Advanced Virtualization and Sharing of IT Resources for Shared Data Centers

Yasuhiro Oonishi†, Takahiro Tamura, Takayasu Suzuki, Kiyoshi Yanagimoto, and Shin-ichi Kuribayashi

Abstract

We present a shared data center approach that facilitates the outsourcing of information technology (IT) systems and describe one of its key components, the shared load balancer, which is implemented by replacing the conventional physical load balancer function with a virtual load balancer function and helps to implement the sharing of servers and firewalls.

1. Background

To stay competitive in the face of strong competitive challenges, companies are increasingly refocusing their resources on areas of core competence and are outsourcing other operations so they can respond quickly to changes taking place in the business environment [1], [2]. One key point in this development is that information technology (IT) resources must ultimately be treated rather like a utility such as electricity or water.

In this article we provide an overview of a shared data center system, which is the first step toward utility computing. We show that virtualization of the data center makes it possible for several companies to share the same IT resources without interfering with one another. This approach can support the same level of services as conventional dedicated data center systems for as little as one-tenth the cost. One key element of the shared data center is a shared load balancer. We developed one because no commercial implementation is available yet. We overview the shared load balancer and describe its main features.

† NTT Information Sharing Platform Laboratories
Musashino-shi, 180-8585 Japan
E-mail: oonishi.yasuhiro@lab.ntt.co.jp

2. Concept of the shared data center

2.1 Shortcomings of existing data center services

Existing dedicated hosting services involve substantial startup and operating costs because separate equipment is set up and operated for each company. This is prohibitively expensive for many small to medium-sized companies and smaller municipal entities. It can also take a substantial time from when the need emerges until hosting services can actually be provided, so resources cannot be obtained as soon as they are needed. This makes it very difficult for businesses to respond quickly and agilely to changes in business environment. These systems are also generally designed for peak traffic conditions, which means that their usage efficiency tends to be low under ordinary non-peak conditions. Unused resources cannot be readily allocated to another company. Data center services offered by application service providers (ASPs) are relatively inexpensive, but they have several disadvantages such as uniform "one-size-fits-all" menu options and interference between traffic from different companies.

2.2 Overview of shared data center services

Shared data center services provide a better alternative in many cases because they offer the best of both approaches [3]. As **Fig. 1** shows, the shared data center opens up a fundamentally new domain that is roughly the same as ASP-based services in terms of cost, yet provides dedicated hosting services in terms

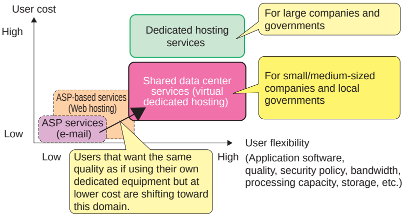

Fig. 1. Domain of shared data center services.

of flexibility (the type of applications that can be used, quality, security policy, bandwidth, processing capacity, memory, and so on). By virtualizing all the key components (server, storage, firewall, load balancer, access bandwidth, etc.) shared data center services allow multiple corporations and municipalities to share the same resources without interfering with one another. This permits services to be made available very quickly on demand at less than one-tenth the cost of conventional dedicated hosting services by sharing devices among at least ten companies (**Fig. 2**). This not only opens up data center services to a far broader range of smaller to medium-sized companies and municipalities, but also accelerates the overall shift toward IT solutions that until now have mostly been exploited more by larger companies and government bodies.

2.3 Steps toward utility computing

The emergence of shared data center services is essentially the first step toward utility computing. This transformation can achieve substantial cost savings by virtualizing and sharing data center equipment and will lead to more advanced services. Eventually, it will lead to the following service capabilities: (1) immediate provisioning of data center services, (2) dynamic allocation of resources reflecting service quality and resource availability, (3) measures to prevent interference among companies sharing resources, and (4) resource switching to ensure that services are not interrupted when equipment fails.

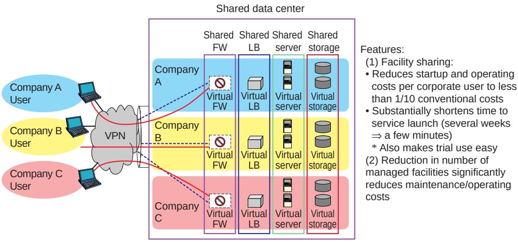

FW: firewall, LB: load balancer, VPN: virtual private network

Fig. 2. Concept of shared data center services.

3. Overview of the shared data center system

3.1 Virtualization and sharing of resources

Our primary objective is the virtualization and sharing of the resources that make up the data center system. Virtualization of access bandwidth and firewalls can be achieved fairly easily using IP-VPN (Internet-protocol-based virtual private network) technology and other virtualization/sharing techniques that have been developed for network resources [4]. For storage and servers, several commercial virtualization/sharing solutions are already available on the market.

The load balancer is a key element of the shared data center system that has not yet been implemented as a virtual/shared commercial product. Moreover, there are hardly any commercial systems without performance degradation that can handle applications based on packet fragmentation*1 for streaming delivery. We therefore decided to develop a shared load balancer (shared LB) to enable the construction of a complete shared data center system.

3.2 Overview of the shared load balancer

(1) Service concept

As **Fig. 3** shows, the shared LB is implemented by replacing the conventional physical load balancer function that is usually set up for each company with a virtual load balancer function. A virtual LB interface is implemented between the shared LB and the server and/or firewall using a virtual LAN (VLAN). Note that all interference between companies is effectively prevented by the ability to manage bandwidth, internal queues, and session numbers in virtual LB units. The ability to provide traffic shaping*2 in virtual LB units is also provided.

Commercial server and firewall equipment that allows multiple companies to share the same equipment without interference is generally quite expensive. On the other hand, there is a wide range of commercial systems designed just to be shared, and this equipment is fairly inexpensive. Essentially we combined a virtual LB with this generic sharing equipment to implement a low-cost server and firewall that support the shared data center services (**Fig. 4**). More specifically, we implemented a function for controlling the number of shared LB sessions that sets the maximum number of sessions (i.e., the total number of sessions and number processed per unit time can both be set at the same time) that can be used by each company for each server and server group. This effectively prevents any particular company from monopolizing the server resources. The same function also prevents an excessive number of access attempts from reaching the server, as illustrated in **Fig. 5**.

At the same time, the shared LB traffic shaping function ensures a fair allocation of bandwidth resource for each company, and again prevents any one company from occupying more than its share of firewall or access bandwidth (toward the network) resources.

In short, the shared LB is more than just a virtualized/shared load balancer. It also helps implement the sharing of servers and firewalls, so it is a key element enabling the construction of an overall-cost-effective shared data center system.

*1 Packet fragmentation occurs when a packet that is too large for the next link is broken up into a number of smaller packets that can be received.

*2 Traffic shaping is a function that can output traffic at a specified uniform rate even when the incoming traffic is bursty by controlling its maximum rate, average rate, and the maximum burst length.

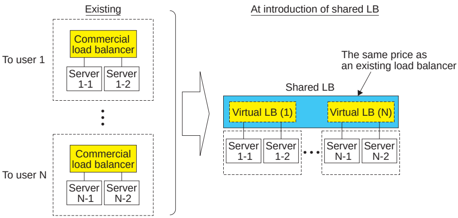

Fig. 3. Effects of introducing shared load balancer (1).

(2) Features of the shared LB system

The shared LB system has seven main features.

1) One implementation supports up to 100 companies (virtual load balancer function)

Each implementation will support load distribution for up to 100 servers without interference. Equipment costs are substantially reduced because multiple companies can share the same equipment. Efforts were made to hold down the cost of the equipment itself to about the same level as the conventional standalone load balancer, so even if the system is just used by a single company, it is still no more expensive than using the conventional physical load balancer. VLAN functionality is supported as noted earlier, so it is easy to interconnect the system to virtual private networks.

2) High-speed processing

Gigabit-per-second-class (wire-speed) processing is achieved using high-speed content addressable memory (CAM) implemented in hardware and network processors and dedicated hardware using network processors. CAM includes comparison logic with each stored bit. The network processor is a programmable device designed specifically for processing network traffic quickly.

3) High-speed streaming delivery processing

In conventional load balancers, packet fragmentation (an essential aspect of streaming deliv-

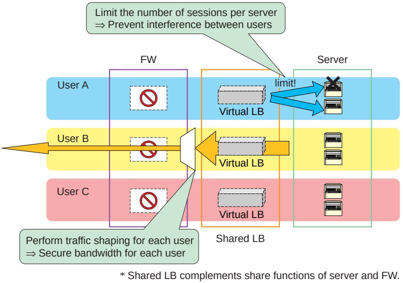

Fig. 4. Effects of introducing shared load balancer (2).

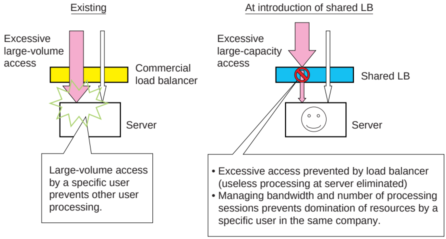

Fig. 5. Effects of introducing shared load balancer (3).

ery) is implemented in software, and this results in substantial degradation of performance. In the shared LB, the streaming delivery is processed at wire-speed using network processors.

4) Support for many distribution capabilities
In addition to round robin, least connection, and other basic server load distribution functions, the shared LB system supports additional functions such as bandwidth control, session number control, and traffic shaping of each virtual LB, server, or server group.

5) Enhanced reliability with VRRP
It has a redundant configuration based on virtual router redundancy protocol (VRRP) for enhanced reliability.

6) Dynamic traffic control and filtering linked to a quality management system and an intrusion detection system. Quality management is provided by a quality management system that monitors the user-perceived-quality of each server and can dynamically select servers for the shared LB.

7) Management on per-virtual LB basis
Equipment and service management are independently implemented for each virtual LB using a graphical user interface-based management system.

4. Future work

We will continue to investigate technologies contributing to the eventual transformation to utility computing. Areas we are studying include fast provisioning of data center resources, dynamic allocation of resources based on service quality and resource demands, and resource switching to keep services running when equipment fails.

References

[1] IBM autonomous computing: http://www.ibm.com/jp/services

[2] Y. Oonishi, K. Yanagimoto, and S. Kuribayashi, "Proposal of the service quality management and control system for making IT resource utility," Technical Paper of the Inst. of Electronics and Communication Engineers of Japan, TM2003-23, Jul. 2003 (in Japanese).

[3] J. Watase, S. Date, T. Tamura, and S. Kuribayashi, "A Study on Adaptive Data Center System," Technical Paper of the Inst. of Electronics and Communication Engineers of Japan, IN2002-156, Jan. 2003 (in Japanese).

[4] H. Hara, "Access system and edge node configuration method for next-generation computer networks," NTT R&D, Vol. 47, No. 4, 1998 (in Japanese).

[5] T. Tamura, J. Watase, M. Tanaka, and S. Kuribayashi, "A Study on a Load Balancing Method in Shared-Data-Center System," Papers from the Annual Meeting on Information Systems of the Inst. of Electronics and Communication Engineers of Japan, B-7-29, Mar. 2003 (in Japanese).

[6] T. Tamura, J. Watase, and S. Kuribayashi, "A Study on a QoS-Adaptive Load Balancing Method," Papers from the Annual Meeting on Information Systems of the Inst. of Electronics and Communication Engineers of Japan, B-7-73, Sep. 2002 (in Japanese).

Yasuhiro Oonishi
Senior Research Engineer, Communication Platform SE Project, NTT Information Sharing Platform Laboratories.

He received the B.E. and M.E. degrees in engineering from Ehime University, Matsuyama, Ehime in 1987 and 1989, respectively. He joined NTT in 1989. He has been engaged in research on IP networking structure and a shared load balancer.

Kiyoshi Yanagimoto
Research Engineer, Communication Platform SE Project, NTT Information Sharing Platform Laboratories.

Since joining Nippon Telegraph and Telephone Public Corporation (now NTT) in 1983, he has been engaged in research on IP networking structure and service quality measurement. He is a member of IEICE.

Takahiro Tamura
Engineer, Secure Communication Project, NTT Information Sharing Platform Laboratories.

He received the B.E. degree in precision engineering from the University of Tokyo, Tokyo in 2001. In 2001, he joined NTT Information Sharing Platform Laboratories, Tokyo, Japan. He has been engaged in research on IP networking structure and network security. He is a member of the Institute of Electronics, Information and Communication Engineers (IEICE).

Shin-ichi Kuribayashi
Project Director, Communication Platform SE Project, NTT Information Sharing Platform Laboratories.

He received the B.E., M.E., and Ph.D. degrees in engineering from Tohoku University, Sendai, Miyagi in 1978, 1980, and 1988 respectively. He joined NTT Electrical Communications Laboratories in 1980. He has been engaged in the design and development of DDX and ISDN packet switching, ATM, PHS, IMT2000, and IP-VPN systems. He researched distributed communication systems at Stanford University from December 1988 through December 1989. He participated in international standardization activities on ATM signaling and IMT2000 signaling protocols at ITU-T SG11 from 1990 through 2000. He is now responsible for the next-generation IP system in NTT Information Sharing Platform Laboratories. He is a member of IEICE.

Takayasu Suzuki
Engineer, Communication Platform SE Project, NTT Information Sharing Platform Laboratories.

He received the B.E. degree in production engineering from Nihon University, Narashino, Chiba in 1990. He joined NTT in 1990. He has been engaged in research on IP networking structure and a shared load balancer.