BELLE2-NOTE-TE-2025-004

Computing Accounting for JFY 2024

The Belle II Computing Steering Group

July 1, 2025

1 Introduction

In this note we present the accounting information for the Japanese Fiscal Year (JFY) 2024, covering the period from April 2024 through March 2025. In the first section, the computing activities originally planned for 2024 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 2024

During 2024, SuperKEKB and Belle II were expecting to collect data during two run periods, one in the fall and one in the spring. Due to various constraints, the spring run period was canceled.

The main activities planned for 2024 were

- Completion of data reprocessing (proc16).
- Completion of MC production (MC15/MC16). 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 2023, we estimated to need, for JFY 2024, the amount of computing resources listed in Table 1. The rightmost column has the total amount of storage and CPU power, in HEPScore23 (HS23), actually used for different activities in 2024. Note that the values in the "Used" column of the table correspond to the maximum usage during the period of interest and may not match the total disk usage, since samples are staged to disk and removed as needed. Additional details related to the requested resources are available in the *Belle II Computing Resource Estimate for 2024-2027*[1]. At the time the resource estimate was prepared, the total integrated luminosity of the data sample was expected to be 0.9 ab^{-1}. The actual delivered luminosity was about 0.64 ab^{-1}, while the recorded luminosity was about 0.58 ab^{-1}.

The difference between the estimated and used resources is primarily due to the lower than expected integrated luminosity of data collected by Belle II. Some of the discrepancies between disk and CPU usage is due to changes to the way the resources are estimated. For example, the data processing buffer was increased in recent estimates, based on the experience and adjustments made to handle data processing requirements. Similarly, the understanding of CPU needs for analysis has evolved as larger data samples are made available and as changes are made to the software and to the skimming scheme.

3 Data Production Activities in 2024

In 2024, the major activity for data processing was the full re-processing of the run 1 data, named proc16 (the number for the reprocessing was advanced to match that for the most recent MC campaign). The re-processing also included a partial recalibration. Some new calibrations were included in the basf2 release that was used (release-08), and some resulted in improvements that were included in the reprocessing. The recalibration was performed using KEKCC resources. Due to the limited amount of luminosity and the complexity of re-calibration, the data from very first experiment, collected in 2019, were not re-calibrated nor reprocessed.

3.1 Data processing and MC production

The processing for all data collected in run 1 was completed and the corresponding MCrd was generated as well. Thanks to a more streamlined process to collect the needed payload, we were able to produce MCrd along with data, with limited delay between the two. The usual mix of generic MC was generated, with equivalent integrated luminosities four times that of data for $q\bar{q}$ and τ pairs sample, so about 1.6 ab^{-1}, and smaller for low multiplicity final states.

A sample of MCri, with background level corresponding to run 1 condition, was also generated and simulated, for a total integrated luminosity of 0.5 ab^{-1}. A similar sample corresponding to the run 2 condition will be generated in 2025.

The old experiments before March 2021 were collected with the HLT in monitoring mode, flagging the events without filtering them to study the trigger performances. This resulted in a very large size for the corresponding raw data, mostly containing events which pass L1 trigger but not any HLT one. An offline filtering of raw events has been completed while the raw data was staged on disk for reprocessing, to be used as input for future re-processing, thus reducing the disk space needed. The un-skimmed raw data has been retained.

[1] https://indico.belle2.org/event/15670/contributions/94921/attachments/35073/51953/BELLE2-NOTE-TE-2023-004.pdf

	Estimated	Used
Tape for raw data (PB)	7.34	6.07
Tape for hRaw data (PB)	2.13	2.24
Tape for cosmics (PB)	0.15	0.0
Total tape (PB)	**9.62 (estimate)**	**8.31**
Disk for data processing buffer (PB)	0.58	2.24
Disk for reprocessing buffer (PB)	0.50	1.90
Disk for data mDST buffer (PB)	0.72	0.24
Disk for data mDST (PB)	1.07	1.03
Disk for random trigger files for MC (PB)	0.16	0.04
Disk for MC buffer (PB)	0.41	—
Disk for MC mDST buffer (PB)	1.35	0.00
Disk for MC mDST (PB)	4.72	3.68
Disk for MC cDST (PB)	0.40	0.16
Disk for data and MC uDST (PB)	8.36	3.05
Disk for data and MC skim buffer (PB) Disk for data and MC uDST buffer (PB)	0.73 2.43	1.29
Disk for user ntuples (PB)	0.89	0.70
Disk for calibration	2.70	0.66
Total disk (PB)	**25.04 (estimate)**	**14.99**
CPU for data processing (kHS23) CPU for data reprocessing (kHS23)	10.6 25.3	21.2
CPU for calibration (kHS23)	16.0	4.1
CPU for MC production (kHS23)	289.2	112.8
CPU for skimming (kHS23)	43.3	8.2
CPU for analysis (kHS23)	136.0	61.8
CPU for local resources (non-production) (kHS23)		160.7
Total CPU (kHS23)	**520 (estimate)**	**204 + 165 (local)**

Table 1: The resource estimate for year 2024. The rightmost column has the total amount of storage and CPU power used for different activities in 2024.

Prior to the full reprocessing, several spacial calibrations and processings of a small sample of data, corresponding to about 30 fb^{-1} collected at the end of run 1, were performed, in order to test and validate release-08.

Another important task has been the prompt calibration and processing of data collected in 2024 data taking. The calibration was performed at BNL. This prompt calibration suffered from several delays, so that it ended only at the end of calendar year 2024 for data taken before summer. In the fall, we started prompt processing of 2024 data, and the corresponding MCrd samples have been produced as well. At the beginning of 2025, when part of the data and MC processing was complete, a problem with tracking reconstruction was found, requiring an update of the software. After this, the full run2 data and MC were re-processed with a patched software version, followed by a skimming campaign. The re-processing was completed by end of fiscal year 2024.

3.2 Skimming

Significant effort has been made to optimize the skim production of data and MC samples, moving from the old schema in which each skim had its own set of output files to a new "flagged-skim" approach. In the latter, skims with significant overlap were produced together and the output saved in the same files, adding a flag to each event to signify which skims it passed. This allows for an offline selection to be performed quickly. The flagged-skim approach has the advantage of significantly reducing the number of productions. Previously one production was required per experiment per skim, now only one is needed per experiment per skim-group. This approach also reduces the overall size of the skim output, since events passing multiple skims are now saved once per group. The drawback is that analysts have to run over a larger dataset, though they are able to skip uninteresting events quickly using the flag variable. Overall, the number of jobs to process the full dataset using skim is reduced by a factor of about 10-40, depending on the skim type, with respect to what is needed to access the unskimmed dataset.

3.3 Signal MC

A significant amount of signal samples have been produced for the MC16 campaign, starting with those for active or planned analyses using the proc16/MC16 dataset. All signal samples have been produced as run-dependent MC. Signal samples with fewer than 10 million events were produced locally at KEKCC, while larger requests were fulfilled using the grid. The signal samples produced locally will be uploaded to the grid for replication analysis use.Last but not least, a non-negligible number of signal samples for MC15 were produced to cope with new requests for ongoing analyses. The total number of MCrd signal samples for MC15 amounts to almost 600, which corresponds to about 13 billion events. For MC16, about 200 have already been produced, corresponding to about 1 billion events.

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 2024 are shown in Table 2. The column "Pledged CPU" shows the amount of CPU power (in kHS23) that the different countries have pledged, according to the June 2024 site report. The pledges are calculated according to the PhD fraction of each country, excluding those with small PhD counts and an inability to provide computing resources.

The column "Provided CPU" reports the amount of CPU (in kHS23) made available to Belle II at the different sites at the time of the last yearly survey (June 2024). 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 kHS23) used on average in 2024. These numbers are obtained by taking the average of the HS23 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 HS23 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 2024 through March 2025, Belle II has used an average of 204 kHS23 of CPU power on the grid. Prompt calibration done at BNL and recalibration done at KEK required on average 4 kHS23. Some detector studies and some physics analysis are done on local resources and 161 kHS23 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 2024 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, excluding those with small PhD counts and an inability to provide computing resources. 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 2024). 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 2024.

The total amount of disk storage used by Belle II during 2024 was 15.4 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 total amount of disk storage indicates the potential usage if data management were handled in a less optimal way.

The amount of tape storage (in PB) that the different Raw Data Centers have actually

Country	Site	Pledged CPU	Provided CPU	Used CPU
Armenia	-	0	0	0
Australia	-	8.4	28	1.8
Austria	HEPHY	7.4	4.8	3.3
Canada	Uvic	19.2	75	33.9
China	Fudan			0
China	USTC			0
China	Shandong			0
China	IHEP			5
China	Total	31.6	15	5
Czech	CESNET	5.3	17.4	6.9
France	IN2P3CC			5
France	IPHC			0
France	LAL			0.2
France	Total	22.3	29	5.1
Germany	DESY	31.5		17.1
Germany	KIT	31.5		20.6
Germany	MPPMU			4
Germany	Total	67	210.1	41.6
India	-	15.8	20.7	0
Israel	-	5.3	2.7	1.1
Italy	CNAF			16.4
Italy	Cosenza			0
Italy	Frascati			1.2
Italy	LNL			0.2
Italy	Napoli			8.8
Italy	Pisa			3.3
Italy	Roma3			0.2
Italy	Torino			5.1
Italy	Total	64.9	161.1	35.2
Japan	KEK	60.3		28.6
Japan	NDU			0.5
Japan	TMU			0.2
Japan	Niigata			0
Japan	KMI			0
Japan	Total	105.8	120.4	29.2
Korea	KISTI			2.7
Korea	SSU			0
Korea	Yonsei			0
Korea	Total	18.9	10.9	2.7
Malaysia	-	0	0	0
Mexico	-	6.3	0	0
Poland	CYFRONET	3.2	2	0.1
Russia	BINP			0.2
Russia	MIPT			0
Russia	Total	28.4	7	0.2
Saudi Arabia	-	0	0	0
Slovenia	SIGNET	12.6	38.5	10.6
Spain	-	0	0	0
Sweden		0	0	0
Taiwan	NTU	7.4	18.3	0
Thailand	-	0	0	0
Turkey	ULAKBIM	2.1	2.8	0.2
UK	UKI-LT2-QMUL	0	21	0.5
Ukraine	-	0	0	0
USA	BNL	66.5	140.5	26.4
VietNam	-	0	0	0
Total All Countries	-	**498.3**	**925.2**	**204**

Table 2: CPU power (in kHS23) that the different countries participating in the Belle II collaboration have pledged, have actually provided, and that has been used in 2024.

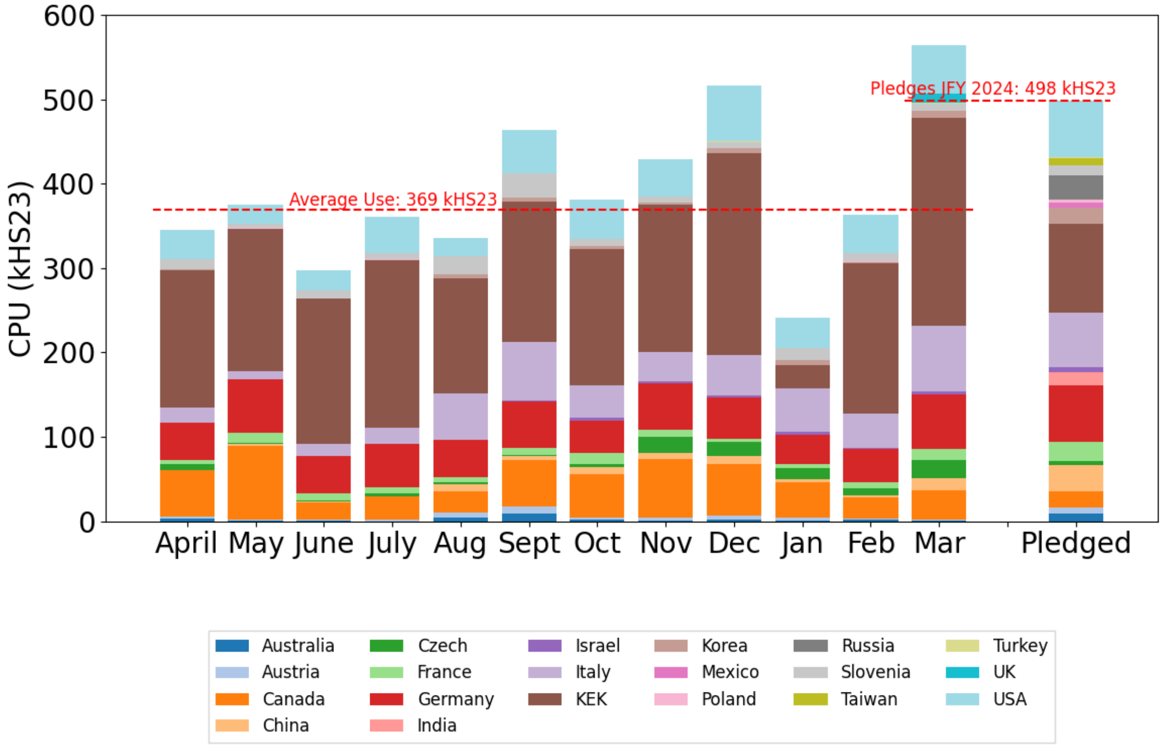

Figure 1: CPU power per month (in kHS23) used in different countries during 2024, including both grid and local resource usage.

provided, along with its use, is shown in Table 4.

A Glossary

Some common abbreviations and jargon used in this document are defined here.

- JFY: Japanese Fiscal Year
- MCri: run-independent MC, produced with fixed conditions
- MCrd: run-dependent MC, produced with real conditions during data taking
- HS23: HEPScore23
- PB: 10^{15} Bytes
- HLT: high-level trigger
- L1: level-1 trigger

Country	Site	Pledged 2024	Pledged 2025	Provided Disk	Used Disk
Armenia	-	100	0	0	0
Australia	Melbourne	380	250	400	12
Austria	HEPHY	330	300	310	281
Canada	Uvic	710	660	987	816
China	IHEP	1430	250	280	200
Czech	CESNET	240	170	440	273
France	IN2P3CC			763	601
France	IPHC			88	68
France	LAL			50	302
France	Total	860	790	901	971
Germany	DESY	2040	2040	2310	1890
Germany	KIT	2040	2040	1940	1480
Germany	MPPMU	260	260	200	93
Germany	Total	4340	4340	4450	3463
India	-	720	690	0	0
Israel	-	240	0	0	0
Italy	CNAF			1320	971
Italy	Frascati			55	9
Italy	Napoli			860	639
Italy	Pisa			209	170
Italy	Roma3			14	1
Italy	Torino			350	322
Italy	Total	2720	2360	2808	2111
Japan	KEK	5240	5000	4095	3320
Korea	KISTI	860	640	200	53
Malaysia	-	50	0	0	0
Mexico	-	290	260	0	0
Poland	CYFRONET	140	130	10	5
Russia	-	1290	0	0	0
Saudi Arabia	-	140	0	0	0
Slovenia	SIGNET	570	560	1210	1070
Spain	-	50	0	0	0
Sweden	-	190	0	0	0
Taiwan	NTU	330	340	869	79
Thailand	-	190	0	0	0
Turkey	ULAKBIM	100	90	143	41
UK	QMUL	0	90	1000	21
Ukraine	-	190	0	0	0
USA	BNL	2720	2470	2900	2690
VietNam	-	100	0	0	0
Total All Countries		**24520**	**19390**	**21003**	**15406**

Table 3: Storage (in TB) that each country has pledged and that has actually provided in 2024. The last column shows the amount of used storage. The pledges for 2025 are also shown for reference.

Site	Pledge	Raw	hRaw	Total
KEK	4.81	3.020	1.250	4.270
Canada - Uvic	0.37	0.039	0.071	0.110
France - IN2P3	0.37	0.186	0.090	0.276
Germany - DESY	0.24	0.128	0.061	0.189
Germany - KIT	0.24	0.144	0.062	0.206
Italy - CNAF	0.49	0.260	0.088	0.348
USA - BNL	3.1	2.290	0.619	2.909

Table 4: Tape storage (in PB) provided by different Raw Data Centers in 2024. For the time being, the University of Victoria is storing raw data on disk with backup on tape.