BELLE2-NOTE-TE-2025-003

Belle II Computing Resource Estimate for 2026-2029

The Belle II Computing Steering Group

July 1, 2025

1 Introduction

In this note we present an estimate of the computing resources needed by the Belle II collaboration in Japanese fiscal years (JFY) 2026 to 2029 and describe the planned activities and the assumptions that lead to this estimate. This estimate supersedes the one presented at the February 2024 BPAC review. Compared with the previous version, this estimate takes into account the need to retain the prompt datasets while the multi-year reprocessing campaigns are ongoing. This results in an increase to disk and CPU requirements, but better anticipates the needs for production of physics results. The luminosity estimate is also revised. In the following, dates refer to JFY unless stated otherwise.

An outline of the Belle II computing model is presented, followed by a description of the planned computing activities. The input parameters used for the resource estimate are discussed and the resource estimate for years 2026 to 2029 are presented. Appendix A shows a breakdown of the resource estimate contribution for each of the planned activities.

2 Outline of the Belle II computing model

Belle II uses a three-level hierarchical structure of computing sites: raw data centers, regional data centers, and MC production centers. Raw data centers are also regional data centers and MC production centers. Regional data centers are also MC production centers. A schematic of the distributed computing system at Belle II is given in Figure 1.

The raw data coming out of the online system are permanently stored, calibrated and processed at raw data centers. The fully reconstructed events, produced from the raw data processing step, are stored in mDST format, which includes the minimum information necessary for analysis. Monte Carlo (MC) events are simulated and reconstructed using the same software used to process detector events and then also stored in mDST format. Detector and MC mDSTs are stored in regional data centers.

Detector and MC mDST samples are then “skimmed” at regional data centers to create a sample of events that suit a specific set of physics analyses. The output of the skimming step is the full mDST information for a selected subset of events. Starting in 2024, the skimmed samples no longer contain additional analysis-related objects, except for the Full Event Interpretation (FEI) skim, which requires extra information determined

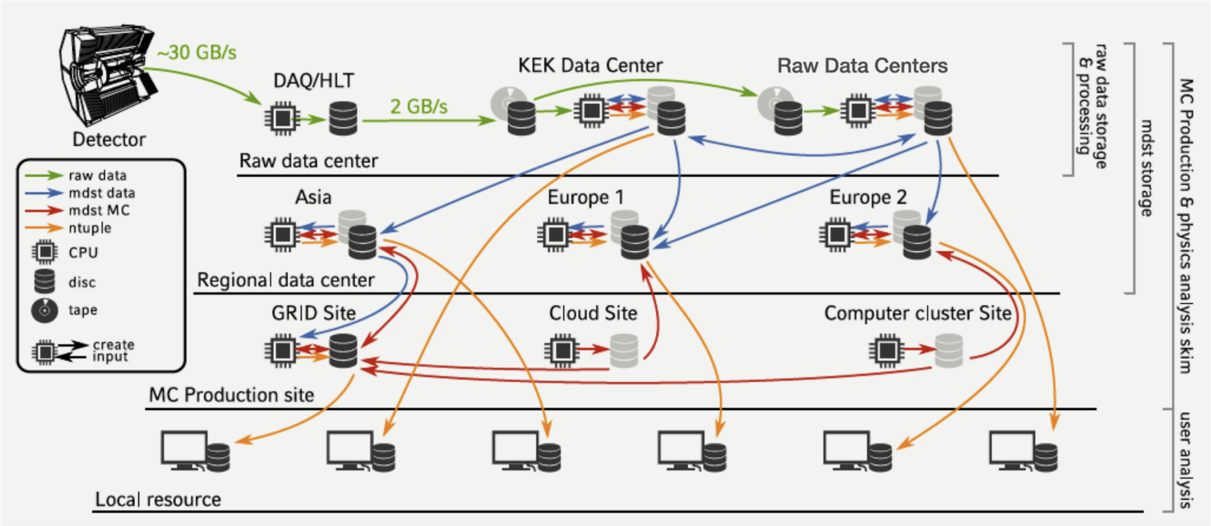

Figure 1: A schematic diagram of the Belle II distributed computing system.

during the skimming step. Reproducing the analysis-related objects rather than storing them results in a marginal increase in processing time while reducing the size of the skimmed sample. The analysis-related objects are also necessary for determination of systematic uncertainties anyway. The objects produced from the FEI, however, are much more time intensive to produce, so they are added to the mDST information for the FEI skim samples.

The understanding of the detector and the quality of the software will improve over time, resulting in new releases of the software and better knowledge of the detector conditions and calibrations that will require reprocessing of the data to exploit the improvements. Reprocessing of detector data with a different software release is expected to trigger the re-creation of the corresponding MC samples and of the skimming of the new data and MC samples.

User analysis will run on skimmed events, unless a suitable skim is not available. Detector data processing, simulation, and skimming will be centrally managed, while physics analysis will be the responsibility of users.

3 Outline of the data flow

A schematic for Belle II data preparation is given in Figure 2.

Two trigger levels are employed in the online system. The level 1 (L1) trigger includes multiple lines and prescales. Events that pass at least one L1 trigger line are processed by the High Level Trigger (HLT), which provides a different set of trigger lines and prescales. Events that pass at least one trigger line are kept, while those that pass no HLT trigger lines are not recorded if the HLT is running in filtering mode, which is the default.

Raw data are transferred from the online system to KEKCC by the core computing team and then registered to the grid for processing at raw data centers. The core computing team also uses the trigger decisions made online by the HLT to produce HLT-skimmed raw data samples (hRaw) that are registered to the grid and replicated to raw data cen-

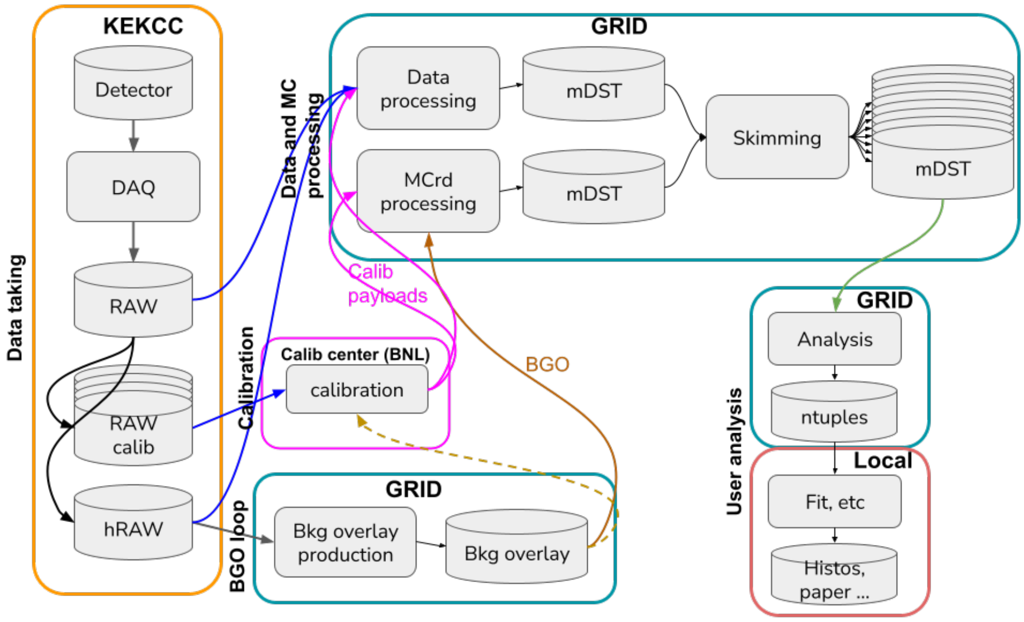

Figure 2: Schematic for Belle II data preparation.

ters for processing. The hRaw samples that are useful for calibration are replicated to BNL for prompt calibration. Only HLT calibration skims, HLT delayed bhabhas used to produce the background overlay files needed for MC production, and the HLT hadron skim samples are produced as hRaw.

Data calibration proceeds according to the following scheme:

- Detector-based calibrations are performed using dedicated "local runs" during data taking.
- Concurrently, tracking and alignment calibrations are performed using the hRaw data.
- When the required payloads are available, the raw data are partially processed to add tracking information and produce the cDST (calibration DST) samples, which include the raw data and tracking objects. This reduces the amount of time required to run calibration algorithms, but retains the raw data needed for calibration.
- All other calibrations use the cDST files as input to produce the necessary payloads.

The cDST samples produced during prompt calibration at BNL are replicated to KEK for subsequent recalibration campaigns.

Raw data prompt processing of the HLT hadron sample begins immediately after the prompt calibration is finished in order to provide samples for analysis as quickly as possible. Most hadronic analyses can use the HLT hadron sample, Prompt raw data processing of the full dataset begins immediately thereafter.

Year	2025	2026	2027	2028	2029
Number of raw data reprocessings	0	0.5	0.5	0.5	0.5
Number of MC reproduction campaigns	0	0.5	0.5	0.5	0.5

Table 1: Number of raw data reprocessing and of MC Campaigns per year. From 2026, we plan on a full data reprocessing and a full MC data reproduction every 2 years.

Analysis skimming begins as soon as the data is ready, with the highest priority skims (systematics) produced concurrently with data processing. Systematics skims consist of decay modes that are used to determine correction factors and are otherwise used for physics performance studies that produce required input for physics results.

Run-independent MC samples, using simulated backgrounds and static conditions, are produced as soon as the major software release is ready and beam background overlay files are prepared and distributed on the grid. Run-dependent MC samples, using random trigger background overlay files and real conditions, are produced on the grid soon after random trigger overlay files are prepared from the HLT delayed-bhabha-trigger hRaw files. Skimming of MC samples begins as soon as the samples are available on the grid.

During early data taking, when significant improvements to software and detector configurations were expected, major reprocessing campaigns happened each year. As the software has become more stable, staring in 2026 the reprocessing is performed once every two years. The planned number of reprocessing campaigns per year are listed in Table 1. The reprocessing of detector data with new software releases will trigger the reproduction of corresponding MC samples, to have consistent data sets, and the re-skimming of detector data and MC samples.

A new major software release is typically tagged in July and tested in August and September to be ready for the prompt processing of the new data collected from the fall and to recalibrate and reprocess the data previously collected. The reprocessing efforts should start late in the calendar year, while prompt processing efforts begin as new data are collected. Since no improvements were planned for the software in 2024, the next reprocessing campaign will begin in 2026.

Only the data for which improvements are expected from new calibration algorithms or for which problems were discovered with prior calibrations will be recalibrated. Otherwise, the existing payloads will be reused when the data is reprocessed with a new release. Data are made available as soon as they are reprocessed, meaning that the reprocessed data set does not become available all at once upon completion.

In parallel with the reprocessing, prompt calibration and processing is performed using the same major release. In this way, the datasets available in the summer of each year are processed with the same major release. This allows for consistent use of performance corrections and systematic studies.

4 The schedule of SuperKEKB

The planning of computing activities depends on the SuperKEKB schedule. The integrated luminosity collected in 2018 Phase 2 data taking was 0.5 fb^{-1} and that delivered by SuperKEKB in Phase 3 since 2019 is shown in Table 2. The Long Shutdown 1 (LS1) for the installation of the PXD began in June 2022 and continued through summer 2024. The data taking period prior to LS1 is referred to as "run 1" while that collected since is called "run 2." The winter/spring 2025 run period was canceled due to various factors. Data taking will resume early in calendar year 2026.

The SuperKEKB luminosity profile for years 2026 to 2029 is presented in Table 2. Note that this shows the delivered luminosity, while the recorded luminosity used for the resource estimate assumes a data taking efficiency of 88%, based on phase 3 data taking. Tables 3 and 5 show, for comparison, the luminosity profiles that were used for the computing resource estimate in February 2024 and February 2022, respectively. The current luminosity profile is based on the run 2 experience with additional assumptions on the evolution of the machine parameters and running time.

JFY	2019	2020	2021	2022	2023	2024	2025	2026	2027	2028	2029
Luminosity (ab^{-1}/year)	0.02	0.1	0.23	0.14	0.02	0.14	0.31	0.75	1.06	1.32	1.58
Integrated Luminosity (ab^{-1})	0.02	0.12	0.35	0.49	0.51	0.65	0.96	1.71	2.77	4.09	5.67

Table 2: The SuperKEKB (delivered) luminosity profile: February 2025. Delivered luminosities are cited through JFY 2024, while projections are made through JFY 2029.

Calendar Year	2019	2020	2021	2022	2023	2024	2025	2026	2027	2028
Luminosity (ab^{-1}/year)	0.01	0.09	0.2	0.19	0.06	0.35	0.5	0.7	1.2	1.2
Integrated Luminosity (ab^{-1})	0.01	0.1	0.3	0.49	0.55	0.9	1.4	2.1	3.3	4.5

Table 3: The SuperKEKB (delivered) luminosity profile: February 2024. Delivered luminosities are cited through calendar 2022, while projections are made through calendar 2028.

Calendar Year	2019	2020	2021	2022	2023	2024	2025	2026	2027
Luminosity (ab-1/year)	0.01	0.09	0.2	0.19	0.06	0.35	0.5	0.7	1.2
Integrated Luminosity (ab-1)	0.01	0.1	0.3	0.49	0.55	0.9	1.4	2.1	3.3

Table 4: The SuperKEKB Luminosity profile: February 2023. Delivered luminosities are cited through calendar 2021, while projections are made through calendar 2027.

5 Size of the MC samples

The B factories experience has demonstrated that it is important to have a MC sample that is many times the detector data sample. MC studies of Belle II performance have shown that the MC sample for hadronic and τ events should be at least 4 times the detector data sample up to an integrated luminosity of 1 ab^{-1} to avoid cases in which the dominant systematic uncertainty is the one coming from limited MC statistics.

The expectation is that this requirement may be reduced to twice the detector data sample up to an integrated luminosity of 5 ab^{-1} and to an equivalent size of the detector data sample for higher integrated luminosities. It is possible that significantly larger generator-level skimmed samples may be necessary to produce analysis-specific generic

Calendar Year	2019	2020	2021	2022	2023	2024	2025	2026
Luminosity (ab-1/year)	0.01	0.09	0.2	0.18	0.19	0.44	0.55	0.69
Integrated Luminosity (ab-1)	0.01	0.1	0.3	0.48	0.67	1.11	1.66	2.35

Table 5: The SuperKEKB (delivered) luminosity profile: February 2022. Delivered luminosities are cited through calendar 2020, while projections are made through calendar 2026.

Class of events	2025	2026 - 2027	2028 - 2029
hadronic	4	4	4
τ	4	4	4
$\mu^+\mu^-(\gamma)$	4	4	4
$\gamma\gamma(\gamma)$	2	2	2
Bhabha	0.1	0.1	0.1
$e^+e^-e^+e^-$, $e^+e^-\mu^+\mu^-$	1	1	1
Other e^+e^-X	1	1	1

Table 6: Ratios of the sizes of the run-dependent generic MC event samples to the detector event sample per year. Data reprocessing and MC reproduction will be performed in two-year periods starting in 2026. No reprocessing is expected in 2025, though prompt data will be processed and the corresponding run-dependent MC produced as usual.

MC samples that are manageable in size. The amount of MC for low multiplicity events can be smaller without significantly impacting the quality of the physics results. Both run-dependent and run-independent MC samples will be produced.

Run-dependent MC samples match the detector conditions on the basis of each run and use background events from delayed bhabha triggers, collected during the data taking, overlaid on the generated events to simulate the effect of machine background. The size of the run-dependent generic MC samples for each event type relative to the size of the detector data sample is summarized in Table 6. Run-independent MC samples are generated using average detector conditions and simulated background events. Table 7 shows the luminosity corresponding to the size of the run-independent MC samples that are planned. In addition, samples of signal events for specific studies will be centrally produced upon request of the physics analysis groups.

6 Inputs to resource estimates

The following details are used to determine the computing resources required for Belle II.

6.1 Accepted cross section

Table 8 shows the accepted cross sections for different classes of events. We plan to accept all of the hadronic and τ events, while the other low multiplicity events will be pre-scaled. Studies done on Belle II skimming code show that more than 90% of hadronic and τ events will be selected by at least one skim.

Class of events	2025	2026	2027	2028	2029
hadronic	0.25	0.25	0.25	0.25	0.25
τ	0.25	0.25	0.25	0.25	0.25
$\mu^+\mu^-(\gamma)$	0.25	0.25	0.25	0.25	0.25
$\gamma\gamma(\gamma)$	0.3	0.3	0.3	0.3	0.3
Bhabha	0.04	0.04	0.04	0.04	0.04
$e^+e^-e^+e^-$, $e^+e^-\mu^+\mu^-$	0.2	0.2	0.2	0.2	0.2
Other e^+e^-X	0.2	0.2	0.2	0.2	0.2
Non-Υ(4S)	0.5	0.5	0.5	0.5	0.5

Table 7: Equivalent luminosity (ab^{-1}) of the run-independent generic MC event samples, which are only produced when a new software release is created.

The low multiplicity events will be used to extract physics results and for detector and calibration studies. An open trigger for low multiplicity events increases the sensitivity for new physics searches like searches for dark matter candidates. On the other hand, the increase in luminosity will require a more restricted trigger and a smaller accepted cross section for low multiplicity events. The optimization will be done during data taking, considering the peak luminosity and the machine background levels.

Table 8 shows the cross section accepted by the current HLT menu. Reductions in low multiplicity event rates are expected for future running periods, but are not factored into the current model. In particular, efforts are ongoing to develop better Bhabha vetoes for the HLT. This should result in a reduction for the accepted Bhabha cross section starting in 2026. Data taking during 2024 showed some disagreements in the expected cross-section, for example for Bhabha events, but the total amount was close to the expected 29 nb. Studies of run-dependent MC samples with different background conditions suggest an uncertainty of 7% on the predictions for the accepted cross sections. The level 1 trigger pass-through, which is used for random trigger overlays and trigger diagnostics was not accounted in the previous accepted cross section, but has been accounted starting in 2025.

6.2 Simulated cross sections

MC events are generated according to the cross sections modeled in the generators. To optimize the usage of computing resources, a set of preliminary cuts is applied to the generated events before passing them to the Belle II detector simulation and reconstruction. The goal of these preliminary cuts is to remove, immediately after generation, those events that will be discarded later by the reconstruction and analysis software. Table 9 shows the cross sections of the different processes after applying the preliminary cuts that were utilized in the MC15 campaign, which was completed in JFY 2023. The cross sections in Table 9 will be used to estimate the computing resources required for MC production during and after JFY 2025.

Class of events	2025	2026	2027	2028	2029
$\Upsilon(4S)$	1.05	1.05	1.05	1.05	1.05
$c\bar{c}$	1.3	1.3	1.3	1.3	1.3
uds	2.12	2.12	2.12	2.12	2.12
$\tau^+\tau^-$	0.8	0.8	0.8	0.8	0.8
$\mu^+\mu^-(\gamma)$	0.83	0.83	0.83	0.83	0.83
$\gamma\gamma(\gamma)$	1.96	1.96	1.96	1.96	1.96
Bhabha	7.31	7.31	7.31	7.31	7.31
$e^+e^-e^+e^-$	0.18	0.18	0.18	0.18	0.18
$e^+e^-\mu^+\mu^-$	0.32	0.32	0.32	0.32	0.32
Other e^+e^-X	0.34	0.34	0.34	0.34	0.34
Other low multiplicity processes	11.7	11.7	11.7	11.7	11.7
L1 passthrough	1.1	1.1	1.1	1.1	1.1
Sum (all)	29.01	29.01	29.01	29.01	29.01

Table 8: Accepted cross sections (nb) for different classes of events. The accepted cross section for 2026 and later is the proposed one, subject to potential revisions.

Process	cross section (nb)
$\Upsilon(4S)$	1.05
$c\bar{c}$	1.3
uds	2.42
$\tau^+\tau^-$	0.92
$\mu^+\mu^-(\gamma)$	1.15
$\gamma\gamma(\gamma)$	3.52
Bhabha	74.5
$e^+e^-e^+e^-$	5.88
$e^+e^-\mu^+\mu^-$	5.87
Other e^+e^-X	2.2
Sum (all)	98.81
Non-$\Upsilon(4S)$	5.8

Table 9: Generated cross sections (nb) for different classes of events. Preliminary cuts at generator level are included.

Site	2025	2026	2027	2028	2029
BNL	0.3	0.3	0.3	0.3	0.3
Italy	0.2	0.2	0.2	0.2	0.2
Germany	0.2	0.2	0.2	0.2	0.2
France	0.15	0.15	0.15	0.15	0.15
Canada	0.15	0.15	0.15	0.15	0.15
hRaw Data (one copy)	1.1	1.5	2.1	2.9	3.9
hRaw Data (including replicas)	2.1	3	4.3	5.8	7.7
Raw Data (one copy)	3.6	5.2	7.3	10.1	13.3
Raw Data (including replicas)	7.3	10.3	14.7	20.1	26.6

Table 10: Sharing of the copy of raw data and size of the raw data sample (PB $= 1 \times 10^{15}$ bytes).

6.3 Raw data size

The raw data are transferred from the online to the offline system and permanently stored on mass storage. The size of one event in ROOT format was measured in run 1 data taking, with incomplete PXD, to be around 60 kB/event. Increased luminosity can possibly result in an increase of the machine background despite the foreseen mitigation actions. Higher background levels result in higher occupancy in the detector and larger raw data size.

To take into account the impact of the full PXD installation and of the uncertainties in the level of machine background, we add 20 kB to the raw data size, resulting in an estimated size of 80 kB/event. We assign a 25% uncertainty on this value and take 80±20 kB/event as the raw data size to estimate the amount of offline storage needed for raw data.

In addition to the raw data corresponding to the all event sample, a subset of events, selected by specific HLT lines, are also stored in raw data format. These data are called hRaw data and account for roughly the 30% of the size of the total raw data sample. They include the raw data needed for prompt calibration, hadronic events that are processed with higher priority, and the delayed Bhabha trigger events used to create the background files for run dependent MC simulation.

One copy of the raw data will be permanently stored at KEK. To avoid catastrophic loss of data, a second copy is stored elsewhere, now shared among different countries according to Table 10. The funding agencies of Canada, France, Germany, Italy, and the US have accepted the sharing according to the top-half of Table 10.

6.4 Calibration step

A subset of events in raw data format, selected according to the requirements of the calibration algorithms, is stored separately and is centrally processed to produce events in cDST format. The separately stored raw data are named "skimmed raw data" and are a subset of the hRaw data discussed previously. Events in cDST format consist of raw data and reconstructed tracking objects that are used to compute the calibration constants for

Year	2025	2026	2027	2028	2029
Storage for prompt calibration (PB)	0.2	0.2	0.2	0.2	0.2
CPU power for prompt calibration (HS23)	4	4	4	4	4
Total cDST data (PB)	0.75	1	1.25	1.5	1.75
CPU power for recalibration (HS23)	10	10	10	10	10
MC cDST (PB)	0.5	0.6	0.7	0.7	0.7

Table 11: Parameters of the calibration model

the various Belle II subdetectors. The skimmed raw data samples are deleted after cDST production.

The prompt calibration is performed using a dataset with an equivalent of 12 fb^{-1}, sampled from the data taken during a period with consistent conditions. These calibration periods are referred to as "buckets" and currently correspond to two weeks of data taking. As the instantaneous luminosity increases, we expect to increase the duration of a bucket to three weeks. Apart from such a change, the amount of computing resources required for cDST production and prompt calibration does not scale with luminosity. They are estimated to be 0.2 PB of storage and 4 kHEPScore23 of CPU power.

The cDST sample is stored and will be used in the recalibration step. The total size of the cDST sample scales with the number of weeks of data taking, but does not depend on luminosity. It is estimated to be 0.5 PB/year. The prompt calibration is done at BNL and, through JFY 2023, recalibration was performed at DESY. Starting in calendar year 2024, the recalibration is instead handled at KEK. In addition, a sample of MC events in cDST format is produced for calibration and detector performance studies. The parameters of the calibration model are shown in Table 11.

An internal review of the Belle II calibration scheme was completed in 2023[1]. Several options for recalibration were considered, as well as the estimated resource needs, projected into future data taking periods. The conclusion of the review was that recalibration at a single site is preferable, at least through 2028.

6.5 Processing Power and buffer space for raw data processing

The CPU power required for processing one event and writing the reconstructed event in mDST format was measured with software release 06-00-08 for different classes of simulated events and different background levels. The software group prepared a note[2] that describes the procedure in detail, including a breakdown of the processing power used in the different framework modules. The average processing power, weighted by the accepted cross sections, is 27.4 HEPScore23 * s / event for the reconstruction of an event simulated with nominal background. This value is 10% lower than the corresponding value of 29.8 HEPScore23 * s / event measured with release 05-01-08 in February 2021.

The processing power necessary to reconstruct a detector event with releases 7 and 8, based on preliminary measurements, are given in Table 12. The resource estimate is based on the processing power for release 8.

---

[1] BELLE2-NOTE-TE-2023-013: `https://docs.belle2.org/record/3579`

[2] BELLE2-NOTE-TE-2016-011: `https://docs.belle2.org/record/371`

Class of events	Release 5	Release 6	Release 7	Release 8
Υ(4S)	64.5	64.8	62	64.8
$c\bar{c}$	53.3	55.4	51.8	57.4
uds	46.8	42.6	40.2	45.9
$\tau^+\tau^-$	36.2	28.2	27.5	38.4
$\mu^+\mu^-(\gamma)$	26.1	23.4	23.4	30.4
$\gamma\gamma(\gamma)$	26.1	23.4	23.4	30.4
$e^+e^-(\gamma)$	26.1	23.4	23.4	30.4
$e^+e^-e^+e^-$	26.1	23.4	23.4	30.4
$e^+e^-\mu^+\mu^-$	26.1	23.4	23.4	30.4
e^+e^-X	26.1	23.4	23.4	30.4
Average	29.8	27.4	27.6	33
Scale factor for software upgrade	1.05	1	1	1
Average including software upgrade	31.3	27.4	27.6	33
Scale factor for background uncertainty	1.1	1.1	1.1	1.1
Average with all scale factors	35 ± 2	30 ± 2	30 ± 2	36 ± 2

Table 12: Processing power (HEPScore23 * s / event) needed to fully reconstruct a detector event.

The experience gained during run 1 data taking has significantly improved the understanding of the machine background and many mitigation actions have been implemented. Increasing the peak luminosity of SuperKEKB is still expected to increase the machine background levels in the Belle II detector and to take into account this possibility we include a 10% increase of the CPU power and quote half of the corresponding increase as background uncertainty. The errors from the foreseen software improvements and from the background uncertainty are summed in quadrature.

The estimated CPU power required to fully reconstruct a detector event is 36 ± 2 HEPScore23 * s / event (see Table 12).

The calibration scheme requires a set of events equivalent to an integrated luminosity of 9 fb^{-1} for each calibration period. The calibration and reconstruction is performed with this data set within eight weeks, which requires a buffer space equivalent to twice the size of the raw data collected in eight weeks of data taking (see Table 17), since two replicas are stored for raw data. Buffer space is also required for raw data staging from tape during reprocessing. Based on the average size of data samples used during recent reprocessing campaigns, equivalent to about 50 fb^{-1}, a buffer space of 500 TB is necessary for this purpose. The total buffer space for raw data processing is accounted in Table 17 and in the appendix.

Additional buffer space is needed for intermediate files produced during mDST production, which cannot be removed until the final merge step is completed and verified. The buffer space for intermediate files is estimated to require approximately 30% of the final mDST disk space.

In principle, it is instead possible to store the raw data on tape and stage them from tape to disk. However, some disk storage will be necessary anyway and additional tape drives will be needed. The savings will be marginal and the operations will be more

Class of events	Release 5	Release 6	Release 7	Release 8
$\Upsilon(4S)$	85	93.6	89.5	93
$c\bar{c}$	71.1	83.1	77.7	86.7
uds	62.5	65.7	62.1	74.3
$\tau^+\tau^-$	45	40.5	39.3	43.2
$\mu^+\mu^-(\gamma)$	35.1	29.9	29.9	30.9
$\gamma\gamma(\gamma)$	35.1	29.9	29.9	30.9
$e^+e^-(\gamma)$	35.1	29.9	29.9	30.9
$e^+e^-e^+e^-$	35.1	29.9	29.9	30.9
$e^+e^-\mu^+\mu^-$	35.1	29.9	29.9	30.9
e^+e^-X	35.1	29.9	29.9	30.9
Average	48.6	47.2	42.5	49.7
Scale factor for software upgrade	1.05	1	1	1
Average including software upgrade	52.9	48.8	42.5	49.7
Scale factor for background uncertainty	1.1	1.1	1.1	1.1
Average with all scale factors	58 ± 4	53 ± 4	47 ± 3	55 ± 3

Table 13: Processing power (in HEPScore23 * sec /event) for Monte Carlo production.

complicated and reduce the speed at which prompt data become available.

6.6 Processing power and buffer storage for Monte Carlo production

The processing power to perform the full simulation and reconstruction of one event and write the event in mDST format has been measured on software release 06-00-08 for different classes of events and background levels. The processing power necessary for simulation and reconstruction with releases 7 and 8, based on estimates, are given in Table 13. The resource estimate is based on the expected processing power for release 8. The average processing power, weighted by the accepted cross sections, is 55 ± 3 HEPScore23 * s / event).

To take into account the possibility of higher backgrounds when increasing the peak luminosity, we include a 10% increase of the CPU power and quote half of the corresponding increase as background uncertainty.

The CPU power required to fully simulate and reconstruct one event is 48 ± 3 HEPScore23 * s / event (see Table 13) where the errors from the foreseen software improvements and from the background uncertainty are summed in quadrature.

The event generation and event reconstruction steps of MC production are performed in a single job, without writing a raw-data-like intermediate output. In principle, writing the intermediate raw data after the generation step will allow to perform only the reconstruction step when reproducing the MC samples. However the saving will be marginal at best, because the reconstruction step requires on the order of 70% of the total CPU power and large storage for intermediate data will be required. In addition, in that case, it will be possible to run simulation only at sites with large storage resources.

It is worth noting that the statement that the reconstruction step requires on the

Class of events	Release 5	Release 6	Release 7	Release 8
Y(4S)	17.8	14.9	15	15.2
ccbar	16.3	13.4	13.5	13.6
uds	15.3	12.4	12.5	12.7
t+t-	13.7	10.2	10.4	10.8
m+m-(g)	12.9	9.3	9.5	10.3
gg(g)	12.9	9.3	9.5	10.3
e+e-(g)	12.9	9.3	9.5	10.3
e+e-e+e-	12.9	9.3	9.5	10.3
e+e-m+m-	12.9	9.3	9.5	10.3
e+e-X	12.9	9.3	9.5	10.3
Average	13.4	9.9	10.1	10.8
Scale factor for software upgrade	1.05	1.05	1	1
Average including software upgrade	14	10.4	10.1	10.8
Scale factor for background uncertainty	1.02	1.02	1.02	1.02
Average with all scale factors	14.3 ± 1.3	10.6 ± 1.1	11.0 ± 1.1	11.0 ± 1.1

Table 14: Sizes of detector events in mDST format (kB/event).

order of 70% of the total CPU power, is only apparently contradicted by using the values of 55 HEPScore23 * s / event for MC production and 36 HEPScore23 * s / event for detector data processing. These values are averages over different classes of events and the detector data sample has a much larger fraction of low multiplicity events, which are less CPU intensive then hadronic and τ events, than the MC sample.

The MC production results in a set of mDST files that are too small for convenient handling. These intermediate files are merged, in a multi-step process, to create the final mDST files. Thus, the MC production generates an amount of intermediate files that is typically twice the size of the final mDST data-set. These intermediate files are deleted in bulk during the MC production, but some buffer space, estimated to be of the order of 30% of the final MC mDST size is needed (see Table 17).

In addition, files with random trigger events collected during the data taking to perform background overlay need to be stored. We estimate the need to be 60 GB/fb^{-1} (see Table 17).

6.7 Event size in mDST format

The mDST event sizes have been measured on software release 06-00-08 for different classes of events and different background level. The software note[3] describes the procedure in detail, including a breakdown of the sizes of different branches. An estimate for the event sizes produced with releases 7 and 8, based on estimates, are given in Tables 14 and 15 for data and MC events, respectively.

To take into account the possibility of having a higher background when increasing the peak luminosity, we include a 2% increase of the event size and quote half of the corresponding increase as background uncertainty. The errors from the foreseen software

---

[3] BELLE2-NOTE-TE-2023-013: `https://docs.belle2.org/record/3579`

Class of events	Release 5	Release 6	Release 7	Release 8
Y(4S)	20.9	18.4	18.5	18.7
ccbar	18.7	16.1	16.2	16.3
uds	17	14.4	14.5	14.7
t+t-	14.7	11.3	11.4	12
m+m-(g)	13.2	9.6	9.8	10.6
gg(g)	13.2	9.6	9.8	10.6
e+e-(g)	13.2	9.6	9.8	10.6
e+e-e+e-	13.2	9.6	9.8	10.6
e+e-m+m-	13.2	9.6	9.8	10.6
e+e-X	13.2	9.6	9.8	10.6
Average	15.5	12.1	11.9	12
Scale factor for software upgrade	1.02	1.05	1	1
Average including software upgrade	15.8	12.7	11.9	12
Scale factor for background uncertainty	1.02	1.02	1.02	1.02
Average with all scale factors	16.1 ± 1.4	13 ± 1.2	12.1 ± 1.2	12.2 ± 1.2

Table 15: Sizes of MC events in mDST format (kB/event).

	Release 5	Release 6	Release 7	Release 8
Average processing power for skimming data (HS23*s/evt)	2.9	4.4	4.6	4.5
skim detector event size (kB)	16.9	12.8	12.6	12
(Events in skim)/(events in mDST) data	1.2	1.2	1.2	1.2
Average processing power for skimming MC (HS23*s/evt)	4.6	5.9	5.6	5.4
skim MC event size (kB)	21.3	17	15.1	14.5
(Events in skim)/(events in mDST) MC	1.2	1.2	1.2	1.2

Table 16: Parameters of the skimming model for Detector and MC events.

improvements and from the background uncertainty are summed in quadrature. Finally, we get 11.0 ± 1.1 kB/event and 12.2±1.2 kB/event for detector and MC events, respectively. Tables 14 and 15 show the sizes of detector and MC events in mDST format.

6.8 Skimming

The skimming software has continually evolved as new skims have been added and some of the skims that selected largely overlapping samples have become single skims. Different skim output files can be created from a single job. Overall there is an optimization process ongoing that will continue in the future as experience is gained. The skimmed samples use mDST format, since reproducing the analysis-related objects rather than storing them results in a marginal increase in processing time while reducing the size of the skimmed sample. The objects produced from the FEI algorithm, however, are much more time intensive to produce, so they are added to the mDST information for the FEI skim samples.

The parameters of the skimming model are summarized in Table 16. The detector mDST data sample will include a larger fraction of low multiplicity events then the MC mDST data sample. The effect of these different sample compositions is included in the

	2025	2026	2027	2028	2029
Disk for raw data processing	1.01	1.73	2.24	2.66	3.09
Disk for data mDST in production	0.04	0.11	0.24	0.22	0.46
Disk for MC in production	0.14	0.28	0.58	0.57	1.14
Disk for random trigger files	0.15	0.27	0.44	0.65	0.9
Disk for skimming in production	0.28	0.68	1.43	1.39	2.81

Table 17: Buffer spaces (in PB) for raw data processing, MC production and skimming. Space for random trigger files for MC production.

estimate of the processing power and the skim sample size in Table 16.

The skimming produces a set of files that are too small for convenient handling. These intermediate files are merged to create the final skimmed files. Thus, the skimming generates an amount of intermediate files that is equivalent in size to the final skimmed dataset. These intermediate files are deleted in bulk during the skimming, but some buffer space, estimated to be on the order of 30% of the final skimmed data-set size is needed (see Table 17).

6.9 mDST and skim replica

We know that some level of replication of mDST and skimmed data will be required for safety and to have reasonable performance of the analysis jobs. For the time being we assume that two replicas will be sufficient.

Until the production with release X is completed, we plan to keep on disk the replicas of release X-1 and X-2. When production with release X is completed, the replicas of release X-2 will be removed from disk. Since replication is handled automatically by Rucio, the space required for two replicas of files with release X is included but no additional buffer for files in production. Additional details about the mDST and skimmed sizes are accounted in Table 18, for reference. Previous estimates assumed that prompt samples would be produced and removed on an annual basis. However, with the reprocessing happening once every two years, the prompt samples that are not part of the reprocessing must be retained.

A key concept in this context is the ”popularity” of the data that helps in optimizing the usage of storage. Data sets that are rarely accessed do not need to be replicated and data that are not accessed at all can be removed from disk storage. As soon as the appropriate tools will become available, we plan to move to a replica policy based on data popularity.

6.10 Analysis

User analysis jobs should be run on skimmed samples to produce ntuples that will be used for offline analysis and systematic studies. The uncertainties associated with analysis jobs are much larger and therefore more difficult to estimate in advance.

The amount of CPU power used for analysis in 2022 and 2023 was significantly larger than originally estimated from the resource estimates. The expected number of concurrent

	2025	2026	2027	2028	2029
Size of prompt data mDST (latest version)	0.24	0.39	0.95	0.69	1.52
Size of reprocessed data mDST (latest version)	0	0.31	0.62	0.79	1.57
Size of prompt MC mDST (latest version)	0.92	1.26	2.62	1.99	4.03
Size of reprocessed MC mDST (latest version)	0	0.62	1.23	1.78	3.56
Size of prompt data mDST (all versions)	0.34	0.73	1.29	1.88	2.71
Size of reprocessed data mDST (all versions)	0.75	1.06	1.38	1.79	2.58
Size of prompt MC mDST (all versions)	3.34	3.02	4.38	4.96	6.99
Size of reprocessed MC mDST (all versions)	1.31	1.93	2.55	3.86	5.65
Size of skimmed data (latest version)	0.32	0.94	2.09	1.97	4.12
Size of skimmed MC (latest version)	1.57	3.62	7.44	7.27	14.63
Size of skimmed data (all versions)	1.45	2.39	3.55	4.9	7.04
Size of skimmed MC (all versions)	7.93	9.54	13.36	17.01	24.36

Table 18: Size of mDST and skimmed samples (all replicas), in PB, for data and MC samples, both for the latest version (X) and for all stored versions (X, X-1, and X-2).

analyses was underestimated. The parameters of the model for the analysis step were adjusted to better reproduce the observed amount of CPU power. Thus far in 2024, the amount of CPU used for analysis expectations is about 130 HS23, which is in good agreement with the expectation of 136 HS23.

The values used in the resource estimate are:

- The average CPU power to run one analysis on one event is 0.05 HEPScore23 * s / event.
- The number of concurrent analyses in 2023 was about 300. We expect this to remain fairly stable for 2025 and later.
- The number of analysis cycles (defined as the number of times an analysis is run on the data and MC samples) is set to 2 and is expected to remain constant. Even as analysis experience grows, analyses that become limited by systematics may require subsequent processing as each analysis evolves.
- The uncertainty is at the 50% level.

The total processing power to analyze one event (all analyses, all analysis cycles) is shown in Table 19 together with the parameters of the analysis model. The user ntuple size scales with the number of concurrent analyses and of analysis cycles. We assume that running one cycle of one analysis on the skimmed data sample will produce on average 0.02 ± 0.01 kB/event of ntuple data.

6.11 Cosmic Data

We conservatively assume that the processing and reprocessing of the cosmic data will require a CPU power of 0.5 kHEPScore23 for 10 months assuming a trigger rate of 100 Hz for 3 months of cosmic data taking. In addition, we assume that the production of the corresponding sample of MC cosmic events will require 1 kHEPScore23 for 10 months and that another 1.5 kHEPScore23 will be needed for cosmic data analysis.

	2026	2027	2028	2029
CPU power to run one analysis (HS23*s/evt)	0.05	0.05	0.05	0.05
Number of concurrent analyses	300	300	300	300
Number of analysis cycles	2	2	2	2
CPU power to run all analyses (HS23*s/evt)	30	30	30	30
ntuple size for all analyses (kB/evt)	12	12	12	12

Table 19: Parameters of the analysis model.

	2026	2027	2028	2029
Total tape (PB)	13.5	19.1	26.1	34.4
Total disk (PB)	24.9	35.8	44.4	63.8
Total CPU (kHS23)	369	507	675	850

Table 20: The resource estimate for years 2026 to 2029. The storage units are PB = 10^{15} Bytes.

6.12 Efficiency

We assume to be able to run the code for 10 months every year because time will be lost due to site downtimes, operational mistakes, software bugs discovered only after the start of the data processing that will require to restart it from scratch, etc. This corresponds to an efficiency in the use of CPU power of 83% that compares well with the 80% efficiency experienced by the different LHC experiments at the beginning of data taking.

7 The Resource Estimate

The input parameters that have been discussed in the previous section are used to obtain the resource estimate summarized in Table 20. CPU and storage for processing and reprocessing of raw data will be shared among the countries hosting raw data, according to the fraction of raw data hosted. CPU and disk storage for MC production and analysis, and storage for mDST and uDST data will be shared among the different countries according to the PhD count. For comparison Table 21 shows the pledged and approved computing resources for 2019-2024. Approval by the Belle II Financial Oversite Panel occurs in October of the preceding JFY (for example, JFY 2025 resources were approved in October 2024). For this resource estimate, the storage units are PB = 10^{15} Bytes.

A detailed breakdown of the different activities and data samples for the years 2026 to 2029 is included in the Appendix A.

8 Conclusion

We have presented the estimate of the computing resources needed by the Belle II collaboration in the years 2026 to 2029 and described the planned activities and the assumptions that lead to this estimate. This estimate has been produced by the Computing Steering

	2019	2020	2021	2022	2023	2024	2025
Total tape (PB)	10	1.4	3.2	8.8	8.6	9.6	11.8
Total disk (PB)	11	11	11	16.5	19.6	25	19.4
Total CPU (kHS23)	299	207	498	385	404	520	247

Table 21: The resources approved and pledged for 2019 through 2025.

Group and approved by the Executive Board.

Appendix A: Resource Breakdown for years 2026 to 2029

Tape (PB)	2026	2027	2028	2029
Tape for raw data	10.34	14.69	20.1	26.58
Tape for hRaw data	3	4.26	5.83	7.71
Tape for cosmic data	0.15	0.15	0.15	0.15
Total tape	13.49	19.1	26.08	34.43

Table 22: Tape resources by data type (in PB).

Disk (PB)	2026	2027	2028	2029
Disk for data processing buffer	1.73	2.24	2.66	3.09
Disk for data mDST buffer	0.11	0.24	0.22	0.46
Disk for calibration	1.2	1.45	1.7	1.95
Disk for data mDST	1.79	2.66	3.67	5.29
Disk for random trigger files for MC	0.27	0.44	0.65	0.9
Disk for MC production buffer	0.28	0.58	0.57	1.14
Disk for MC mDST buffer	0	0	0	0
Disk for MC cDST	0.6	0.7	0.7	0.7
Disk for MC mDST	4.95	6.93	8.82	12.64
Disk for data and MC uDST	11.93	16.9	21.9	31.41
Disk for data and MC skimming buffer	0.68	1.43	1.39	2.81
Disk for data and MC uDST buffer	0	0	0	0
Disk for user ntuples	1.34	2.24	2.12	3.45
Disk for cosmic	0.01	0.01	0.01	0.01
Total disk	24.9	35.81	44.42	63.85

Table 23: Disk resources by data type (in PB).

CPU (HS23)	2026	2027	2028	2029
CPU for data processing	26.6	37.5	46.7	56
CPU for data reprocessing	20.3	20.3	51.2	51.2
CPU for calibration	14	14	14	14
CPU for MC production	151.4	183.6	304.5	331.5
CPU for skimming	22	26.7	44.6	48.6
CPU for analysis	134.9	225.2	213.7	348.3
Total CPU	369	507	675	850

Table 24: CPU resources by usage (in HS23).