



Feedback from Tier-1 Evaluation Committee

VSC Users Day 2017

Walter Lioen walter.lioen@surfsara.nl

Chairman Tier-1 Allocation Board



Tier-1 Allocation Board – TAB

Started 2014, composition:

- Sadaf Alam (CSCS, Switzerland)
- Nicole Audiffren (CINES, France)
- Derek Groen
(Computer Science, Brunel University London, UK)
- Walter Lioen, **chairman**
(SURFsara, The Netherlands)
- Gavin Pringle (EPCC, UK)

Evaluation Guidelines 2017

1. Technical criteria will continue to receive priority in the evaluation of Tier-1 applications.
2. Applications should justify the use of Tier-1, i.e. only for job volumes that cannot be executed on Tier-2.
3. HTC and HPC projects will be treated on an equal footing.
4. For the time being, no restrictions will be placed on the number of projects per sub-discipline, but a number of aspects such as distribution over sub-disciplines, universities, types of jobs, etc. will be monitored in 2017 and re-evaluated later.
7. Care will be taken to ensure that the final decision states the exact grounds on which it is based, to be backed up by a technical evaluation, especially for projects that are found to be technically sound but that cannot be allocated any computing time due to the limited availability of node days, and that this is transparently communicated to the applicants. Regular feedback will be provided to the VSC Users Committee.

TAB Meeting

- Every proposal is reviewed by two board members doing an independent rating
- During the review meeting, the first and second reader discuss their findings with the full board
- The coordinators from participating universities are present as observers and can/will provide additional information on request
- The full board decides on a final rating

Acceptance Rate and Quality – Evolution

- March 2014
 - 9 proposals
 - Acceptance rate: virtually 100%
- February 2017
 - 30 proposals
 - Acceptance rate: some 50%
 - Call is very successful!
 - Proposals are rated/ranked based on quality
 - Highest ranked proposals are awarded until we run out of available resources
 - Quality of proposals clearly improved over the years
 - Only small number of rejects because of quality issues
- June 5, 2017
 - Upcoming cut-off date!

Tier-1 Application Form

- Fill out form
- Completely answer all questions / provide all requested information
- Motivate / substantiate, especially
 - Need for Tier-1
 - Amount of node days / disk storage
 - # nodes / run
- That's basically all
- Stop here?

Tier-1 Application Form Q6

- *Persons mandated by the Applicant to compute on the Tier-1 – experience of using HPC resources in the past (Tier-0/Tier-1/Tier-2 infrastructure in Belgium and abroad)*
 - What infrastructure
 - Number of years
 - Programming experience (if relevant) / user experience
 - If you are keen to mention experience of the supervisor / group as a whole mention this separately

Tier-1 Application Form Q8

- *Justify the **number of node days** requested.*
 - Provide all requested information (including memory usage)
 - Don't mix up nodes / CPUs / cores
 - If possible, use the example Table 1 as a template
 - Clarify the computation of the total number of node days relate data in table to description (don't make it a puzzle)
 - Check the math ...

Tier-1 Application Form Q8 (continued)

- For estimates based on other clusters specify the system:
 - Name (in most cases sufficient for VSC Tier-2 clusters)
 - Vendor name / System type
 - # sockets per node / CPU type
 - Memory size per node (and type)
 - Interconnect
- Example:
 - BrENIAC; NEC HPC1816Rg; 2 Intel Xeon E5-2680 v4 CPUs/node; 128 – 256 GB (DDR4 2400 MT/s) memory/node; EDR InfiniBand interconnect
- In this case explain how you come to you estimates (e.g. based on: clock speed / core count / instruction set / memory speed)

Tier-1 Application Form Q9

- Describe the *software* (...)
Provide the results of *scaling* tests that were conducted with this software, *preferably on the current VSC Tier-1* (using, e.g., a *Starting Grant*) for *system/problem sizes that are on par with those of the intended computing tasks* (e.g., same mesh sizes, actual molecular system, ...). If not run on the current VSC Tier-1, provide the name, architecture, #cores, memory, etc. of the machine that was used to obtain these results and how you think this compares to the current VSC Tier-1. If a different system/problem size is used provide some guidance how it relates to the problem size in the application.
Provide both a table and scaling plot such as table 2 and plot 1 below (...)

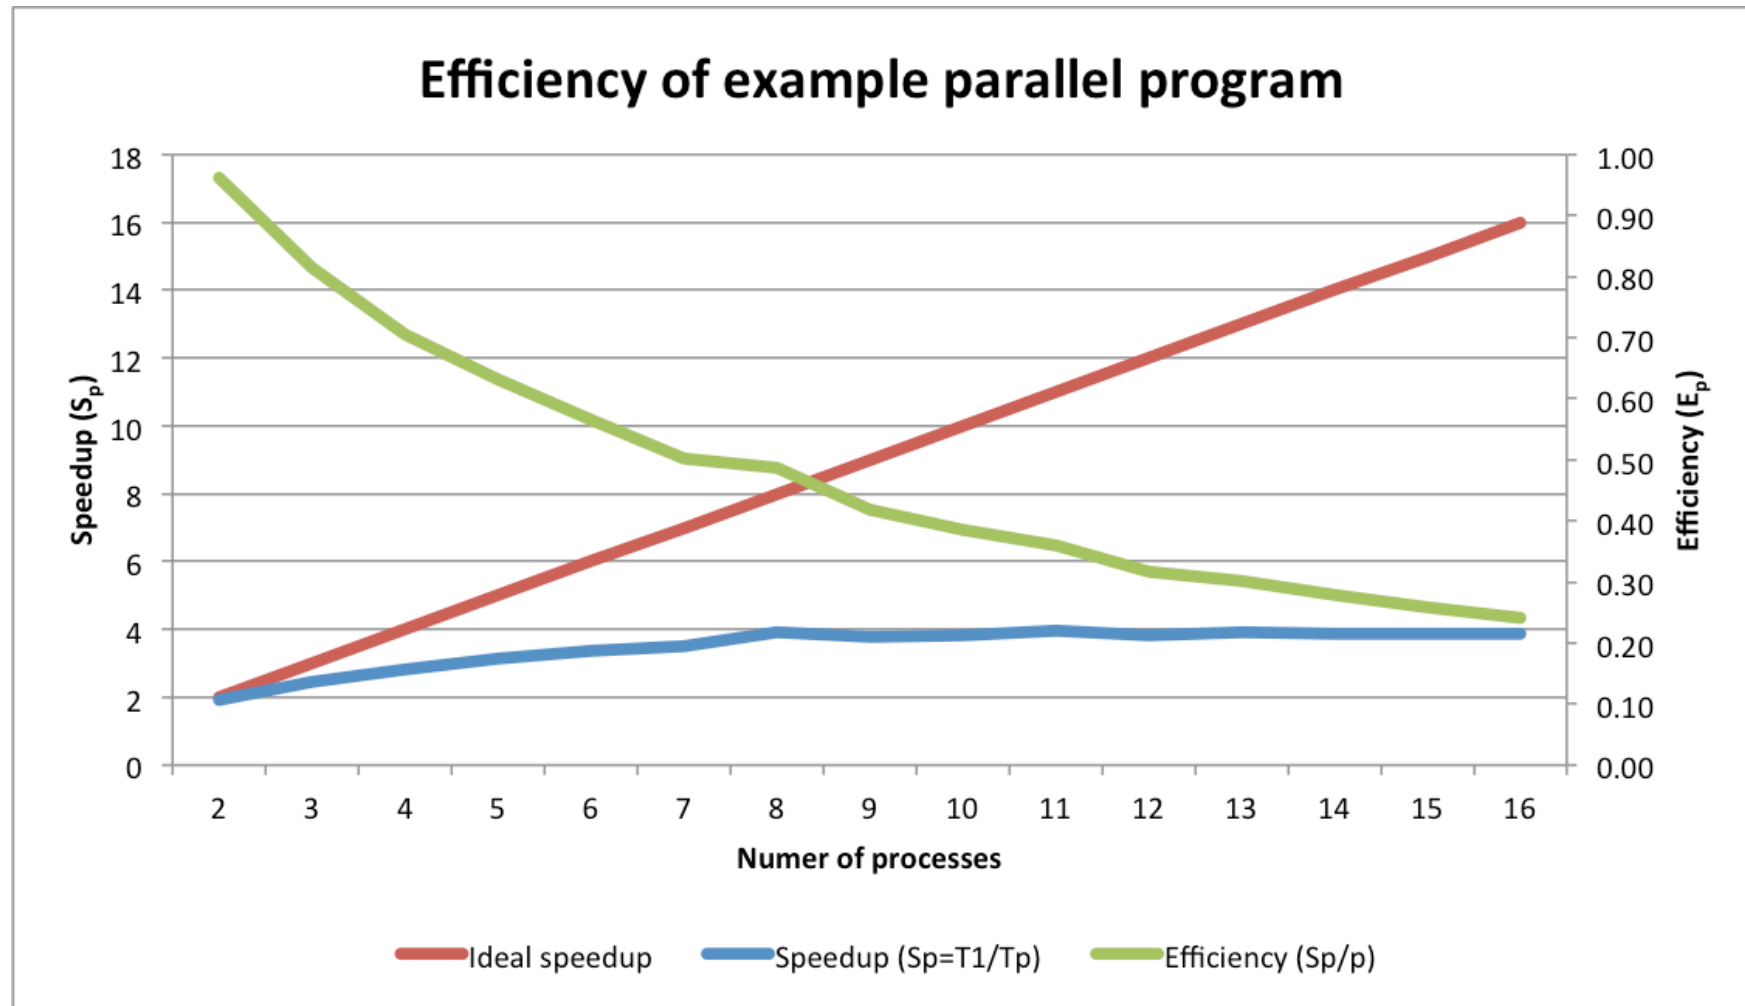
Tier-1 Application Form Q9 (continued)

- Ideally show scaling of relevant problem(s) on Tier-1 (If needed, use a Starting Grant)
- Only use scaling data from another cluster if the code scales (very) well (preferably / if possible show scaling for a wider range of node counts)
- Also here: specify the system (as for the estimate for the number of node days).
- Scaling data can only be skipped for HTC projects but make sure to **state this explicitly**.
(Normally, HTC can scale linearly: running many single node jobs, possibly using the worker framework.)

Tier-1 Application Form Q9 (continued)

- Start (strong) scaling data at 1 node or at least at a minimal number of nodes (based on the memory requirements)
- Use at least 3 core counts
- Provide scaling in the range of the node counts that you target
- Provide legible numbers in graphs (corresponding with the table)
- Preferably use (relative) parallel Efficiency graphs instead of speed-up graphs

Speed-up vs Parallel Efficiency



Tier-1 Application Form Q9 (continued)

- Scaling is important, however
 - Use at least 70-90% parallel efficiency if possible
 - Maximize throughput: use lower node counts / higher parallel efficiency (if you are not bound by the time to solution / maximizing memory usage)
 - HTC and HPC projects will be treated on an equal footing
 - Use the worker framework for parallelism across the problem (perfect scaling!)
- Eric Lindahl (chairman of PRACE Scientific Steering Committee) on embarrassingly parallel problems: this in fact is very smart parallel usage

Tier-1 Application Form Q10

- (Describe how you will manage the resources)
- Provide a data management plan (transfer of files to/from Tier1)
 - Mandatory for I/O intensive applications, i.e. using one or more of:
 - Large amount of disk space (GBs)
 - Large number of inodes: many (small) files and/or directories (e.g. use tar and compression)
 - Bandwidth to disk
 - Number of iops (I/O operations per second)
 - Large volume offloading / stage out to Tier-2

General Recommendations

- Proposals should be self-contained!
 - Do not refer to previous proposals or external documents without providing any context
 - The TAB can only evaluate what is in the proposal and is not supposed to use any other information or knowledge
- Remove unused example tables and graphs
- Check the math (computation of node days)
- Use Starting Grants
- Complete the form
- Clearly motivate the number of node days (~ 100 k€)
- Demonstrate efficient use of Tier-1
- Ask your VSC coordinator for assistance with your proposal