1

Next-Generation Content Distribution Technology in the Broadband Age

ATSUSHI OGASAWARA *(Affiliated Fellow)*

1 Introduction

In recent years, Japan saw the rapid development of its communications infrastructure under the e-Japan Strategy. As a result, ADSL (Asymmetric Digital Subscriber Line) and fiber-optic communications have become widely available, making Japan the country where a broadband network environment is available at prices lower than anywhere else in the world and where more than half (54.5% as of the end of December 2002) of the total population have access to the Internet.

These changes are leading to a sharp growth in the demand for viewing streaming video or other types of motion pictures through the Internet. It is expected that as the fusion of communications and broadcasting advances, this demand will further increase and so will the demand for the distribution of large-volume, high-resolution images (rich content) such as high-definition video. To meet such demands, clearing technical hurdles is an urgent necessity.

This is particularly important to Japan, which owns the world's first-class broadcasting and video technologies as well as world-renowned image content ranging from films to animations to video games.

The e-Japan Strategy II, announced on July 2, 2003, and the e-Japan Priority Program, published on August 8, 2003, also emphasize the importance of the utilization of content, along with digital home appliances, as one of the key elements that will enable Japan to assume the leadership in the forthcoming network society. There are, however, problems that Japan needs to resolve before it can take advantage of its comprehensive strength covering from distribution technologies used upstream in the network to terminal devices located downstream.

While mentioning the possibility that such innovation in content distribution could transform the traditional concept and structure of the network, this article discusses how Japan could buildup technical and industrial competitiveness through this possible transition.

2 Problems in large-volume data distribution

The major problem in distributing information-rich content, or rich content, such as high-resolution motion pictures is the bottlenecks in large-volume data traffic.

In the past, the problem did not draw much attention since the gap in speed between the Gbps-level backbone networks and the "last one mile" to households was extremely wide. However, the gap narrowed as household computers became more sophisticated with the advent of CPUs with up to 3 GHz processing power, and as the communications speeds of the last one mile increased to a range of a few to 100 Mbps, thanks to the widespread availability of ADSL and fiber-to-the-home (FTTH). As a result, the relatively low capacity on the content server side has become highlighted.

Figure 1 is a conceptual image of the bottleneck problem. It shows that attention has mostly been paid to improving the situation in the last one mile, while the server-side bottleneck, which exists upstream in the networks, has been neglected.

The major causes of user complaints in the current network environment are an absolute insufficiency in the performance of servers, a superficial lack of server capacity due to the poor network infrastructure around servers, and the fluctuation of server capacity caused by varied access volume. Internet users are also frustrated even when viewing streaming videos, which are, unlike broadcast programs, accessible whenever they like, because people tend to find spare time during the same time periods of a day, resulting in an overload on the servers and thus an unsatisfying viewing experience for many. This server overloading issue, along with the problem concerning poor access speeds due to insufficient capacity, has emerged as a serious challenge.

For example, in the jurisdiction of the Ministry of Education, Culture, Sports, Science and Technology (MEXT), the National Space Development Agency of Japan (NASDA, or Japan Aerospace Exploration Agency (JAXA) since October 1, 2003) has been offering live pictures of the H-II rocket launches through the Internet. Visitors to their website have been experiencing difficulty accessing the images because of extreme congestion, and even those who did gain access were only given low-frame-rate images that were far from satisfactory. While this is an exceptional case where the overwhelming peak in access number occurs at the moment of the launch, it shows the need for efficient techniques to provide for the required server capacity and to balance the load on servers.

These two issues will become increasingly critical in years to come, when content to be distributed will be larger in volume and higher in resolution.

If the issues of insufficient server capacity and fluctuations in server load were to be resolved solely by installing servers with large enough capacity and scale, only a limited number of businesses that could afford such a large investment would be entitled to enter the content distribution market.

If Japan wants to develop the content business as a new pillar of the Japanese industry, it must encourage the entry of venture firms and small to medium-sized companies. In addition, reducing their investment requirements is yet another critical factor of the development in this area.

One of the reasons that the initial Japanese IT business, which had been emerging in the latter half of the 1990s, burst like a bubble is that fluctuations in server load required companies to build facilities beyond their investment capacity. This is another hurdle that must be overcome.

Figure 1: Network bottlenecks (conceptual image)

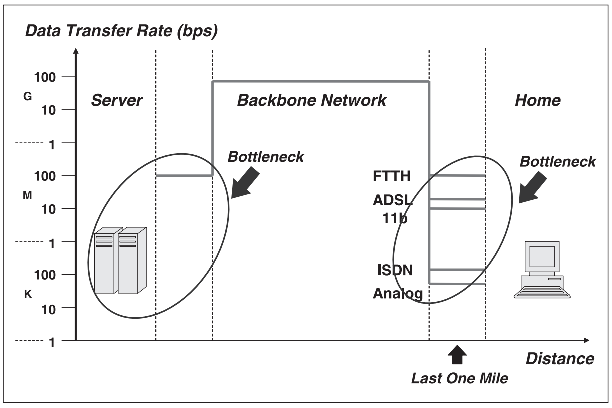

3 Technologies to eliminate the bottlenecks

3.1 Enhancing server performance

(a) Increasing the processor's speed and memory bandwidth

The most essential approach to improving server performance is to increase the speed and memory bandwidth of the processor. Compared with general-purpose processors used for regular PCs, the processors for 3D graphics are required to provide higher operating frequency, wider memory bandwidth, and additional vectorization capability. Japan's vector processor technologies are superior to those of the U.S., as demonstrated in the Earth Simulator, the world's fastest supercomputer. Furthermore, current Japanese video game systems incorporate processors that operate at as high as the 128-bit level (as opposed to 32-bit level Intel Pentium processors for standard PCs), with some even enabled for vector processing. Japan is probably the most advanced country with respect to processor technologies applicable for 3D graphics.

Japan could further advance these technologies and leverage them to construct servers at affordable costs.

(b) Parallel/distributed processing

Another feasible approach to enhancing server performance is the introduction of parallel/distributed processing.

Although increasing the parallel/distributed processing power of a single processor is a possible strategy, there is a more cost-efficient method in which multiple processors are connected to form a cluster that provides greater processing power.

3.2 Coping with load fluctuations

(a) Content delivery network (CDN)

Content delivery network (CDN) services have been developed as one of the technologies that help businesses cope with fluctuations in server load, a problem as serious as the issue of sever performance. As described by Naoki Nishikado of Mitsubishi Research Institute, Inc. in the September issue of Science & Technology Journal, a CDN consists of cache servers (usually mirror servers) that serve as geographically-distributed content sources, and controls them to route users to the nearest server to ensure faster access. This technology is advocated by Akamai Technologies, Inc. of the U.S.

(b) The grid

As compared to content delivery networks, in which data is processed by a group of servers on an independent basis, there is another approach in which multiple computers or processors are connected through a network to scalably perform parallel/distributed processing in accordance with the required computing power. This technology is known as grid computing, an environment in which numerous computers are connected to a network to operate like a single massive-scale computer system through the sharing of CPU resources and memory space. For the successful implementation of a grid computing project, not only computers but also a high-speed network infrastructure is indispensable.

Figure 2:Content Delivery Network (CDN)

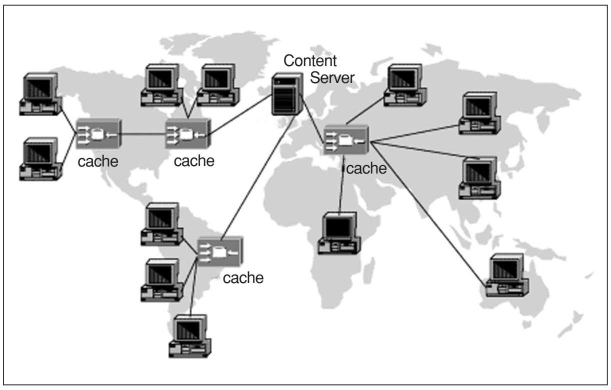

Source: Excerpts from ATT material

In Akamai's content delivery network, which is based on distributed processing on cache servers, computing resources available to a user will never exceed the cache server's processing capacity, that is to say there are as many independent computing spaces as the number of cache servers located across the world. For example, a network that can simultaneously deliver content to hundreds of thousand of people worldwide is actually divided into independent spaces each accommodating a few thousands of people. In a distributed processing environment created through grid computing, by contrast, the computing resources recognizable to a single user would be exactly the same as the entire network resources, which the rest of the hundreds of thousand of people also have access to.

Whether to use cache servers or a grid to distribute content produces a significant difference in the results of load balancing. While the difference may be limited as far as the content is not time-sensitive or interactive, if the content

Figure 3:Grid server for gaming (IBM)

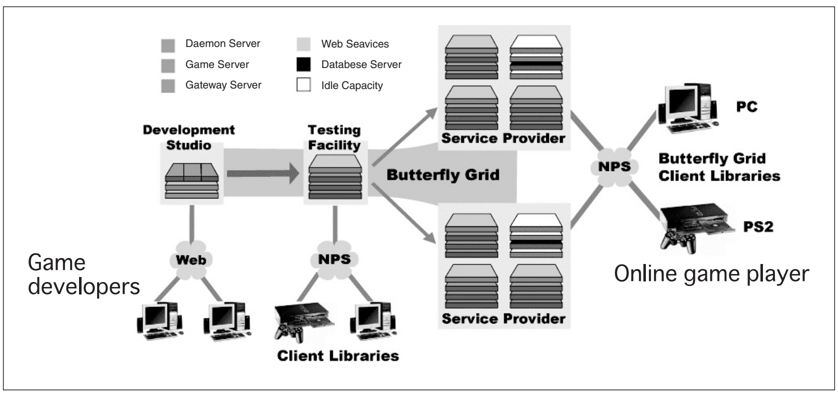

requires users to participate and substantially interact with each other (e.g., participation in politics in an electronic government, educational applications, and network gaming or other personal entertainment applications), the difference is considerable.

Another advantage of grid computing is that it frees content providers from the need for excessive server investment to ensure accessibility even during peak periods. They can instead purchase computing resources as required in the same way as they purchase electricity. This would also contribute to facilitating the entry of venture or small to medium-sized companies to the content industry.

4 Background of grid computing technology

Grid computing has been developed mainly for the purpose of scientific high-performance computing (HPC), for which mainframes and supercomputers at universities and national research institutions are connected through high-speed networks.

More recently, encouraged by the launch of large-scale projects in the U.S. such as TeraGrid of the NSF and the Biomedical Informatics Research Network (BIRN) of the NIH (reported by Masao Watari in the sixth issue of this bulletin), grid research in the field of high-end HPC systems that use high-speed communications lines is making rapid progress.

In 2002, under the Bush administration, the U.S. government began integrating federal IT research and development initiatives such as HPCC, NII, NGI, and IT2 into the collective program called the "Networking and Information Technology Research and Development (NITRD)." This move could be considered as a symbol of a growing fusion between communications network technologies and information systems technologies, a fusion in which grid computing is expected to serve as a major pillar.

The U.S. is also actively promoting standardization and advocating the Globus Toolkit, middleware for distributed heterogeneous environments.

The wave of grid computing is even sweeping across the business sector, in addition to the academic sector.

IBM, which has been making earnest efforts for the application of grid technologies to B2B since around 2000, proposed the Open Grid Services Architecture (OGSA), which combines grid computing with Web services, as an extension of the Globus Toolkit.

The proposal was initially met with great excitement among experts, while it received cooler response from corporations, because even if the OGSA could be implemented relatively early within the corporate intranet framework, further extensions would require an additional network infrastructure, an obstruction that corporations thought would not be eliminated soon.

However, as broadband communications quickly became widely available and as the development of the content distribution business increased the need for solutions to the server-side bottleneck problem and to server load balancing, the demand for grid computing dramatically rose.

A new movement in the U.S. is that Oracle announced "Oracle Database 10g" ("g" represents "grid"), a new database based on grid computing, at OracleWorld 2003, which started on September 8 in San Francisco. At the event, Hewlett Packard (HP) Chief Executive Officer Carly Fiorina mentioned in her speech that HP would actively move toward grid servers in the next three to five years. These facts underscore that grid computing is rapidly being accepted not only in IBM but also across the entire IT industries of the U.S.

In Japan, the opening ceremony of the "Center for the Grid Research and Development (NAREGI)," an institution established by MEXT, was held on July 1, 2003. The Center was built as a new R&D base for NAREGI (National Research Grid Initiative), a five-year project that started in April 2003, based at the National Institute of Informatics (Director General: Yasuharu Suematsu). The NAREGI project aims at "the creation of an environment for next-generation computing systems, with the view of strengthening Japan's international competitiveness in the information and communications field." Through collaboration between industry, academia and government, NAREGI intends to develop a high-performance computing environment on the basis of grid computing in which multiple computer resources are effectively combined to substitute for a single supercomputer. A major objective of the project is the research and development on application software for use in the nanotechnology and biotechnology fields, which are expected to be critical fields of science in the next generation. The scientific fields the NAREGI project actually embraces are so extensive as to include even areas that are closely related to industrial applications, such as new communications principles, electronic devices, optical devices, molecules, and the development of new drugs.

The Ministry of Economy, Trade and Industry (METI) also launched its own initiative, "Business Grid Computing Project," on July 15, 2003. Chaired by vice president of Waseda University Yoichi Muraoka, the project seeks to support the development of middleware that can primarily be used for grid computing.

What made METI take this step was the awareness that grid computing technology will be the key to acquiring international competitiveness in the information systems sector. METI also recognizes that this sector is currently dominated by the U.S. from upstream to downstream network elements, under the influence of such industry leaders as Intel in the MPU (Micro Processing Unit) domain, Microsoft, with Windows, in operating systems, IBM and HP in data distribution servers, and Oracle in databases.

When grid computing becomes widely applicable for industrial purposes in addition to scientific computations, it will be the time when a transition starts in computer and network architectures. Having realized that this will present a good opportunity for Japan to assume the world leadership in the information systems sector, METI made the strategic decision of launching the project.

5 Directions of Japan's research and development

While deploying IBM, HP, or other American-made servers in combination with Akamai's content network services may be one solution, there are other options Japan could consider. As mentioned in the beginning of this article, Japan is armed with sophisticated 3D graphic technologies that could lead the global IT industry and unique video content such as animations and video game software, as it faces the forthcoming full-scale broadband network age. Japan is also the world leader in the development of digital home appliances and in optical communications technology. The grid is a highly potential solution that could help Japan take advantage of these technological strengths in the information, communications, and electronics sectors in an integrated manner.

There are two major methods to create a grid computing environment. In the first method, a grid is constructed by interconnecting the company's servers with a view to balancing load. This could be implemented on different scales, ranging from a grid created on a network or a cluster formed within a data center. The other method is to build a large-scale grid embracing user-side (client-side) computing resources.

In the latter case, in particular, processors embedded in digital home appliances and the device proposed as the "home server" may comprise a grid. By linking up these appliances and devices via high-speed fiber-optic cables, in an extreme case, even a massive-scale fast grid that covers all of Japan might become a reality.

Such a concept deserves serious attention from the viewpoint of delivering content, because it suggests the possibility that Japan as a whole could function as a huge server that offers to the world its exclusive high-resolution image processing technologies, animations and game software. This would be feasible if the nation could find an appropriate solution to interconnect processors in video game systems and digital home appliances, Japan's two areas of strength, instead of U.S.-dominated PC processors, and to build high-performance server systems by exploiting its world-class optical technologies, such as semiconductor lasers, and fiber-optic networks developed through the e-Japan program.

If this becomes a reality, Japan will come to play a much more important role as the hub of the high-speed networks linking the three continents, namely, North America, Europe and Asia.

A major problem in the delivery of high-resolution video content is delays that occur in long-distance data transmissions. For example, it is difficult to ensure Quality of Service (QoS) for high-resolution motion pictures transmitted to somewhere in Southeast Asia from the U.S. via Japan, because of signal delays.

Equally distant from North America, Europe and Asia (provided that an eastbound route from Europe is established), Japan is most conveniently located to play the key role in content distribution. Japan could largely benefit in economic terms from playing such a part, because it is obvious, even without citing examples in the U.S., that information critical to business activities, as well as spillover information, converges to the hubs of information and traffic. Acting as such a hub would have a significant impact on the domestic business community not because the illegitimate collection of personal or corporate information might be possible, but because even overall statistics concerning what type of information is being sent to which region would be valuable and useful for business.

6 Research and development on constituent technologies for grid computing

The MEXT-led NAREGI project, mentioned earlier, pursues the following themes as the constituent technologies of grid computing, or high performance computing as they mention it.

(1)Research and development of grid infrastructure software

- Research and development of resource management in the grid environment
- Research and development of grid programming modules in the grid environment
- Research and development of upper layer grid software/environment and grid application development tools
- Research and development of integration and operation technology for grid software
- Research and development on the adaptation of nano-simulation software to the grid environment

(2)Development of networking technology

- Research and development of network communication infrastructures

These themes are studied at the following institutions.

Industry: Fujitsu, Hitachi, NEC, etc.
Academy: National Institute of Informatics, Institute for Molecular Science, Tokyo Institute of Technology, Osaka University, Kyushu University, etc.
Government: National Institute of Advanced Industrial Science and Technology, ITBL Project, etc.

Among them, Fujitsu, Hitachi and NEC also participate in METI's Business Grid Computing Project. The project draws a fast-paced schedule and intends to build prototypes of grid computing software by the end of fiscal 2003, followed by field testing in 2004 and commercialization in 2005.

In addition to the software technologies being developed through these projects as elements directly comprising grid computing, namely, middleware and application software technologies, there are other areas where research and development efforts should be intensified. They are device technologies and hardware systems, in which areas Japan is already particularly competitive.

To develop grid computing environments even more efficiently, optical technologies need to be advanced. This involves not only the construction of additional backbone infrastructures, such as optical routing systems, to allow for large-volume long-distance communications, but another leap in fiber-to-the-home (FTTH) solutions to initiate a move from the current 100 Mbps range to the Gbps range.

Yet another critical element of grid computing is the promotion of the use of optical wiring for the processor's internal data transfers and the computer's internal connections between processors and memory, between chips, and between boards.

The U.S. took action ahead of Japan and its Defense Advanced Research Project Agency (DARPA) announced on September 11, 2003 that it would fund $30 million over four years to support the joint research project of IBM and Agilent Technologies (a spin-off from HP, providing testing and measuring instruments and semiconductor products), which pursues "terabit per second optical interconnect technology for multiprocessing servers." DARPA says the effort is part of the High Productivity Computing Systems (HPCS) initiative, which aims to develop interconnect technology that delivers 40-Tbps performance by 2010. Such technology, if used for connections much longer than chip-to-chip or board-to-board distances or for consumer products, will contribute to improving to a large extent the performance of cluster servers and grid servers.

Japan, as the pioneer and the world leader in fiber-optic communications technology and optical device technologies such as the lasers for long-distance communications systems and the vertical-cavity surface-emitting laser (VCSEL), should undertake greater research and development efforts than the U.S. is making, in order to maintain the current dominance in these optical fields.

Other constituent technologies that need to be developed include a home information appliance network based on, like IPv6 networks, Japan's original technologies, middleware to integrate information appliances and game systems, and operating systems suitable for these terminals. In other words, the key constituent technologies of grid computing will be achieved by pioneering the development of new architectures that combine software and hardware.

7 Conclusion

The above chapters discussed the importance of distributed processing and optical technology, with an emphasis on grid computing technologies as a tool to meet the growing demand for the distribution of rich content such as high-resolution motion pictures. The most remarkable point is that these technologies could trigger a major change in the 30-year-old concepts of the Internet and the network, both of which originate in ARPANET of the U.S.

The current Internet structure is by far dominated by the U.S., including TCP/IP and other protocols (procedures), fundamental technologies, the top-level route servers that manages IP addresses and domain names, server systems that deliver data, operating systems running on terminals, and the computer as the basic hardware component.

If the network architecture moves from the traditional server-client model to the grid model of shared computing resources, however, Japan will have a chance to pave the way for next-generation networks and computers that are of unconventional designs incorporating Japan's original technologies for IPv6 and digital home appliances. This could allow Japan to build a competitive advantage with its collective strength in the content and other industries. From a technical point of view, it is still a very challenging task to collect computing resources of an indefinite number of users (client computers) to construct a grid, while ensuring delivery control. This hurdle, however, could be overcome by first resolving technical difficulties in a small-scale cluster server system and then gradually expanding the range of the target.

Such an effort would benefit Japan not only technically, but in a sense that it might help transform the business structure of Japanese companies. The electronics industry, which is, along with the automobile industry, Japan's most internationally competitive sector, is suffering a serious decline in profitability in home appliances and audio-visual equipment, mainly because of the rise of China.

Furthermore, once-hopeful Internet businesses are in most cases having difficulty making a profit, as a result of the failure to earn wide acceptance from consumers about the concept of the B2C (Business-to-Consumer) business, especially that of the model where services are offered online for a fee. Their profitability has been also hurt by excessive investments to meet the growing demand for server capacity, as mentioned earlier.

By contrast, U.S. firms, such as IBM and HP, which have strength in information distribution functions upstream in networks, are enjoying high profitability.

What these facts suggest is that while Internet businesses are essentially fee-based business, customers in the B2C sector are reluctant to accept such a business model except for e-commerce in which purchased products are physically delivered to them.

A possible action for Japan to take under the circumstances is to actively move into B2B (Business-to-Business) commerce, which exists upstream in networks and has been dominated by the U.S. By combining B2B with the C (Consumer) sector, the basic business model for consumer electronics and other traditional industries, Japan could establish what one might call B2B2C, which would be a model unlike either the traditional Japanese style or U.S. style.

Here is a possible scenario to roll out B2B2C commerce. While Japanese consumer electronics manufacturers are currently proposing home servers to households, it is usually not very easy for consumers to understand the features of the equipment and benefits of spending a considerable amount of money to purchase it. To address this problem, the industry can launch a program in which the unused power of home servers may be loaned to those who are in need, and the owners of home servers can receive compensation for the loaned computing resources in the form of electronic money. Such an incentive to consumers will accelerate the proliferation of home servers.

MEXT's NAREGI project seeks for the development of middleware to construct grids and new application software for the nanotechnology and biotechnology fields, while METI's Business Grid Project aims at the development of middleware and application programs for the B2B sector. A possible addition to them will be a project that focuses on the C sector with emphasis on content distribution.

Similarly important as such comprehensive initiatives is the research and development of, not to mention software technologies to facilitate these network revolutions, detailed constituent technologies including optical technologies (communications, optical wiring, optical interconnection and optical routing), integrated computer architectures incorporating digital home appliances, and new business models.

(Original Japanese version: published in October 2003)