B-factory Programme Advisory Committee
Full report for
the review meeting on the Phase 3 operation

27–28 October 2019 at KEK

A. Andreazza* (Milano), D. Cassel (Cornell), P. Collins&* (CERN), G. Corti (CERN),
S. Gori (UCSC), W. Hulsbergen* (NIKHEF),
M. Ishino (Tokyo), Z. Ligeti (LBL), F. Meijers (CERN)*,
N. Neufeld& (CERN), B. Ratcliff* (SLAC),
M. Sullivan (SLAC), H. Tajima (Nagoya), M. M. Titov* (Saclay)
and chaired by T. Nakada (EPFL)
* Expert members, & Remotely participated in part

17 January 2020

1 Short Summary

The SuperKEKB machine and Belle II detector are now entering the second phase of commissioning, where the machine should achieve luminosities well above those provided by the KEKB machine with a background condition that allows the Belle II detector to perform safe data taking. Similarly, the Belle II experiment should become capable of fully exploiting the delivered luminosities for physics analysis.

The SuperKEKB machine team has shown that the machine could be operated with a luminosity well above 10^{33} cm^{-2}s^{-1} and a background condition acceptable for the Belle II to operate. The background level became too high when the machine operated above $\sim10^{34}$ cm^{-2}s^{-1}. In the meeting, it was reported that the newly installed collimators reduced the background as expected, which shows that the simulation is providing improved understanding of the machine behaviour, and the simulation effort should continue. A very small beam spot size at the interaction point, which is one of the key ingredients to achieve luminosities well above those of the KEKB with acceptable background levels, has not been fully demonstrated yet. The proposed working plan to understand the machine properties with $\beta^* = 2$ mm first, then to proceed reducing it to 1 mm and below is very appropriate. For optimising the amount of useful data for physics before the summer shutdown in 2020, giving a priority for the machine work during this year is sensible and strongly supported. The current effort on understanding the background sources and devising mitigation plans is essential in order to progress toward the design luminosity. The committee congratulates the successful collaboration between the SuperKEKB machine operation team and the Belle II background study group and encourages their collaboration to continue and be strengthened.

It was reported that there were many repairs and improvements in the machine during the summer shutdown. The committee is pleased to hear that now a fast beam abort signal can be generated by the power supply of the superconducting final focusing quadrupole magnets (QCS). This clearly improves the protection of the Belle II vertex detector system. Introduction of further protection systems such as a very fast beam abort system is recommended. The committee is pleased to learn that the power supply of the Pixel Detector (PXD) is being modified to allow fast ramping down for protection from beam background.

The goal of the Belle II collaboration to collect 200 fb^{-1} of data by the time of the summer conferences in 2020 appears to be extremely ambitious. However, the committee fully encourages the collaboration to make every effort for this goal, since this would allow the collaboration to start providing physics results comparable to the Belle experiment in the core physics programme.

The overall performance of the Belle II detector is good. Although the second layer of the PXD is not fully installed, there should be no real degradation in physics performance with the currently expected background level. However, the PXD readout in a gated mode has not been fully demonstrated and remains a concern. Construction of new PXD ladders, needed for the new PXD with two full layers, is progressing well with a plan to complete the production by the summer 2020, if the current high yield is maintained. Assembly of the new PXD integrated with the rest of the vertex detector will require expertise and skill. The collaboration should ensure that the necessary personnel will be retained for the work as well as those needed for the operation and maintenance of the system. The dark current issue of the Central Drift Chamber (CDC) is under control for the moment, after various interventions. Test chamber studies for understanding the long term behaviour of the CDC should continue. In parallel with the careful monitoring of the CDC and the test chamber, the Belle II management should work out a plan for a possible major repair. It could even become necessary to construct a new CDC. An algorithm to filter the cross-talk signals based on the TDC information has been validated in software. This must be implemented as firmware in the ASIC of the CDC frontend card for the Level-1 trigger to work efficiently at high luminosities. Acquisition of radiation hard photon detectors (MCP-PMT) for the barrel particle identification system (TOP), to replace those which are not expected to last long due to the beam background radiation, is in progress. Ensuring the radiation hardness of the delivered devices is essential. The $\mathrm{K_L}$-Muon detector (KLM) has made good progress by merging the barrel and endcap groups together. However, further effort is needed for long term stable operation of the device and fully exploiting its performance. The committee recommends the KLM group producing a work plan for this.

The currently functioning trigger system is sufficient for the experiment for the moment. However, several subsystem trigger components are missing. The committee would like to see a clear plan for the completion of the full trigger system. The trigger performance should be monitored in real time to detect quickly any problem with the detector or machine background that may arise.

In the Belle II operational system, the committee is pleased to see many routine actions have been automatised in data taking and processing. Further effort should be invested to make those automatised procedure more complete and reliable, which is essential for sustaining a long period of data taking. Both online and offline computing as well as the control system of the experiment are working well. Dedicated human resources given to the online and offline teams using the maintenance and operation fund of the collaboration are showing a very positive impact. The committee encourages continuing this practice to ensure smooth running of the experiment. For the offline computing resources, the committee encourages the collaboration to utilise resources in small computing centres more effectively by finding suitable roles appropriate for their capabilities. The committee would also like to congratulate the collaboration for making a decision on the choice of technology for the new readout card, which will replace the current COPPER board.

Finally, the committee appreciates the effort to explore physics results with the early Belle II data and encourages the collaboration continuing to do so. It is also important to plan for a transition from the Belle analysis work. The committee is looking forward to seeing progress in all aspects of the experiment during the next annual review meeting in February 2020.

2 Physics plan of Belle II

2.1 Status

The committee would like to congratulate the Belle II collaboration for the successful start of Phase 3 data taking, and for completing the first papers that utilise the 0.5 fb^{-1} Phase 2 data. In particular, the collaboration has finished a paper on the luminosity measurement, using both the $e^+e^- \to e^+e^-$ and $e^+e^- \to \gamma\gamma$ channels. The paper has been submitted to Chinese Physics C. Furthermore, a new search for an invisibly decaying Z' boson produced in association with $\mu^+\mu^-$ or $\mu^\pm e^\mp$ has been completed. The paper is close to submission. These channels were not previously investigated by any other experiment. One additional paper using Phase 2 data on the search for axion-like particles decaying into photons, $e^+e^- \to a\gamma,\ a \to \gamma\gamma$, will be completed soon.

The collaboration has collected 6.5 fb^{-1} of Phase 3 data in the spring 2019, with a peak luminosity around 5×10^{33} $\mathrm{cm}^{-1}\mathrm{s}^{-1}$ during the physics run. Of this data, 5.6 fb^{-1} was collected on the $\Upsilon(4S)$ resonance and 0.83 fb^{-1} off-resonance. These data have enabled the collaboration to measure the B meson oscillation frequency and lifetime, and to reconstruct several B meson decays, such as $B \to DK$, $B \to K\pi$, $B \to K^*\gamma$, $B^0 \to J/\psi K^{(*)}$. Studies of hadronic B decays are in good shape, and results for the semileptonic decays are expected by next summer. For these measurements, one could envision a combination of Belle and Belle II analyses. A large array of additional analyses are also in the pipeline and are planned for the winter of 2019–2020. In particular, new results on the measurement of the branching fraction for $B^0 \to D^*\ell\nu$ are expected soon. Furthermore, the analysis of just 340 pb^{-1} data might lead to a measurement of the D lifetime, with uncertainty close to the world average. This will be a clear demonstration for the capability of the Belle II vertex detector. Time-dependent CP violation measurements will also be possible with a working flavour tagger.

The collaboration has ambitious goals for the coming winter, and especially for the summer of 2020, and beyond. The goal is to collect about 200 fb^{-1} data by the ICHEP 2020 conference. The 200 fb^{-1} data would enable the collaboration to start setting new milestones in the core physics program, as well as achieving the most stringent bounds on many dark sector models. A data set of up to 1 ab^{-1} is envisioned by summer 2021, which would allow many measurements to go substantially beyond past Belle analyses. Several triggers are not ready for the autumn run, and some trigger and detector efficiencies are below expectations. In particular, the single photon trigger is under investigation.

2.2 Concerns

- Several triggers were reported not to be ready for the autumn run. It is important to ensure that they are working for the spring 2020 data taking.
- The performance of some detector components were reported to be below expectations or not well reproduced by simulation.

2.3 Recommendations

- The committee encourages the collaboration to make every effort to achieve the goal of collecting 200 fb^{-1} data by the 2020 summer conferences. This would allow starting to obtain physics results close to the ones of the Belle experiment in the core physics program. This is important to ensure competitiveness with the LHCb experiment and also essential for providing Ph.D. students and postdocs the best possible career opportunities.
- The committee also encourages systematically developing and pursuing a program to study low multiplicity events, that could lead to several world leading bounds on light dark sector particles with a relatively modest amount of data.

3 Machine related issues

3.1 Status

The committee was pleased to hear the report of the experience gained with the beginning of the 2019 run and appreciates how the accelerator team continues to make improvements in the operation of SuperKEKB. By the end of the initial 2019 runs, the accelerator team achieved a $\beta_y = 2$ mm, which is a new record for lepton colliders, and continuous injection was also commissioned. Following the July report by the Accelerator Review Committee (ARC), the KEK accelerator team and the Belle II collaboration developed a common plan aiming to achieve an integrated luminosity of 200 fb^{-1} by mid-2020. Significant progress has been reached in understanding the machine background, which has so far prevented Belle II from fully exploiting the highest luminosities.

A dispersion mismatch discovered in the LER energy compression system and the subsequent correction has led to a much better beam injection to LER that is closer to the design.

A significant effort has gone into reducing the time between an abort signal generation and the actual abort of the beams. The transportation time of the abort signal to the abort system has been shortened and there are now two abort gaps in the bunch pattern allowing for a quicker beam abort. In addition, the time for the generation of an abort signal from the QCS power supplies has been greatly reduced and the PXD power supply has been altered to speed up ramp down times. This will ensure that the beams are out of the machine before the magnetic fields in the final focus magnets start to significantly decay due to a quench event.

Damaged cables were found for the beam position monitors just inside of the cryostats. Rerouting these cables has eliminated future possible damages. This should improve the diagnostics for seeing where the beams are at the collision point.

Slightly damaged bellow-fingers at the QC1L HER position just inside of the cryostat have been replaced and the condition of the fingers was carefully watched during the reinsertion and connection of the cryostat to the central chamber.

The two collimator heads that had been damaged during beam loss events in June and July have been repaired. The faster beam abort should reduce the chances of collimator damage in the future.

The background studies have continued to shed light on various background sources. In addition, mitigation plans for these sources have been implemented and studies are underway. Currently, the largest background is coming from the LER beam-gas events. This background should slowly improve as the LER ring is scrubbed.

The background simulation qualitatively reproduces these features, but the amount of beam gas events seen in data is about a factor ten larger than the predictions. This may mean the vacuum pressure at the location of the beam is higher than measured. Also the Touschek background is observed to be similar for LER and HER, while the simulation expects the background in HER to be few order of magnitude smaller.

Additional collimators in the LER are being considered and a new collimator (D06V1) will be installed. This will have the same betatron phase as D02V1 and therefore should reduce the collimator loss rate on D02V1. D02V1 is the last vertical collimator before the IP and it was one of the damaged collimators in the spring run. D06V1 should also reduce the chances of damaging D02V1. For this reason, and because D06V1 is far from the IP, the collimator head might be made with a low Z material (C or Ti) to improve the robustness of the collimator head to intense beam losses.

There has been an unexpected increase in the HER Touschek background. The exact reason for this is still unknown but adjusting the tune of the HER has been proven to greatly reduce the overall HER background rate and this may be related. If the HER coupling is too small, then the HER vertical size is too small and the HER Touschek lifetime would become very short causing an increase in the background rates.

The background studies have also found that the LER background rate in the detector is sensitive to the LER orbit just upstream of the detector.

3.2 Concern

- Accelerator backgrounds continue to be high and efforts to take data at high luminosity and/or high beam currents has proven to be difficult.

3.3 Recommendations

- Although the reduction of the delay time between an abort signal and the actual beam abort is excellent news, efforts to further reduce this delay are strongly encouraged.
- Further studies of collimator placement and collimator robustness against beam instabilities should be carried out.
- This close collaboration between the background study group and accelerator team has been the most efficient way for the accelerator and the detector to reach the common goals of high luminosity and low backgrounds and should continue.

4 Issues related to the Belle II detector

4.1 Pixel Detector (PXD)

4.1.1 Overall status

The PXD group reported significant progress on understanding the performance of the current PXD and on the status of the effort to install a new PXD in 2021.

The PXD suffered temporary damage after radiation accidents due to "dust events." Test are planned in November at the Mainz Microtron, in order to reproduce the failure mode, which is likely to be due to the SWB chip. Both the affected modules were eventually recovered, as well as a broken optical link. The ROI selection could be tested only on about 1/2 of the modules due to the requirement on DHE ordering. ROI operation was verified by comparing the Belle II ONSEN output with local DAQ data on which ROI software emulation was applied, showing a perfect matching between ONSEN output and expectation from the ROI simulation.

Rate capability has been improved by new DHH firmware, distributing the data load from one DHH to four ONSEN. As a result, the PXD is not a bottleneck for running at 30 kHz trigger rate. DHE-DHP desynchronisations are observed in high multiplicity events, like the ones with high background due to HER injections. The problem is currently handled by a trigger veto, waiting for the commissioning of the gated mode in the autumn run.

The status report for the PXD replacement includes an updated schedule targeting the delivery of the new PXD at KEK by October 2020. That would allow the installation of the detector, now named PXD2021, to be performed in summer 2021. The target is to have 12 (+6 spare) ladders for the outer layer (L2) and 8 (+4 spare) ladders in the inner layer (L1).

A critical issue has been the ladder assembly procedure. A first review of the assembly procedure was held in February 2019, with a follow up review in July 2019. These resulted in a simplification of the procedures and improvements in the ladder design, with additional ceramic stiffeners. The cleanliness of the assembly laboratory has been increased and the four-eyes principle is now implemented on the whole assembly chain. Since then the ladder assembly yield has dramatically improved with 11 assembled ladders of which only one failed. A weekly schedule for ladder assembly has been developed and it is planned to be followed routinely. It includes testing and gluing two ladders per week, operating on batches of delivered tested modules. Currently there are ten ladders ready for L2 and 2 ladders ready for L1.

Modules to build the remaining ladders are being collected from the three sources: the leftovers of the previous PXD production, the recovery of modules with the SWBv2.2 switcher chip, and new sensor production. In particular, a bug of SWBv2.2 output lines reported in the June BPAC review prevented the usage of this ASIC. All of the L1 modules already assembled with the SWBv2.2 have been recovered by replacing SWBv2.2 chips with the previous SWBv2.1 version. Already diced SWBv2.1 chips were retrieved from AMS and single-chip bump deposition performed by IZM. Testing of these chips in KIT will start in December. As to sensor production, the PXD9-20 batch was delayed by a metal contamination in the etching process, resulting in an incomplete removal of the handle wafer. The issue has been resolved and now the first half of the batch is expected to be delivered in November, followed by the remaining half in February 2020. These deliveries should provide enough components for completing PXD2021 if the yield is kept at the 100% level for L1. The remaining PXD9-21 batch is therefore kept on hold. No lack of other ASICs and components is expected.

The committee commends the PXD2021 collaboration effort and it is impressed by the progresses achieved since the June review. The path to the PXD detector completion is now well delineated and with reasonable time profile and contingency.

4.1.2 Concerns

- While the schedule is now realistic, there are some critical steps on the way. The number of good SWBv2.1 chips available will only be known after testing in December. The number of L1 sensors in the PXD9-20 batch is barely enough to provide the 9 L1 pairs missing to achieve the target ladder production. Therefore, for the completion of the L1 production, 100% yield in the module and ladder assembly is required. This is a concern because no L1 ladder has been assembled since the reviews and therefore the impact of the improved assembly procedure has not been verified for L1.

- There are some open items in the schedule, like whether flip-chip can be performed at HLL or at IZM, as well as the effectiveness of the distribution of module tests across several labs.

- If the schedule is delayed and installation is postponed further after the 2021 summer shutdown, the capability of the collaboration to maintain the expertise both on the detector assembly and installation becomes also a critical issue. This is in addition to the fact that a long term commitment for the detector operation and maintenance after the replacement of the current PXD with PXD2021 is also required.

4.1.3 Recommendations

- An analysis should be performed to evaluate the impact on the schedule of potential yield issues in L1 module and ladder assembly that may require the processing of the PXD9-21 sensor batch. A similar analysis is needed to assess the impact of moving the flip-chip from HLL, which is the baseline, to IZM.

- An effort should be made to guarantee that expertise and resources are available to complete the PXD2021 construction and installation, and for the long term operation and maintenance of the detector. This planning will be helped by a decision by the Belle II collaboration about the start of the long shutdown needed for the PXD replacement.

4.2 Silicon Strip Detector (SVD)

4.2.1 Overall status

The SVD performed well in Phase 3 and cosmic ray runs. In a cosmic ray run, the expected cluster properties were confirmed, with a hit finding efficiency $> 99\%$. During operation in the previous data run, there was no major issue except for some concern about the data bandwidth limitation due to less-than-ideal trigger timing and higher occupancies. The SVD group fixed a few remaining minor issues during the summer shutdown and is ready for the next physics run.

4.2.2 Concerns

- Concerns on occupancies, which are higher than the expected values, remain.

- The bandwidth limitation will be another concern at higher background conditions, if the trigger timing is not improved.

4.2.3 Recommendation

- The SVD group should continue to work with the background group to reduce the SVD occupancies.

4.3 Central drift chamber (CDC)

4.3.1 Overall status

The committee was pleased to learn that a variety of measures put the CDC dark current issue under control and CDC operation has been stable. Water was added to the gas and its content was maintained at about 2000 ppm. The gas flow circulation rate was doubled from 0.5 to 1 volume/day. HV was reduced somewhat for layer 53 and 55, while layer 54 was put at a yet lower voltage. A new protection procedure is implemented where the HV is cut when the current stays at a predefined level for a few tens of second. High leakage currents are still a matter of major concern and are being monitored carefully with updated online reference plots to track long-term behaviour. Improved monitoring tools show that there was no significant drop in tracking efficiency in z, and dE/dx measurements per layer were reasonable.

Several field wires in a 5-cell by 5-cell test chamber built before CDC construction were replaced with wire left over from the construction. A series of tests involving irradiation were conducted with HV and gas conditions similar to those of the CDC. After irradiation of as much as 40mC/cm/wire, some degradation of gas gain and some instabilities in the current behaviour were observed, where the latter could be due to the Malter effect. Following further tests, there will be an attempt to remove affected field wires to look for deposits in order to confirm the Malter effect. The next priority would be continuation of ageing tests with more precise H_2O/O_2 control.

Simultaneously, computer simulations of CDC performance with a He/CO_2 mixture were performed. This is to investigate the possibility of replacing the current He/C_2H_6 mixture with one without hydrocarbons. The results so far were as expected: i.e. a higher voltage will be required to achieve the same gain and the drift velocity will be lower. The He/CO_2 studies are ongoing with plans to include magnetic field effects and with different CO_2 fractions. The studies will then address the resulting tracking performance.

The CDC was able to operate at a 30 kHz trigger rate in the global cosmic ray test. An algorithm based on TDC information to reduce cross talk has been validated in software. Implementation of this algorithm in the ASIC firmware of CDC frontend cards is underway to improve the Level-1 trigger at high luminosities.

Studies of tracking efficiencies with various event types, such as $e^+e^- \to \tau^+\tau^-$, $e^+e^- \to e^+e^-(\gamma)$ and slow pion from D^* decays, are advancing well. Overall agreement between the data and simulation is good. However, there still exist small discrepancies that need to be understood.

4.3.2 Concerns

- Although stable CDC operation has been achieved for the time being, there is no guarantee that the measures adopted so far will be adequate to maintain CDC performance for the rest of the Belle II program.

- There has been significant progress in attempting to determine if the high current problems are due to the Malter effect, but these studies have not yet led to firm conclusions or clear mitigation strategies.

- Studies of CDC performance with a particular He/CO_2 mixture, as an example of a gas without hydrocarbons, have begun but still much work is required to establish the feasibility of such a change in the gas mixture.

- It is not obvious that sufficient personnel are already involved in the efforts to understand and mitigate the CDC high current problem.
- The essential elements of how a replacement chamber should be designed or built differently are not yet addressed clearly.

4.3.3 Recommendations

- The efforts to understand the reasons for the high CDC currents and possible mitigation strategies must continue with very high priority.
- Every effort should be made to ensure that the level of personnel engaged in CDC hardware is adequate to understand and resolve CDC high-current issues before more drastic action might be required.
- The Belle II management should be prepared for more drastic action, e.g. a major repair or even construction of a new CDC.

4.4 Barrel particle identification (TOP)

4.4.1 Overall status

The committee commends the TOP group for their hard work and progress on several fronts since the June meeting. They continued to make substantial progress in solving operational issues during the summer and appear to be ready for the autumn running. Further results of TOP PID performance studies remained similar to those reported in June, and substantial differences between simulation and data remain. The conventional MCP-PMTs will not survive beyond summer of 2021, and could limit the attainable luminosity even accepting some degradation in performance. All 224 conventional tubes must be replaced during the 2021 summer shutdown.

Going into the summer, about 96% of the system channels were functional. There was substantial summer work inspecting the fibre, data, and trigger links to the front end modules. There are a number of problems of varying severity with boardstacks and more broadly with DAQ that are not yet fully resolved. In most cases, these can be reasonably accommodated at the system level during running. In more specific cases, e.g. Boardstack 6a, the link is permanently lost, and repair must await full access that will come when the MCP-PMTs are replaced in the summer of 2021. Other summer work included improvements to the user interface in the operating system to ease operational loads and lower barriers to becoming an expert operator, and streamlining calibration, trigger matching tuning, and PMT gain monitoring. More work appears necessary to understand and mitigate the rather frequent single event upsets. All calibration data is in hand. Further calibration runs are being done during daily local runs for stability monitoring.

The overall TOP performance remains to be better understood. Kaon identification efficiencies and pion fake rates have been studied using identified kaons and pions from

D^* tagged D^0s and pions from K_S decays. Kaon efficiencies match with the simulation results, but the pion fake rates do not.

An active background simulation effort is continuing to understand how machine backgrounds affect MCP-PMT lifetimes, and how to cope with backgrounds as the run proceeds and the luminosity increases. At the operational level, the TOP group has a liaison with the beam background group. Members of the TOP group take background shifts and provide constant feedback. New simulation now not only deploys a better characterised background model, but also implements a more realistic spectrum of degraded QE. Results from these studies require that the average beam backgrounds (Touschek plus beam gas) be less than 2 MHz/PMT in order to be able to defer MCP-PMT replacement until summer 2021. The rate should drop to less than 1 MHz/PMT after the summer of 2022 at the latest so that the ALD tubes can survive until 2027 when 50 ab^{-1} is expected to be attained.

The production of the new MCP-PMTs has proceeded in a timely manner at Hamamatsu, and it seems to be straight forward to complete the testing and assembly well in advance for the 2021 summer installation. Unfortunately, the lifetimes for more than 70 of these new tubes are expected to be shorter than 10 $\mathrm{C/cm^2}$. The cause of the short lifetimes in many tubes from some batches is still under investigation with testing both at Hamamatsu and Nagoya, but it is probably related to the thickness of the ALD layer. The team needs to replace these bad tubes and is in negotiations with Hamamatsu. A solution may require additional funds. A meeting with Hamamatsu is scheduled for the end of November.

4.4.2 Concerns

- As already noted in the June meeting, accepting a running condition with a 2 MHz/PMT rate for Touschek plus beam-gas backgrounds from the LER may have a serious impact on the performance of bar box modules with conventional tubes even before the summer 2021 shutdown if the integrated luminosity goals are met.
- It is unclear that production issues for the present production run of MCP-PMTs with ALD are fully resolved, and substantial contracting issues remain open.
- Intense effort is needed to understand the TOP performance and tune the simulation.

4.4.3 Recommendations

Most of the recommendations from June remain open and relevant.

- Understanding and reducing LER backgrounds are very high priorities for the TOP and background groups, impacting both the short-term and long-term health of the TOP system.

- Continuing to closely monitor the present round of ALD MCP-PMT production and testing, while quickly deciding how to obtain further tubes from Hamamatsu are both crucial in order to be able to replace the conventional tubes during the summer shutdown of 2021.
- Understanding the discrepancies in TOP's particle identification performance between data and simulation is another high priority for the TOP group. With sufficient data now becoming available, it should be possible to bin the PID data in different regions of phase space, and/or physical or component space, in order to identify the origins of the differences.

4.5 K-Long Muon system (KLM)

4.5.1 Overall status

The committee is pleased with the reports of significant progress on hardware and software issues made by the KLM group since the previous BPAC meeting.

Critical known hardware problems were addressed: all 16 6U-VME crates for the barrel RPC readout were replaced to deal with power limitation and reliability issues, five Data Concentrators with issues were replaced, a problematic SCROD and several of RPC HV cables were replaced, and a number of dead scintillator channels were repaired. Although a few hardware problems remain, as a result of these improvements a total of 99.5% of all KLM channels are operational: 96.7% of the BKLM scintillator channels, 99.9% of the BKLM RPC channels, and 99.6% of the ELKM scintillator channels.

Integration of the BKLM and EKLM systems has progressed well on many fronts, including: merging slow-control systems, unifying data quality monitoring, developing a common expert shift team, and unifying μ and K_L identification. Common alignment and calibration software are almost complete, but efficiency and background monitoring remain to be unified. Timing information is useful for checking consistency of 2D tracks and suppressing background tracks. After calibration, the timing information for the BKLM is very good, i.e. σ is 7.5 ns for the scintillators and 8.9 ns for the RPCs. The timing codes for the BKLM and EKLM are now merged into a unified system and the result will be tested with forthcoming data.

Weekly KLM meetings were established to coordinate the efforts for the BKLM and EKLM, addressing committee concerns about setting priorities for work on a unified KLM system

So far the KLM firmware only utilises hits on detector elements. There has been significant progress in developing wave-form readout, but this upgrade is not yet ready for routine operation. Meanwhile, the simplified firmware works well with stable operation up to 50 kHz with occupancies of 2–3%. The KLM DAQ operates relatively trouble-free at processing rates somewhat above 30 kHz. The triggers generated by the BKLM and EKLM are still not usable due to large latency jitter. An ongoing effort has addressed this problem by delaying triggers to a fixed latency, but not all bugs in the system have been identified and eliminated.

The efficiencies of the RPC layers in the BKLM are near 90%, with some troubling variations. However, the efficiencies of the scintillator layers are only about 10%. BKLM stand-alone and extrapolated-track efficiencies are consistent. In the EKLM, stand-alone efficiencies are about 80%, but extrapolated-track efficiencies are lower.

Muon identification is well-advanced, but there are some discrepancies between data and simulation. K-Long identification is less well-developed. Although much work to refine the KLM firmware and analysis software remains to be done, observation of a very clean $J/\psi \to \mu^+\mu^-$ peak is very encouraging.

4.5.2 Concerns

- Uncertainties in understanding KLM efficiencies remain. The very low efficiency for BKLM scintillators is particularly troubling. The difference between standalone and extrapolated-track efficiencies for the EKLM, which was said to be not related to the alignment, is also worrying.
- The calibrations of some of the SiPMs still require verification.
- The lack of the KLM trigger remains a matter of some concern.
- There has been some increase in the personnel committed to the KLM, but the actual number of available FTEs may still be marginal.

4.5.3 Recommendations

- There should be concerted efforts to understand the outstanding efficiency issues and resolve them.
- The KLM trigger should be completed and thoroughly tested.
- The collaboration should examine whether the number of FTEs committed to the KLM project is sufficient for the required effort. If not, collaboration management should attempt to identify additional personnel for the KLM project.

4.6 Detector performance

4.6.1 Overall status

The committee commends the collaboration for forming and rapidly implementing the Performance Group in June 2019. This provides a coherent approach to understanding and improving all aspects of physics performance as the experiment proceeds. The organisation seems complete, well structured, and coordinated, with many opportunities for interactions between the subgroups. It also can make collaboration wide recommendations for physics analysis, provides detailed ongoing internal reviews of performance, and will lead to performance publications. The associated personpower continues to grow.

The committee was excited to see performance results on neutrals, trigger, beam performance, tracking, charged hadron identification, and lepton identification using much of the 6 fb^{-1} of data collected during the Phase 3 run of 2019, plus earlier Phase 2 data. In general, detector performance is quite encouraging at this stage, though performance with data often lags significantly behind that of simulation. Calibration and alignment for all detectors is one primary area of focus. Another is to study the response properties of the individual sub-detectors. Systematic errors are also a crucial area of study, with evaluation of PDFs, backgrounds, triggers, run dependence, etc. Understanding and resolving these issues and data/simulation discrepancies will be central to the collaboration's work as they progress towards more physics results.

Studies of detector performance for photons and π^0s are advancing rapidly. In general, performance at high energy seems to agree rather well with simulation, while the low energy studies are in an early phase. Energy resolution seems to be worse than expected and timing resolution in data is also worse. The latter is especially important to understand as background levels increase.

Trigger studies are at a rather early stage, and a number of L1 trigger components are still not available. The calorimeter L1 trigger is stable, but the tracking triggers are less so, and have lower efficiency than expected. There was an angle dependent two track efficiency drop during the Phase 3 run, which is under study. High Level trigger signatures for low multiplicity events are reasonably stable but there are some run dependent losses (at about the 10% level) for some Bhabha channels and single photons that need to be understood.

Beam studies are underway. The offline luminosity in Phase 2 has been determined to a precision better than 0.5%, and published as the first Belle II physics paper. Work continues for Phase 3. Other beam related issues such as beam spot position and size, beam energy, the boost vector, and B counting are under active study. There are 6% fewer B pairs than expected in the reprocessed data set, while there are 7% more B pairs than expected in the prompt data set. Those discrepancies could be due to the uncertainties in the subtraction of continuum events. Continuum data samples taken under different beam conditions are needed to understand the discrepancy.

A wide variety of tracking studies are in progress or planned. A number of these analyses are just getting started. So far, there appears to be rather good agreement between data and simulation for tracks from three-prong τ decays, although these have limited momentum and pointing information. The slow pion analysis is very early but encouraging to date.

The hadronic particle identification group continues to work on methodology, but baseline likelihood methods are in place for both binary (e.g. π/K) and global PID. Kaon identification efficiencies and π fake rates have been studied using identified kaons and πs from D^* tagged D^0s and πs from K_S events. Reasonable π/K separation is being achieved globally with substantially higher fake rates in data than in simulation at the same efficiency, and with large differences between specific regions of detector phase space. These differences will most likely need to be explored at the subsystem level (TOP, CDC, ARICH).

A number of lepton identification performance studies using a wide variety of channels are underway by several different groups and are being documented. Initial performance seems reasonable at this stage, though efficiencies for both electrons and muons are a few percent low, compared to simulation. Hadron fake rates measured in data with πs from K_Ss are typically around 2–5 times higher than in simulation. Global particle identification methods (such as MVA, BDT, and NN) are being developed.

4.6.2 Concerns

- While there are many groups associated with this effort across the collaboration, it is unclear if the human resources are fully adequate for all tasks. It is also unclear how effort is being balanced between the competing needs of the sub-detectors and higher level developments.
- The list of priority for the short term seems appropriate. The committee agrees that a detailed understanding of the relative performance of Belle/Belle II is crucial for physics in the short to intermediate term. Systematic comparison of performance with BaBar, especially in the particle identification sector, could also be very helpful for understanding performance, and might suggest directions for developing the best methodologies.
- The stability of performance versus backgrounds for both tracking and particle identification needs to be better understood.

4.6.3 Recommendation

- Many features of the present overall performance are most likely driven by anomalous behaviour in sub-detectors and must be studied and resolved at this level first. It is crucial that the individual sub-detector analysis efforts be kept strong or even strengthened, both to deal with alignment and calibrations, and to understand the origin of the performance anomalies and how they can be resolved.

5 Trigger and DAQ

5.1 Trigger

5.1.1 Overall status

The trigger has been working with good hadron efficiency during Phase 3. The nominal L1 trigger rate was typically around 2.5 kHz, dominated by beam background (about 90%). The main trigger was based on the CDC 2D track and ECL cluster sub-trigger. For 2020, it is expected that the background rates could increase by a factor 5 with increased currents. This would correspond to about 12 kHz trigger rate, without modifications to the trigger.

A lot of progress has been made since the last BPAC in June to prepare for the autumn run to implement fixes, mitigation actions for detector readout issues, trigger enhancements and new features. These are not crucial for the physics data taking during the 2019 autumn run, but it will be very beneficial to commission and validate them with beam as soon as possible. In particular, mitigation of CDC noise, the CDC timing, the ECL with tighter thresholds and the CDC-3D Neural Net are ready to be tested with beam.

5.1.2 Concerns

- The CDC 2D track trigger suffers from low efficiency due to noise. A TDC based crosstalk filter on the CDC front-end and a reduced track counting of tracks with similar parameters is expected to reduce the rate substantially without affecting efficiency. This still remains to be validated with beam.
- The CDC-3D and CDC-3D Neural Net are still missing, although progress has been made.
- The only timing source in Phase 3 was the ECL. The CDC timing is just ready to be tested with beam.
- The KLM trigger is generated by both BKLM and EKLM. However, it suffers from large jitter. The solution to delay triggers to a fixed latency has been implemented and is still being debugged.
- The TOP trigger has not been ready yet. The reasons for the low efficiency and low resolution are not understood.
- Substantial progress has been made with the trigger simulator, but the firmware simulation is not complete and there are inconsistencies between simulation and data.
- Although correct priority has been given to the CDC trigger, the human resource situation remains very worrying, both for the global trigger as well as sub-detector trigger systems.

5.1.3 Recommendations

- Getting one of the CDC-3D-trigger systems fully commissioned should be the highest priority of the team.
- The remaining timing sources from TOP and CDC should be added.
- The commissioning of the KLM trigger with fixed latency should be pursued.
- The trigger performance should be monitored in realtime to detect quickly any problem with the detector or machine background that may arise.
- The committee would like to see a clear plan for the completion of the full trigger system.

- The committee felt that an adequate L1 trigger simulation should be prepared as it is important to estimate the trigger rate in advance. This could be an opportunity for a new group/person to join and become familiar with the Belle II trigger.

- New persons, ideally a group with a institutional commitment, should strengthen the central trigger team, in particular for long-term operation.

5.2 DAQ-online

5.2.1 Overall status

For Phase 3, an overall availability of the DAQ of up to 80% has been achieved. Also the recovery time after failures has been greatly reduced. This is an achievement for the early phase of data taking for an essentially new experiment.

The HLT became the active rejection mode during Phase 3, rather than tagging events. The data reduction by HLT selection was about a factor of ten. The Phase 3 DAQ throughput was limited by the HLT processing. The HLT capacity has been almost doubled during summer by adding a further four units to the five HLT/storage units deployed in Phase 3.

The robustness of the slow-control in Phase 3 has been significantly improved and numerous bugs have been corrected. The introduction of modern software management methods has contributed significantly to this. A powerful monitoring of dead time and downtime has been established. Work has started to automate recovery procedures.

The personpower situation has been somewhat alleviated by adding DAQ experts from the sub-detectors to the core team. They contribute to on-call expert shifts and started contributing to development work. New personpower for the HLT has already resulted in the adaptation of more modern software technologies, such as ZeroMQ for the HLT framework. The collaboration agreed to fund three FTEs from FY2020 onward. They have been allocated to specific areas already.

5.2.2 Concerns

- The DAQ design foresees an ultimate 30 kHz L1 rate. This is not yet reached for all sub-detectors. This holds in particular for the SVD, although options are identified to reach that value.

- The first approach for fast reconstruction in the HLT with simplified tracking has not yet reached the desired performance in tracking and trigger efficiency. Hence the autumn run will use the same reconstruction as offline.

- The number of HLT units will remain 10 for a number of years (till 2024). Although there is a plan to increase the number of cores by 40% in 2020, it is not clear that this is sufficient with an increase of beam related background.

- The problem of the mismatch in the number of events between several information sources in the online-offline data transfer is disturbing. Resolution of this issue is still pending.

- Although the system to monitor the BUSY state of individual sub-detector readout channels is operational and has a good time resolution, there are still a few issues at the level of the generation of the BUSY signal that makes the monitoring not yet optimal.
- Downtime monitoring seems not to take into account the ultimate usability of the data for physics analysis. This might result in a skewed perception of the actual efficiency of the experiment, when reviewed during running on a daily or weekly basis.
- In a non-negligible number of failure scenarios, the collaboration depends on the availability of two or even a single expert from central or sub-detector teams. This is not a long-term sustainable situation.
- The personpower situation in the area of system administration remains precarious. This problem will increase with the expansion of the HLT farm.

5.2.3 Recommendations

- Data reduction by sending the HLT decision to the ONSEN and subsequent ROI selection has been verified for half of the PXD modules. Full implementation and validation should be pursued.
- The efforts to attract and secure "junior" experts should be continued. These are likely to play a major role as well for the development and commissioning of the DAQ upgrade in the coming years.
- The efforts to improve and to expand the coverage of the system documentation should be continued. This is important for training expert shifters and new experts more quickly.
- The good practices of modern software management should be continued.
- All incidents creating data taking downtime should be carefully followed and the root cause should be fixed. Inefficiencies and downtime due to slow-control and system failures should be carefully monitored and broken down by source, to direct the priority for improvements.
- The efforts to reduce the recovery time of the system should continue.
- Wherever possible recovery procedures should be automated to increase the lifetime of the DAQ and relieve both shifters and experts.

5.3 Decision on the new DAQ card

5.3.1 Overall status

The DAQ upgrade concerns the replacement of the COPPER boards, currently used to read out the front-end electronics of all sub-detectors via about 800 Belle2link optical connections. The replacement is foreseen during the summer shutdowns in 2020 - 2021. A DAQ upgrade committee has examined two proposals and submitted its conclusions to management in October 2018. Subsequently, it has been decided to select the PCIe40 proposal. The BPAC committee would like to congratulate the collaboration for making a decision according to a well-defined process.

5.3.2 Concern

- There is no concern.

5.3.3 Recommendations

- The committee would like to see a plan with milestones for the firmware and software development and installation and commissioning schedule for each sub-detector.

6 Offline software and computing

The committee is pleased to note that the Belle II Software and Computing groups are addressing many of the concerns and recommendations expressed in previous reports. The positive impact of additional human resources in computing are clearly visible. Implementation of automated procedures for calibration and computing operation has started and is progressing nicely.

Considerable experience has been gained with the *2019a* and *2019b* data processing. Processed data have been provided in a timely manner to the collaboration and the data processing schema is evolving although the processing schema for the *2019c* autumn run will be same as that of the spring data taking. In the future data processing schema, the offline KEKCC resources will be reserved for prompt processing to produce cDST files for calibration and mDST files of offline HLT selected skims for the quality assurance monitor. Full data processing will only be performed on GRID sites. New offline skims targeted for analysis will be produced on GRID sites; they will be provided already from the next reprocessing, *proc10*. The full reprocessing of the whole 2019 Phase 3 data for the winter conferences after the end of data taking in mid December is to be achieved by late January. This will provide a crucial scale test for the readiness of calibration and processing for the 2020 spring run.

Intensive studies are being carried out under newly appointed coordinators that include those of the recently formed Performance group. The collaboration has formalised the mandate of the new group that is responsible to provide reconstruction recommendations for physics analysis, provide data versus simulation (MC) corrections and systematic uncertainties, as well as evaluating the Belle II performance in degraded background conditions. The Performance and Data Processing groups are also expected to provide feedback on prompt processing before a major reprocessing takes place. The collaboration is ensuring good communication between all groups involved via bi-weekly coordination meetings in which the Belle II management and all relevant groups (Data

Processing, Software, Computing, Performance) participate. It has been reported that many new institutes will be joining in the performance studies.

6.1 Offline software and calibration

6.1.1 Overall status

First analyses on the performance of tracking, particle ID, neutrals, trigger and beam parameters for Phase 3 have been carried out and are under internal review. Performance of tracking and particle ID with single as well as combined detectors were presented to the committee. Data from the latest state-of-the-art *proc9* (full) and prompt processing have been analysed and compared to the latest MC samples produced, MC12, with simulated background. Run dependent studies with beam background overlaid MC have also started with preliminary results of CDC standalone tracking efficiency showing an improved agreement between data and MC with respect to run independent MC with simulated background. Agreement at better than 1% has been found between data and Monte Carlo for tracking efficiency in Phase 3 with fast tracks from $e^+e^- \to \tau^+\tau^-$ events. Particle identification is worse than expected in all sub-detectors, with lower identification efficiency of electrons and muons with good purity, and π/K separation with larger fake rates in data than in MC for the same efficiency. Reconstruction efficiency of high energy photons and π^0 is in agreement between data and MC, but both energy and timing resolution are worse than expected. L1 and HLT performance studies are ongoing and trigger lines where efficiency is lower than expected, e.g. L1 track trigger and HLT single photon, are under investigation. Finally beam energy and luminous region parameters were expected to be soon available for analysis at the time of the review.

Software release 3 was used for *proc9* and prompt processing as well as for MC12 production. A new release 4 was deployed at the beginning of October and performance studies comparing it to the earlier version were carried out with its pre-release. Release 4 provides new and improved reconstruction and analysis algorithms as well as improvements in the interfaces to the conditions database and computing system. A new global tag replay feature has been introduced to automatically use the information stored in the file metadata and ensure the default use of the same conditions used during the production data processing. This release also enables run dependent simulation of all detectors. Release 4 will be used for data taking in the *2019c* autumn run, for the full re-processing (*proc10*) of the ~ 6 fb^{-1} of Phase 3 *exp 7+8* data and for the next MC13 production campaign. A new release 5 is in preparation for spring 2020. It will include a framework for systematic uncertainties and improvements in the simulation of the detector response. A major change in this release will be the update to a more recent Geant4 version and to a Belle II custom-made physics list: a new group has been identified to work on the validation of the changes and carefully check the effect on the quality of the simulation. The HLT poses strict limits on software performance. Effort to improve it has increased and there is a new institutional commitment for benchmarking and performance optimisation.

Procedures for the production and handling of calibration constants have been put in place. For all calibrations, responsible groups have been found. There will be two levels of calibration. About once a week a prompt calibration will be performed on recent data and used for the reconstruction of the prompt dataset. These data are used to monitor detector performance. The prompt calibrations will be superseded by a full recalibration that takes place only once every few months and is used for a full reprocessing. Only the reprocessed data will be used for actual publications.

The calibration procedures have been successfully applied to all 2019 data, leading to a large improvement in detector performance between the prompt and the reprocessed data. However, the calibration tasks were manually operated, which required a great effort, and introduced delays. To minimise the delay between data dating and prompt processing in the future, an effort has started to fully automatise the prompt calibration using a system called Apache Airflow. First tests of the system look promising. The plan is to make the system operational by the end of November.

The committee was pleased to note the introduction of a new run registry that stores basic run information and quality flags. Run information logging has greatly improved, and there is no need anymore to manually gather information from various sources.

The committee was informed that in order to foster collaboration between the Software group and the newly formed Performance group, the Software coordinator and/or deputy regularly attend the weekly performance meetings and joint workshops on Software and Performance are planned.

Advantages and disadvantages of making the Belle II software public have been identified. Consensus is emerging in the Executive Board to license the BASF2 code under the LGPL licence and discussion on an initial proposal for a change of the Belle II ByLaws will continue within the Belle II Executive Board in the next months.

Precision benchmarking of the Belle II performance requires identical analysis procedures for Belle and Belle II data. The committee was informed that this is already technically possible and a framework allowing to use Belle II software, as e.g. the Full Event Interpretation tool, to analyse Belle data has already been exercised.

6.1.2 Concerns

- While the sub-detector groups have strengthened their software involvement in the calibration software with a few people active in most of them, the CDC tracking is at the moment with only one person. In spite of the automation of the procedure, lack of redundancy in expert knowledge constitutes a potential risk in the timely deployment of calibration for unexpected situations.

- A list of known software open tasks has been prepared to attract new contributors and about five FTEs have been identified. The group should continue motivating skilled contributors and attracting new ones, in order to retain the low level software knowledge essential for the long term maintenance and development of core elements.

- For deep understanding of the Belle II detector performance, actual behaviour of the sub-detectors and experimental conditions must be fed back accurately to the simulation. It is not clear to the committee if a plan with all groups involved is in place for this task.

6.1.3 Recommendations

- Differences between data and MC are still present and a through investigation of detector response and/or calibrations is needed to accurately quantify the detector contribution to systematic errors. The committee recommends to continue having good communication and foster common studies with the detector experts.

- The migration to the new Geant4 version will require careful validation. In addition to low verification, high level reference distributions should be identified with detector experts and the Performance group to quickly identify the potential impact on physics measurements.

6.2 Computing

6.2.1 Overall status

The committee was pleased to note the progress on automation and monitoring of computing operation tasks. The whole chain of data transfer from online to offline and from offline to distributed GRID resources is operating smoothly and is practically fully automated. The automation of the only remaining manual task of transferring data from the offline core computing pipeline to the new dedicated *BelleRawDirac* service, presented in June, was in progress at the time of the review.

The origin of the remaining discrepancy in the number of events registered online and offline is under investigation and seems to stem from the online system. Monitoring of raw data registration is working nicely and monitoring of data copy is being improved. The level of the data transfer activity since June is similar for all offline nodes with regular spikes corresponding to the daily transfers that have been occurring. It is expected that multiple transfers a day will occur with the increase of luminosity.

Six new front-end servers have been procured and deployed at KEKCC to match the HLT increase in resources and supplement core computing operations. A *Fast-lane* RAW data copy tool with a seven days data persistence has been setup for the detector experts enabling them to copy and analyse data immediately after a run is taken for fast checks. The Belle II existing custom distributed data management framework has been consolidated and is now in maintenance mode with all development efforts migrated to Rucio. The replica policy system has been improved to enhance scalability and was being commissioned at the time of the review. An experienced Rucio developer has been hired at BNL to work on the Rucio implementation for Belle II starting from November 2019. Tests have been successfully carried out at BNL on a dedicated Rucio-DIRAC testbed and Rucio servers showing that the data registration and transfer performance exceed the Belle II requirements. Discussion on a new BelleDIRAC system with the improved understanding of DDM API requirements has started. The timeline for Stage 1 of the migration to the Rucio-based DDM has been revised with the API now expected to be ready in January 2020. A set of checkpoints in the Stage 1 migration is foreseen at the next Belle II General Meeting. The Stage 2 migration to Rucio as file catalog discussion is also progressing and with good communication being established with the developer at Imperial College working on the file catalog plugin.

Valuable experience has been gained with data (re-)processing of Phase 3 *2019a* and *2019b* running simultaneously at KEKCC and BNL. It was confirmed that raw data processing and re-processing jobs can run at BNL by artificially throttling them at KEK. Overall, it took slightly more than a month to complete the whole data processing with a share of 60% - 40%, respectively at KEKCC and BNL, and the reprocessing was done totally at BNL. The major issue encountered with the *proc9* on the GRID resource was the staging of files with jobs completing before new staged files become available at BNL. This was solved in processing of *bucket7* by checking the status of the staging before job submission.

Use of the CPU distributed computing resources provided to the Belle II collaboration has decreased in 2019 with respect to 2018.

MC production is proceeding smoothly and a new MC13 campaign is in preparation aiming to produce about 1 ab^{-1} of generic events. An improved workflow for requests of production of signal samples by the Physics Working Group liaisons will be put in use. Production of run-dependent MC has been exercised at KEKCC in the current MC12 campaign. Ongoing effort to carry it out on GRID resources have identified the size of the background files as the main outstanding issue. Methods to reduce file size and/or optimal distribution are under discussion with detectors and distributed computing experts. An analysis of the resource requirements to produce run-dependent MC samples on the Belle II distributed computing sites has been carried out with a conclusion that each site has to have a proper balance of CPU and storage capacity. Three categories of sites (GRID sites with much more CPU than storage, those with much more storage than CPU and non-GRID sites without storage) cannot be used at the moment. Streaming data from sites with storage resource to those nearby with CPU resource for processing may solve the issue but requires good network bandwidth.

MC skimming is fully operational on the GRID sites and Physics Working groups are encouraged to use exclusively skimmed samples. A new dataset searcher tool to replace the manually populated address tables will be made available with MC13 to ease the retrieval of MC samples.

The KEKCC resources are scheduled to be replaced in summer 2020 with capacities expected to increase to 15,000 CPU cores and 17 PB and 100 PB respectively of disk and tape storage. A detailed plan and impacted services is being prepared based on the previous exercise that took place in 2016 to minimise the downtime. An overlap period of two months for the current and new KEKCCs is foreseen.

6.2.2 Concerns

- Although all types of jobs, from MC and skims production to raw data processing and users' analysis are running on distributed computing sites, the average CPU usage in 2019 has been only 25% of that of the previous year and concentrated in the period of April to June. While this can be understood to reflect the commissioning status of the Belle II experiment, the adverse impact on many remote sites should not be underestimated.

- Production of run dependent MC is critical for a deep understanding of the Belle II performance. It is important to understand how background file size needed at each site will increase with the increase in background sample collected with increased luminosity. Using efficiently all distributed computing sites may require dedicated development efforts that would put additional strain on the group woking on the distributed computing.

6.2.3 Recommendations

- The positive impact of the dedicated new personpower in core computing hired using the maintenance and operation fund of the collaboration is clearly visible. The committee encourages the Belle II collaboration to continue to support computing activities in this manner and further secure sufficient human resource as needed to ensure smooth running of the experiment.

- A training course and a detailed documentation for computing operation expert shifters has been introduced. Nevertheless more personpower is needed to ensure appropriate coverage for the operation throughout the year.

- With the current understanding of CPU and storage requirements for producing run-dependent MC samples at distributed computing sites, 15% of the sites would not be able to be used for these productions. The committee encourages the collaboration to utilise resources in small computing centres more effectively by finding suitable roles for their capabilities, e.g. giving them priority for user analysis jobs and/or production of simulated background MC samples.

- Having an experienced Rucio core team member at BNL should enhance the commissioning of Rucio as the Belle II DDM and help in speeding up the migration process. Nevertheless, deployment should be carried out carefully after ensuring all necessary functionality is proven. The committee supports setting up of specific checkpoints for the different steps of the migration phases as planned. The final step of moving to use Rucio-DDM for data replication should be done in the least intrusive way and the committee feels the plan to do it outside of data taking periods very sensible.