A Large-scale Power-saving Cloud System Composed of Multiple Data Centers

Toshiaki Suzuki, Tomoyuki Iijima,
Isao Shimokawa, and Toshiaki Tarui
Central Research Laboratory
Hitachi, Ltd.
Yokohama, Kanagawa, Japan
{toshiaki.suzuki.cs, tomoyuki.iijima.fg,
isao.shimokawa.sd, toshiaki.tarui.my}@hitachi.com

Shinichi Kuwahara, Hidenori Takagi,
and Tomohiro Baba
Telecommunications & Network Systems Division
Hitachi, Ltd.
Kawasaki, Kanagawa, Japan
{kuwahara_s, takagi_hide}@itg.hitachi.co.jp,
tomohiro.baba.mn@hitachi.com

***Abstract*— A large-scale power-saving cloud system-composed of multiple data centers (DCs) and a wide-area network (WAN) connecting them is proposed. In this system, to reduce power consumption of the DCs and the WAN, virtual machines (VMs) are migrated and data routing paths are optimized under the condition that quality of service (QoS) is maintained by simultaneously providing necessary CPU resources and network bandwidth for services by a VM. To address the issue of excess VM migration (causing network congestion) due to separate control of "server resource" and "network resource" by a conventional power-saving scheme, the proposed system controls power consumption by cooperation between an inter-DC management server and a WAN management server. To determine an appropriate resource allocation, conditions for various resources (such as CPU loads and bandwidth consumed by network switches) are monitored in real time. In addition, future loads for the resources are periodically predicted. An appropriate VM reallocation is only executed when necessary resources after the reallocation can be guaranteed. A prototype system comprising 200 VMs, 200 servers, and four DCs was developed and evaluated. The evaluation results indicate that the system can achieve power saving by VM migration between DCs under the condition that the necessary CPU resource and network-access bandwidth for providing services by a VM are maintained.**

Keywords- power saving; QoS; cloud system; virtual-machine migration; network congestion; resource allocation

I. INTRODUCTION

Lately, the amount of electric power consumed by information and communication technology (ICT) systems has been dramatically rising [1] in conjunction with the increasing number of data centers (DCs) being constructed. As one of the biggest issues concerning ICT systems, including DCs, power-saving measures have therefore been attracting lots of attention. [2].

Subsequently, to address the above-mentioned power-consumption issue, technical developments and standardizations aiming to make ICT systems more power efficient have been actively promoted. For example, "server-resource virtualization" (that is, saving power consumed by a server on the basis of optimization of necessary resources) has been under research and development [3], [4]. In addition, many standardization activities, such as those undertaken by the Energy Management Working Group (EMAN) in the Internet Engineering Task Force (IETF) [5], the Institute of Electrical and Electronics Engineers (IEEE) [6], the International Telecommunication Union - Telecommunication Standardization Sector (ITU-T) [7], and the Distributed Management Task Force, Inc. (DMTF) [8], are continuing.

Although the above-mentioned activities have aimed at reducing the electric-power consumption of ICT systems, the power consumptions of the "server resource" and "network resource" are controlled separately. Power-saving control has therefore been optimized for each resource, and total optimization of electric-power consumption while maintaining service quality provided by a large-scale cloud system comprising multiple DCs and a wide-area network (WAN) to connect them has not been focused on by conventional activities. Besides, if electric-power saving for one resource is conducted separately, it might cause a serious problem for other resources. For example, an excessive aggregation of servers by virtual-machine (VM) migrations might degrade access quality to a VM since data flows are also aggregated to the same routing path; as a result, network link bandwidth is exceeded, and network congestion occurs.

We are aiming to develop efficient power-saving control scheme for both network and server resources while guaranteeing network and server "quality of service" (QoS), such as bandwidth and CPU power, by integrated power-consumption management of both network and server resources. In a previous work [9], we proposed a power-saving cloud system managed by one control system. In the present work, aiming at total electric-power saving for both WAN and DCs resources, we propose a large-scale power-saving cloud system managed by cooperation between a WAN management server and integrated DC management servers.

The rest of this paper is organized as follows. Section II explains the requirements of a power-saving cloud system. Section III proposes a large-scale power-saving cloud system managed by a WAN management server and integrated DC management server. The proposed system simultaneously saves electric power and guarantees access bandwidth to a VM. Section IV describes a prototype system and presents some results of a performance evaluation. Related works are shown in section V and section VI concludes the paper.

 ISBN: 978-1-61208-259-2

II. REQUIREMENTS OF POWER-SAVING CLOUD SYSTEM

A power-saving cloud system provides various services and resources, such as application software, CPU processing power, and storage, via a network. To create a power-saving cloud system and to reduce total electric-power consumption during off-peak hours (such as late evening), only the minimum resources required for providing cloud services should be activated.

To control the power of a target system, average loads on physical servers, VMs on those servers, and network nodes should be monitored. In addition, VMs should be appropriately reallocated according to the future loads on servers and VMs predicted under a predefined threshold during off-peak hours. After the appropriate reallocation of VMs, unnecessary physical servers should be turned off. The nodes or ports on the nodes that transmitted data to unnecessary physical servers should also be turned off or switched from active mode to sleep mode. Furthermore, service quality (such as access bandwidth to a VM) should be guaranteed before, as well as after, the power-saving control by VM migration. In addition, a power-saving scheme should be applied to not only small cloud systems comprising a single DC but also large-scale systems comprising multiple DCs.

To summarize the above-mentioned requirements, the power-saving control should be executed according to the following procedures, namely, four power-saving policies.

- Policy 1: Power consumption of the DC can be reduced by turning off unnecessary physical servers that are no longer used after an appropriate reallocation of VMs by VM migration in the DC.
- Policy 2: Power consumption of the DC can be reduced by turning off unnecessary physical servers and network nodes that are no longer used after aggregation of running physical servers and data transmission routes by VM migration in the DC.
- Policy 3: Power consumption of the DC can be reduced by turning off unnecessary physical servers and nodes in the DCs that are no longer used after aggregation of physical servers and data transmission routes by VM migration between DCs based on cooperation between DC management and WAN management.
- Policy 4: Power consumptions of the DC and WAN can be reduced by turning off nodes or their ports in the WAN that are no longer used after VM migration between DCs and aggregation of data-transmission routes.

The above four policies are resource-control procedures from the viewpoint of power saving. In addition, resources should also be controlled from the viewpoint of service quality. More specifically, power consumption of the system should be reduced by aggregation of both server resources and network resources while service quality of a network path between an end user and the VM providing application services is maintained.

III. PROPOSED POWER-SAVING CLOUD SYSTEM

A. System Architecture

The typical structure of the proposed power-saving cloud system is shown schematically in Fig. 1. The system is composed of multiple DCs and a WAN connecting them. More specifically, the DC consists of multiple switches (SWs) for transmitting data, servers for providing various services, a DC management server for controlling resources in the DC, and an inter-DC management server for controlling multiple DC management servers. The WAN consists of multiple SWs, an "integrated-mining-of-flow" (IMF) apparatus for monitoring network conditions, and a WAN management server for controlling resources in the WAN.

In the power-saving cloud system, the DC management server monitors the loads of servers, VMs, and SWs in the DC in real time. In addition, it predicts future loads of these resources according to statistical analysis (such as autoregressive model analysis [10]) based on the past history of loads. Besides, to reduce electric-power consumption on the DC side, it determines and controls an appropriate reallocation of resources such as VMs and routing paths.

To reduce power consumption of the WAN, the IMF monitors loads of SWs in the WAN. The WAN-management server receives statistical-monitoring data and predicts future loads on each SW. Electric power consumed by the WAN is saved by optimizing data-routing paths and turning off SWs or their ports that are no longer used.

In summary, power consumption of the total system is reduced by reallocating VMs between the DCs appropriately on the basis of cooperation between multiple DC-management servers and the WAN-management server.

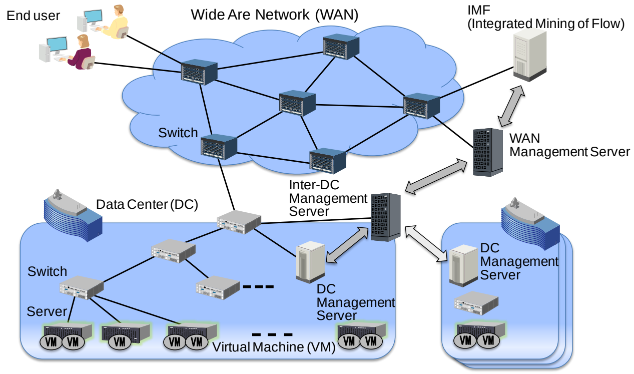

Figure 1. Proposed power-saving cloud system

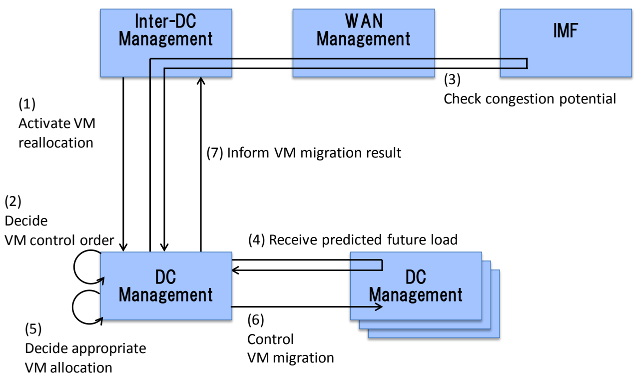

Figure 2. Process steps of proper VM resource reallocation

B. Overview of Power-saving Scheme by VM Migration

The process steps of a typical power-saving scheme by VM migration based on multiple management servers are shown schematically in Fig. 2. In the proposed system, the inter-DC management server activates power-saving control according to the loads on the physical servers and VMs [step (1)]. The DC management server determines the order of VM migration [step (2)]. The DC management server checks a “congestion potential” via the inter-DC management server, the WAN management server, and the IMF [step (3)]. For VM migration between DCs, the DC management server receives predicted loads on alternative physical servers from other DC management servers [step (4)]. To move the VM according to the predicted loads on the outside servers and the effectiveness of the power saving, the DC management server determines one alternative server [step (5)] and triggers an actual VM migration [step (6)]. The VM-migration result is transmitted from the DC management server to the inter-DC management server [step (7)].

C. Detailed VM Resource Reallocation

The seven above-mentioned steps for resource reallocation are explained in detail in the following:

1) VM reallocation trigger by inter-DC management server: The inter-DC management server starts or stops optimizing reallocation of VMs to each DC management server when the loads on servers and VMs are low (such as late evening).

2) Determination of VM reallocation order by DC managemnet server: The DC management server determines the order to reallocate running VMs for each virtual local area network (VLAN). The reallocation order is determined according to (i) decending order of idle power, (ii) ascending order of number of running VMs on a server, (iii) ascending order of assigned CPU resources, and (iv) ascending order of assigned memory resources.

3) Checking of congestion potential in WAN by DC management server: To maintain access quality to the VM after VM migration to another DC, the DC management server receives the congestion potential concerning the WAN from the IMF via the inter-DC management server and the WAN management server. The congestion potential is evaluated according to the history of the monitored data and predicted future loads in the case of fluctuation of bandwidth for each port of the switch. If there is any posibility of network congestion in the future, data-routing paths including the congestion point are not used for VM migration.

More specifically, the IP address of the VM to reallocate, the identifier of the source DC, and the identifier of the VLAN to which the VM belongs are transmitted from the DC management server to the inter-DC management server. A list of alternative DCs that can accommodate the migrated VM and above-mentioned information from the DC management server is then transmitted from the inter-DC management server to the WAN mananagement server. After that, information about a routing path (from a WAN edge point connecting a user to another WAN edge point connecting an alternative DC) and the above-mentioned information from the inter-DC management server are transmitted from the WAN management server to the IMF. The congestion potential at the routing path between the user and the altenative DC is sent from the IMF to the DC management server via the WAN management server and the inter-DC management server.

4) Determination of target server for VM migration by DC management server: The DC management server detemines the appropriate VM reallocation by considering all alternative DCs. Specifically, all severs that can provide enough resources to run the intended VM in the future and guarantee access qualtiy to the VM at the same time are selected as alternative servers for the reallocation of the VM. The most effective server for power saving is then selected as a final target server for the VM migration.

More specifically, the DC management server predicts future loads on the CPU and consumption of the bandwidth resource by the intended VM. In addtion, it gets information concerning predicted available future resources (such as CPU and bandwidth) for the alternative servers in other DCs from other DC management servers. It finally determines one target server to which the intended VM is reallocated by comparing the received available future resources for all alternative servers in other DCs and the amount of necessary resources for the intended VM in the future.

5) Determination of VM reallocation by DC management server: The DC management server determines whether target servers can provide enough resources (such as CPU processing power and memories) for running the intended VM in the future. The only servers that can provide enough resources are registered as alternative servers for VM migrations. In addition, the DC management server determines whether switches on the routing path between the entrance of the DC and the

 ISBN: 978-1-61208-259-2

alternative server in another DC can provide enough bandwidth for the intended VM after the VM migration. It checks the congestion potential for the routing path between the WAN edge connecting the DC and another WAN edge connecting an end user on the basis of the monitored information from the IMF. To determine the most appropriate alternative server, the DC management server checks all the above-mentioned evaluation points, i.e., CPU load, network congestion, and bandwidth. The most appropriate server that can meet the requirements stated in Section II and has the most effective power-saving advantage is then selected as the target server for the VM migration by the DC management server.

6) VM migration by DC management server: The VM migration is executed according to the trigger of the DC management server. As for VM-migration methods, various technologies have been developed by several organizations [3], [4], and these technologies can be used for an alternative VM-migration scheme by combining them with the proposed power-saving cloud system. After the VM migration, the DC management server updates stored topology information. In addition, to predict future load, when the VM has been migrated to a server in another DC, the history of the VM's resources (such as CPU load) is moved to another DC management server.

7) Information about VM-migration completion sent from DC management server to inter-DC management server: After all VM migrations have been executed, the completion of all VM reallocations is transmitted from the DC management server to the inter-DC management server. In addition, the migration histories from the source servers to destination servers are transmitted from the DC management server to the inter-DC management server. The inter-DC management server receives the migration histories and stores them. These histories are used when the migrated VMs are returned to the original allocated servers when CPU load increases.

D. Power-consumption Model

A power-consumption model for the proposed cloud system is defined as follows. The amount of power (P_{All}) consumed by the cloud system is given by formula (1), where P_{IT} means power consumption of IT equipment, and P_{NET} means power consumption of network nodes. Formula (2) indicates P_{IT} is calculated by a summation of power consumption (P_{SV}) of servers since the proposed system includes multiple servers as IT equipment. Here, i (i = 1, 2, 3, . . . , N) mean the number of the server. In addition, n means CPU load (%) on the server. P_{SV} is given by formula (3). $P_{idle(i)}$ means the power consumption of the ith server during idle time, and $P_{max(i)}$ means power consumption during maximum load. Formula (4) gives P_{NET} of a network calculated by the summation of the power consumption of each node. Here, k (k = 1, 2, 3, . . . , M) mean the number of the node. In addition, m means load (%) on a node in terms of bandwidth. The power consumption of the node (P_{NODE}) is given by formula (5). $P_{idle(k)}$ means power consumption by the kth node during idle time, and $P_{max(k)}$ means power consumption during maximum load. Here, P_{SV} and P_{NODE} are assumed to fit a linear function, as shown in Fig. 3. The relations between the power and CPU loads and between the power and traffic are independently evaluated in advance. According to that evaluation, the relation between power consumption and load (traffic) fits a linear function well (as shown in Fig. 3).

$$P_{All} = P_{IT} + P_{NET} \tag{1}$$

$$P_{IT} = \Sigma_i P_{SV(i)}[n] \tag{2}$$

$$P_{SV(i)}[n] = P_{idle(i)} + (P_{max(i)} - P_{idle(i)})(n/100) \tag{3}$$

$$P_{NET} = \Sigma_k P_{NODE(k)}[m] \tag{4}$$

$$P_{NODE(k)}[m] = P_{idle(k)} + (P_{max(k)} - P_{idle(k)})(m/100) \tag{5}$$

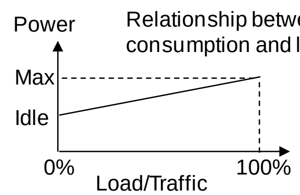

Figure 3. Assumed power consumption based on load/traffic

IV. Evaluation of Proposed System

A. Evaluation System

The evaluation system is shown schematically in Fig. 4, and the number of pieces of ICT equipment is listed in Table I. In the system, switches, servers, and VMs in the DCs are emulated by open-source software, while switches in the WAN and management servers are real apparatuses. The performances of the power-saving control of the DCs and WAN are evaluated in detail in the following sections.

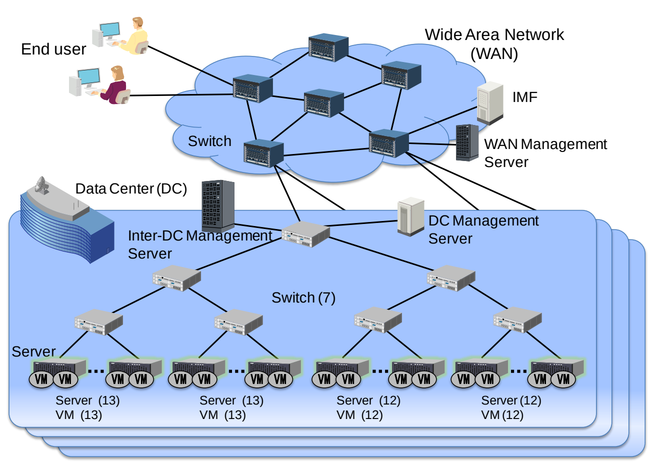

Figure 4. Evaluation system

 ISBN: 978-1-61208-259-2

TABLE I. NUMBER OF PIECES OF ICT EQUIPMENT

	Item	Number of pieces of ICT equipment in 4 DCs
1	WAN management server	1
2	IMF	1
3	SW in the WAN	6
4	DC	4
5	Inter-DC management server	1
6	DC management server	4
7	SW in the DCs	28
8	Server in the DCs	200
9	VM on the servers	200

B. *Evaluation of Power-saving Control for DCs*

The effectiveness of the power-saving control for DCs per day was evaluated. First, a CPU load model of a VM in the DC is assumed. The bandwidth consumed by the VM for one day is also assumed. The power consumption by DCs for one day is evaluated according to these assumptions.

1) Workload model for a VM per day

The assumed loads on the CPU as well as the incoming flow to and the outgoing flow from a VM are schematically shown in Fig. 5. As depicted in the figure, the peak load is set only one time (around noon), and the loads during business hours are high while the loads during the night for the CPU, incoming flow, and outgoing flow are low. The effectiveness of the power-saving control scheme is evaluated by comparing two cases: executing appropriate VM reallocations and not executing them.

The topology of the DC is shown in the lower part of Fig. 4. The specifications and number of pieces of each apparatus in the DC are listed in Table II.

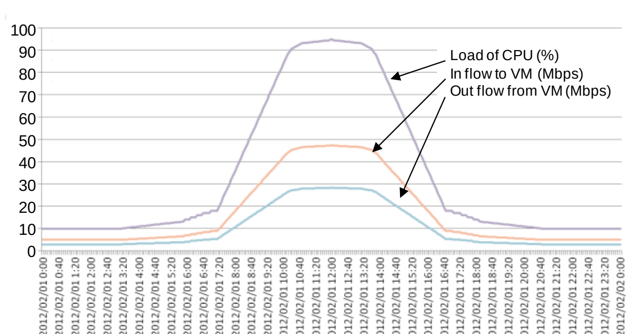

Figure 5. Load model of VM (CPU and in/out flow)

TABLE II. SPECIFICATIONS AND NUMBER OF PIECES OF EACH APPARATUS IN ONE DC

	Apparatus	Idle power	Max. power	Number
1	Server (Model 1)	120 W	170 W	17
2	Server (Model 2)	110 W	150 W	17
3	Server (Model 3)	177 W	251 W	16
4	VM	—	—	50
5	Switch	350 W	450 W	7

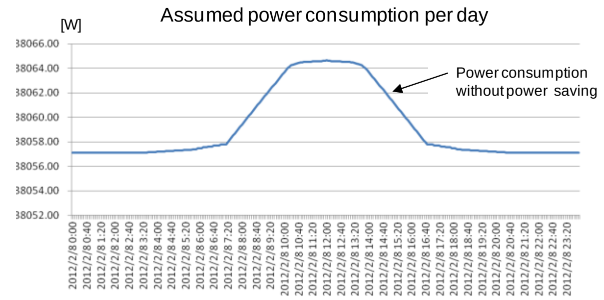

Figure 6. Electric power consumption of a DC per day

2) Electric-power consumption of a DC per day

Electric-power consumption of a DC for one day (under the assumed loads for each server shown in Fig. 5) is shown in Fig. 6. The effectiveness of the power-saving control scheme under the following three conditions was evaluated. In the first condition, the VM is reallocated when the CPU loads are less than 75%. In the second and third conditions, reallocations are executed under CPU loads of 50% and 25%, respectively. On the other hand, when the load on the CPU is over these thresholds, reallocated VMs are returned to the original locations to guarantee service quality.

3) Energy consumption of DCs per day

The evaluated fluctuations of power consumption of all DCs for the three above-mentioned power-saving control conditions (CPU loads of 75%, 50%, and 25%) are shown in Fig. 7. The result in the case of no VM reallocation is also shown in the figure for comparison. As shown in the figure, the effectiveness of the power-saving control scheme under the three conditions is verified.

In addition, the results for VM reallocation keeping VM access quality and energy consumption per day are listed in Table III. The number of VMs is shown in the upper row, while the number of servers (SV) is shown in parentheses in the lower row. According to the table, some VMs are migrated between DCs (since the number of VMs in the DC is changed after appropriate VM reallocations). In addition, the number of running servers is dramatically reduced after the VM migration. Here, the CPU resource for a server is assumed to be enough for six VMs with CPU loads of 50%. The reductions in energy consumption at CPU loads of 25%, 50%, and 75% are 45.2%, 45.7%, and 47.6%, respectively.

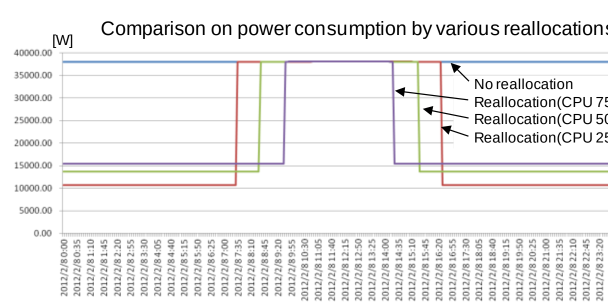

Figure 7. Electric-power consumption of a DC per day

TABLE III. ELECTRIC-ENERGY CONSUMPTION OF DCS PER DAY

	Optimization timing	DC1	DC2	DC3	DC4	Energy consumption per day (kWh)
		VM (SV)	VM (SV)	VM (SV)	VM (SV)	
1	No optimization	50 (50)	50 (50)	50 (50)	50 (50)	913.421
2	CPU load: 25%	60 (5)	48 (4)	48 (4)	44 (4)	500.388
3	CPU load: 50%	54 (9)	48 (8)	48 (8)	50 (9)	496.395
4	CPU load: 75%	52 (13)	48 (12)	52 (13)	48 (12)	478.323

C. Power-saving Evaluation for WAN

Power-saving control for a wide-area network (WAN) for one day was evaluated. In particular, the effectiveness of the power-saving scheme based on bandwidth control by link aggregation was evaluated. The topology of the evaluated WAN is shown in Fig. 4. In addition, the specifications of the switches in the WAN are the same as those listed in Table II.

Power-saving control by appropriate data routing (including link-aggregation control) was executed after appropriate VM reallocation between DCs. The fluctuation of power consumption of the WAN is shown in Fig. 8. In addition, energy consumptions under the three types of control are compared in Table IV. According to these results, the reductions in energy consumption achieved by the power-saving control scheme under CPU loads of 25%, 50%, and 75% are 10.4%, 12.0%, and 13.7%, respectively.

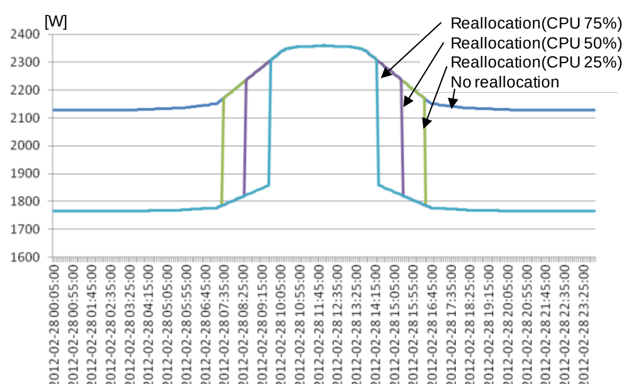

Figure 8. Electric power consumption of WAN per day

TABLE IV. ELECTRIC-ENERGY CONSUMPTION OF WAN PER DAY

	Optimization timing	Electric-energy consumption per day (kWh)	Reduction (%)
1	No optimization	52.470	–
2	25% CPU load	46.989	10.4
3	50% CPU load	46.194	12.0
4	75% CPU load	45.265	13.7

TABLE V. ELECTRIC-ENERGY CONSUMPTION OF ENTIRE CLOUD SYSTEM PER DAY

	Optimization timing	Optimization term	Electric-energy consumption (kWh)	Reduction (%)
1	No optimization	–	965.891	–
2	CPU load: 25%	15h00m	547.377	43.3
3	CPU load: 50%	17h00m	542.589	43.8
4	CPU load: 75%	19h10m	523.588	45.8

D. Power-saving Effect for Entire Cloud System

The effectiveness of the power-saving control scheme for an entire cloud system is shown in Table V. The reductions of energy consumption achieved by the power-saving control scheme under CPU loads of 25%, 50%, and 75% are 43.3%, 43.8%, and 45.8%, respectively. As shown in the table, energy consumption is reduced by approximately 40%. In addition, the highest reduction is accomplished under CPU load of 75%.

E. Discussion of Power-saving Effect

According to the results of this evaluation of a large-scale power-saving cloud system composed of multiple DCs and a WAN, energy consumption of the entire system is reduced by about 40%. With regard to only the power saving for the DCs, energy consumption is reduced by over 45%. On the other hand, energy consumption of the WAN is reduced by only about 10%.

The reason that the reduction of energy consumption of the DCs is high is the effectiveness of turning off unnecessary servers after appropriate VM reallocation. On the other hand, the reason that the reduction of the energy consumption of the WAN is low is that unnecessary switches were not turned off (since turning off unnecessary links (network ports) is only possible for the assumed system).

With regard to power saving for the entire system, the reductions in energy consumption achieved by the power-saving control scheme under CPU loads lower than 25%, 50%, and 75% are 43.3%, 43.8%, and 45.8%, respectively. On the other hand, the periods for the resource optimization under the three above conditions are 15 hours, 17 hours, and 19 hours and 10 minutes, respectively. When the power-saving control is executed under a CPU load of 75%, the period for the optimization is the longest, and reduction in energy consumption is the highest. These results verify the effectiveness of the proposed power-saving control scheme.

V. RELATED WORK

In previous researches, many power-saving schemes for ICT systems have been proposed. For example, power-saving schemes for node and link levels [11], [12] have been proposed. These schemes are useful for our proposed system to reduce power consumption for the link level. In addition, power-saving schemes [13]-[16] for the network level have been proposed. Power-saving schemes for the DC/server level [17]-[20] have also been proposed.

In conventional power-saving researches like those mentioned above, network and DC/server resources are controlled separately. Therefore, integrated management for maintaining network QoS and reducing energy consumption of servers is addressed in the current study.

VI. CONCLUSION

A large-scale power-saving cloud system comprising multiple DCs and a WAN is proposed. As for this system, VMs are reallocated under condition of guaranteeing necessary CPU resources and network bandwidth for providing cloud services. Energy saving for the entire system is executed by cooperation between a DC management server and an inter-DC management server. The energy saving for the DCs is executed by VM migration between DCs and aggregation of running servers under maintained network-access quality to a VM. On the other hand, the energy saving for the WAN is executed by controlling link aggregation according to the required bandwidth for transmitting user data between a first WAN edge connecting the user and a second WAN edge connecting the DC.

A prototype system, composed of 200 VMs, 200 servers, and four DCs, was developed and evaluated. The evaluation results verify that the functions for reallocation of VMs between DCs and control of link aggregation can reduce power consumption of the DCs and WAN under maintained service quality. In addition, they show the possibility of energy saving by approximately 40% (under the conditions assumed in this evaluation). Moreover, they also show that power-saving control should be executed when CPU load is 75%, i.e., not when CPU load is 50% or 25%.

The power-saving cloud system will be further evaluated in the case that switches in the WAN under various CPU loads and consumed bandwidths are turned off. In addition, while aspects of QoS concerning VM access are partially evaluated in [9], power-consumption control while keeping QoS should be evaluated in detail. Besides, the prototype power-saving cloud system will be enhanced so that it can handle multiple-use cases, i.e., multiple users.

ACKNOWLEDGMENT

Part of this research was supported by the MIC (The Japanese Ministry of Internal Affairs and Communications) projects “Research and development on signaling technology of network configuration for sustainable environment” and “Research and development on power-saving communication technology – Realization of the Eco-Internet”.

REFERENCES

[1] C. L. Belady, Microsoft Corporation, “Projecting Annual New Datacenter Construction Market Size,” Mar. 2011 http://cdn.globalfoundationservices.com/documents/Projecting_Annual_New_Data_Center_Construction_PDF.pdf [retrieved: Jan., 2013]

[2] GreenTouch, “Our Mission,” http://www.greentouch.org/index.php?page=about-us/ [retrieved: Jan., 2013]

[3] VMware, Inc. Homepage, http://www.vmware.com/ [retrieved: Jan., 2013]

[4] Xen Homepage, http://www.xen.org/ [retrieved: Jan., 2013]

[5] The Internet Engineering Task Force (IETF), “Energy Management Working Group (EMAN WG) Charter,” http://datatracker.ietf.org/wg/eman/charter/ [retrieved: Jan., 2013]

[6] Institute of Electrical and Electronics Engineers (IEEE), “IEEE802.1,” http://www.ieee802.org/1/ [retrieved: Jan., 2013]

[7] International Telecommunication Union - Telecommunication Standardization Sector (ITU-T) Study Group 13 (SG13), “Question 21/13 – Future Networks,” http://www.itu.int/ITU-T/studygroups/com13/sg13-q21.html [retrieved: Jan., 2013]

[8] Distributed Management Task Force, Inc. (DMTF), “Cloud Management,” http://www.dmtf.org/standards/cloud [retrieved: Jan., 2013]

[9] T. Suzuki et al., “Power-saving ICT platform that guarantees network bandwidth for cloud-service systems,” TS-A4: Cloud Computing Technical Session, World Telecommunications Congress, Mar. 2012

[10] International Business Machines Corp. Homepage, http://www-01.ibm.com/software/analytics/spss/products/statistics/forecasting/ [retrieved: Jan., 2013]

[11] M. Yamada, T. Yazaki, N. Matsuyama, and T. Hayashi, “Power efficient approach and performance control for routers,” Proc. of IEEE International Conference on Communications (ICC Workshops 2009), pp.1-5, June 2009

[12] Y. Fukuda, T. Ikenaga, H. Tamura, M. Uchida, K. Kawahara, and Y. Oie, “Performance Evaluation of power saving scheme with dynamic transmission capacity control,” Proc. of IEEE Globecom Workshops 2009, pp.1-5, Nov./Dec. 2009

[13] J. Baliga, R. Ayre, K. Hinton, W. V. Sorin, and R. S. Tucker, “Energy consumption in optical IP networks,” Journal of Lightwave Techology, vol. 27, no. 13, pp.2391-2403, July 2009

[14] C. Lange, D. Kosiankowski, R. Weidmann, and A. Gladisch, “Energy consumption of telecommunication networks and related improvement options,” IEEE Journal of selected topics in quantum electronics, vol. 17, no. 2, pp.285-295, March/April 2011

[15] Y. Zhang, P. Chowdhury, M. Tornatore, and B. Mukherjee, “Energy efficiency in telecom optical networks,” IEEE Communications surveys & tutorials, vol. 12, no. 4, pp.441-458, Fourth quarter 2010

[16] R. Bolla, R. Bruschi, F. Davoli, and F. Cucchietti, “Energy efficiency in the future internet: A servey of existing approaches and trends in energy-aware fixed nework infrastructures,” IEEE Communications surveys & tutorials, vol. 13, no. 2, pp.223-244, Second quarter 2011

[17] J. Baliga, R. W. A. Ayre, K. Hinton, and R. S. Tucker, “Green cloud computing: Balancing energy in processing, storage, and transport,” Proc. of IEEE, vol. 99, no. 1, Jan. 2011

[18] S. Yang, L. Chen, H. Tseng, H. Chung, and H. Lin, “Designing automatic power saving on virtualization environment,” IEEE International Conference on Communication Technology (ICCT 2010), pp.966-970, Nov. 2010

[19] Y. Yao, L. Huang, A. Sharma, L. Golubchik, and M. Neely,”Data centers power reduction: A two time scale approach for delay tolerant workloads,” Proc. of IEEE Infocom 2012, pp.1431-1439, Mar. 2012

[20] T. Imada, M. Sato, and H. Kimura, “Power and QoS performance characteristics of virtualized servers,” Proc. of IEEE/ACM International Conference on Grid Computing, pp.232-240, Oct. 2009

 ISBN: 978-1-61208-259-2