

## Discussion Document

# Expectations for of DUNE Physical Computing Resource Provision by Collaborating Nations

Draft 5 – September 2019

## 1. Purpose of document

The DUNE (+protoDUNE) experiment is under construction. Data taking preparation will start in earnest in ~2024 with the advent of the first far detector module. Preparatory studies and Monte Carlo production are already underway, and protoDUNE data already exists. A second protoDUNE run is planned for 2021-22.

DUNE needs to formalize how computing capacity will be provided across its international partners.

FNAL management seeks to internationalise computing for the next generation experiments such as DUNE. FNAL expectation is that ~ 25% of computing resource should be provided by FNAL as the host lab, and the remaining 75% provided by partners external to FNAL.

This document provides a discussion of the means to achieve this in the short and long term.

*Status: This is a discussion document. It has at this point:*

- (i) Been iterated within Computing management*
- (ii) Discussed with the Spokespersons*
- (iii) The next stage is to informally discuss the ideas with countries*
- (iv) Following this the matter will be presented to the IB*

*The reader should therefore take it as implicit that any statements or principles are only “proposed” at this stage.*

## 2. Scope

**DUNE needs several forms of computing resources. They are:**

- 1) CPU resources for simulation, reconstruction and analysis of DUNE data.
- 2) Storage resources dedicated to DUNE data
- 3) Professional personnel for development and maintenance of core computing infrastructure
- 4) Dedicated personnel to manage and operate DUNE computing

This document is mainly concerned with the first two “physical” resource types.

### 3. Principles

#### **Principle 1: Physical Computing resources are a collective responsibility**

Given the ever more “internationalised” nature of DUNE, we propose to follow the practice of the LHC experiments, with physical computing resource provision being a collective responsibility shared between collaborating countries subject to some form of “fair share” expectation.

There is already an assumption in this statement: namely that the unit of aggregation for determining a fair share is a country (as opposed to an institute). We believe this is the correct thing to do.

*The alternative would be for computing to be considered like a sub-detector, and as such signed up to only by those who want to work on that sub detector. This would result in computing being provided only by a subset of institutes (not countries) on a voluntary basis. We do not consider this viable, and it has rarely (never?) worked elsewhere for experiments of the scale of DUNE. We do not consider this further here.*

#### **Principle 2: DUNE does not seek to formally oblige partners/nations to contribute at this stage. The matter is left as a point of good citizenship.**

We feel that it is too early to seek to make computing provision the subject of formal obligation. It is far better at this stage to make reasonable expectations known, and then rely upon each country to do its best to play its part. For this reason, this document only refers to *expectation of national contribution*. This is often found helpful by national contingents in order to demonstrate a “fair ask” into their respective funding processes.

The position can be revisited nearer to data taking and could become a “harder” expectation.

### 4. Computing Tiers

It is probably useful to follow LHC “tiers” convention and recognise different types of site. We provide a rough guide here; detailed criteria will be enumerated in due course.

- **Tier-0** is the Host lab (FNAL), which a priori has unique functions (e.g. first destination and responsibility for custodial storage of raw data)
- **Tier-1** sites offer CPU, Disk and tape services. They satisfy a minimum level of service agreement and response times to be defined (similar to LHC). Tier-1 resources will be promised (or “pledged”). These sites are expected to have of order of several 10s Gb/s network connectivity now rising to ~ 100 Gb/s by the start of data taking.
- **Tier-2** offer CPU and Disk. A Tier-2 will provide a specified service level and response time, but which will be less onerous than a Tier-1. Tier-2 resources will be promised (or

“pledged”). These sites are expected to have of order of several 10s Gb/s network connectivity by the start of data taking, and in some cases ~ 100Gb/s

- **Tier-3** any other site which will contribute but does not meet Tier-1 or Tier-2 criteria. Tier-3 sites may, for example, provide a lower service level and may focus on CPU capacity only. Tier-3 sites may not be able to pledge resources, and may be considered to be opportunistic.

Having defined these Tier levels, the current thinking is that DUNE need not differentiate strongly between Tier-1 and Tier-2 contributions from partner countries. Thus, the likely stratification will be

- Tier-0 as defined above
- Tier-1 and Tier-2 provide promised (pledged) resources external to FNAL
- Tier-3 for any other contributions not in the above and probably not pledged.

The DUNE production computing team can plan production computing based on this stratification.

## 5. Who sets and agrees the DUNE resource request.

DUNE will estimate its computing resource requirements according to its computing model. These should be formal documents that will be reviewed annually.

Naturally it cannot be that DUNE may request an arbitrary unlimited amount of resource. Some form of independent scrutiny process is needed to “*keep DUNE honest*”, i.e. to ensure that the DUNE resource requirements are reasonable. FNAL or DUNE management should probably instigate some form of scrutiny body.

The agreed output of this body becomes the approved DUNE resource requirements.

## 6. Metric for fair share

The most immediately obvious *Fair share Metric (FSM)* would be based on M&O or common fund share. This is may to be based upon PhD *authors*.

It is not important for this document to specify exactly what this FSM will be, but simply to assume that there will be one.

## 7. Long term expectation of national contribution (for discussion)

The most obvious viable long-term way to provide the required computing capacity is to expect each collaborating country to provide a share based upon the agreed FSM.

The Tier-0 contribution is special as discussed earlier, and may be around ~25%. The residual requirement is then to be provided by all other DUNE members excluding FNAL.

The overall share for country-X,  $F_x$ , would then naturally be defined as”

$$F_x = \text{FSM}_x / \text{SumOf}\{\text{FSM of all of countries, excluding FNAL}\}$$

This would be used to determine the total CPU and disk resource contribution expected from each country.

Then Tier-1 hosting countries may then choose to provide their contribution as Tier-1 services and/or Tier-2 services as they see fit.

Non-Tier-1 hosting countries will provide their contribution as Tier-2 services (or possibly Tier-3).

Simply agreeing and making this policy known now will, we hope, help countries in discussions with their funding agencies in planning for future operations.

## Tape

Tape will be dealt with separately as only Tier-1 hosting countries can provide tape services. It may be necessary to offset this against their CPU and Disk expectation but in practice we expect Tier-1 countries will be happy to provide this in addition. Further discussion needed here

## Open questions

- **USA:** This raises the question of whether US institutes obligation should be offset by the FNAL host lab contribution, at least in part ? Would this then have an unintended consequence of inhibiting US Tier-2s wishing and able to contribute?.
- **CERN contribution:** The above does not naturally apply to CERN. We do not have a proposal at this point, other than to suggest bi-lateral discussions between DUNE/FNAL and CERN to determine that which CERN wishes to contribute. This, along with the FNAL contribution, is removed from the total requirement to set the residual Tier-1 and Tier 2 requirement. Of course, one way might be to simply use CERN authors, This needs further discussion with CERN.
- **Regional consortia:** Assuming there will always be a set of very small countries (e.g. < 10 authors) it may not make sense to expect each of these to provide a separate computing centre. We may therefore expect groups of very small, or regionally related, countries to form a consortium for the purposes of computing provision. This should be the basis of common-sense discussion.

## 8. Short term pragmatic expectation

The proposal of the previous section may well be the correct thing to aim for when DUNE is in steady data taking in several years (operations mode). It is certainly important to establish this principle early on so that all partners will have this in mind when negotiating with their funding agencies. Also, so that this expectation is made manifest to new institutes joining.

However, it has two problems:

- (i) We are not in steady state. We are in construction. This means the number of PhD authors is low and probably imbalanced (i.e. not reflecting that which it will be in steady state). It may therefore be difficult to know the appropriate FSM at this stage.
- (ii) Based upon current DUNE members in the Database, this would mean that some countries who have already indicated their willingness and ability to provide a substantive share of computing would be asked to reduce this (as they have a very small number of people in the DB). It makes no sense to turn such countries away.

Therefore, we propose a much more pragmatic approach in the short term as an interim measure for the next few years:

Countries with a large pre-existing computing infrastructure (e.g. that in place for the LHC), are asked to seek to be able to provide at least  $>\sim X\%$  of the overall DUNE computing requirement at Tier-1 and/or Tier-2 level. This is regardless of their actual FSM calculated at this point in time.

The value of X needs to be agreed, but we imagine this to be at least 10% for Tier-1 hosting countries (but may be more). At this level, this will not require a huge commitment from each country in the next few years.

Such a minimum threshold is suggested to enable the production team to work with a limited number of organisational entities for the bulk of the production work. The load on the team scales with number of organisational entities.

At the same time resources from other countries that do not historically have a large compute infrastructure are welcome, and such may allow such smaller countries/consortia to build their capability towards operations.

To take this forward, informal discussions will take place with countries to seek their opinions and likely positions. Consensus will be sought, and if this is possible a pragmatic proposal will be made.

## 9. Software staff in lieu

It is probably worth discussing a problem that is bound to arise for DUNE. That is the traditional problem that partner countries will not, on aggregate, furnish the necessary number of software engineering staff to construct the Computing project (software infrastructure) or to run computing operations. This tends to be a practical fact for several reasons including that institutes tend to be happy to apply for staff for detector components “close to their heart” as it supports their local expertise, but few find that computing is seen to be to their long-term benefit. It is also often the case that grant applications for software staff are hard to get awarded. Perhaps we will be proven wrong in DUNE, which will be great, but in case not then the following may help.

Some countries/institutions may not have access to resources sufficient to meet their fair share obligation of hardware resources. However, it may be that they find it easier to obtain staff support. In this case contributions of software developer/operations staff in lieu of physical hardware may be an appropriate substitute.

## 10. Computing Board.

The LHC experiments typically constituted a *Computing Board* (called different things in different experiments). The requirement is that the members are national liaisons who are in a position able to interact effectively with their respective funding agencies. Naturally there will also be additional members or in-attendance people. The Chair should be able to interact with the Experiment and lab management.

DUNE proposes to constitute such a **Computing Contributions Board** (CCB) within the Computing organization.

The prime charge of this body is to ensure that the required computing resources are available to DUNE. It will also deal with other matters as may arise (generally of policy or strategy, but not technical or operations)

This body will need to have an appropriate relation to the DUNE RRB.

The formation of this body should be presented to, and endorsed by the IB. Then each nation (or regional consortium) would nominate its member.