# *Belle II Experiment at the KEK electron-poistiron collider*
## Strategic Plan

### 1. **Belle II Experiment and Scientific Impact**

The Belle II experiment is a next-generation flavour factory experiment will study the fundamental constituents of the universe and search for evidence of new physics such as the source of dark matter. Belle II is a ninternational collaboration with 600 researchers in 23 countries that will study the collisions of high intensity beams of electrons and positrons at the KEK Laboratory in Tsukuba, Japan.

One of the primary goals of the experiment is to further studt the impact of CP-violation, which is believed to be the origin for the observed dominance of matter over anti-mater in our present universe. The results from the earlier BaBar and Belle experiments confirmed the effect of CP-violation as described by the theory of Kobayashi and Maskawa, who shared the Nobel Prize in physics in 2008. Canadians had a leading role in the BaBar experiment at the SLAC National Laboratory and made significiant contribution to the computing.

The measured level of CP-violation by BaBar and Belle is not sufficient to account for the observed matter and antimatter asymmetry in the universe. Therefore, a much deeper understanding of the related phenomena is required. The Belle II experiment will have a sample of data that is approximately fifty times larger than the previous experiments, enabling us to study CP-violation in greater detail and search for new physics with unprecedented precision.

Starting in 2018, the SuperKEKB facility at the KEK Laboratory will collide electrons and positrons in the centre of the Belle II detector (see figure 1). The electrons and positrons in SuperKEKB are circulated in opposite directions (and at different energies) in the storage rings. The two beams will collide in the centre of the Belle II detector. The Belle II detector is cylindrical in shape with layers of sub-detectors that are used to track the charged particles, and measure the energy of both neutral and charged particles from the collisions. The scale of the Belle II detector is optimized for the energy of the collisions and designed to record all collisions in a high intensity environment.

The Belle II Canada group includes ten faculty members from Montreal, McGill, UBC and Victoria. The Belle II Canada group, as part of its international obligation to the Belle II Collaboration, is expected to contribute a proportionate level of computation resources to the project. The level of the contribution is determined by a number of factors including the size of the Canadian group and the other contributions to the project. It is expected that the Institute of Particle Physics (Canada) will sign an international Memorandum of Understanding that will give our commitment to provide our share of computing resources and manpower to the Belle II Collaboration.

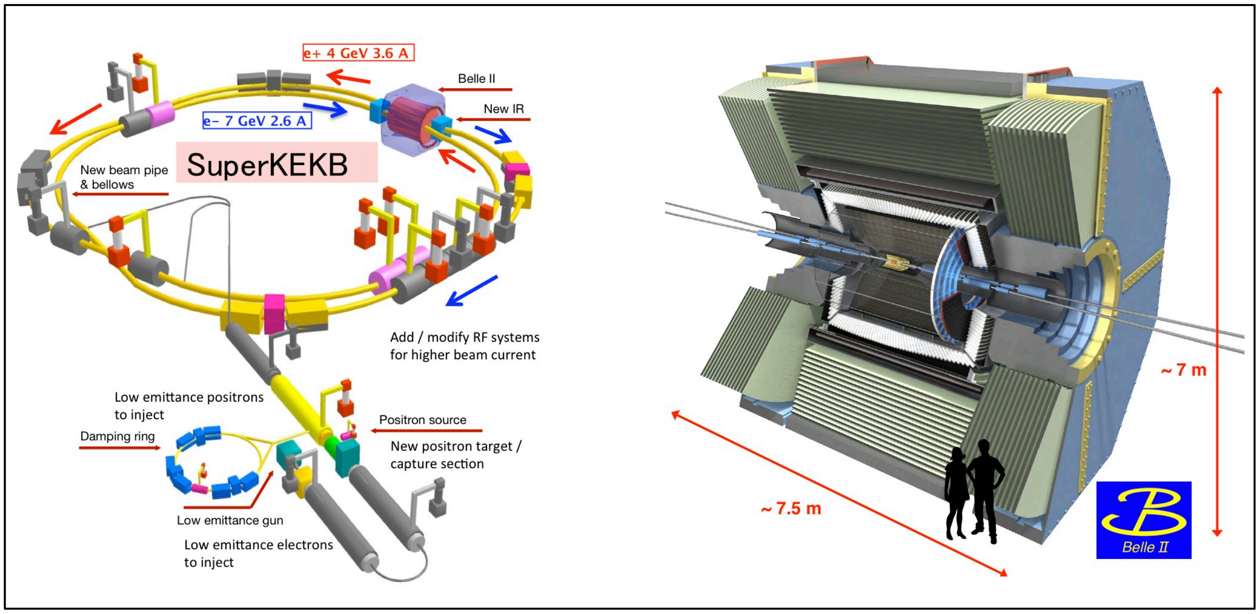


*Figure 1*
*The left image shows the SuperKEKB accelerator complex. Positrons are created by colliding electrons with a metal target and collected in a damping ring where the transverse energy of the particles is reduced or damped in order to make a low emittance (small) beam that is then injected into the positron storage ring (red arrows). The electrons (blue arrows) are directly injected into the electron storage ring. The two beams will be collided in the centre of the Belle II detector.*

The Belle II project is expected to continue for the next 15-20 years and its computing requirements will continue to grow by significantly over this period. The scale of the Belle II computing will be comparable to a small LHC experiment at CERN. Canada is expected to begin storing the raw data of the Belle II experiment in 2021.

## 2. Belle II Computing

Belle II is using computing resources distributed around world integrated into a hierarchical structure, similar to the ATLAS computing system[1] (see figure 2). The main computing centre for the Belle II experiment is located at the KEK where the raw data is received from the detector in near-real time. A second copy of the raw data will initially be sent to the Brookhaven National Laboratory[2] in the US for the 2018-2020 period as the size of the raw data sample is modest (12 PB). The numbers listed in Table 1 for the raw data are for two copies.

[1] We use the word "system" rather than "platform" to describe the Belle II computing due to its distributed nature.
[2] In 2017, the US moved the primary computing site from the Pacific Northwest National Laboratory (PNNL) to the Brookhaven National Laboratory (BNL).

In 2020-2021, the increasing size of the raw data sample will make it difficult for the KEK and BNLsites to host the entire samples. To distribute the data, Germany, Italy, Korea, India, Australia and Canada will each take a fraction of the raw data sample; for example, Canada will be responsible for 10% of the total raw data sample. The resources needed in Canada in 2021 are outside the period being considered in this RAC request, however, we need to start, planning as the storage needed in those years is significant. It is possible, that slower disks or possibly active tape could be used for the raw data storage. We expect that the data will be accessed a few times per year for reprocessing.

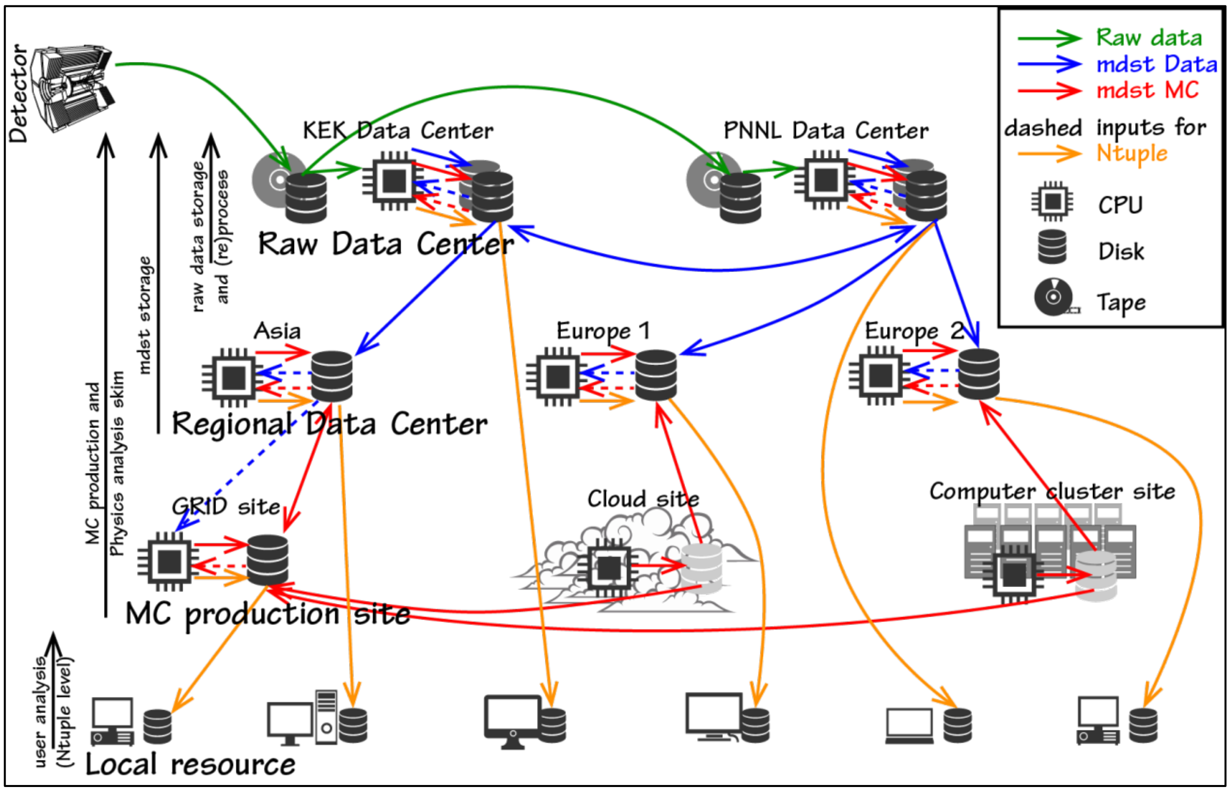


*Figure 2*
*The Belle II computing model for period up to 2020. Prior to 2020 the raw data will be copied to the BNL Data Centre (shown as PNNL) and after 2020 the raw data will be distributed to sites around the world (including Canada). Canadian resources are considered a "Regional Data Centre".*

In addition to raw data storage and processing, there are MC production samples and derived data sets for user analysis. These sample are expected to reach 20 PB by the end of 2020. The MC generation and user analysis will use the same computing centres plus smaller centres distributed around the world (see Figure 2).

The 2018-2020 computing requirements for Belle II Canada are modest with storage going from 400 TB to 800 TB and processing going from 400 to 800 cores 2018 to 2020 (see Table 1). The request takes a step upward in 2021 when Canada starts storing and processing the raw data.

| | 2018 | 2019 | 2020 | 2021 | 2022 |
|---|---|---|---|---|---|
| **Total collaboration** | | | | | |
| Raw Data (PB) | 2 | 4 | 12 | 28 | 50 |
| Derived data (PB) | 4 | 9 | 18 | 20 | 38 |
| Total storage (PB) | 6 | 13 | 30 | 48 | 88 |
| | | | | | |
| CPU (kHEPSPEC06) | 175 | 250 | 340 | 440 | 620 |
| CPU (kilo-cores) | 12 | 17 | 23 | 29 | 41 |
| | | | | | |
| **Canadian assignment** | | | | | |
| Raw data (PB) | 0 | 0 | 0 | 3 | 5 |
| Derived data (PB) | 0.1 | 0.4 | 0.8 | 0.8 | 1.4 |
| | | | | | |
| CPU (kHEPSPEC06) | 4 | 6 | 12 | 15 | 17 |
| CPU (cores) | 400 | 600 | 800 | 1,000 | 1,200 |

Table 1
Storage and processing requirements for the Belle II experiment for 2018-2012. We list the total resources and the expectation of resources from the Canadian group. Our measurements show that each core on the CC clouds is approximately 15 HEPSpec06. Some of CC-West cores are hyperthreaded and the performance is dependent on the other activities in the node.

Currently the Canadian processing is provided using the compute clouds of Compute Canada and other organizations. We managed the resources with our VM provisioning service (called CloudScheduler) (see Figure 3). Jobs from KEK are submitted to our HTCondor batch system in Victoria and CloudScheduler makes requests for the instantiation of virtual machine (VM) instances on the remote clouds. Currently we use between 6-10 clouds for Belle II.

Note that the UVic group has been funded with a CFI Cyberinfrastructure Competition I award for 2016-2019. The goal of the project is to enhance the distributed cloud system by adding the ability to run data-intensive applications on all clouds. We are establishing a federated storage system and are using in production for Belle II input data files. In addition, we are incorporating object storage systems into the production evironment and are using object storage at the University of Victoria centre.

Currently Belle II Canada has a disk storage allocation of 200 TB. We have been given access to 70 TB, which has been sufficent for our needs. Our processing allocation is 425 CPUs on the CC clouds. We have fully utilized these resources and have been able to opportunistically use idle processing.

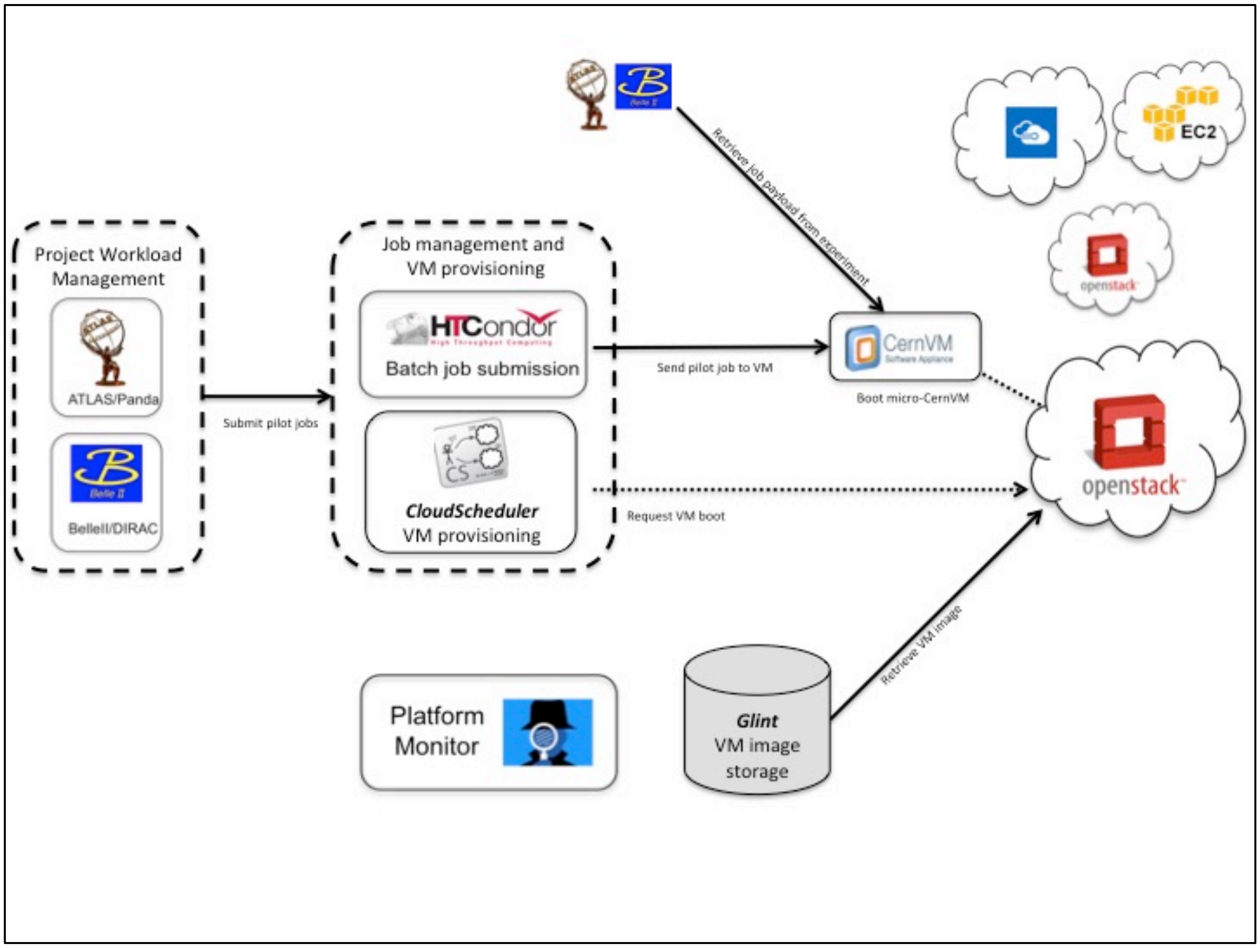


*Figure 3*
*Overview of the distributed cloud computing system. The project workload management system submits a pilot job to HTCondor. CloudScheduler monitors the HTCondor queue and requests a VM image be booted on one of the remote clouds. The cloud boots a CernVM image from the local image repository. Once the VM instance is running and registered with HTCondor, the jobs are submitted to the VM instance. Both ATLAS and BelleII submit pilot jobs that connect back to a central server to retrieve the production job.*

## 3. Belle II Canada Group

Belle II Canada consists of ten faculty members at the Universities of Montreal, McGill, British Columbia and Victoria. The Principle Investigator (PI) of the Belle II Canada group is Dr. C. Hearty at the University of British Columbia. The PI of this proposal is Dr. R. Sobie at the University of Victoria. The CCV's of Hearty and Sobie are included with this application.

The Belle II Canada group has a track record of accomplishments on many experiments including the previous flavor factory BaBar experiment. The group had leading roles in the BaBar experiment (Prof. J.M. Roney of the University of Victoria is the BaBar spokesman) and has been leaders in the analysis of the physics data. Currently the group leading the development of detector components for the Belle II experiment and is developing reconstruction software for the data from those detectors.

The team has a track record of expertise in HEP computing in BaBar and ATLAS, and is already making a significant contribution to the Belle II project (as described earlier). The Belle II Canada is one of the top sites in the collaboration for the generation of simulated (MC) data sample.

The Principal Investigator of this project is Dr. Randall Sobie who is a Research Scientist of the Institute of Particle Physics and Adjunct Professor at the University of Victoria. He is a member of the ATLAS and Belle-II experiments. He has also been involved in computing for HEP applications for 25 years. He is Director of the VISPA (HEP) Research Centre at the University of Victoria. He is Director of HEPNET/Canada, which is responsible for national and international network connectivity for the Canadian particle physics community ($1.5 million in NSERC support). Sobie is the PI of a CFI Cyberinfrastructure Competition I award ($1.648 million). He has received $2 million in support from CANARIE for the development of the compute-intensive distributed cloud computing system used by ATLAS and Belle II. In 2002, he was awarded a CFI grant ($7 million); it was the first CFI-funded project to focus on data storage and management. He is a member of the CANARIE Research Advisory Committee and is Chief Editor of the SoftwareX Journal.

## 4. Highly Qualified Personnel (HQP)

The Belle II Canada is committed to the training of undergraduate and graduate students, post-doctoral fellows and technical staff. Currently the Belle II Canada group has approximately 10 students and staff and is expecting this number to grow to 30 as the project starts to record collision data. The group has a track record (spanning many decades) of training HQP. Many of the graduate students are now faculty members in Canadian universities or research scientists in laboratories around the world. Former technical staff are employed by D-Wave (quantum computing) in Vancouver, a Silicon Valley start-up company, and another is head of research computing at the University of Victoria.

In addition to the physics training of HQP by the Belle II Canada group, our computing activities foster the training of HQP in novel areas of cloud computing, high-speed networks and high-performance storage systems. For example, the Victoria group typically employs 3-6 undergraduate science or computer-engineering students each year. These students are highly motivated and excited to work on cloud computing projects. They are employed for one or two 4-month terms before returning to complete their studies and typically move on to industry. Some of our former students are employed at Amazon, Microsoft, Intel, DELL and small Canadian high-tech companies.

## 5. Summary

We have described the computing and storage requirements that are required to meet our international commitments to the Belle II experiment at the KEK Laboratory in Japan. The resources will serve both the Canadian and international community of over 600 researchers on a project with the potential to discover new physics.

The experiment has been fully vetted by the funding agencies of the participating countries (including NSERC and CFI). The project meets the strategic goals of the Federal Government of Canada and its funding agencies (including Compute Canada). We expect that the experiment will publish over 500 journal papers over the lifetime of the project covering a broad range of particle physics topics. Further, it is an excellent training ground for both undergraduate and graduate students as well as technical staff.

The Canadian group has a track record of excellence in cloud computing, high-speed networks and high-performance storage systems. The group currently holds a CFI Cyberinfrastructure award that is being used to develop the software infrastructure. The group has already demonstrated its ability to full utilize its allocated resources and exploit opportunistic resources when available. Further the group has been able to augment its resources with grants from Amazon and Microsoft, and has access to other research clouds in Canada, Europe and the United States. In addition, the group has a history of strong relations with the IT industry.

## Appendix A Additional contributions

A number of our computing projects have been highlighted by various organizations:

The 2016 CFI Cyberinfrastructure Competition I Award:
https://www.uvic.ca/communicationsmarketing/media/news/2016+cloud-computing-physics-randall-sobie+ring
http://www.bctechnology.com/news/2016/10/7/UVIC-Research-Scientist-Leading-Project-to-Develop-a-Data-System-for-the-Worlds-Largest-Physics-Experiment.cfm?j=0

CANARIE article on our cloud computing software:
https://www.canarie.ca/software/platforms/hep/

In 2014, we helped CANARIE, Canada's national research network organization, test the new 100 gigabit/second transatlantic link.
http://www.canarie.ca/canadian-physicists-achieve-100-gigabitsecond-transatlantic-transmission-enabled-by-canarie-and-its-global-partners-2/

Our activities in computing have been highlighted by industry:

In 2017, Microsoft made a video on our use of the Azure cloud resources:
https://youtu.be/5Kixh66AJNM

Brocade published an article on our work with 100G networks and fast storage systems:
https://www.brocade.com/content/dam/common/documents/content-types/success-story/brocade-uvic-ss.pdf

San Disk published a web story and made a video of our project:
http://www.sandisk.com/about-sandisk/press-room/press-releases/2014/sandisk-helps-accelerate-high-energy-physics-research-on-the-origin-of-the-universe-supercomputing-2014/
https://www.youtube.com/watch?v=CiUZITuYjoQ