

# Understanding your Allocation

## On this page:

- [Introduction](#)
- [Job Priorities](#)
- [Service Units Explained](#)
  - [Magnus](#)
  - [Zeus](#)
  - [Topaz](#)

## Introduction

On all Pawsey supercomputers, compute time is budgeted in quarterly periods, irrespective of the duration of the project. Typically this is achieved by uniformly distributing the Service Unit (SU) allocation across relevant periods. For example, a project that is awarded 1,000,000 SUs during January--December in a calendar year would be credited with 250,000 SUs in each of the four quarters.

This breakdown of time is implemented in order to promote a more uniform demand on our systems. The supercomputers are uniform in capacity over time, and time is linear. Past experience tells us that supercomputers are subject to excess levels of demand, and consequently significant delays in running jobs, at the end of the calendar year. With this in mind, unspent SUs from one quarter are not carried over into the subsequent quarter. However, a project can continue to run computations on our resources even after its SU budget is exhausted if resources are available. We implement a fair-share system that means such jobs run, though are queued with a lower priority than is normally the case.

Time awarded on Pawsey supercomputers include a factoring in of four weeks of downtime per year. Since downtime affects all projects, *individual projects cannot be compensated for any additional downtime*. Pawsey aims to minimise the downtime of all its services through careful planning and execution of scheduled maintenance.

Our intent is to achieve fair access to Pawsey supercomputers for all of our users, and to promote access to resources in a manner that makes effective use of the available computing capacity.

Commercial projects have some minor differences. Only usage is charged, not allocation, so downtime does not affect charging. Unused time in each quarter is also not charged for. Additional quarterly allocation may be granted if it is available.

See [Project Accounting](#) for how to see usage and allocation information of your projects.

## Job Priorities

Pawsey uses the fairshare algorithm in its job prioritisation. This takes into account the quarterly usage relative to the quarterly allocation. A project that is close to using up its allocation is a lower priority than one that is just starting. Fairshare aims to have all projects be at 100% usage of their allocation for the quarter. At the end of each quarter, usage and therefore fairshare are reset.

To ensure that Pawsey infrastructure is not idle (to maximise science outcomes), merit projects that have exceeded their quarterly allocation enter a low priority mode, so will run if there are no higher priority jobs. In this mode the projects can use more than their quarterly allocation. The low priority mode is reset at the start of every quarter. Low priority mode is only applicable to merit projects.

The usage of Director's Share projects cannot exceed their allocation. If more time is required, contact the Pawsey [Helpdesk Service](#).

## Service Units Explained

All projects that use the Pawsey compute resources are allocated Service Units (SU) through one of the several merit allocation based schemes. How service units are defined and consumed will depend on which system you are using. In general, 1 SU is equal to 1 core hour.

### Magnus

SU on Magnus are based on "one core per hour of walltime". However, you should note that the small job that can run on Magnus requires using the entire node and each node on Magnus has 24 cores. Any job that runs on Magnus will consume SU based on the usage of the entire node regardless, as it is not possible to share a node on the Cray with different users. Ideally all jobs on Magnus should use some multiple of 24 cores per job. A single Magnus service unit is defined as:

```
1 node per hour of walltime = 24 SUs
```

SU are charged based on the fraction of actual walltime used not on the wall time requested in your job scripts. If you specify a walltime of 5 hours on 96 cores (4 nodes) but the job finishes in 2 hours 30 minutes and 10 seconds then your allocation is charged 240.27 SU.

## Zeus

The workq contains CPU-only nodes, and these nodes can be shared. You are only charged for the CPUs allocated to the job, rather than the whole node. There are 28 CPU cores per node, so if you do use the full node this is slightly higher than the Magnus charging rate.

```
1 cpu core per hour of walltime = 1 SU
1 node per hour of walltime = 28 SUs
```

## Topaz

The gpuq and gpuq-dev contains nodes which have two Volta GPUs, and these nodes are shared. Each Volta GPU is charged the equivalent of 36 CPU cores. The number of CPU cores used in these jobs is not included.

```
1 GPU per node per hour of walltime (gpuq, gpuq-dev) = 36 SUs
1 GPU node (with 2 GPUs) per hour of walltime (gpuq, gpuq-dev) = 72 SUs
```

The nvlmq and nvlmq-dev contains nodes which have four Pascal GPUs, and these nodes are not shared. Each Pascal GPU is charged the equivalent of 36 CPU cores. The number of CPU cores used in these jobs is not included.

```
1 GPU node (with 4 GPUs) per hour of walltime (nvlmq, nvlmq-dev) = 144 SUs
```

Please contact the Pawsey [Helpdesk Service](#) if you have any questions, comments, or suggestions.

**Pages in this section:**