Proposed new NCI Allocation Model

Background

To ensure a fair utilisation of the HPC resources between research projects, NCI uses the concept of a Service Unit (SU) to parcel up the time available on the supercomputer. Research projects are then allocated an account containing a number of SU to be used within a three-month period (quarter).

When a researcher within a project submits a job to the supercomputer, the job scheduler calculates the size of the job in SU and checks if the project has sufficient SU within their account. If they do, the job is scheduled to run. At the end of a job, the total resources used by the job is calculated and decremented from the project's allocation. When a project runs out of SUs, researchers within that project are no longer able to run jobs on the supercomputer.

Problems with the past allocation model

Over the past few years, it has become apparent that there needs to be greater flexibility in the accounting system.

For scheme managers, trying to maximise the usage of their SU allocation, they need to be constantly checking usage across multiple projects and then requesting transfers of SU via the help desk from one project that is underutilising to another project that needs more. This process can be very time consuming and potentially frustrating for researchers if a project runs out of SU before additional SU can be put into their account. If they fail to perform this juggling act, they run the risk of the resources of the organisation being underutilised and wasted.

For researchers, the allocation process can be problematic particularly if the research project requires just a little more storage or needs access to some cloud computing resources for a short period of time as part of their research workflow. Without access to these resources, work has to stop and the time-consuming process seeking approvals or requesting the scheme manager eek resourcing from other projects begins. The net result is that SU that had been allocated to the project goes unused and potentially lost at the end of the quarter while they wait.

The proposed new accounting model attempts to address these issues by creating an SU value associated with all resources and then allowing scheme managers and researchers some limited flexibly to allocate or redistribute resources subject to availability automatically.

To explain how the new proposed allocation scheme will work, the model is described in 4 parts:

- Scheme Manager allocation

- Lead CI allocation
- Storage allocations and accounting
- Cloud Allocations and accounting

Scheme Manager Allocations

The new proposed scheme offers several additional features. The first if these is the ability to purchase additional temporary storage for projects by converting SU resources. This now means that a Scheme manager can use SU resources to allocate to projects, keep some unallocated in reserve and use them for project storage.

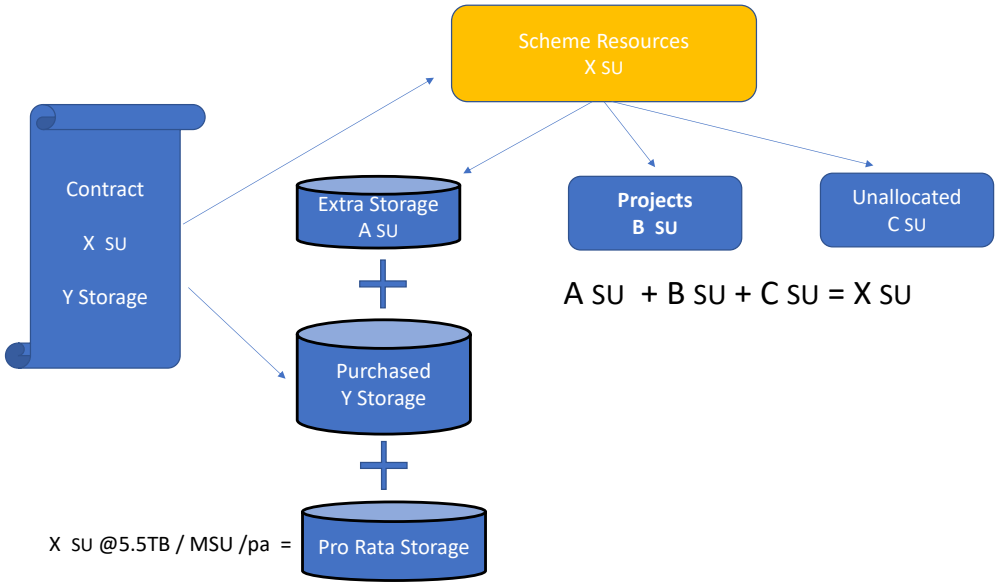

Figure 1 Institutional contract breakdown into resources

Allocating to projects

The proposed changes will allow Scheme managers to provide a fixed and a variable SU allocation to a project as well as direct what the SU resources may be used for within a project.

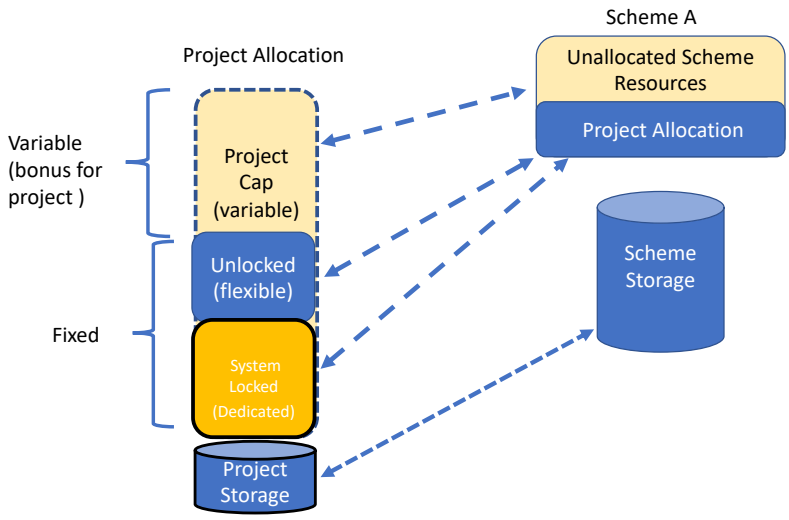

Figure 2 Allocation of resources to a project

The fixed allocation is like the current allocation procedure with each project assigned a share of the total SU pool. If 100% of the scheme resources are allocated to projects, then the situation is identical to the current scheme allocation process.

If, however not all of the scheme allocation has been allocated to projects then, a variable allocation can be made to projects allowing them to go beyond the fixed allocation and use the shared pool of resource up to the limit of the variable allocation IF it is available.

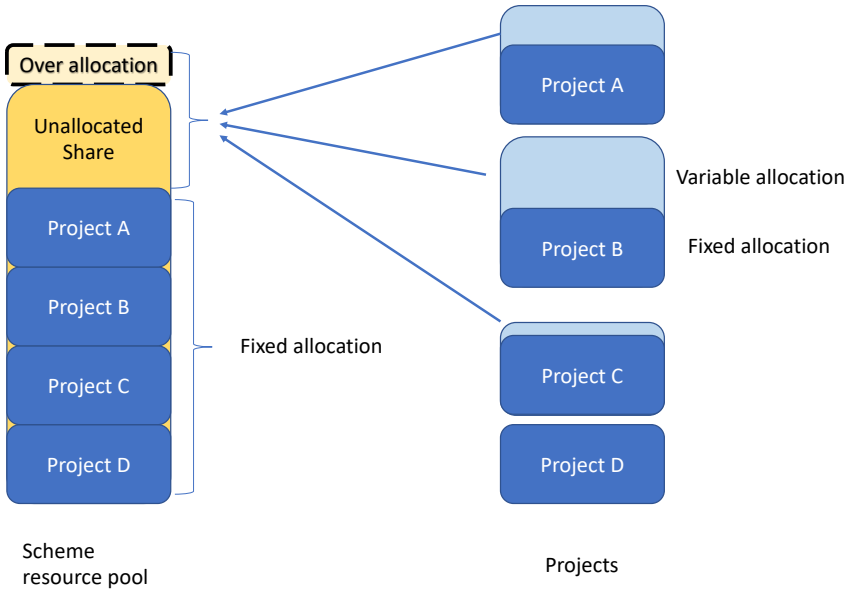

Figure 3 Scheme SU resource pool available to be allocated out to projects

The figure above shows the total SU resource pool for a scheme manager. They have allocated out resources to 4 projects but have not used all the allocation. As shown, Project D in the figure has only a fixed SU allocation and has not been given a variable allocation so will need to contact the scheme manager if they need more resources

The remaining projects A to C have been given a variable allocation and are able to draw on the remaining unallocated SU share from the scheme pool. From the diagram not all projects have been given the same variable allocation and the total of variable allocations assigned to all the projects and the fixed allocations exceeds the total Scheme allocation.

Access to the unallocated pool is on a shared "first come, first served" basis. This means that although some of the projects have been given a variable allocation, the projects are not guaranteed to be able access any of the additional resource if the scheme runs out of unallocated SU resources before they access it. Projects are only guaranteed to get access to their fixed allocation.

Projects A - C will only get access to the shared pool once they have exhausted their own fixed allocation. If project B for example does not use all its fixed allocation, then it will not draw down from the shared pool so leaving it available for Projects A and C to use.

As can be seen in the diagram Project C has a small variable allocation that will cap how much the project can use of the unallocated share. Once Project C reaches this cap, researchers in this project will not be able to submit any more jobs unless granted more SU by the Scheme Manager.

Scheme managers can add more SU to a project through a fixed allocation if there is still SU available from within the unallocated share. The SU is then added to a project and is deducted from the unallocated share. The project is then guaranteed to use this SU within the normal quarterly constraints.

Scheme managers can add a variable allocation at any time since this will just increase the level of over commitment and does not guarantee the project will be able to access additional SU.

It should also be noted that the "use it or lose it" rule still applies to Schemes on a quarterly basis. This means that if a Scheme manager fails to allocate out all the SU or the projects A – C don't use all the unallocated share with the quarter, then unused share remains unused and does not roll forward into the next quarter.

Over Allocation

To prevent underutilisation, Scheme manager have been encouraged to over allocate on the assumption that all projects won't use their full share. This works when the assumption holds true but becomes problematic for NCI and other users of the system when multiple projects use their full allocation taking the Scheme in total over its allocation.

Under the new proposed allocation model, Projects that have been granted the right to access additional SU will be able to use it until the total Scheme share has been used. When

the Scheme has reached their fixed allocation, then projects using this unallocated share will then be blocked from running new jobs. Projects that have not reached their allocation will continue to operate normally.

Also note that under the new proposed model, scheme managers will not be able to over allocate fixed allocations. In other words, the fixed allocations of all projects under the scheme cannot be greater than the total allocation for that scheme.

SU Transfers
As with the existing allocation process, Scheme managers can request transfers of SU between projects. For example, if Project B was not expected to use all its SU resources, the scheme manager could arrange for part of the fixed allocation from Project B to be transferred to the fixed allocation of another project, thereby guaranteeing the project that SU. Alternatively, the scheme manager can transfer this time back into the unallocated share and make the time available to any project that has a variable allocation to use when the have exhausted their fixed allocation, including Project B from which the SU were taken in the first place.

Scheme Allocation Strategy
The flexibility that a fixed and variable allocation provides is effectively a base allocation of SU for a project with the option of some bonus SU if available. To minimise the shuffle of SU from one project to another, Scheme managers can now choose to allocate a percentage share of the total to projects, then provide all projects with a variable cap and allow the unallocated share to act as buffer filling in automatically when projects need more time.

For example, a Scheme manage has 100 kSU available to distribute for the quarter to the four projects A-D. For simplicity assume all projects want 25 kSU for the quarter. Instead of allocating each project 25 kSU and then trying to claw back SU from one project to give to another, the Scheme manager can allocate out 20 kSU fixed allocation with a cap of 30 kSU to the 4 projects (i.e. 20 kSU fixed allocation + 10 kSU variable allocation = 30 kSU cap). This would mean that each project will be guaranteed for getting 20 kSU and then competing for the remaining 20 kSU based on who needed it. By placing a cap of 30 kSU (a variable allocation of 10 kSU), no one project can take all the remining time automatically.

Multi-Stake holder projects

On Gadi there are several projects that receive SU allocation from more than one scheme. These projects fall into two major categories

1) Merit Projects that receive a Merit Allocation (ALCG, NCMAS, etc) and also receive additional SU from their home institution
2) Collaborative projects where the project receives time from one or more institutional schemes

Merit Projects
Projects that receive an allocation from a Merit scheme such as NCMAS are awarded resources for use on the system to which they were awarded. If a project receives an

allocation of 100 kSU from NCMAS to use on GADI, that project is not able to use these SU on any other system.

These projects are often eligible for a "top up" of SU awarded under an institutional scheme. As an example, consider a project that receives 100 kSU from NCMAS and requests 30 kSU from their home institution. Assuming sufficient capacity is available, the Scheme manager can transfer the 30 kSU as either a fixed or a variable allocation. As a fixed allocation, the project is guaranteed to get 130 kSU but as a variable allocation, the project is guaranteed to get 100 kSU and can use up to 30 kSU from the home institution's unused share.

With all joint projects there is a priority order or a set order in which allocations are used.

For Merit based projects, the SU allocated from the Merit share are used before any other allocation. Assume after 2 months of the quarter, the above example project with a fixed allocation of 130 kSU (100 NCMAS + 30 home institution) had used 90 kSU. Behind the scenes this would mean 90 kSU of the 100 kSU from the NCMAS allocation had been used and 0 kSU of the 30 kSU from the home institution scheme had been used. Since none of the home intuition allocation had been used, a scheme manager could legitimately pull part or all that 30 kSU allocation back off the project to redistribute to other projects if needed

Collaborative Projects

With collaborative projects the order in which SU are used depends on how the projects are set up. Unless otherwise indicated, it is assumed as an even split between schemes, but it can be a 60:40 or 70:20:10 etc depending on the number of contributing schemes.

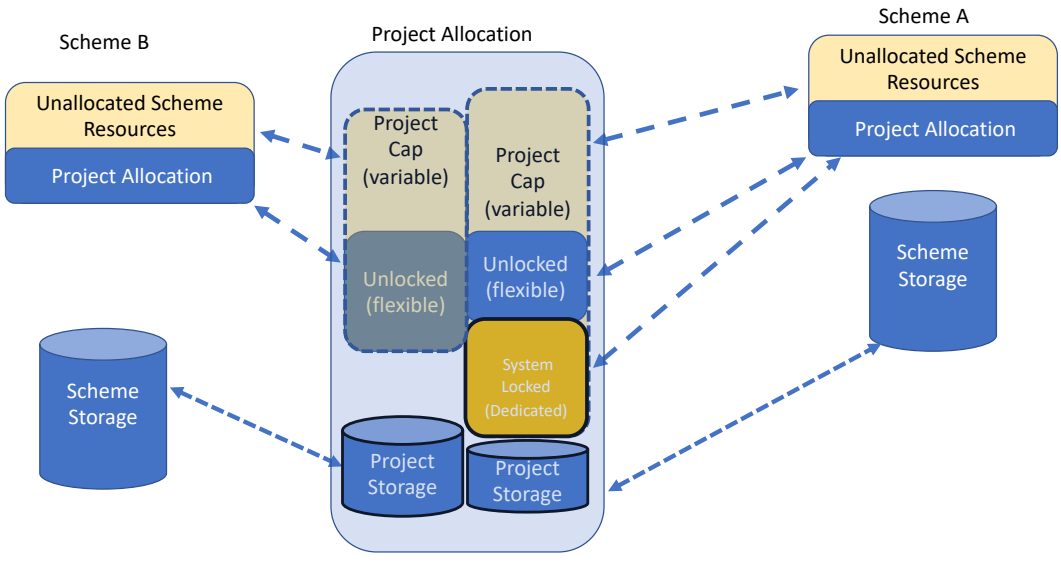

Consider a project that received a fixed allocation from Scheme A and Scheme B and has a 60:40 split. To keep the example simple, the project is given 60kSU from Scheme A and 40

kSU from Scheme B. After a period, the project has used 80 kSU. This would mean that 48 kSU had been used out of Scheme A and 32 ksu from Scheme B.

If it became necessary for the Scheme manager of Scheme B to pull back time, then they could recover a maximum of 8 kSU. The project would then be left with 12 kSU and could continue to operate.

Alternatively, if Scheme manager B added 50 kSU extra, then the project now has 12 kSU left from Scheme A and 58 kSU from Scheme B giving a total of 70 kSU to the project. The project would continue to use allocations in a 60:40 split until all 60 kSU from Scheme A had been used. After this point in time, the project would then use 100% of the remaining SU from Scheme B until it had all been used.

Both Scheme A and Scheme B could also provide the project with a variable allocation. In this case the fixed allocations would still be used first then the variable in the same ratio.

A slightly more complex case might occur where the joint project needs more time, Scheme A provides another 50kSU fixed allocation and Scheme B provides a 50 kSU variable allocation. In this case the contribution from A and B would be decremented in 60:40 split while time was available in both schemes. When one scheme became unavailable, the projects would continue to use time from the remaining scheme.

A twist on this example is where a project receives both a Merit allocation and a top up from two or more separate schemes. The priority order remains the same with 100% of the SU merit used first, then the SU from the joint schemes at their agreed ratio.