

Utility Computing Technology for IDC Operation

● Motomitsu Adachi ● Hideaki Orikasa

(Manuscript received June 23, 2003)

Internet data centers (IDCs) have been providing various business and services in recent years, and utility computing is expected to become the new operating technology for IDCs of future outsourcing businesses. By using this technology, users can receive stable, high-quality services without having to make needless investments, even when there are unexpected increases in the volume of access or business. This is achieved by sharing IT resources at the center and assigning the necessary servers and network resources to heavy-load services. This paper describes the benefits that utility computing technology can provide to center users and the basic requirements for IDCs. It also describes the components of utility computing technology, including the technologies for system monitoring, IT resource reallocation, and system control.

1. Introduction

With the popularization of the Internet in recent years, the way in which IT systems are used has changed greatly and there is a growing need to outsource some business tasks. The likely reason for this is that in the current situation, in which data processing requires ever more specialized knowledge due to advances in IT technology, companies have even greater expectations of entrusting system operations to an Internet Data Center (IDC) to stabilize the quality of their business operations. Against this background, future outsourcing businesses cannot simply take on a company's processing tasks, because they will also have to develop new technologies that boast high value-added functions.

This paper explains utility computing, which is a new technology for operating an IDC that meets the above-mentioned requirements. The following section describes the basic requirements of an IDC and explains the advantages of an IDC that makes use of utility technology. Finally, this

paper explains the element technologies required for implementing utility computing.¹⁾⁻⁶⁾

2. Basic requirements of an IDC

An IDC can provide services in various ways, from the so-called housing service whereby the IDC undertakes to manage a user's facilities and operate them in a stable environment, to something akin to a hosting service whereby the IDC runs the business tasks itself and provides services to customers. The actual services are also diverse, from so-called mission-critical core tasks such as inventory control to WWW-based information services such as homepage administration.

Thus, an IDC provides various services, and these services are used in different ways and have various requirements. For example, while mission-critical tasks require high reliability and stable service quality, information services require greater economical efficiency, even though they have lower priority than the mission-critical tasks. It is important for an outsourcing business oper-

ating in such an environment to provide stable services at the quality level required by each task. The following summarizes the basic requirements for constructing and operating an IDC (**Figure 1**).

1) Availability

Provide services continuously without stopping the system.

2) Expandability

Be able to construct a flexible system that can be expanded to a large-scale system as conditions dictate (in contrast to a business network that is constructed according to a fixed plan and remains effective over many years).

3) Security

Prevent illegal access to the IDC and data tampering.

4) Quality

Provide high-quality services while keeping the response time within a specified limit.

Naturally, these requirements must be met in a balanced manner and at reasonable cost.

So-called utility computing technology is gaining attention for its ability to meet the demand for high availability and stable quality. This technology dynamically changes the allocation of IT resources to suit the user's demands and the system operating conditions. In practice, it sequentially reinforces servers and network resources that cause bottlenecks in order to satisfy user demands to keep the service response time

within a specified limit. It also looks for unexpected loads and errors and, when they are detected, flexibly positions IT resources so that services do not stop or are otherwise degraded.

3. Advantages of an IDC that uses utility technology

This section explains the advantages to be gained from using an IDC that has adopted utility technology.

As was explained above, IDC users benefit because, when access to a service increases and the service becomes overloaded, a resource such as a server can be acquired from the shared resource pool owned by the IDC and the resource relocated, thereby preventing a deterioration in response time. Furthermore, an IDC enables users to start small with low initial investment costs, because they do not need to install large facilities to start services run by the IDC. Even if a user prepares the required resources for planned services at the start of system operation, it is difficult for the user to predict and prepare the volume of resources that will be required in the future. However, by using utility technology, a user need only prepare the minimum required facilities, and if there is an overload, the technology prevents a deterioration of service by re-allocating resources from the resource pool. Utility technology also records a log of all IT resources that each user has used, making it possible to apply a quantity-based accounting system whereby the user is charged for the amount used. This approach reduces operating costs. Utility technology offers the following advantages:

- 1) It limits initial and additional investments by making it possible to start small,
 - 2) reduces operating costs by calculating charges using a quantity-based accounting system,
 - 3) flexibly handles load increases and seasonal fluctuations that are difficult to predict, and
 - 4) enhances reliability through maintenance and resource relocation when an error occurs.
- When access to service A increases, this tech-




Figure1
Basic requirements of an IDC.

nology allocates the required volume of servers and network resources from standby resources to prevent service deterioration (**Figure 2**). The standby resources can be located in the same IDC or in a remote center connected by a broadband network. For example, assume that several IDCs located in different companies and prefectures are linked to an IDC operated by Fujitsu as shown in **Figure 3**. In this case, if there is an unexpectedly high volume of access to a service provided by a local IDC, the IDC can still provide a stable ser-




Figure 2
Concept of utility computing.




Figure 3
Linkage between local IDCs and central IDC.

vice by immediately allocating standby resources of the linked central IDC. Therefore, a corporate or local IDC can start a service without requiring large facilities.

4. Technology enabling utility computing

The technology map for utility computing is shown in **Figure 4**. The platform layer consists of a server made up of blade servers, storage, and network devices. On top of the platform layer, there is an organic layer. This layer uses load analysis to perform server resource allocation control and uses organic computing technology to perform on-demand server assignment. The organic layer also uses grid computing technology to hide the operation of individual devices from higher layers. These functions of the organic layer form the basis for utility computing. The functions of the utility computing layer are described below.

1) Policy management

This function translates policy from an abstract level such as the SLA (Service Level Agreement) to executable rules. It creates several commands for a platform by checking the platform status and operation rules. This function also prevents conflicts between different policies.




Figure 4
Technology map for utility computing.

2) Total service quality analysis

This function analyzes the total service quality by measuring various parameters of the Web front servers, application servers, and database. Then, it predicts the IT resources that will be required in the future.

3) Provisioning functions for IT resources

These functions provide the IT resources that are required for stable, high-quality services without bottlenecks. These functions include resource reservation functions that operate according to the results of long-term access analysis.

Moreover, the technology for measuring the platform status can also be accessed from each layer. Each layer can access information about the servers, storage, and network load status. It is also important to consistently control IT equipment such as servers and network devices distributed over a wide area without conflict.

Next, we describe the on-demand IT resource control function, which is a core technology of utility computing and operations.

This function checks the operation status of the system, prevents system overloads, and provides steady service quality. Utility operation is achieved through the dynamic allocation of IT re-

sources using the following technologies (**Figure 5**):

1) Monitoring and measurement technology

This technology monitors and measures the operating status of IT resources.

2) Resource relocation technology

This technology detects the location of system bottlenecks from the results of the monitoring and measurement technology and calculates how to relocate resources to avoid them.

3) Control technology

This technology controls the addition of servers when relocating resources and controls packet destinations in order to distribute the load.

4.1 Monitoring and measurement technology

The monitoring and measurement technology checks the status of servers and networks. It measures a server's instantaneous load factor against its maximum capacity and the available bandwidth of a network being used. The main characteristic of this technology is that it can measure these values without adding a measuring agent to the server or sending a large measurement packet, which would disrupt the system. The following paragraphs introduce server measurement technology and the method of measuring a network's used bandwidth.

To illustrate the technology for measuring a server's operating status from a remote point, we next describe an example in which the target server's response time is analyzed during HTTP communication with the server (**Figure 6**). This




Figure 5
Dynamic function loop of on-demand type control.




Figure 6
Measurement technology for server status.

technology measures the target server's basic performance and then estimates the load based on this value.

The server's basic performance is measured as follows:

- 1) Send get requests to the server in quick succession,
- 2) measure the average data arrival interval x , and then
- 3) obtain the reciprocal $1/x$ as a measure μ of the system's basic performance (requests/sec-ond).

The server's load factor is estimated as follows:

- 1) Measure the HTTP delay average T and ping delay d and
- 2) use $T = (1/(\mu(1-\rho)) + d)$, d , μ to estimate the server load ρ .

The following explains the method for measuring a network's used bandwidth. The method sends a measuring packet (probing packet) to the target network and path and measures the performance of the path. A classic, commonly used measurement method that uses probing is ping-
ing, which uses an ICMP (Internet Control Message Protocol) echo packet.^{7),8)} Pinging can measure the reachability of a host on the path and the packet's round trip time (RTT). Increasing the probing frequency makes it possible to find the lowest RTT value. However, the asymptotic improvement method that is used to increase the probing count in order to improve the measurement accuracy naturally has limits and the influence of error cannot be avoided.⁹⁾

To get around this problem, developers at Fujitsu considered that if they did not evaluate each probe independently but focused on the correlation that could be observed between probes, even if individual probes contained errors, they could eliminate the effect of the errors and thereby enhance the measurement accuracy. They then developed a technology that detects the correlation between probes that arises when probing packets are queued in a bottleneck and measures the per-

formance without increasing the probing count.¹⁰⁾ The correlation between probes characteristically arises when multiple probing packets are queued in a bottleneck at the same time. This occurs because, if probing packets are not queued at the same time, each probe becomes an independent event. However, if probing packets are queued at the same time, a correlation necessarily arises between the earlier and later probes because the later probing packet is sent after the earlier one. The measurement accuracy of the conventional method is directly affected by the difference in the RTTs of the probing packets. In contrast, the method proposed here detects the probing speed at which the rate of positive RTT differences becomes discontinuous. Then, it uses the probing speed and not the RTT value to calculate the bandwidth of the bottleneck (**Figure 7**). The measurement accuracy is therefore not affected by the disparity in the RTTs. This means that even in a congested network that contains other traffic, values can be measured with good accuracy by sending just a small number of probing packets. This method can be used to detect the performance of links and any bottleneck.

4.2 Resource relocation technology

Based on the results of the measurement technology explained above, resource relocation technology determines the location of bottlenecks in the system, eliminates them, and calculates the resource volume and location of the contents required to provide stable services. For example, it calculates the required number of servers from the current server status and growth in the access volume and calculates the most suitable packet route based on the measured used bandwidth. It also calculates the optimal system configuration. For these calculations, it creates a model of the interdependency between the resources that make up the service to be provided. This technology also provides a regulator function for preventing struggles and incompatibilities in standby resources between multiple services.




Figure 7
Relation between RTTs when probing speed is changed.

4.3 Control technology

The control technology controls the servers, the load distribution facilities, and the network facilities of the IDC to realize the optimal configuration calculated above. In concrete terms, it controls the platform and applications in order to allocate services from standby resources. It also checks whether service provisioning is stable as a result of the control and, if not, provides feedback, examines the cause, and isolates the server that encountered the problem.

5. Conclusion

Research and development of the technology presented here is being carried out at the IDC Laboratory established by Fujitsu Laboratories in November 2002. This laboratory has nearly 1000 servers and storage units set up for operation and a mock regionally distributed environment. It is promoting research and development of the construction and operation of IDCs. The utility

computing technology introduced in this paper is undergoing experimentation and evaluation at the IDC Laboratory and is scheduled for practical application at Fujitsu's IDC in the near future.

References

- 1) Sun Microsystems: N1 : Revolutionary IT Architecture for Business.
<http://www.sun.com/software/solutions/n1/index.html>
- 2) Sun Microsystems: N1: Revolutionary IT Architecture for Business Technology Brief - The I-Fabric Architecture.
http://www.sun.com/software/solutions/n1/tb_ifabric.html
- 3) IBM: Managed Hosting - Linux virtual services.
http://www-3.ibm.com/services/e-business/hosting/managed_hosting.html
- 4) IBM: On demand World. 2002.10.
<http://www-6.ibm.com/jp/e-business/pdf/>

- e-business_ondemand.pdf*
- 5) Hewlett-Packard: resource virtualization: utility data center.
<http://h30046.www3.hp.com/solutions/utilitydata.html>
 - 6) Ejasant: UpScale Data Center: Dynamic & Flexible Resource Management.
<http://www.ejasent.com/platform.shtml>
 - 7) IETF IP Performance Metrics WG.
<http://www.ietf.org/html.charters/ippm-charter.html>
 - 8) V. Paxson et al.: Framework for IP Performance Metrics. Request for Comments: 2330, May 1998.
 - 9) V. Jacobson: Pathchar - A Tool to Infer Characteristics of Internet Paths. MSRI, 1997.
<ftp://ftp.ee.lbl.gov/pathchar/msrj-talk.pdf>
 - 10) M. Adachi et al.: NEPRI: Available Bandwidth Measurement in IP Networks. ICC, 2000.



Motomitsu Adachi received the B.S. and M.S. degrees in Electronic Engineering from Waseda University, Tokyo, Japan in 1977 and 1979, respectively. In 1979, he joined Fujitsu Laboratories Ltd., Kawasaki, Japan where he was engaged in research and development of service communication systems. From 1987 to 1989, he was engaged in research and development of heterogeneous communication network architec-

tures and services as a visiting researcher at Bellcore, USA. He is currently researching Internet operation. He is a member of the Institute of Electronics, Information and Communication Engineers (IEICE) of Japan and the Communication Society of the IEEE.



Hideaki Orikasa received the B.S. degree in Mechanical Engineering from the University of Tokyo, Tokyo, Japan in 1976. He joined Fujitsu Ltd., Tokyo, Japan in 1976, where he was engaged in development of operating systems and systems integration. Since 1999, he has been engaged in development of IT service management methods. He is currently engaged in development of a new IT infrastructure technology for

Internet Data Centers (IDCs). He is a member of the Information Processing Society of Japan.