BELLE2-NOTE-TE-2024-003
Version: 1.6

Computing Accounting for JFY 2023

The Belle II Computing Steering Group

May 27, 2024

1 Introduction

In this note we present the accounting information for the Japanese Fiscal Year (JFY) 2023, covering the period from April 2023 to April 2024. In the first section, the computing activities originally planned for 2023 are summarized. In the second section, the activities undertaken are described, along with the accounting information. In the following, dates refer to JFY unless otherwise stated.

2 Planned activities and resource estimate for year 2023

During 2023, SuperKEKB and Belle II were not expecting to collected data, being in the process of various updates both for the accelerator and the detector during Long Shutdown 1 (LS1). Therefore, no prompt processing was planned. Furthermore, the previous major reprocessing (proc13) was near completion at the end of 2022, as was the corresponding run-dependent MC production.

The main activities planned for 2023 were

- Completion of data reprocessing (proc13).
- Completion of MC production (MC15). The ratio of the generic run-dependent MC (MCrd) event sample to the detector event sample is 4 for hadronic and τ events and smaller for low-multiplicity events. Another 1 ab^{-1} equivalent of run-independent MC (MCri) is produced. In addition, samples of signal events for specific studies are centrally produced upon request of the Physics Analysis Groups.
- Skimming of detector and MC data.
- Physics analysis based on MC and reprocessed recorded data.

Taking into account the planned activities and the best knowledge of the parameters of their modeling available in February 2022, we estimated to need, for JFY 2023, the amount of computing resources listed in Table 1. The rightmost column has the total amount of storage and CPU power actually used for different activities in 2023. Note that 0.2 PB of disk was requested for prompt calibration at BNL, while the 1.5 PB of

space required for recalibration was provided as tape at DESY. The recalibration center at DESY also includes 451 TB of disk space, which was purchased for this purpose.

The discrepancy in the pledged and used CPU for user analysis in Table 1 comes primarily from an underestimation in the CPU required per analysis cycle and the number of concurrent analyses. The CPU estimate was revised in the resource estimate created in 2023 to better reflect the actual use in previous years.

3 Data Production Activities in 2023

The planned data reprocessing (proc13) was completed by the beginning of 2023 as well as the MCrd campaign. This was the second full MCrd campaign, and it went much smoother than the first one, thanks to the experience gained both on setup and preparation as well as for grid operations.

The major activity for 2023 was the skimming campaign. For the first time, a full skim was run for data, MCri, and MCrd. About 70 different skims were produced, grouped together in about 15 "combined-skims". Each combined skim runs over a single input, but produces the multiple output files, one for each skim. The full skim campaign lasted several months, mostly due to the complexity to run on the MCrd samples.

Another major activity was the production of signal samples, both for MCri and MCrd, the latter for the first time. These productions took place on grid. We produced about 700 different signals for MCri, and about 300 for MCrd. Despite the reduced number, the MCrd signal production was significantly more complex than for MCri, involving a complex setup phase and operation on the grid.

In addition, several dedicated productions of MCri (equivalent to approximately 1 ab^{-1}) were performed to test improvement in background modeling as well as improved decays tables.

Additional productions were performed to validate release-07 (s-proc4). The new release was used to produce mDST samples for a fraction of the data collected in the last experiment, in which the beam-related background was higher. For the first time, we also produced the corresponding MCrd sample. This process was repeated three times due to issues found in the release. In order to test different configurations of tracking reconstruction software, four different processings was performed for both data and MCrd.

As soon as the latest software version, release-08, was available, we started the full validation process (s-proc5). This includes production of dedicated samples of MCri, both generic and signal, for a first evaluation of the performance. Then we produced mDST for data and MCrd as we did for s-proc4, in this case with just one tracking configuration. The production was repeated twice due to some serious problem found in the software.

The other major activity for 2023 was the review of the calibration for reprocessing. A task force was set up along with a review panel. A thorough assessment of the resource needs for this re-calibration was completed, based on the experience gained so far. Two different re-calibration scenarios were considered: one fully local, as was done so far at DESY/NAF, and one using resources in multiple sites via the grid. The review showed that the amount of resources needed was not large enough to justify the additional complexity of a grid-based calibration. In addition, the available resources at KEKCC are

	Estimated	Used
Tape for raw data (PB)	6.70	5.75
Tape for hRaw data (PB)	1.90	0.97
Tape for recalibration (PB)	2.00	0.92
Total tape (PB)	**10.59 (estimate)**	**7.64**
Disk for data processing buffer (PB)	0.07	2.07
Disk for data mDST buffer (PB)	0.26	0.13
Disk for data mDST (PB)	0.93	0.60
Disk for data cDST (PB)	0.0	0.08
Disk for random trigger files for MC (PB)	0.12	0.03
Disk for MC buffer (PB)	0.18	0.63
Disk for MC mDST buffer (PB) Disk for MC mDST (PB)	0.60 4.22	2.86
Disk for MC cDST (PB)	0.30	0.14
Disk for data and MC uDST (PB) Disk for data and MC uDST buffer (PB)	7.61 2.04	2.47
Disk for data and MC skim buffer (PB)	0.60	0.9
Disk for user ntuples (PB)	0.84	0.84
Disk for calibration	0.20	0.12
Total disk (PB)	**19.64 (estimate)**	**11.69**
CPU for data processing (kHS06) CPU for data reprocessing (kHS06)	7.01 21.70	7.20
CPU for calibration/recalibration (kHS06)	13.00	4.05
CPU for MC production (kHS06)	248.99	72.44
CPU for skimming (kHS06)	36.67	16.55
CPU for analysis (kHS06)	76.32	143.68
CPU for local resources (non-production) (kHS06)		108.43
Total CPU (kHS06)	**404 (estimate)**	**352**

Table 1: The resource estimate for year 2023. The rightmost column has the total amount of storage and CPU power used for different activities in 2023. Note that the total disk use corresponds to the maximum actual use after optimization and does not correspond to the sum of maximum values for each use case (see the text for more discussion).

well matched for the near term. Therefore, the decision was made to move re-calibration to KEKCC from DESY, which will no longer provide those resources.

4 Accounting

The amount of CPU power that the different countries participating in the Belle II collaboration have pledged, have actually provided, and that has been used on average in 2023 are shown in Table 2. The column "Pledged CPU" shows the amount of CPU power (in kHEPSpec06) that the different countries have pledged. The pledges are calculated according to the PhD fraction of each country. The column "Provided CPU" reports the amount of CPU (in kHEPSpec06) dedicated to Belle II at the different sites at the time of the last yearly survey (June 2023). While some CPUs are accessible only to Belle II, shared CPUs are used by different experiments according to a fair share algorithm that takes into account the amount of pledged CPU. It is important to note that the resources that a site provides, on top of "pledged" resources that are guaranteed to the Belle II collaboration, may include opportunistic resources or other resources that are only temporarily available and that will not necessarily be available in the future. The column "Used" shows the amount of CPU (in kHEPSpec06) used on average in 2023. These numbers are obtained by taking the average of the HS06 per core times the number of cores times the number of hours they were used for each site and then dividing by the average number of hours per month (720) to get the average HS06 value that can be compared with the pledged and provided CPU per site. Fig 1 displays the CPU power used in different countries per month. Note that for some sites the amount of CPU used is less than that pledged amount due to the fact that there were not enough Belle II jobs to make full use of the available resources.

The used CPU power is extracted from the EGI portal for the sites who export their accounting to it. It is extracted from the KEK Dirac portal for the other sites. From April 2023 through March 2024, Belle II has used an average of 240 kHEPSpec06 of CPU power on the Grid. Prompt Calibration done at BNL and recalibration done at DESY required on average 4 kHEPSpec06. Some detector studies and some physics analysis are done on local resources and 110 kHEPSpec06 have been devoted to that.

The amount of disk storage (in TB) that the different countries participating in the Belle II collaboration have pledged, have actually provided, and the maximum amount in use during 2023 are shown in Table 3. The column "Pledged Disk" shows the amount of disk storage that the different countries should have provided. The pledges are calculated according to the PhD fraction of each country. The column "Provided Disk" reports the amount of disk storage made available to Belle II at the different sites at the time of the last yearly survey (June 2023). The storage available only to local Belle II users and not configured as a storage element is not included. The last column shows the maximum amount of storage in use during 2023.

The total amount of disk storage used by Belle II during 2023 was 11.69 PB. This amount does not correspond to the sum of maximum use for each individual category in Table 1, since actual usage varies throughout the year. The sum of maximum values gives a total of 12.95 PB. This amount of disk storage was never required at a given point during 2023, but indicates potential usage if data management were handled in a less

Country	Site	Pledged CPU	Total CPU provided	Total used
Armenia	-	1.6	0	0
Australia	Melbourne	6.6	28	3.1
Austria	HEPHY	5.7	4.8	3.3
Canada	Uvic	12.7	100	59.7
China	Fudan	0	0	0
China	USTC	0	0	0
China	Shandong	0	0	0.1
China	IHEP	0	0	0.3
China	Total	25.4	15	0.4
Czech	CESNET	4.1	16.4	5
France	IN2P3CC	0	0	3.5
France	IPHC	0	0	0
France	LAL	0	0	0.2
France	Total	17.6	21.6	3.7
Germany	DESY	0	0	23.9
Germany	KIT	0	0	20.7
Germany	MPPMU	0	0	2.1
Germany	Total	62.3	207.7	46.7
India	-	9.8	24.7	0
Israel	-	3.3	2.7	0.8
Italy	CNAF	0	0	15.2
Italy	Cosenza	0	0	0
Italy	Frascati	0	0	0.4
Italy	LNL	0	0	0.2
Italy	Napoli	0	0	3.7
Italy	Pisa	0	0	1
Italy	Roma3	0	0	0.6
Italy	Torino	0	0	3.7
Italy	Total	50	150.5	16.5
Japan	KEK	0	0	46.2
Japan	NDU	0	0	0.2
Japan	TMU	0	0	0.1
Japan	Niigata	0	0	0
Japan	KMI	0	0	0
Japan	Total	71.5	90.3	46.2
Korea	KISTI	14.8	10.9	0.3
Malaysia	-	1.6	0	0
Mexico	-	4.9	2.4	0
Poland	CYFRONET	3.3	2	0.4
Russia	BINP	0	0	0.1
Russia	MIPT	0	0	0
Russia	Total	25.4	7	0.1
Saudi Arabia	-	2.5	0	0
Slovenia	SIGNET	9	38.5	11
Spain	-	0.8	0	0
Sweden		0	0	0
Taiwan	NTU	5.7	18.3	0
Thailand	-	3.3	0	0
Turkey	ULAKBIM	1.6	0.9	0.2
UK	UKI-LT2-QMUL	0	0	0
Ukraine	-	3.3	0	0
USA	BNL	0	0	33.4
USA	Univ of Miss	0	0	0.4
USA	Total	54.9	155.4	33.8
VietNam	-	1.6	0	0
Total All Countries	-	403.7	897.1	239.9 (108.4 local)

Table 2: CPU power (in kHEPSpec06) that the different countries participating in the Belle II collaboration have pledged, have actually provided, and that has been used in 2023.

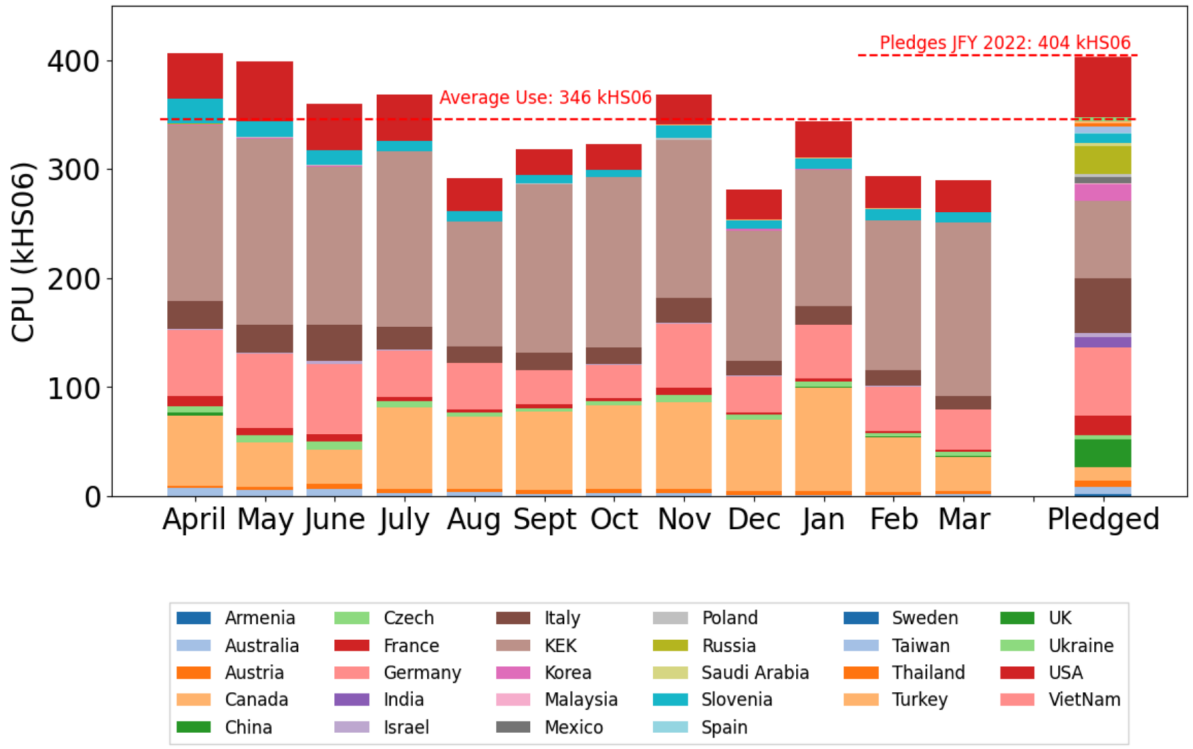

Figure 1: CPU power per month (in kHEPSpec06) used in different countries during 2023, including both grid and local resource usage.

optimal way.

The amount of tape storage (in PB) that the different Raw Data Centers have actually provided, along with its use, is shown in Table 4.

Country	Site	Pledged Disk	Provided Disk	Used Disk
Armenia		68	0	0
Australia	Melbourne	267	380	6
Austria	HEPHY	169	310	247
Canada	Uvic	320	1080	695
China	IHEP	1017	280	49
Czech	CESNET	136	238	204
France	IN2P3CC	320	595	421
France	IPHC	100	88	70
France	LAL	0	50	20
France	Total	420	733	512
Germany	DESY	1720	2310	1650
Germany	KIT	1720	1940	1230
Germany	MPPMU	220	200	93
Germany	Total	3660	4450	2973
India		407	0	0
Israel		136	60	0
Italy	CNAF	820	819	750
Italy	Frascati	0	11	9
Italy	Napoli	0	590	310
Italy	Pisa	0	209	43
Italy	Roma3	0	2	1
Italy	Torino	0	350	313
Italy	Total	1810	1981	1425
Japan	KEK	0	3650	2986
Japan	KMI	0	300	42
Japan	Total	2910	3950	3028
Korea	KISTI	678	200	40
Malaysia		34	0	0
Mexico		203	0	0
Poland	CYFRONET	237	11	8
Russia		881	0	0
Saudi Arabia		68	0	0
Slovenia	SIGNET	339	1210	1090
Spain		34	0	0
Taiwan	NTU	237	869	45
Thailand		136	0	0
Turkey	ULAKBIM	68	143	41
Ukraine		102	0	0
USA	BNL	1970	2900	2100
USA	Univ of Miss	0	0	0
USA	Total	1970	2900	2100
VietNam		68	0	0
Total All Countries		16373	18795	12463

Table 3: Storage (in TB) that each country has pledged and that has actually provided in 2023. The last column shows the amount of used storage.

Site	Pledge	Raw	hRaw	Other	Total
KEKCC	4.40	2.66	0.398	0	3.058
Canada - Uvic	0.34	0.122 (test)	0	0	0.122 (test)
France - IN2P3	0.34	0.12	0.029	0	0.149
Germany - DESY	0.22	0.089	0.018	0.921 (calib)	1.027
Germany - KIT	0.22	0.103	0.021	0	0.124
Italy - CNAF	0.45	0.208	0.042	0	0.249
USA - BNL	2.83	2.17	0.461	0	2.631

Table 4: Tape storage (in PB) provided by different Raw Data Centers in 2023. For the time being, the University of Victoria is storing raw data on disk with backup on tape.