to two types of managers: an RS agent and an LB agent. If a host acts as a load balancer, the LB agent runs on this host. If a host runs as a real server, the RS agent is executed. Every host has the potential to be a load balancer and a real server, and all agents should be able to run in all hosts. The role of each host in a load cluster group is assigned by the manager through an

SNMP message. The agent should be able to set its local host as a load balancer and a real server.

From the consideration of these special features in a load cluster system, the LCMS MIB is defined. Figure 7 shows the LB MIB in the LCMS MIB. When the LC manager assigns a role to a host, the agent sets the *agentType* field in the *generalInfo* category with the value of 0, 1, 2 and 3. If the *agentType* field is set to 1, the OIDs under the load balancer are only meaningful and the host is running as a load balancer. In the same manner, if a host is running as a real server, the OIDs under the real server are meaningful, the backup server and the value of *agentType* is 2.

In one cluster group, one or more services can be created, as in Figure 4. The load balancer in a cluster group can provide more than one service, such as web service, FTP service, and streaming service. Therefore, the real server group can differ from each service. We defined two MIB tables to store service information and real service information in a single load balancer. By the attribute *rsSrIndex* in the *realServerTable* the relationship between real server and service can be found, which is illustrated in Figure 7.

In the current version of LCMS MIB, the security check between the manager and the agent is performed by SNMP v1 community name.¹⁹ This is

intended to be replaced to the user-based security model in SNMPv3¹⁸ in our next version of LCMS.

—High Availability and Resource Management—

The RS manager and the LB manager are responsible for processing faults in the cluster group. The RS manager takes action when a real server fails to perform its assigned service, as presented in the previous section. Next, we discuss in more detail how the RS and the LB manager operate when a fault occurs in a host.

First, when a real server has a fault, the RS manager detects it and removes this host among a request-forwarding list in a load balancer, and reports to the LB manager. The LB manager updates IR that this host is down. Afterwards, the RS manager and the LB manager periodically check this host until this host wakes up again. When this host wakes up, the RS manager reports to the LB manager. Then the LB manager reads the role assigned to this host from IR and send SNMP message for the job assignment. Also, the LB manager can detect that this host wakes up again. If the RS manager has sent a message already, then the LB manager does nothing. But if the RS manager has not sent a message the LB manager

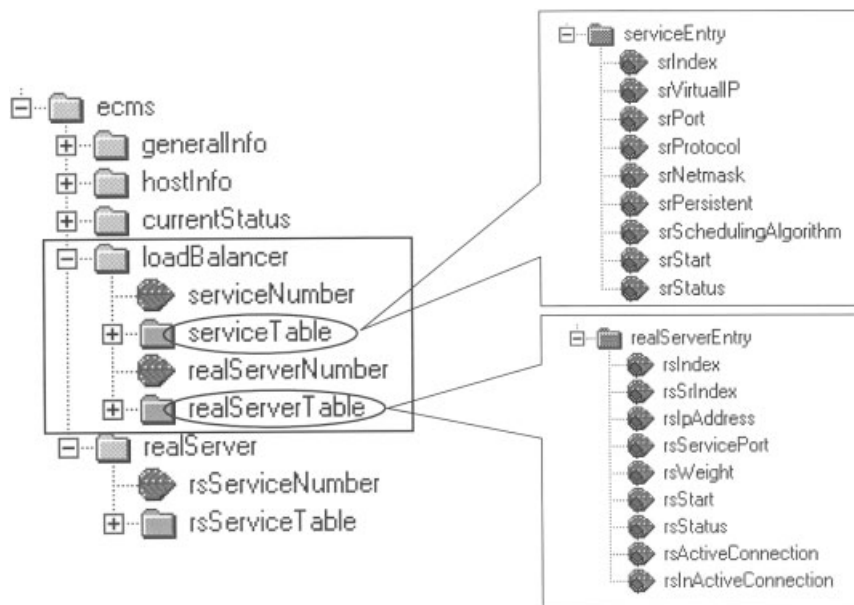


Figure 7. LB MIB in LCMS MIB

performs the job assignment action on this host. In this way, we increased the HA by two managers checking the status of a dead real server.

Second, when a load balancer is dead, the LB manager selects a host that does the least services and has the smallest workload among all real servers in the cluster group. Then the LB manager assigns the role of a load balancer to that real server. During that time the real server works as a load balancer not as a real server. When the original real server wakes up again, the role is assigned back to the original. Among a cluster group, there is no redundant backup server. We increased the efficiency of the resource usage. In our cluster composition architecture, there is no dedicated backup server for a load balancer. That means any real server can function as a backup server when a load balancer fault occurs.

LCMS Implementation

Using this design architecture, we have developed LCMS, which is deployed on a commercial ultra-dense server called Netstech EnterFLEX 2100.²⁹ The software is called EnterFLEX Cluster Manager.

For LCMS, we used MySQL DB²⁵ as an information repository, which runs in the same machine as the LC manager. The UCD-SNMP and its extension are used to make an LB and RS agent at each host. The SNMP agent is programmed using C language. All managers and user interfaces are developed with Java and JSP. We used the joesNMP library²⁶ provided by OpenNMS for SNMP communication between manager and agent. The LC manager is made usually in the form of JSP, and RS manager and LB manager is made of a Java application and JSP.

We used JSP, Applet, and Servlet for the user interface. Tomcat 3.2.1²⁷ is used as a JSP engine, and Apache web server is used as front-end web server. We have selected Java/JSP for LCMS because it is flexible, portable, easy to develop, and has a rich class APIs.

To validate our LCMS, we first developed a streaming service cluster system using a Linux Virtual Server. The streaming service on the Internet is performed by two server types. One is by a web server and the other is a specialized streaming server. The Windows Media Player from

Microsoft,²² Real Server from Real Networks²³ and Quicktime Server from Apple²⁴ are generally used in the current Internet environment. We set up a streaming service cluster supporting these three kinds of media formats and two types of servers in the Linux environment. We used a direct routing method of LVS load balancing. We used an NFS file server to store all stream media files, which is shared among all real servers. The Apache server and the streaming server are used on the real server side. Streaming service web page has been developed using JSP in JAVA.

The screen shots in Figure 8 show selected user interfaces of the streaming service we developed, and the screen shots in the bottom are LCMS user interfaces for information repository operation and new cluster generation.

Conclusion and Future Work

The load cluster system is a cost-effective solution to construct a high-performance Internet service. As deployment ratios of the load cluster system increase year after year, we need a good management solution. In this paper, we proposed a new management method for a load cluster system, which is based on SNMP, the Web and DB.

The load cluster system is a cost-effective solution to construct a high-performance Internet service.

The contributions of this paper are as follows. The first contribution is management workload distribution. In this type of dispatcher method, a load balancer can be easily overloaded and can be a single point of failure. Our proposed LCMS reduces the HA workload in a load balancer by separating HA functionality into an RS manager and an RS agent. The second contribution is that LCMS minimizes the network bandwidth needed for management functions using SNMP. Our last contribution is to efficiently manage the host and load cluster information using an information repository. LCMS follows the current trend of web-based user interface for manager convenience.

Our future work is as follows. It is necessary to measure the exact workload of our LCMS, and

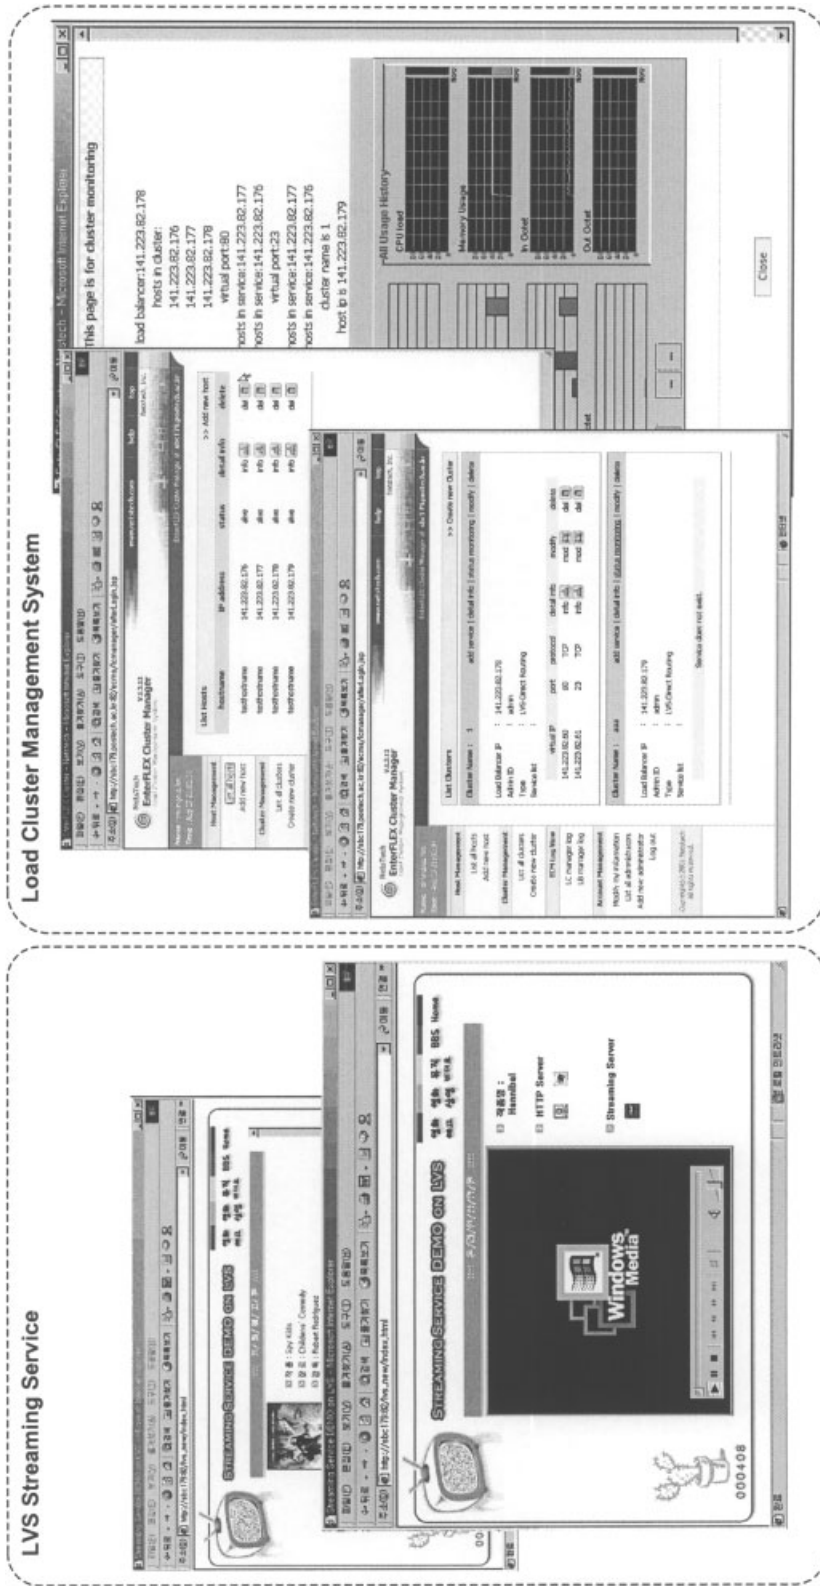


Figure 8. LCMS and streaming service pages

compare it with other load cluster management systems. This LCMS is designed and developed over the dispatcher approach of load balancing, so we plan to apply this LCMS to other types of load balancing technique. We also plan to apply the management information into a request packet-scheduling algorithm. We anticipate that the load balancer will use this information as CPU load, memory usage, and current network bandwidth, that a client requests and service workload will be more equally distributed among real servers.

References

1. Cragle J. Load balancing web servers. *Windows 2000 Magazine* June 1998.
2. Brisco T. DNS support for load balancing. *RFC 1794*, April 1995; <http://www.ietf.org/rfc/rfc1794.txt>
3. Kwan TT, McGrath RE, Reed DA. NCSA's World Wide Web server: design and performance. *IEEE Computer* November 1995; **28**: No. 11, 68–74.
4. Katz ED, Katz ED, McGrath R. A scalable HTTP server: the NCSA prototype. *Computer Networks and ISDN Systems* Nov. 1994; **27**: No. 2, 155–163.
5. CISCO Systems. Cisco IOS server load balancing and the catalyst 6000 family of switches. 1999; http://cisco.com/warp/public/cc/pd/si/casi/ca6000/tech/ios6k_wp.pdf.
6. Foundry Networks. Cutting through Layer 4 hype, White paper, http://www.foundrynet.com/whitepaper_layer4.html.
7. Aron M, Sanders D, Druschel P, Zwaenepoel W. Scalable content-aware request distribution in cluster-based network servers. *USENIX Annual Technical Conference* June 2000.
8. Tewari R, Dias D, Mukherjee R, Vin HM. High availability for clustered multimedia servers. *Proceedings of International Conference on Data Engineering*, February 1996; IEEE Press, 345–354.
9. Robertson A. High-availability Linux project. May 1998—now; <http://www.linux-ha.org/>
10. Zhang W. Linux Virtual Server project. May 1998—now; <http://www.linuxvirtualserver.org>.
11. Zhang W. Linux Virtual Server for scalable network services. *Ottawa Linux Symposium 2000*, July 2000.
12. Mack J. LVS-HOWTO. April 2001; <http://www.linuxvirtualserver.org/Joseph.Mack/HOWTO/LVS-HOWTO.html>.
13. Netscape. *Netscape Proxy Server Administrator's Guide*, Netscape Inc. <http://developer.netscape.com/docs/manuals/proxy/adminnt/revpxy.htm>.
14. Buyya R, Cortes T, Teixeira O. Cluster environment observer. Monash University, Australia, <http://www.csse.monash.edu.au/~rajkumar/ClusterObserver/>.
15. Fischer M. A cluster management software for Windows 9X, NT, 2000, 1999; <http://mufasa.informatik.uni-mannheim.de/lrsra/persons/markus/jobcontrol.htm>.
16. The Symera Team. NCSA Symera, 1997–1998; <http://symera.ncsa.uiuc.edu/>
17. MPI Forum. The Message Passing Interface (MPI) Standard. <http://www-unix.mcs.anl.gov/mpi/>.
18. Blumenthal U, Wijnen B. User-based security model (USM) for version 3 of the simple network management protocol (SNMPv3). April 1999; *RFC2574*, IEEE. <http://www.ietf.org/rfc/rfc2574.txt>
19. Stallings W. *SNMP, SNMPv2, and RMON 1 and 2*, Addison-Wesley, Reading, MA, 1999; 163–203.
20. TurboLinux. Turbo Linux Cluster Server 6 user guide. <http://www.turbolinux.com>, 2000.
21. RedHatLinux. Piranha white paper. <http://www.redhat.com/support/wpapers/piranha/index.html>, 2001.
22. Microsoft, Microsoft Windows Media. <http://www.microsoft.com/windows/windowsmedia/en/default.asp>.
23. Real Networks, Real media technology. <http://www.realnetworks.com/>.
24. Apple, QuickTime. <http://www.apple.com/quicktime/>.
25. Net-SNMP Project Team. Net-SNMP, Sourceforge, <http://net-snmp.sourceforge.net/>.
26. OpenNMS. joeSNMP Library 0.2.5. <http://www.opennms.org>.
27. Jakarta Project. Tomcat. <http://jakarta.apache.org/tomcat/index.html>.
28. Colajanni M, Cardellini V, Yu PS. Dynamic load balancing in geographically distributed heterogeneous web servers. *Proceedings of the 18th International Conference on Distributed Computing Systems, ICDCS*, 1998.
29. Netstech. EnterFLEX 2100 user guide 2.0. <http://www.netstech.com>, July 2001. ■

If you wish to order reprints for this or any other articles in the *International Journal of Network Management*, please see the Special Reprint instructions inside the front cover.