Hitachi’s Involvement in Networking for Cloud Computing

Chikako Tanaka
Takashi Ezaki
Yasuhiro Uchiyama, Dr. Eng.
Kazuma Yumoto
Takeki Yazaki

OVERVIEW: The trend toward use of the cloud for IT systems is bringing major changes to the environment in which networks operate. The cloud handles large quantities of data and, in addition to high-speed and high-capacity networks, also requires higher levels of stability and reliability. There is a need to conduct overall optimization of IT resources in customers’ information system platforms, including their use of the cloud. With network virtualization as the key technology, Hitachi is proceeding with the supply of operational management systems and other network products and their installation at customers’ sites, with the aim of making improvements in operations and utilizing IT resources at data centers. These enhancements to networking capabilities are intended to contribute to Hitachi’s cloud business through the provision of highly reliable cloud platforms. For carrier networks that support cloud services, Hitachi is also working on next-generation transport that provides higher reliability and availability in addition to satisfying basic requirements such as responding to traffic growth and achieving flexibility in service provision. While Hitachi is currently considering entering overseas markets, this article describes its activities in Japan, which will form the base for any such initiatives.

INTRODUCTION

AS advances are made in cloud computing, corporate data is increasingly being consolidated in data centers and there is a trend toward making active use of cloud services based around these data centers. This is driving demands for data centers to achieve greater reliability and operational performance in their role of providing the infrastructure for corporate systems.

The basic requirements for networks include the following:

(1) Handle the rapid growth in traffic resulting from the consolidation of computing.

(2) Ensure the reliability needed for using services over a network.

(3) Improve data center operations to maintain reliable operation, including fault response, operational monitoring, and the introduction, migration, and reconfiguration of cloud services.

To cope with the growth of traffic concentrated around data centers in which IT (information technology) equipment is increasingly being consolidated, a bandwidth of 10 Gbit/s has become widely used for internal data center communications, and there is demand for even faster speeds. To ensure trouble-free operation, a level of reliability (delay, response performance, and operational characteristics) equivalent to that of a leased line is typically demanded for services accessed over a network. Also, to improve operation, there is a need to implement flexible system configurations with improved efficiency of use that are able to respond promptly and appropriately to extensions to or reconfiguration of the IT equipment in data centers, including network equipment. It is also necessary to optimize equipment based on business continuity considerations, and to take account of the formulation of operating procedures.

This article describes what Hitachi is doing to satisfy the basic requirements for cloud networks and its work in the field of data center and carrier networks.

SATISFYING BASIC REQUIREMENTS FOR CLOUD NETWORKS

Data Center Internal Networks

For the LANs (local area networks) at data centers, progress is being made on standardizing low-latency, lossless Ethernet communications and LAN/SAN (storage area network) integration techniques with the aim of improving reliability and achieving better maintenance and operational characteristics through measures such as reducing the number of IT devices and cables. In response to this trend toward integration of switches in data centers, Hitachi is working on the supply of network equipment, including the next generation of Ethernet switches.

Hitachi is also working on supplying network solutions to help improve data center maintenance and operation, including the efficient collection of network configuration data, minimizing the scope of faults, and reducing the cost of installation or migration from existing systems. In particular, Hitachi is strengthening its support for network virtualization (networking between virtual systems), including the supply of auto-provisioning functions that automate the setting and management of network configuration data in order to reduce system migration, installation, and other related costs. In addition to this, Hitachi is also enhancing functions for interoperation between data centers to allow for data center backup and situations in which IT equipment from a number of different data centers is used.

Carrier Networks

The growth of the cloud requires that carrier (telecommunication company) network systems provide communication services with even higher levels of reliability than before. The provision of communication services faces the following two challenges:

(1) Network services such as the Internet and telephony are becoming more complex as each of them has system-specific transport systems.

(2) IP (Internet Protocol) networks are largely autonomous and distributed best-effort systems in which communication routes are determined autonomously. This limits the ability to guarantee service quality.

The MPLS-TP (Multi-protocol Label Switching—Transport Profile) packet-based transport method has attracted interest as a means of overcoming these challenges. Unlike conventional routers and switches, networks that use MPLS-TP aim to provide an end-to-end service quality guarantee. A feature is the provision of a route control function for managing and maintaining communication routes over the packet network to allow control and management of the entire network. Hitachi is working on developing and implementing this technology.

It is anticipated that the trend toward cloud computing will expand the use of IT systems globally by improving the efficiency with which they are used and creating new business opportunities. Data centers provide the infrastructure for cloud services and Hitachi is working on the networks they use internally as well as the carrier network systems that support them.

DATA CENTER NETWORKS

Individual System Optimization to System-wide Platform Optimization

In the past, most enterprise information system platforms have had silo-type designs in which the system's servers, network, storage, and administration operated independently of other systems. Recent years, however, have seen growing optimization of enterprise information system platforms in response to strong pressure for cost reduction in the user companies.

The following describes what Hitachi is doing to provide network solutions for companies that want to optimize their information system platforms.

Examples of companies embarking on system upgrades aimed at optimizing their overall infrastructure by using virtualization are becoming more common.

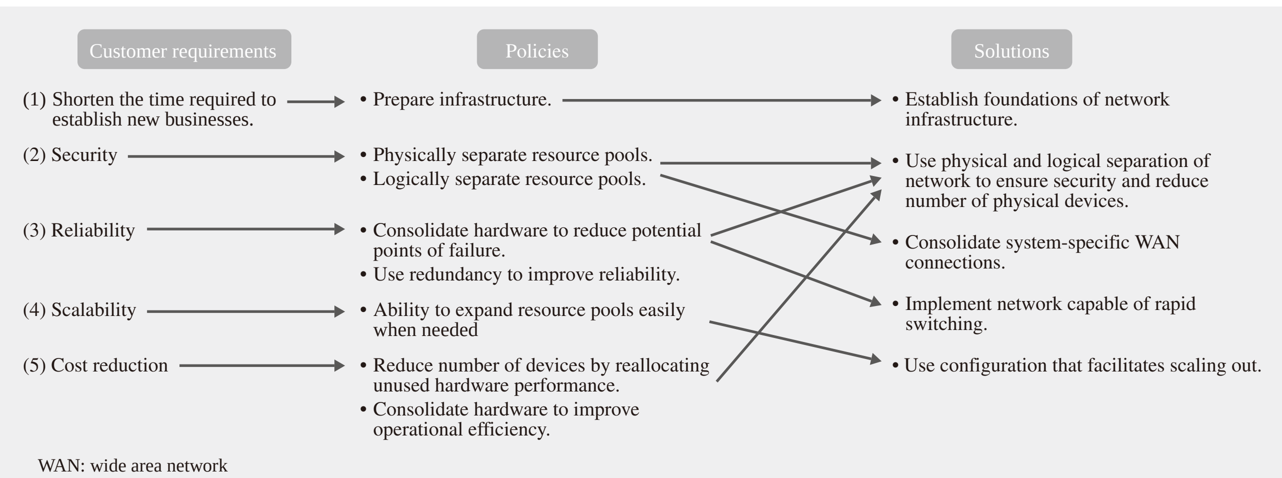

Fig. 1—Policies for Satisfying Customer Requirements.
The policies and solutions for satisfying the typical requirements of customers who want to optimize their network infrastructure.

An analysis of their typical requirements broadly divides them into the following five:

(1) Shorten the time required to establish new businesses.

(2) Security

(3) Reliability

(4) Scalability

(5) Cost reduction

An analysis of the requirements for infrastructure optimization shows that a common feature is that companies are shifting their focus from individual system optimization to system-wide optimization. Based on these requirements, Hitachi has taken the following view on how to deal with networks (see Fig. 1).

Advances in virtualization have made it possible to provide resources for servers easily and quickly. Similarly for networks, to keep up with faster resource provision, it is necessary to put in place highly flexible networks that are able to respond to changing server resource allocation. A way of achieving goals such as security and minimizing the risks associated with configuration changes is to split up the resource pool physically for production, development, and other systems, and to logically isolate communications between resource pools so each can respond as a network. For reliability, potential points of failure are kept to an absolute minimum through hardware consolidation. Reliability is also enhanced by using simple redundancy functions. For expandability, networks are designed in a way that makes it easy to add hardware when expanding the resource pool becomes necessary in the future. Finally, in terms of cost, the past practice of configuring systems independently and on-premises meant that there was no way of making use of any spare capacity in hardware resources. In contrast, sharing a single hardware resource across a number of systems reduces both the number of devices and their operating costs.

For corporate customers, Hitachi has implemented data center networks with the following features in line with these policies (see Fig. 2).

Racks have edge switches fitted to allow connection of all servers in the rack, and lower-level core switches are provided for consolidating the edge switches. When the number of racks increases, additional capacity can be installed simply by connecting cables to these lower-level core switches. Furthermore, a multi-ring network can be formed between these core switches by providing higher level core switches to consolidate the lower-level core switches.

Ring networks can handle future network expansion and can be switched over quickly in the event of a failure. Similarly, they can also be used to implement large and highly reliable layer 2 networks and can cope with situations that require seamless

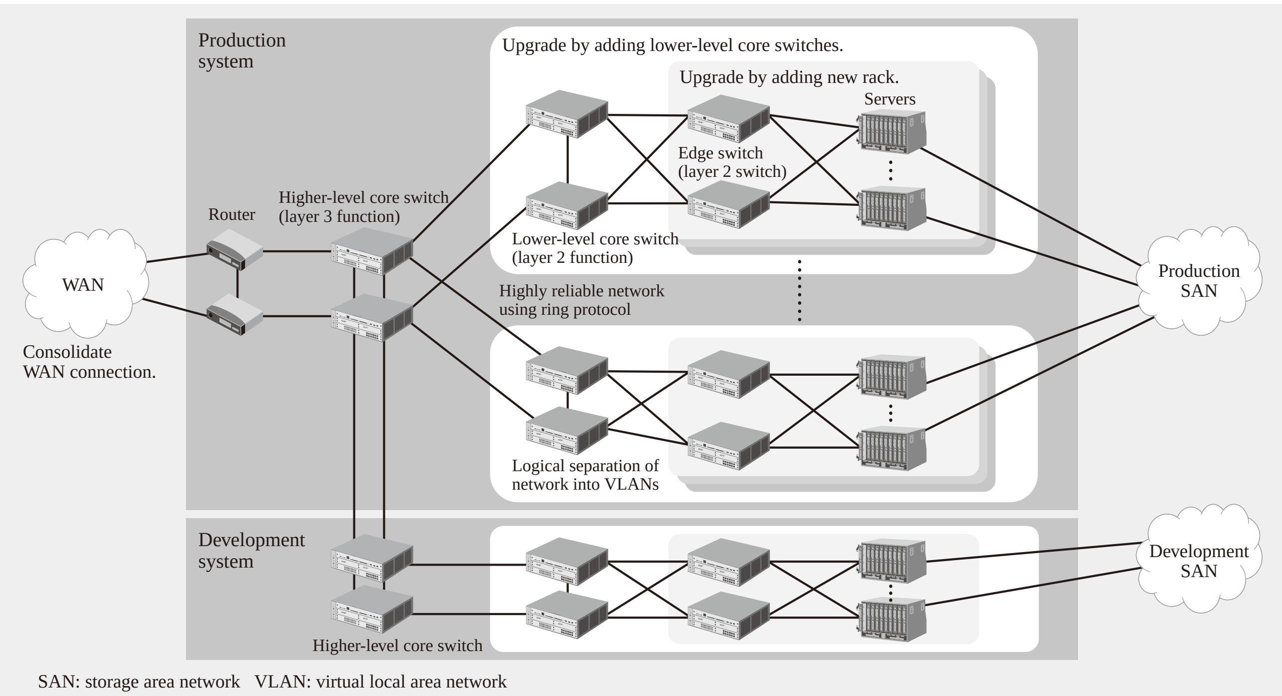

Fig. 2—Example Configuration of Data Center Network.

This example shows a highly reliable network for data centers designed to satisfy customer requirements.

networks such as in distributed systems or live migration of servers.

In the case of networks used for mission-critical systems, in particular, it is necessary to split the network physically. This is implemented by replicating the configurations under the high-level core switches and using a ring connection between the high-level core switches when communication between systems is required (such as for operation and monitoring communications). Similarly, when communication between systems (segments) is required, this is made possible by using a high-level core switch with a routing function. Alternatively, if inter-system communications is not a requirement, this can be achieved using a configuration that does not permit inter-segment communications.

As technology advances further, it is anticipated that future data center networks will develop along the following three lines:

(1) Distributed networks integrating LAN/SAN
(2) Unified management of integrated networks
(3) High-speed networks between data centers

For operational reasons, there will be a need for simple and easy ways of expanding capacity as LAN/SAN integration progresses. In response, Hitachi is working toward making possible unified management of multiple network devices and is seeking to implement integrated distributed networks that make it easy to add hardware with zero configuration. Meanwhile, demand for disaster recovery is also anticipated as companies formulate BCPs (business continuity plans). Accordingly, networks will also require seamless connectivity, including high-speed backup and synchronization of data between geographically separated data centers. While 10-Gbit/s NICs (network interface cards) are now in wide use, it is anticipated that even faster networks such as 40 Gbit/s or 100 Gbit/s will become the standard in the future.

Improving the efficiency of connections between data centers is an important factor as demand for overseas connections is anticipated to grow from, for example, companies setting up networks between their overseas offices as they become increasingly global. In existing TCP (transmission control protocol) networks, throughput is influenced by physical distance. This means that resources may not be used effectively even if the speed of the network infrastructure is increased. Therefore, Hitachi is undertaking steps to achieve higher throughput (more efficient use of bandwidth) between data centers.

For the future, Hitachi intends to propose and supply optimum networks for its customers that take account of developments in network standardization and are designed to handle future requirements.

Reducing Operating Costs by Automating Provisioning

On systems such as those used for cloud computing that require frequent changes to the network configuration, it is desirable that the various services can be provided easily using automated procedures. The cloud pools servers, storage, and networking as common resources that can be allocated on demand. Cloud systems require the design and configuration of mechanisms that can deliver services to users promptly, and with a minimum of intervention by the department responsible for management and operation.

However, when the system operation is based on use of virtualization, there are a number of issues that are over-complex if conventional operating methods are used for the servers, network equipment, and other hardware. There is also a considerable risk that a single operational error might affect the entire cloud system.

Although server virtualization makes the allocation of resources easier than in the past, network provisioning is required when making changes such as adding or moving servers to provide users with resources, which are part of a larger cloud system including a network. Example operations that can be performed on virtual servers include adding new or additional servers and moving or removing servers. Such operations create a need to control the information for associated resources, including network equipment and virtual switches.

Accordingly, it is necessary to maintain a database for the unified management of information about servers, networks, and other resources, and to provide a way of managing resources from a GUI (graphical user interface) based on the different use cases. Hitachi intends to provide software with functions that include automatically maintaining consistency across a wide range of different devices, automating operations, and preventing operational errors (see Fig. 3).

CARRIER NETWORKS

POTS Support

The movement of more IT systems to the cloud is causing major changes to the carrier network environments that provide the infrastructure of the Internet. With the consolidation of computing leading to dramatic growth in traffic levels, ways of ensuring high quality and highly reliable transport are essential

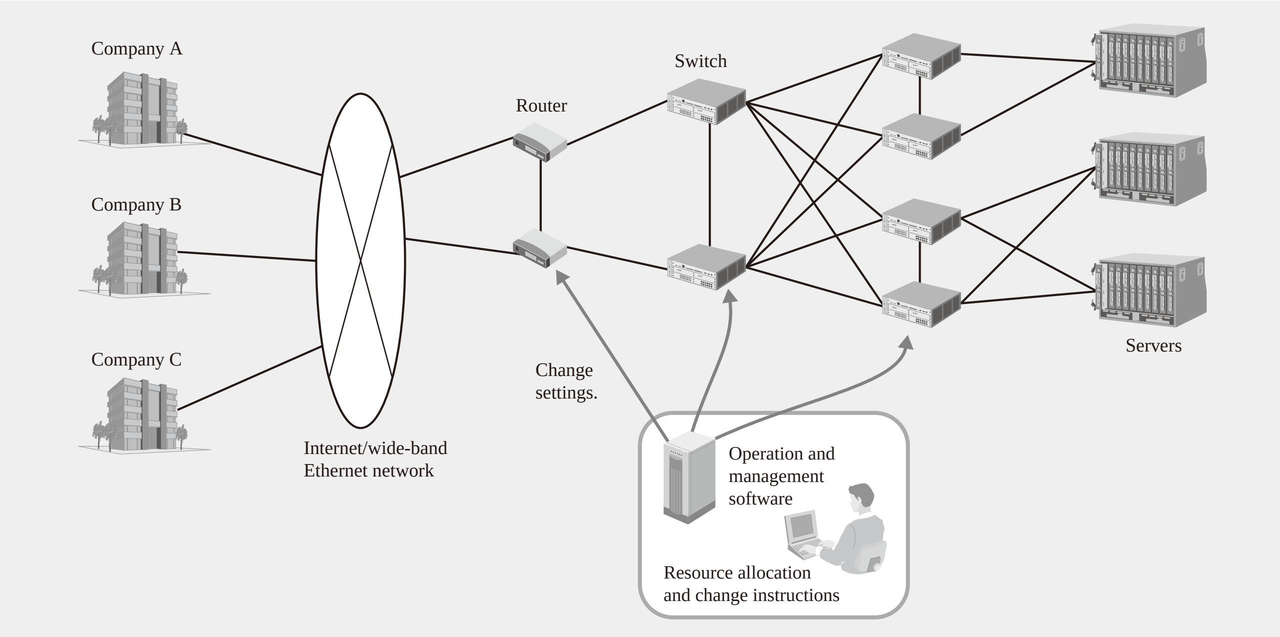

Fig. 3—Overview of Network Operation and Management Software.
The network administrator uses the GUI (graphical user interface) to allocate resources to suit different use cases and to update the settings on the network equipment.

for making services available via networks. Carriers also need the ability to incorporate a wide range of services into their networks flexibly. This applies not only to the cloud, but also extends from the best-effort Internet services they already provide through to legacy services such as telephony and leased lines, which require high reliability and a bandwidth guarantee.

The requirements of future networks are: (1) High capacity, (2) High quality and reliability, and (3) Flexibility in service provision. To achieve this, Hitachi is developing a POTS (packet optical transport system), which incorporates a family of POTS platforms, as well as an integrated EMS (element management system) able to perform management operations remotely, including monitoring equipment for faults and adding or removing services (see Fig. 4).

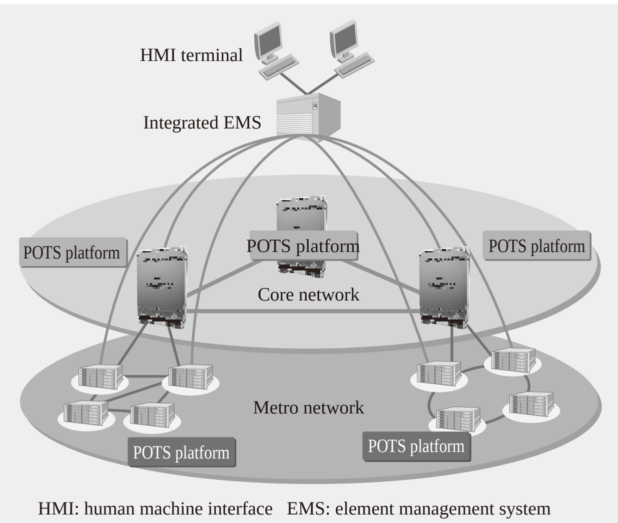

Fig. 4—POTS.
POTS (packet optical transport system) combines wavelength division multiplexing and electrical multiplexing switch functions.

The POTS includes a wavelength division multiplexing function for high-capacity transport and an electrical multiplexing switch function that ensures high quality and reliability despite being packet-based. The following sections describe the features of the POTS intended to satisfy the network requirements described above.

High-capacity Transport

The use of 100-Gbit/s optical coherent technologies with high-density wavelength division multiplexing provides a maximum transport capacity of 8.8 Tbit/s using 88 wavelengths. Each wavelength of the wavelength-division-multiplexed signal can be remotely switched through up to eight paths, thus supporting a range of different network configurations, including mesh and ring topologies. The system can also use electrical switches with a switching capacity per shelf of more than 1 Tbit/s to ensure that the electrical switching capacity can keep up with the greater transport capacity.

Ensuring High Reliability and High Quality

POTS uses the MPLS-TP transport protocol to ensure that Ethernet and other packet-based services can have levels of reliability and quality equivalent to those of legacy services, such as voice or leased lines, while still being provided efficiently. It also incorporates QoS (quality of service) functions to allow packet services to be provided with quality equivalent to that of legacy services. These include a priority control function, which transmits packets based on the relative priority assigned to each service, and a traffic shaping function, which performs peak rate control of packet transmission based on the service's bandwidth usage. The system also features extensive OAM (operation, administration and maintenance) functions, including alarms and continuity checking for MPLS-TP services, which can detect failures and identify the failure generation point. Not only wavelength-division-multiplexed optical signals, but these packet services are centrally managed by an integrated EMS, which can provide a graphical display and other service provision information. This integrated EMS can make it easy for the operator to identify the failure generation points and to perform failure recoveries when failures occur. To improve the reliability of the network, the integrated EMS has a configuration that allows redundancy along with comprehensive database backup functions to provide service continuity and recovery in the event of a fault or disaster. The system also achieves high quality services with high reliability by supporting failure recovery (protection) at the service level.

Flexibility of Service Provision

A diverse range of services are delivered using adaptation technologies that provide interfaces for legacy services such as voice and leased lines, as well as packet communications, and allow a wide range of services to be provided in a transparent way. It is possible to combine a number of networks on the same platform, including not only newly configured networks, but also the migration of existing networks that require upgrading.

To keep up with the evolution of requirements for future carrier networks, this system has adopted an architecture that ensures the expandability required for functions to be provided in the future and also allows these to coexist with existing functions. Hitachi also intends to continue working on measures aimed at further enhancing POTS, including doubling the capacity of interfaces used to provide services and expanding the types of services that can be supported.

POTS consists of a family of POTS platforms and an integrated EMS, and thus satisfies the carrier network requirements associated with greater adoption of the cloud, namely greater capacity, high quality, high reliability, and flexibility of service provision.

RESEARCH AND DEVELOPMENT

Next-generation Data Center Network Architecture

Hitachi is working on devising and implementing a network architecture for the next generation of data centers (described below), which satisfies the above requirements for data center networks and keeps up with developments in technology. With the adoption of virtualization in data centers, it is possible to achieve flexible system configurations while also improving the efficiency of IT resource use by extending or reconfiguring the pooled IT resources (servers and storage) from multiple racks as required. The rack-to-rack networks (aggregation networks) that support these IT resource pools must be capable of flexible expansion (scaling out) in step with the addition of IT resources, and must allow end-to-end connections as required, both within and between racks. A problem with these network requirements, however, is that scaling out is not possible when using configurations like those in the past based on large switches because

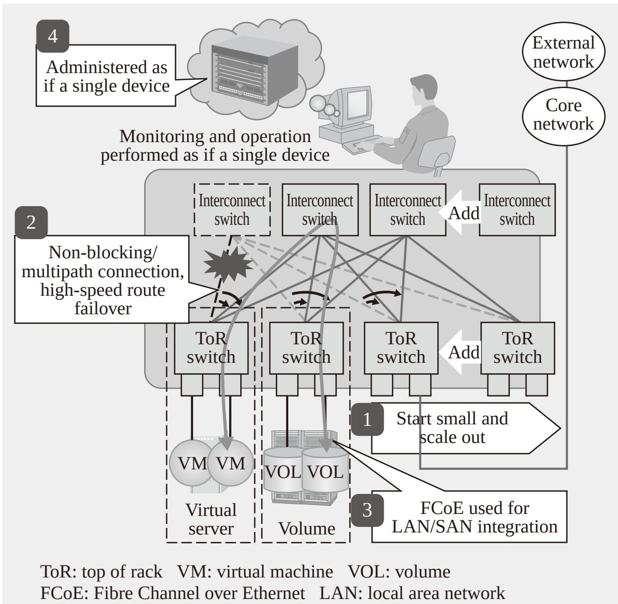

Fig. 5—Scalable Network.
The number of ports and total bandwidth can be increased by separating the switches that provide connection ports to servers and other resources from the switches that provide interconnection between these so that additional switches of each type can be added separately.

they result in significant oversubscription (ratio of total port capacity and processing capacity), which risks performance degradation when the volume of inter-server communications increases.

To resolve these problems, Hitachi is developing technology for data center networks with the following characteristics (see Fig. 5):

(1) Use of a fat-tree architecture, which can be upgraded by adding small box-mounted switches. This allows users to start small and then scale out.

(2) High-speed route failover combined with full use of the available bandwidth is achieved through the use of multipath connections and non-blocking communications between any two end points with an oversubscription ratio of 1:1.

(3) The cable count and administration workload is reduced by integrating the SAN into the Ethernet network using FCoE (Fibre Channel over Ethernet).

(4) A mechanism is adopted in which multiple switches can be administered as if they belong to the same device. This provides the ability to configure multiple switches at a time, support rack-to-rack cabling, and monitor the overall performance/load balance.

Testing of these techniques in actual implementations aimed to demonstrate high-speed route failover with a significant improvement in the failover speed when a fault occurs (5 to 30% better than the standard technology used previously). These network technologies help reduce the overall cost of data center operation and management by eliminating the need to consider network performance when deciding where to locate IT resources.

More Efficient Operation and Management—Automatic Network Design

Data centers use a wide range of networking equipment such as layer 3 switches, firewalls, and load balancers. The virtualization technologies built into these devices are used to implement private and public clouds. As the cloud allows a number of customer or business systems to be implemented virtually, it increases the workload associated with design and configuration because of the need to guarantee that the systems remain independent of each other. In particular, reducing this workload was a challenge for the design and configuration of previous data center networks, which assumed that SI (system integration) would be based on management work sheets maintained by network engineers.

To overcome this challenge, Hitachi is developing technology for automatic network design for data centers that automates the network design of business systems running on the cloud, as well as the design of individual network devices, and also generates the command sequences (configurations) used to configure each network device (see Fig. 6).

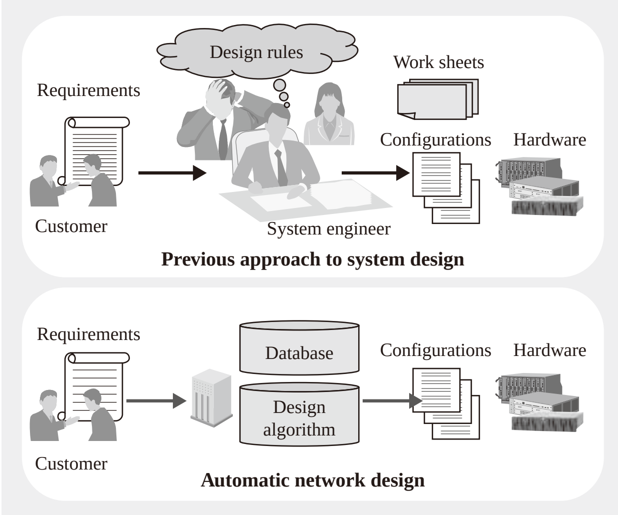

Fig. 6—Automatic Network Design.
Design efficiency is improved by defining design rules so that the configurations for each device can be generated automatically.

The technology for automatic network design defines the design rules used by network engineers and uses defined design algorithms to calculate automatically the design values required for each network device when configuring business systems. This guarantees the independence of each business system. This technology also improves the efficiency of network design and testing of individual business systems, and is anticipated to cut the design and configuration costs for data center networks in the cloud era by about 75%.

Hitachi is also strengthening its involvement with operation and management technologies that can be used on the integrated LAN/SAN networks being increasingly adopted by data centers.

Improving Efficiency of WAN

Hitachi is working on research and development of the cloud networking technologies for reliable business systems, social infrastructure, and other applications. In legacy IT systems, servers, storage, and other IT devices were installed at each user site and most communication took place within the site.

However, as cloud computing is used in wide area, a significant volume of communication data now passes over the WAN (wide area network). This causes communication delays due to data being sent via routers and other network equipment, and results in problems such as a decrease in communication bandwidth or a slow application response.

This decrease becomes an obstacle to activities such as disaster recovery or data backup, which are essential elements for ensuring the continuity of core businesses requiring high reliability. Furthermore, a demand of cloud service for social infrastructure applications, such as disaster monitoring or traffic control, will increase. Because these applications require the real-time processing, it is difficult for today's cloud to support these applications.

To solve these problems of cloud computing, Hitachi is working on research and development of "cooperated TCP" technology, which aims at improving communication bandwidth and a new cloud architecture for improving response time of each application.

Cooperated TCP

TCP is a commonly used protocol for reliable data transmission. In TCP, when the receiver receives the packets, it sends an ACK (acknowledgment) back to the sender. The sender judges that the packets have been successfully received, when it receives the ACK. When it does not receive the ACK, it resends any packets. This mechanism achieves reliable data transmission. To control the transmission rate, the number of bytes that can be sent without receiving an ACK (cwnd) and its maximum value (winsize) is determined by the sender and receiver. A problem is that because the sender is not permitted to send any more packets during the RTT (round trip time) after it has sent the maximum number allowed by cwnd, the bandwidth decreases when the RTT is long. Also, the sender reduces cwnd by a specified proportion if it detects a packet loss. This also causes the bandwidth problem.

For such reasons, we have developed cooperated TCP. It uses the packet loss rate as an indicator of network congestion and controls the transmission rate dynamically without using cwnd or winsize. By not using cwnd and winsize, cooperated TCP prevents bandwidth decrease, which is caused by long transmission delay. Furthermore, it limits the effect of packet loss by control of the transmission rate based on network congestion.

Fig. 7 shows a test configuration for cooperated TCP. It consists of a device for generating delays and packet loss to simulate the WAN, together with a sender [PC (personal computer) 1] and receiver (PC 2) with both conventional and cooperated TCP.

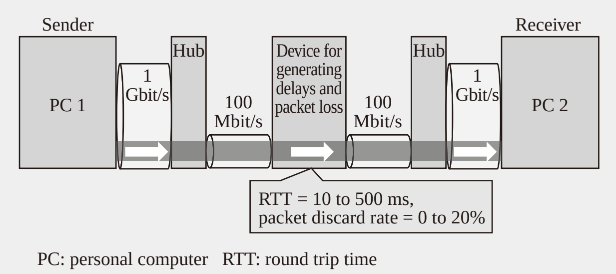

Fig. 7—Test Configuration for Cooperated TCP.
The sender (PC 1), which supports cooperated TCP (transmission control), uses FTP (file transfer protocol) to transmit a file to the receiver (PC 2) over a test network that simulates a WAN.

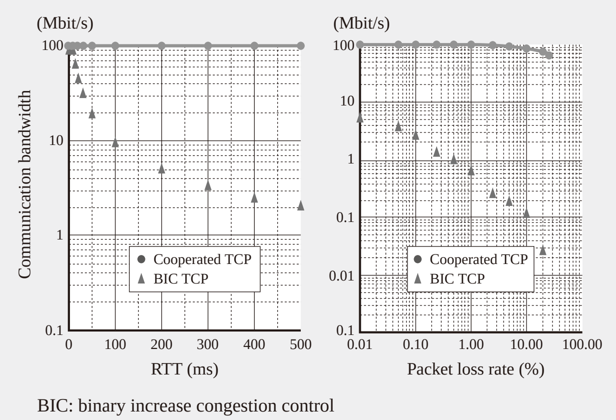

Fig. 8—Results of Cooperated TCP Testing.
The graphs show the results of measuring the communication bandwidth for different values of RTT and packet loss rate. Cooperated TCP achieves a significantly better communication bandwidth than BIC TCP.

Fig. 8 shows the results of a test in which a FTP (File Transfer Protocol) session is established between PC 1 and PC 2.

The left-hand graph shows the relationship between RTT and communication bandwidth when the packet loss rate is zero while the right-hand graph shows the relationship between the packet loss rate and communication bandwidth when the RTT is 200 ms. The left-hand graph shows that the bandwidth of BIC (binary increase congestion control) TCP(3) decreases as the RTT increases. On the other hand, cooperated TCP achieves above 90 Mbit/s. The right-hand graph, meanwhile, shows that a bandwidth of conventional TCP can fall to 5 Mbit/s when the packet loss rate is 0.01%, the loss rate of a high quality network, cooperated TCP achieves a bandwidth of 90 Mbit/s or more on networks with a packet loss rate of 1% or less. This represents an approximate 20-fold increase in efficiency for global data transfer (efficiency of line bandwidth utilization). These results indicate that cooperated TCP can significantly reduce the degradation in communication bandwidth caused by the increasing communication delays.

Actions for future

In the near future, humans, objects and IT services in real world will be connected by network and massive amount of data will be collected to cloud, which we call "Broadgather." Knowledge extracted by analyzing the data stored in IT platform will be returned to humans and objects in real world and advanced value reproduction cycle will be achieved (see Fig. 9).

The value reproduction cycle will not be limited to existing applications such as keyword search or information sharing, but is expected also to help resolve a wide range of issues, including natural disaster control or resource and environmental problems.

However, the degradation in response time will be an obstacle to implementing services such as disaster monitoring applications that require real-time response. Hitachi is working on research and development, because we believe resolving these problems will be a first step for the cloud to become a real information and communications platform for society. The following paragraphs describe a cloud architecture that improves response time of cloud applications by deploying a part of computing resources on a network.

Fig. 10 shows a proposed cloud architecture, which consists of front-side layer, back-side layer, and the management nodes. The back-side layer corresponds to the data centers in existing cloud systems and is used to handle information processing of accumulated

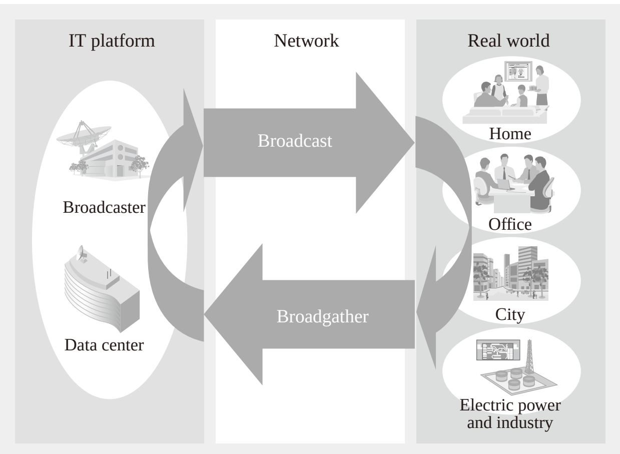

Fig. 9—Future Information and Communication Platform.
A value reproduction cycle is created in which knowledge obtained by analyzing information from the real world is returned back to the real world via networks.

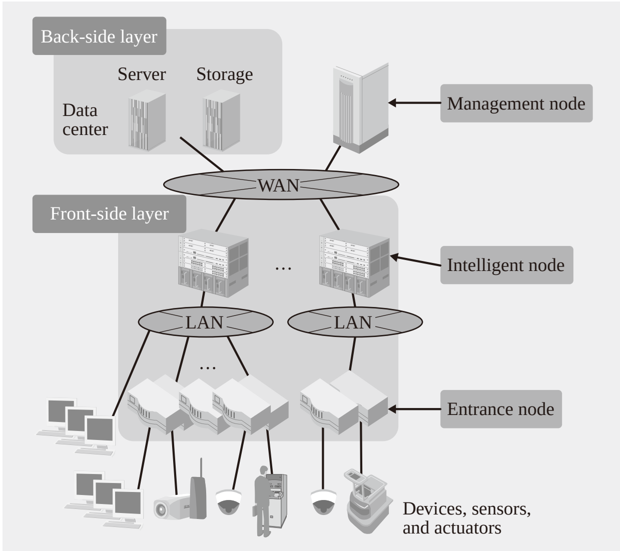

Fig. 10—Proposed Cloud Architecture.
Response time is improved by placing a front-side layer of computing resources between the WAN and real-world sensors.

data, including database management and data mining. The front-side layer is located between WAN and real-world sensors. It also handles real-time processing to consolidate or compress the sensor data flowing into the WAN to reduce its volume. Another task is processing a part of this data in order to send feedback messages to real-world actuators.

Front-side layers consist of entrance nodes and intelligent nodes. Entrance nodes process sensor data and transmit the necessary information to intelligent nodes, and perform control of actuators based on feedback messages from intelligent nodes. Intelligent nodes process the data from entrance nodes in a real-time manner and judge feedback messages to the real world. Sensor data required for processing such as data mining is sent to server and storage systems in the back-side layer via the WAN.

When a trouble of intelligent node occurs, its operation is taken over by another intelligent node to prevent service disruptions and improve the system reliability. Tasks such as allocating functions to these nodes are handled by the management nodes.

Because entrance and intelligent nodes are installed close to the real world, they are able to process sensor data and control actuators in a timely manner. Also, increased WAN power consumption is prevented by limiting the information sent to the WAN.

Hitachi continues developing technology for implementing entrance and intelligent nodes, and accelerates the use of the cloud for social infrastructure services.

CONCLUSIONS

This article has described what Hitachi is doing to satisfy the basic requirements for cloud networks and its work in the field of data center and carrier networks.

It is anticipated that the scope of applications for cloud computing will continue to expand as a way of using IT systems to improve business practices and create new business opportunities.

Hitachi continues to propose and deliver solutions to customers in order to achieve overall optimization of information system platforms, including dealing with the increasingly complex deployment of IT equipment at data centers while cutting operating costs and boosting the efficiency with which this equipment is used. It is also working on the supply of POTS, which provides high reliability and availability for the carrier networks used by data centers.

Hitachi is also continuing its involvement with the networks that form part of the infrastructure of society in order to develop services that offer greater stability, reliability, and efficiency.

ACKNOWLEDGMENT

The part of this research was supported by MIC (Ministry of Internal Affairs and Communications of the Japanese Government) "Research and Development on Secure Cloud Networking Technologies (Intelligent Distributed Processing Technologies)" and "Research and Development on Management Platform Technologies for High Reliable Cloud Services."

REFERENCES

(1) M. Hiraiwa et al., "Hitachi's Involvement in Networking for Cloud Computing," Hitachi Review **59**, pp. 206–212 (Dec. 2010).
(2) Fibre Channel over Ethernet, http://www.t11.org/fcoe
(3) L. Xu et al., "Binary Increase Congestion Control for Fast Long-Distance Networks," IEEE INFOCOM (Mar. 2004).

ABOUT THE AUTHORS

Chikako Tanaka

Joined Hitachi, Ltd. in 1993, and now works at the Network Business Strategy Planning Office, Strategic Business Development Division, Strategy Planning & Development Office, Information & Telecommunication Systems Company. She is currently engaged in the development of network business strategy.

Takashi Ezaki

Joined Hitachi, Ltd. in 2005, and now works at the Enterprise Network Systems Department 1, Network System Solution Operations, Network Systems Solutions Division, Information & Telecommunication Systems Company. He is currently engaged in the development of next-generation data center network solutions.

Yasuhiro Uchiyama, Dr. Eng.

Joined Hitachi, Ltd. in 1999, and now works at the Packet Transport Project, Network System Operation, Telecommunications & Network Systems Division, Information & Telecommunication Systems Company. He is currently engaged in the development of POTS (packet optical transport systems). Dr. Uchiyama is a member of the Japan Society of Applied Physics (JSAP).

Kazuma Yumoto

Joined Hitachi, Ltd. in 1993, and now works at the Social Infrastructure Information Network Research Department, Yokohama Research Laboratory. He is currently engaged in the research and development of DC networks. Mr. Yumoto is a member of the Information Processing Society of Japan (IPSJ).

Takeki Yazaki

Joined Hitachi, Ltd. in 1995, and now works at the Network Systems Research Department, Central Research Laboratory. He is currently engaged in the research and development of cloud networking. Mr. Yazaki is a member of the IPSJ and The Institute of Electronics, Information and Communication Engineers (IEICE).