The Design and Implementation of the NHN Next Generation Service Infrastructure Using Virtualization Technology and Remote Data Centers

NHN (Next Host Network) is IIJ's next generation service infrastructure. When NHN is adopted, service development flow and demand forecasts for equipment upgrades at IIJ are set to change dramatically. Here, we provide an overview of NHN, its technological components, architecture and implementation.

At IIJ we have implemented a service infrastructure that provides flexibility, while keeping data center costs low through the utilization of remote data centers. We designed the "NHN" (Next Host Network) in 2008, with the aim of consolidating onsite work in one place, devising systems for the efficient use of rack space and electricity, and reducing the operating costs. It was prepared as an infrastructure for the migration of both new and existing services. In this whitepaper we explain the circumstances leading up to the introduction of NHN, its technological components, architecture and implementation.

4.1 Background to Adoption of NHN

IIJ has until now operated over 200 racks and several thousand servers distributed over multiple data centers for use in our own services. As equipment was constructed with each host reserving a rack for individual services, the rack space and network devices were designed with headroom to cope with future increases in demand, resulting in lost potential for the infrastructure as a whole. Using a system configuration such as this that is vertically segmented for each service makes it difficult to expand when plans change and also reuse equipment when a service is discontinued. Service infrastructure costs piled up, and as racks used different equipment and cabling, the complexity of operation had also become an issue.

Until now we had taken factors such as support when a failure occurs into consideration, and used a data center with a favorable location in Tokyo because of the ease of access to a physical system, aiming for optimal server operation. However, with the improvement in performance and lower costs of IA servers over the past few years, and power consumption for each server continuing to rise, facility costs such as rack space, air conditioning costs for cooling, and power costs now make up a higher percentage of the total than equipment costs such as networks, servers, and storage. Our Tokyo data center receives a lot of business from customers, and even when there was space to install equipment, we were not able to install enough servers due to limitations such as a shortage in our capacity for cooling, UPS, and emergency generators.

As we saw that a configuration with systems concentrated at the Tokyo data center was reaching its limit, IIJ decided to drastically reduce costs by moving service equipment that was not dependant on location to a low-cost facility on the outskirts of Tokyo. We considered all options that would lead to reduced facility costs and power requirements, such as container-type data centers. Advances in multiplexing technology have made it possible to keep network costs down, and this also motivated our plan to move to a data center out of central Tokyo.

While considering our options, we also moved ahead with transitioning to a configuration that ensures freedom and mobility of operations even when a remote data center is used.

Moving service equipment to a data center away from the city reduces space costs, personnel expenses, and power costs, leading to cost reductions across the entire service infrastructure. We began looking into NHN with the goal of improving operational freedom at the same time.

© 2009 Internet Initiative Japan Inc.

4.2 A Design Plan Optimized Remote Data Center

For NHN, we isolated the following points that must be given priority and points that could be put on hold, based on our experience as a server operator.

- Regarding storage, foreseeable failures such as individual HDD failures are acceptable. However, there is a need to design a configuration with high reliability to prevent failures that lead to service outages.
- Server devices sometimes fail. As there are wide-ranging points of failure, and there is a limit to how much the failure rate can be curtailed, a configuration that suffers little impact when a failure occurs is ideal.
- Failures of network devices used on the edge are not very common in our experience to date. NIC redundancy settings, etc., are only implemented when required.

Next, we will detail the failure rate of devices that IIJ uses, and explain the sequence of events that lead to the plan mentioned above.

4.2.1 Storage Failure Rates at IIJ

After investigating the logs for the approximately 100 DAS*[1] that IIJ uses, we confirmed that despite the number of storage units changing very little, HDD failure rates had dropped over the past few years. Another trend in the last few years is that almost no correlation was observed between storage load and failure rates. It seems that HDD failures occur regardless of load, and it is rare for failures to occur more often in HDDs in a certain storage unit.

Year	HDD Failures	HDD Failures in the Same RAID
2005	32	7 x 1 unit, 5 x 1 unit, 4 x 1 unit, *[2] 2 x 2 units, 1 x 12 units
2006	22	2 x 3 units, 1 x 16 units
2007	16	2 x 1 unit, 1 x 14 units
2008	7	2 x 2 units, 1 x 3 units
2009	4	1 x 4 units (as of September 2009)

As the HDD failure rate is dropping each year, storage failures other than HDD failure, and in particular those that lead directly to service outages, have become more conspicuous. Specifically, this refers to failures such as the following.

- Suspension of function due to the failure of cache memory on a RAID controller
- Unstable operation thought to be caused by a RAID controller failure
- Failure of connection interfaces on SCSI cards, etc.
- Performance degradation due to the failure of a battery backup unit on a RAID controller or it reaching its lifespan

Failure numbers amount to a few times per year overall, but when failures such as those above occur in devices without RAID controller or connection path redundancy, it is not possible to restore operation until the device is replaced onsite, leading directly to a service outage. When considering a remote data center option, we decided that for storage service outages it would take an unacceptably long time to complete the replacement of a device. For this reason, we made RAID controller and connection path redundancy a requirement when selecting devices, although this would make hardware costs higher than at present.

4.2.2 Server Failure Rates at IIJ

Regarding server failure rates, servers are often replaced on the spot as a precaution to prioritize the restoration of service, and in some cases the symptoms cannot be recreated after the device that has failed is subsequently taken back and tested, meaning that there is no statistical information regarding the parts that have failed.

We estimate that the failure rate is approximately 1-2%, based on the number of servers that were repaired in the six months between April and September 2009, and the number of servers that were switched out and are awaiting inspection. However, this figure does not include failures in redundant parts, such as the failure of one of a pair of redundant HDDs. For this reason, we estimate that a significant number of server devices fail.

*1 An abbreviation of Direct Attached Storage. At IIJ we chiefly used devices that connect externally via SCSI I/F until 2008. HDD failures for setups where a local HDD is added to a server are not included.

*2 As the devices for which 7, 5, and 4 units failed in the same RAID in 2005 were the same product type and were installed in the same period, it is highly likely they were part of a faulty batch. We removed them from service use in 2005.

© 2009 Internet Initiative Japan Inc.

Stoppages caused by memory failures and reboots stand out as reasons for failures. There were a wide range of other reasons such as thermal runaway due to fan failure, backplane failures in connectors for devices such as HDDs, power (VRM) failures, and processor failures.

When a failure to boot occurs in a server device fitted with a local HDD, it is necessary to carry out repairs onsite to relocate the HDD from the failed equipment into working server equipment. A great deal of personnel expenses are incurred in order to have operators who can perform server device maintenance such as that mentioned above on duty 24/7 at a data center. When considering a remote data center, costs will increase if operators are on duty 24/7, but if operators are not on the premises it will take a great deal of time to restore service. For this reason, we made diskless operation in which servers have no local data a prerequisite when selecting devices.

4.2.3 Network Device Failure Rates at IIJ

Failure numbers for edge L2 switches that accommodate hosts, etc. are not high in our experience to date. Excluding cases that involve specific faulty batches such as faulty condensers, the figure is less than 1% of the units in operation.

It is also possible to adopt a redundancy configuration using NIC redundancy settings, but this makes the addition of NIC and network devices necessary, increasing costs. Additionally, as adopting a redundant configuration over servers using devices such as a load balancer usually makes operation easier, we went with a design implementing NIC redundancy settings only when higher reliability was necessary, and did not add redundancy to switches accommodating the servers for the basic configuration. Also, to limit the impact of edge L2 switch failures, we designed a system that made it possible to either form a redundant configuration with devices under the control of other server switches using a load balancer, or deal with failures by migrating the contents of server devices controlled by the faulty server switch to a server under the control of another server switch using remote control.

4.3 NHN Configuration

For NHN, we used the following configuration that also supports remote data centers.

- A server pool system is used and large numbers of servers with identical specifications are installed. Onsite work such as device installation and the physical replacement of failed devices can be consolidated into monthly scheduled maintenance.
- Server configuration is standardized and no modifications are made to physical configurations such as changes to the amount of physical memory installed or the installation of a local HDD, in order to curb configuration management costs.
- Energy-saving servers are used to increase the number of servers that each rack can accommodate.
- Racks are not assigned to specific services to increase the number of servers that each rack can accommodate.
- Devices and cabling are standardized to reduce the possibility of mistakes when device installation or replacement work is outsourced on a case-by-case basis.
- Consolidation and degree of freedom is raised through combination with virtualization technologies such as Xen and OpenVZ.
- VLAN is used to make it possible to configure a logical network without changing onsite cabling.
- iSCSI is used to configure a cheap SAN. RAID controller and connection path redundancy are required for storage to prevent service outages as far as possible.
- It is assumed that diskless servers will break down. The goal is to make it possible to complete temporary restoration by switching over devices remotely when a hardware failure occurs.
- If the time it would take for switchover to a backup server to be completed when a failure occurs in a diskless server results in a longer service outage than is acceptable, two servers from separate groups are prepared in advance and a device such as a load balancer is used to create a redundant configuration, making an application-based redundancy configuration possible.
- All work that is required after installation, including OS installation, can be executed remotely.

Next, we explain each technological component.

© 2009 Internet Initiative Japan Inc.

4.3.1 Adoption of iSCSI-based IP SAN

When use in a remote data center is taken into consideration, the service downtime of storage can lead to extended failures, so we sought even higher storage reliability than we had in the past. IIJ has introduced a system using iSCSI, as it excels at IP network technology, and it makes it possible to put together cheaper systems than FC SAN. Also, by making RAID controller and connection path redundancy a requirement, we have reduced the risk of storage failure that may render our services inaccessible to a minimum.

4.3.2 Implementing a Diskless Server Combining iSCSI Storage and Energy-Saving Servers

Few current server devices support iSCSI boot, and installing iSCSI HBA on all server devices is costly. Because of this, we devised a way to make SAN boot with iSCSI possible using PXE boot, which is supported by the onboard NIC of most server devices.

Information such as the iqn*3 and IP address that is necessary for iSCSI boot is passed on using DHCP options. This makes it possible to replace a failed server device with a backup device by overwriting configuration information and booting the backup server.

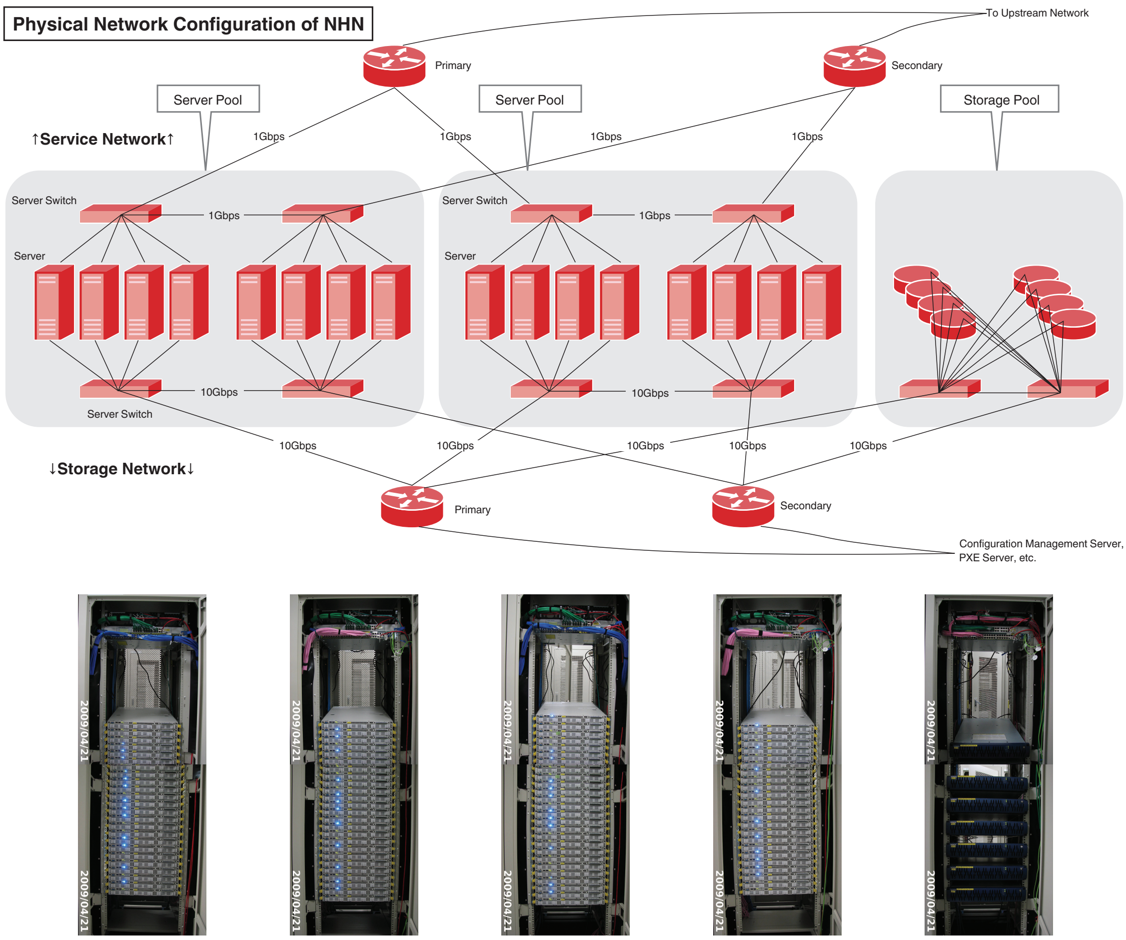

Figure 1: Physical Configuration of NHN

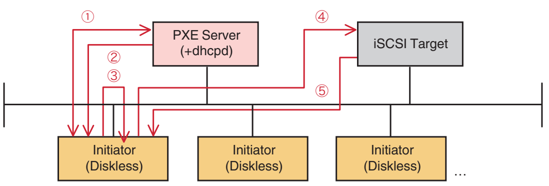

1. Address and parameters are acquired using DHCP.
2. Kernel and initrd images are acquired from TFTP.
3. Kernel is booted in an initrd ramfs root environment.
4. iSCSI initiator is initialized based on DHCP parameters (visible as a local HDD).
5. This is mounted as a block device, and switchroot is used to switch to the mounted location.

Figure 2: Diskless Server (Linux) Boot using PXE boot and iSCSI

*3 An abbreviation of iSCSI Qualified Name. Here this is used for the unique identification of iSCSI targets.

© 2009 Internet Initiative Japan Inc.

To avoid the need for onsite work in the event of a server failure, for NHN we used a diskless configuration for servers so they have no local data. If a failure occurs in one server device, it is possible to boot backup servers that are already in place while still maintaining the storage content of the failed device by overwriting configuration information remotely and rebooting.

4.3.3 Implementing a Virtual Network That Requires No Cabling Changes Using VLAN

For NHN, we rendered onsite cabling changes unnecessary by having multiple networks coexist over a single physical cable, making all maintenance possible remotely. Technically, each server is configured as a network using access VLAN and trunk VLAN, with multiple networks made to coexist over a single physical cable as necessary.

VLAN is not a new technology, but for NHN we implemented a system that made it possible to link it with the configuration information that IIJ manages, and automatically overwrite the VLAN settings of a server switch that accommodates a given host. This reduces VLAN setting errors in network devices, and cuts operating costs.

4.4 Results of Adoption of NHN

IIJ began looking at its future use of remote data centers in 2008, implementing drastic revisions to service host configurations under the keyword "NHN," and preparing an infrastructure for the migration of new and existing services.

We adopted new technology for NHN based on our perspective and experience as server operators, introducing virtualization, diskless servers using iSCSI, the centralized management of host information, and dynamic updates for network device VLAN settings.

When we adopted NHN for our internal system, the differences with previous service host configurations were not immediately apparent, and opinions such as the following were sometimes heard.

- I don't trust virtual servers. I'd rather use physical servers.
- This level of performance is not necessary, so I want cheaper servers.
- There is too much memory, so I'd like it reduced.
- Disk unit prices are higher than local disks.

Opinions such as these were heard at first, but even though individual servers have slightly higher prices, people understood the cost benefits when taking into account factors such as consolidation, facilities, and operating costs after the introduction of virtualization, and currently there is less resistance to in-house introduction.

Additionally, people are beginning to see the merits of servers that can be ready to use a few days after being requested, and can be returned when no longer required without requiring any physical work to be carried out at the data center. Because until now equipment was constructed using different devices for each service, it was difficult to reuse devices in other services when plans changed, and when devices were required urgently there was a delay before they could be used while waiting for their purchase and installation to be completed. Due to the introduction of NHN, service development flow and the system for making demand forecasts for equipment upgrades are gradually changing.

IIJ will continue investigating the security and I/O virtualization aspects of this technology, and polish it for eventual release to IIJ GIO and other customers.

Authors:
Yasumitsu Makino
Section Chief, System Operating Section, System Infrastructure Division, IIJ Service Business Department
Mr. Makino is engaged in the design and operation of server infrastructure for corporate and private services. Since 2008 he has consolidated equipment procurement and onsite construction work for service hosts at the System Operating Section, and worked towards optimizing the consolidation and operation of equipment through infrastructure systemization.

Shinya Hanataka
System Operating Section, System Infrastructure Division, IIJ Service Business Department
Tadashi Kobayashi
System Operating Section, System Infrastructure Division, IIJ Service Business Department

© 2009 Internet Initiative Japan Inc.