RIKEN Accel. Prog. Rep. **38** (2005)

CCJ Operation in 2003-2004

S. Yokkaichi, Y. Goto, H. Hamagaki,[*1] T. Ichihara, O. Jinnouchi, A. Kiyomichi, Y. Watanabe and H. En'yo

1 Overview

The operation of CCJ[1–3)], RIKEN Computing Center in Japan for RHIC[4)] physics, started in June 2000 as the largest off-site computing center for the PHENIX[5)] experiment at RHIC. CCJ performs, in PHENIX computing, the three roles of 1) the simulation center, 2) the Asian regional center and 3) the center of spin physics. Recently, DST (Data Summary Tape) production from raw data has become more important.

Many analysis and simulation projects are being carried out at CCJ, including some PHENIX official projects. They are shown on the web page: `http://ccjsun.riken.go.jp/ccj/proposals/`. The report from each project is described in this volume.

2 Current configuration

We have approximately 180 PC nodes operated using Linux, 166 are calculation nodes and the others are service nodes. Each calculation node has 1 GB of memory, 2 GB of swap area, 10∼31 GB of local work area and dual CPUs (Pentium III 700 MHz ∼ 1.4 GHz, Pentium 4 2.0 GHz). Red Hat 8.0/kernel 2.4.20 is operated on the calculation nodes, the same as at the RCF (RHIC computing facility)[6)], which is the main analysis facility for the PHENIX. The upgrade to Red Hat 8 was carried out in Jan. 2004.

We also have six SUN servers operated using Solaris, one is a service node and five are data servers, which are connected to large RAID systems (38 TB in total) and operated as NFS servers. Service nodes are used for various purposes such as a login server from WAN, a database server for analysis, an AFS client to share the PHENIX computing environment. Calculation nodes are connected with 100Base-T and data servers are connected with 1000Base-SX via a gigabit switch. Gigabit switch Alteon ACE 180 was replaced by Cisco Catalyst 4506 in Feb. 2004.

HPSS[7)] 4.5 is used as a mass-storage system in CCJ, the same as RCF. Approximately 180 TB/650000 files of data have been stored in CCJ-HPSS as of Oct. 2004. Five IBM p630 servers operated using AIX are used as the HPSS core server and data/tape movers. Eight 9940B tape drives (30 MB/s I/O & capacity 200 GB/cartridge) and 3000 tape cartridges in a StorageTek PowderHorn 9310 tape robot are available. The robot can handle approximately 5000 tapes, thus we can extend the total tape capacity to 1 PB. The upgrade to the current HPSS system was carried out in Nov. 2003, as a part of the RIKEN Super Combined Cluster System (RSCC)[8)], which is explained in the following section. We also have a tape duplication facility at BNL which consists of an IBM p630 server and two 9940B tape drives.

On the calculation nodes, the batch queueing system LSF[9)] 4.2 was operated and upgrading to LSF 5.1 was carried out in Apr. 2004.

In Mar. 2004, the AFS client was upgraded from a SUN WS to a PC (Pentium III 1.4 GHz/memory 1 GB) using OpenAFS on Linux. In Jun. 2004, a PC (dual Xeon 2.6 GHz/memory 2 GB) functioning as an NFS server was deployed which was connected to an 8 TB FC RAID system.

The DB system for PHENIX analysis was gradually being changed from Objectivity/DB (OBJY) to PostgreSQL this year. The last OBJY update at PHENIX was in Jun. 2004. The limited support of PostgreSQL in CCJ started in Apr. 2004.

3 Cooperation with RSCC[8)]

RSCC was deployed in Mar. 2004 by the RIKEN Advanced Center for Computing and Communication. Integrated operation between RSCC and CCJ was investigated from 2003, and opened for users in May 2004. RSCC includes 1024 nodes/2048 CPUs of PC clusters and HPSS as a mass-storage device, which is shared with CCJ.

As calculation nodes, 128 PC nodes (dual Xeon 3.06 GHz/2 GB memory/100 GB local disk), forming the 'pc2c' cluster in RSCC are dedicated to CCJ. Our data servers and HPSS can also be accessed by these nodes the same as for our old PC nodes. In particular, the user's home area and the PHENIX computing environment are shared using NFS.

As a batch queueing system, NQSII/ERSII (made by NEC) is operated in the pc2c cluster. For the time being, the performance of the system is insufficient for our usage in experimental nuclear physics, e.g., to process up to 1000 ∼ 10000 jobs per day. The redundancy of the system is also worse than that of LSF. Improvement or replacement should be planned.

To estimate how appealing our CPU power is to users, we examine the 'obsolescence index (OI)', which is defined as the 'total CPU power in the cluster' divided by the 'fastest desktop CPU power available in that year'. In other words, the cluster is equivalent to the fastest desktop PC × OI for that year. If there is no CPU upgrade, the OI decreases because CPUs become faster year by year. Here, the 'effective CPU power' is measured using a typical PHENIX software

*1 CNS, University of Tokyo

and normalized with respect to 1 GHz Pentium III.

The trend of the OI of CCJ (building commenced in 1999) is shown in Fig. 1. It shows that, at the end of JFY 2004, the fastest Intel CPU is estimated to be 4.2 GHz (solid triangles), CCJ (with RSCC pc2c) have calculation power equivalent to approximately 1100 Pentium III 1 GHz CPUs (solid squares) and equivalent to approximately 250 of the fastest available CPUs (open squares) for the PHENIX software. RSCC made CCJ very appealing in terms of CPU power.

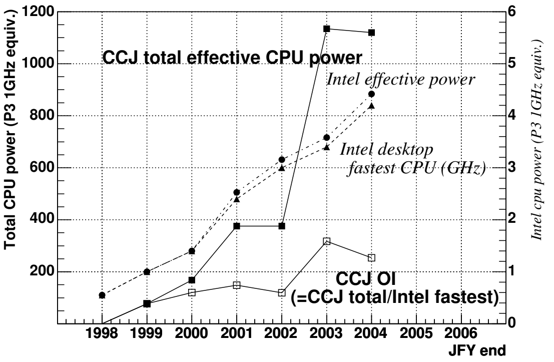

Fig. 1. Obsolescence index (OI) of CCJ

4 Hardware problems and retirement

Here, we report hardware problems that occurred in the term of Dec. 2003 ∼ Nov. 2004. Concerning the PC nodes, 10% of nodes were down due to various problems involving the cooling fan, disk, memory, CPU and power supply. In total, 25% of calculation nodes has been down because 15% were down in the previous year[2)]. The problems occurred in the older (and slower) node group purchased in JFY 1999∼2001. There were almost no problems with the newer (and faster) node group purchased at the end of JFY 2001 including Linux Networx R-Cluster and IBM xSeries 330. Concerning the SUN servers, M/B replacement in Mar. and memory & CPU replacement in Oct. were the only two problems that required the shutdown

Concerning the RAID systems, the controller, cable, FC-hub and I/F card problems which caused a disk-offline are approximately 10 times including two cases that required the shutdown of the servers. Simple disk failure is not included because disks comprising RAID are hot swappable. The old problematic 3.2 TB SCSI RAID purchased in JFY 1999 was retired in Mar. 2004. The disk-failure rate of the old 6 TB FC RAID purchased in JFY 2000 became very high (4 times/month or more) this year. We could not neglect the probability of the simultaneous failure of two disks, which could cause the loss of data, therefore we retired the RAID in Dec. 2004.

The air-conditioning system in the machine room was down on Mar. 9, 2004, which was caused by a mistake by the building maintenance staff. A high room temperature for 6∼8 hours could have damaged the hardware, although data were not lost in this case.

The most serious problem occurred in Feb. 2004, that is, an operational error of the disk caused a loss of data in the user's home area. Although backup data was restored, the data updated in the previous 7 days were lost. (Backup using the *dump* command was operated once a week. Unfortunately, the accident occurred on the 7th day.) Backup using *rsync* once a day was also started after the accident.

5 Data transfer via network

In PHENIX Run4 (2003/11 ∼ 2004/5), approximately 200 TB of raw data for the heavy-ion collision experiment and 20 TB for the p-p collision experiment were recorded. The transfer of the p-p data from BNL to RIKEN via WAN was carried out and a transfer rate of approximately 100 Mbps(12 MB/sec) for 24 hours (or more) was achieved using *bbftp*[10)]. In the test, approximately 70 MB/s was achieved using *GridFTP*[11)].

The project to transfer the Run5 (2004/11∼) data via WAN is planned to produce DST in CCJ, particularly for the p-p collision data. The amount of data of Run5 will be approximately 200∼400 TB for just the p-p data for 2∼3 months. If a transfer rate of 70 MB/s (= 6 TB/day = 180 TB/month) cannot be maintained continuously, transfer using tape will also be required.

6 Outlook

The following tasks should be processed in the next several months: upgrade to Scientific Linux 3 (and LSF 6.0) on the calculation node, reinforcement of the DB server for PHENIX, improvement of the batch queueing system of the pc2c cluster in RSCC, establishment of the (semi-)automatic data-transfer procedure from the PHENIX counting house and the efficient data-I/O procedure for the p-p DST production at CCJ.

References

1) `http://ccjsun.riken.go.jp/ccj/`
2) S. Yokkaichi et al.: RIKEN Accel. Prog. Rep. 37, 257 (2004)
3) Y. Watanabe et al.: RIKEN Accel. Prog. Rep. 36, 262 (2003), T. Ichihara et al.: RIKEN Accel. Prog. Rep. 35, 236 (2002)
4) `http://www.bnl.gov/rhic`
5) `http://www.phenix.bnl.gov`
6) `http://www.rhic.bnl.gov/RCF/`
7) `http://www.hpss-collaboration.org/`
8) `http://rscc.riken.jp`
9) `http://www.platform.com/products/LSF/`
10) `http://doc.in2p3.fr/bbftp/`
11) `http://www.globus.org/datagrid/gridftp.html`