Computing proposal for the T2K and Hyper-Kamiokande experiments in Japan

IRFU/DPhP, LLR and LPNHE Neutrino groups

November 25, 2019

Abstract

This document describes a proposal for computing requests to CC-IN2P3 for T2K and Hyper-Kamiokande, the current and next generations long-baseline neutrino oscillations experiments based on the Water Cherenkov technology. French groups are actively working on the currently running T2K experiment, primarily on the near detector. Hyper-Kamiokande has been recently approved by the MEXT (Japanese Ministry of Education, Culture, Sports, Science and Technology) and will be built in Japan in the next decade. French groups are working on defining potential contributions to this promising experiment. As computing is an important element for the success of these experiments and CC-IN2P3 is playing a key-role for LHC and many other experiments, a French contribution to the T2K and Hyper-Kamiokande computing effort seems natural and desirable.

Contents

1 Introduction 2

2 T2K and Hyper-Kamiokande Computing Model 2

3 Computing Needs 3
- 3.1 T2K Phase II . . . 3
- 3.2 Hyper-Kamiokande first stage: construction . . . 4
- 3.3 Hyper-Kamiokande second stage: exploitation . . . 4
- 3.4 Storage and file replication . . . 5

4 Proposal to CC-IN2P3 6
- 4.1 Goal of the request . . . 6
- 4.2 Scenario 1: T2 site for ND280 data . . . 6
- 4.3 Scenario 2: T1 site for ND280 data . . . 7
- 4.4 Scenario 3: T1 site for all HK data . . . 8
- 4.5 Requested support . . . 8
- 4.6 Costs estimations . . . 9

5 Conclusions 9

1 Introduction

T2K is a long-baseline neutrino oscillation experiment currently taking data in Japan. The composition and energy of a muon neutrino beam produced at J-PARC in Tokai is precisely measured using a suite of near and far detectors, described in Table 1. Hyper-Kamiokande is a next-generation experiment that will be built in Japan starting in April 2020 and will start data taking in 2027 [Abe *et al.*(2018)].

INGRID and ND280 are currently taking data for the T2K experiment and will continue the data taking for the next 20 years. A new near detector, the Intermediate Water Cherenkov Detector (IWCD), will be built at about 2 km from the neutrino target and will consist of a water Cherenkov tank covered with photomultiplier tubes (PMT); this detector can be moved to be on and off the beam axis.

Finally, a new far detector called Hyper-Kamiokande will consist of an underground water Cherenkov detector located 8 km South of the current Super-Kamiokande detector used as far detector for T2K. This detector will be instrumented with at least 20,000 high-efficiency 20" PMTs (rouhgly two thirds more than Super-Kamiokande), providing a 20 % photocoverage. To complement these 20" PMTs, the option to add a few thousands of multi-PMT modules (mPMTs) is being actively studied. Another 8,000 20" PMTs will be used to monitor the outer detector surrounding the main Hyper-Kamiokande vessel. What we will call Hyper-Kamiokande experiment therefore corresponds to the near detector suite and the new far detector.

During the construction and exploitation of the detector suite, several simulations will produce data that will need storage (more details about the expected required space are provided in Sec. 3). Starting from 2027, data files will also be produced by these detectors. Because of the size of the detector and the number of detectors, the Hyper-Kamiokande far detector will produce a significant fraction of data that will need to be stored and processed. It doesn't seem feasible to store and process all the produced data into one single site (e.g. at the Kamioka Tier-0 site as it is the case for Super-Kamiokande).

For these reasons, we propose to use the computing resources of the CC-IN2P3 as part of T2K and Hyper-Kamiokande experiments computing schemes. This document decribes the computing model envisonned by the T2K and Hyper-Kamiokande collaborations (Sec. 2), the expected global computing needs (Sec. 3) and finally the proposal to CC-IN2P3 (Sec. 4). Most of the material of this document is based on [Di Lodovico et King(2019)] and internal discussions.

2 T2K and Hyper-Kamiokande Computing Model

Similarly to LHC, ND280 has used a tiered system since 2010 composed of layers, or *Tiers*, 0, 1, and 2. Figure 1 represents the computing model for the ND280 detector. A T0 site is at KEK where the data are storeed on HPSS/disk combo storage system. The raw data are transferred to RAL and Triumf as T1 sites. Recently, Triumf decided not to contribute to the ND280 computing anymore; as a consequence, a new T1 site is highly desired by the collaboration. Data reconstruction and MC generation are run on T1 and T2 sites and then copied on at least one T1 site and one or more T2 sites. A set of python scripts and DIRAC utilities is used for data transfer and registration.

Hyper-Kamiokande also proposed to use a tiered computing model. Figure 2 shows a scheme of the proposed Hyper-Kamiokande dataflow. Data collected by the detectors in Japan are stored in the two T0 sites (Kamioka and KEK) and copied to T1 sites after reduction. Data calibration and reconstruction along with MonteCarlo simulations will be performed at T1 and T2 sites. Data and workload management be-

Name	Distance along beam	Angle wrt. beam	Expected data rate
INGRID	280 m	$0°$	78 GB/day
ND280	280 m	$2.5°$	214 GB/day
IWCD	2 km	$0° - 4°$	170 GB/day
Far detector	295 km	$2.5°$	5 TB/day

Table 1: List of detectors and primary characteristics. The near detectors (INGRID, ND280 and IWCD) will be used both for T2K and Hyper-Kamiokande. The far detector expected data rate corresponds to the new far detector built for Hyper-Kamiokande.

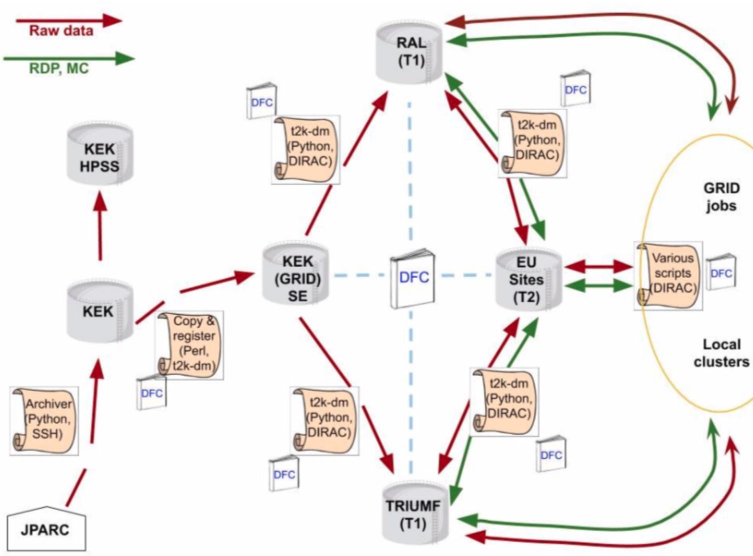

Figure 1: ND280 Data Flow Diagram for raw data and reconstructed data (RDP) and MC. The T0 site at KEK and the T1 sites (RAL and Triumf) are depicted.

tween sites is done using the DIRAC (Distributed Infrastructure with Remote Agent Control) framework [Tsaregorodtsev *et al.*(2008)]. T2 sites should have enough computing and storage resources to proccess and store part of these data or produce events simulation and reconstruction.

In this scheme, we propose to include CC-IN2P3 as a T1 or T2 site, that will therefore store part of data produced by T2K and Hyper-Kamiokande. Let's point out that currenly CC-IN2P3 already host part of the ND280 data but it is not integrated into the T2K tiered model; efforts are currenly being made to have resources allocated to T2K integrated within the T2K DIRAC system. Hyper-Kamiokande could therefore benefit from the momentum of T2K computing-related work.

3 Computing Needs

Given the timescale of each experiment, the computing resources requirements are split into three parts. The T2K experiment will continue taking data until 2026. Starting in 2020, a 7-year long period corresponds to the construction of the Hyper-Kamiokande far detector and IWCD. In 2027, when the Hyper-Kamiokande far detector comes online, part of the T2K resources will be transfered to the Hyper-Kamiokande collaboration, including the data taken so far by the ND280 detector. The second stage for Hyper-Kamiokande with an expected start in 2027 should last about 10 years: during this period, all the detectors will be producing data that will need processing for the Hyper-Kamiokande experiment.

Below we give first estimations of the computing needs during these three items and how they relate to one another. Let's point out that the computing needs will last for more than 17 years: it is therefore expected that frequent usage-based updates on the present request proposal will be made. This initial request is supposed to be conservative though.

Let us also point out that the storage needs described below correspond to a single copy of each file: we describe more precisely the file replication model in Sec. 3.4.

3.1 T2K Phase II

The T2K experiment is curently taking data. The near detector suite (ND280 and INGRID) is taking about 8 TB/year[1] and the total amount of data collected so far is 125 TB. Starting in 2022 following the ND280 detector upgrade, it will record data at a rate of 32 TB/year. Annual raw data analyses and MC production are made: the amount of space and computing resources needed is about 9.6 CPU.hours and 0.27 5GB per 5×10^{17} POT. Each production is expected to be about 110 TB.

[1] This assumes a 150 days of operations each year with 200 GB/day as data rate.

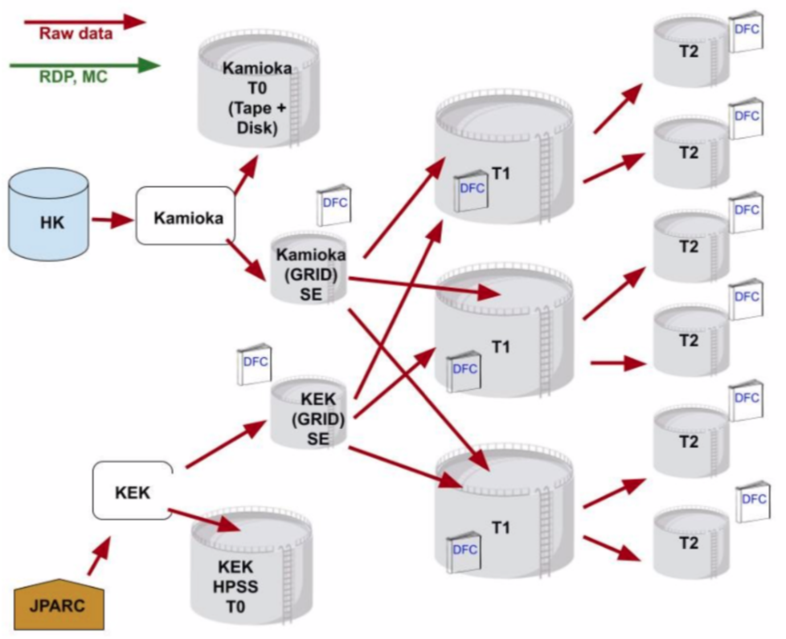

Figure 2: Hyper-K Data Flow Diagram for raw data and reconstructed data (RDP) and MC. The two T0 sites (Kamioka and KEK) are depicted. In this proposal, CC-IN2P3 would be one of the T1 sites.

3.2 Hyper-Kamiokande first stage: construction

During the construction of the IWCD and Hyper-Kamiokande far detector, no data will be collected by these two detectors. Therefore during this period, the only storage needs for HK would come from MonteCarlo studies. Note that the T2K data recorded by working detectors (namely INGRID and ND280) during this period will be stored on T2K dedicated resources.

The estimations in Table 2 were made for MonteCarlo simulations on the existing detectors and conservative projections for the new ones. For the ND280 detector, we expect 5 MC productions over the construction period with constant statistics (20×10^{22} POT per production, so ≈ 110 TB per production). The size of the MC is larger than the raw data, since they correspond to fake raw data but with larger statistics (roughly 20×10^{22} POT per production, compared with 2×10^{21} POT by the end of T2K in 2026)[2].

For the INGRID detector, only 2 MC productions are expected as this detector is not part of the oscillation fits for the sensitivity studies, leading to rather small resources needs compared with the other detectors.

3.3 Hyper-Kamiokande second stage: exploitation

Table 3 shows the estimated needs for storage and CPU during the data taking phase starting in 2027. The main item to Hyper-Kamiokande computing budget is related to the far detector. This estimation is based on Super-Kamiokande current data rate: assuming 50,000 PMTs, the data rate will be 5 TB/day. We can

[2]These MC productions will be different from the T2K ones, and therefore require additional space.

Detector	MC (CPU.hours)	MC Storage (TB)
INGRID	0.13M	7
ND280	19.2M	2,250
IWCD	97M	52
Far detector	20M	500
Total	136.33M	2,824

Table 2: Expected computing resources for the Hyper-Kamiokande experiment during the construction phase, considering only one copy of each file.

expect a decrease of this rate if the number of PMTs used in the tank also decreases. This will be confirmed in the next couple of years as the design of the experiment is finalized.

It is unclear whether the IWCD detector will run all year round or only periodically in coincidence with neutrino beam periods: in the former case, this will also allow IWCD to be sensitive to supernovae neutrinos at a cost of a significantly larger dataset (62 TB/year compared with 30 TB/year). The projected computing needs in this document correspond to the case where data are continuously recorded, leading to 620 TB for 10 years of data taking.

The rate of data collected by ND280 is expected to be 32 TB/year, meaning that the storage space needed to store the data collected by this detector between 2027 and 2037 will be about 320 TB. It is expected that at the end of the T2K experiment, the 445 TB T2K dataset will be shared with Hyper-Kamiokande leading to a total of 669 TB for this detector by the end of the data taking period in 2037. During the data taking period, we expect 7 MC productions with increasing size compared with the construction period.

The storage space needed by ND280 for MC data is about 10 times larger than the actual dataset because the MC data corresponds to full simulations of events in the ND280 detector and with larger statistics. For the far detector, since the amount of raw data is very large, only a small fraction of the simulated events would have a full detector response.

At the end of the data taking period, the amount of produced data (raw, processed and MC) will be about 28.6 PB.

3.4 Storage and file replication

As stated earlier, Tables 2 and 3 show the storage needs for only one copy of each file. Table 4 summarizes the file replication strategy accross the T0, T1 and T2 sites. Replication and data management will be done with custom tools created by the collaboration based on iRODS and the DIRAC framework.

Given the expected amount of data and the need for distinct locations for data replicas, several sites will be needed outside of Japan. Triumf which has been a T1 site for T2K does not intend to contribute as such for Hyper-Kamiokande. However Rutherford Appleton Laboratory (RAL) is requesting significant space and computing resources to be a T1 site, but it seems unlikely all the 28.6 PB could be stored there in several copies. Therefore having a secondary T1 site in Europe to store and process part of the total dataset would be highly beneficial.

The possibility to have two copies of the data at CC-IN2P3 has been discussed and seemed feasible. Such scenario is not detailled at the moment in this document, but it is possible to estimate the cost by roughly doubling the storage capabilities at any given time. To reduce the costs, the second copy should be kept on tape.

Detector	Data Storage (TB)	MC (CPU.hours)	MC Storage (TB)
INGRID	226	0.51M	26
ND280	669	42.2M	4,950
IWCD	620	684M	367
Far detector	18,440	25M	500
Total	19,955	751.71M	5,858

Table 3: Expected computing resources for the Hyper-Kamiokande experiment during the data taking phase, considering only one copy of each file.

Detector	Construction (MC)	Data taking (raw data)	Data taking (MC/proc)
INGRID	$\mathrm{T1_{IN}}+n\mathrm{T2_{IN}}$	$\mathrm{T0_{IN}}+2\mathrm{T1_{IN}}$	$\mathrm{T1_{IN}}+n\mathrm{T2_{IN}}$
ND280	$\mathrm{T1_{ND}}+n\mathrm{T2_{ND}}$	$\mathrm{T0_{ND}}+2\mathrm{T1_{ND}}$	$\mathrm{T1_{ND}}+n\mathrm{T2_{ND}}$
IWCD	$\mathrm{T1_{WC}}+n\mathrm{T2_{WC}}$	$\mathrm{T0_{WC}}+2\mathrm{T1_{WC}}$	$\mathrm{T1_{WC}}+n\mathrm{T2_{WC}}$
Far detector	$\mathrm{T1_{FD}}+n\mathrm{T2_{FD}}$	$\mathrm{T0_{FD}}+2\mathrm{T1_{FD}}$	$\mathrm{T1_{FD}}+n\mathrm{T2_{FD}}$

Table 4: File replication for the T2K and Hyper-Kamiokande raw data, processed data (proc) and MC production for the construction and data taking periods. T0 sites are in Japan, whilst T1 and T2 are outside of Japan. Different labels indicate the tiers for the different detectors.

The current plan for transfering and indexing data is to use iRODS[3]. This data distribution tool is routinely used at CC-IN2P3 and allows to access data stored on disks and on tape. It is therefore possible to store part of the data on differents supports depending on the usage frequency: for example, it is possible to use data stored on tape during annual production analyses even if they are accessed at a slower pace with respect to more expansive disks. Disks would then be used only for data used at higher frequencies (like the results of the most recent production and end-user analyses). Even if this is preliminary estimation, the costs estimation in Section 4.6 will attempt to include this distribution of storage.

4 Proposal to CC-IN2P3

4.1 Goal of the request

In this document, we propose a stronger involvement of CC-IN2P3 into T2K and Hyper-Kamiokande computing. Depending on the level of engagement desired by the IN2P3 directorate, three scenarii are suggested.

The first and minimal one consists in becoming ND280 T2 site which is our current plan of actions. The second Scenario consists in becoming ND280 T1 site. The last Scenario consists in becoming T1 site for all the Hyper-Kamiokande data and MC.

These three scenarii have different timescales of implementation which allow us to start ramping up efforts as experimental phases evolve. Moreover it is possible later to decide to increase our involvement into the overall computing effort, keeping in mind that a strong involvement at the present will make us more visible within both the T2K and Hyper-Kamiokande collaborations. Let us emphasize the fact that these scenarii are based on our current estimations of the computing and storage needed described earlier. **It is expected that we will need to reevaluate these every year or so, but our current and rather conservative estimations should not significantly increase with these reevaluations.**

4.2 Scenario 1: T2 site for ND280 data

Scenario 1 corresponds to our current course of actions which is to become a T2 site for the T2K and Hyper-Kamiokande experiments. Since Super-Kamiokande (far detector for the T2K experiment) has its own computing system, the only data we have to consider for T2K are the ND280 and INGRID data. We would share our resources with the T2K for the computing and storage of a subset of the data[4]. Our current T2K allocation is about 60 TB storage and 20 MCPU.hours computing each year. A limited amount of additional resources would be requested: with approximately an extra 200 TB storage[5] and light storage increment each year (about 40 TB), we could store the raw data and have the resources necessary to contribute to each MC production and locally store the results of the most recent production. Given that this storage will be primarly for storing raw data which are needed only once a year for production, it is possible to put these all on a more cost effective support such as tapes and only keep the current allocation on disk for end-user analyses[6].

On top of this, we request some startup storage and resources for Hyper-Kamiokande that would be included into Hyper-Kamiokande computing model.

Table 5 summarizes the storage and computing resources requested in this Scenario. This contribution is rather minimal with not a lot of visibility within the collaboration. But we emphasize once more that such minimal scenario could be a starting point for another scenario. For instance, this scenario would allow us to later contribute to computing if a stronger commitment to T2K and Hyper-Kamiokande is envisonned by the IN2P3 directorate. However, it is essential that 2020 requests are met in order to continue our current work in integrating our resources into existing collaborative frameworks, both for T2K and Hyper-Kamiokande.

[3] https://doc.cc.in2p3.fr/irods

[4] Let us point out that work at IN2P3 is being done to enable such sharing. Tools like DIRAC allow the shared resources to be managed in coordination with the rest of the collaboration.

[5] As a T2 site, there is no need for backup storage.

[6] Each production weights roughly 100 TB. Therefore we could store only the latest production on disk and keep all the raw data on tape.

4.3 Scenario 2: T1 site for ND280 data

Scenario 2 is an extension of Scenario 1, where additional allocated storage and computing resources would allow us to become a T1 site for the ND280 and INGRID[7] data for both T2K and Hyper-Kamiokande. As stated earlier, such decision could be made within the next year or two, even though an early decision would put us in a better position and increase our visibility.

In this scenario, we request approximately 170 TB in 2020 for the T2K near detector production and 100 TB dedicated to Hyper-Kamiokande as a startup contribution for MC studies for ND280. From there, we will request more space each year for hosting new raw data and the previous and current productions results for T2K (currently, T2K has produced 6 sets of MC dataset) and storage for the results of the Hyper-Kamiokande productions related to ND280. The initial computing request of 5 MCPU.hours should be sufficient to cover a partial contribution to the MC productions; after the end of T2K-II, the T2K resources will be merged with Hyper-Kamiokande's for MC productions and end-user analyses using ND280 data. This scenario is described in Table 6.

These resources will, like Scenario 1, be shared with the rest of the collaboration using DIRAC.

In this scenario, only the most recent MC production results should be stored on disks as they will be frequently used for end-users analyses: this would represent between 110 TB for the first productions and 275 TB for the last Hyper-Kamiokande productions. The raw data and the older productions results could then be stored on tape, reducing the overall cost.

In Table 6, we have separated T2K requests from Hyper-Kamiokande requests as they are viewed by IN2P3 directorate as two different projects. Note however that at the end of T2K Phase II in 2027, the storage and computing resources for T2K and Hyper-Kamiokande will be merged, leading to about 25 MCPU.hours available each year for ND280.

To our mind, this scenario makes a lot of sense given our current involvement in the ND280 detector upgrade and exploitation: the request remains by itself moderate and slowly ramping up. It also gives us more time to refine requests predictions, especially for the new detectors Hyper-Kamiokande and IWCD: in the next couple of years, first productions will be made for these detectors and some data will be collected with mPMTs prototypes for IWCD.

[7] Given the required space and resources for INGRID compared with T2K, we assume that we have enough space to store INGRID data and MC along with ND280's.

	T2K				Hyper-Kamiokande			
Year	Requested MCPU.h	Total MCPU.h	Requested space (TB)	Total space (TB)	Requested MCPU.h	Total MCPU.h	Requested space (TB)	Total space (TB)
2019	0	20	0	59	0	0	0	0
2020	0	20	170	229	3	3	100	100
2021	0	20	40	269	0	3	0	100
2022 → 2036	0	20	40	309 → 869	0	3	0	100

Table 5: Computing and storage requests in Scenario 1. This requests will cover the storage of ND280 raw data and partially cover the storage of the most recent production results. From 2022 until the end of Hyper-Kamiokande data taking period, the request for storage will be of about 40 TB each year.

	T2K				Hyper-Kamiokande			
Year	Requested MCPU.h	Total MCPU.h	Requested space (TB)	Total space (TB)	Requested MCPU.h	Total MCPU.h	Requested space (TB)	Total space (TB)
2019	0	20	0	59	0	0	0	0
2020	0	20	170	229	5	5	100	100
2021 → 2022	0	20	40	269→ 309	0	5	350	450→ 800
2023 → 2026	0	20	40	349→ 469	0	5	500	1,300→ 2,800
2027 → 2036	0	Merged with HK	0	Merged with HK	0	25	540	3,809 → 8,669

Table 6: Computing and storage requests in Scenario 2 for T2K and Hyper-Kamiokande. We separated the construction and data taking periods: at the transition from the Hyper-Kamiokande construction to data taking, the resources dedicated to T2K will be merged with Hyper-Kamiokande's. Starting in 2027, Hyper-Kamiokande requested space includes the 40 TB needed for storing new raw data of ND280.

Therefore we expect that computing and storage needs for Hyper-Kamiokande will be reevaluated within 3-4 years; it will be necessary to then evaluate a possible change to another Scenario e.g. Scenario 3.

4.4 Scenario 3: T1 site for all HK data

Scenario 3 consists in becoming a T1 site for both T2K and Hyper-Kamiokande experiments. This means being able to store volumes given in Tables 2 and 3 for Hyper-Kamiokande and described in Section 3.1 for T2K. The time profile are presented in Table 7.

Several things are to be considered for this Scenario. First, most of the storage support would be tapes: as we have explained above, a majority of the data here corresponds to raw data that will be accessed about once a year during production campaigns. Once the most recent productions for all the detectors should be kept on disk: this represents at most 275 TB for ND280, about 50 TB for IWCD, and 50 TB for the far detector. Additional storage for end-users analyses should also be available. **To summarize, by the end of 2037, the total amount of data on tapes would be about 24.5 PB while about 400 TB of data would be stored on disks.**

Concerning the computing reesources, the projections given in Table 7 assume that data processing and MC productions will be shared between RAL and CC-IN2P3 since no other T1 and T2 sites have been identified so far. It is clear that, as additional international contributions to computing are being identified over the next couple of years, these computing resources will be reevaluated to something more appropriate. We should therefore consider this part of the Scenario as an upper limit on the overall request.

Given the current capabilities of CC-IN2P3 and the fact that our request extends over 17 years, becoming a T1 site for Hyper-Kamiokande is definitely possible in terms of technical capabilities. This in-kind contribution would be highly visible within the collaboration and could be viewed as an equivalent international contribution to Hyper-Kamiokande.

4.5 Requested support

Hyper-Kamiokande collaboration has already started developping tools for transfering data between sites and has cumulated some experience on DIRAC framework. A collaboration between Hyper-Kamiokande software development team and the DIRAC development core team is on going.

Moreover, thanks to the Jennifer-II European funding, Hyper-Kamiokande and Belle-II collaborations are working together to build common tools for distributed computing, storage management and software development and deployment.

Another possible contribution could be to host Hyper-Kamiokande software. T2K software is currently hosted on a GitLab server in Poland. Work is being done for getting Continuous Integration (CI) support by collaborators at LPNHE. Since these two items and associated expertise already exist at CC-IN2P3, we could consider having dedicated servers for both software and CI hosted at CC-IN2P3 with replicas in other countries for redundancy purposes.

	T2K				Hyper-Kamiokande			
Year	Requested MCPU.h	Total MCPU.h	Requested space (TB)	Total space (TB)	Requested MCPU.h	Total MCPU.h	Requested space (TB)	Total space (TB)
2019	0	20	0	59	0	0	0	0
2020	0	20	170	229	5	5	100	100
2021 → 2022	0	20	40	269→ 309	5	10→15	350	450→ 800
2023 → 2024	0	20	40	349→ 389	5	20→25	500	1,300→ 1,800
2025 → 2026	0	20	40	429→ 469	5	30→35	500	2350→ 2,900
2027 → 2036	0	Merged with HK	0	Merged with HK	0	55	2,150	5,519 → 24,869

Table 7: Computing and storage requests in Scenario 3 for T2K and Hyper-Kamiokande. We separated the construction and data taking periods: at the transition from the Hyper-Kamiokande construction to data taking, the resources dedicated to T2K will be merged with Hyper-Kamiokande's. Starting in 2027, Hyper-Kamiokande requested space includes the 40 TB needed for storing new raw data of ND280.

4.6 Costs estimations

TODO: Estimation des couts sur les X prochaines années

5 Conclusions

References

[Abe *et al.*(2018)] K. Abe et al. (Hyper-Kamiokande)*Hyper-Kamiokande Design Report*, arXiv:1805.04163 [physics.ins-det] (2018).

[Di Lodovico et King(2019)] F. Di Lodovico et S. King, *Tiered computing model*, Hyper-Kamiokande Internal Report (2019).

[Tsaregorodtsev *et al.*(2008)] A. Tsaregorodtsev, M. Bargiotti, et al., *DIRAC: a community grid solution*, Journal of Physics: Conference Series **119** (6), 062048 (2008).