

Welcome to GPGPU User Community survey

Who should take part? This survey is open to all scientific communities, who already use or plan to use General Purpose Graphics Processing Units (GPGPUs) for compute intensive applications.

The use of General Purpose Graphics Processing Units (GPGPUs) or co-processor accelerator devices, such as Intel's Xeon Phi co-processor, over the past few years has resulted in significant increases in the performance of certain applications. With three of the top 10 clusters in the current June 2012 Top500 supercomputers list using NVIDIA GPGPUs, we would expect the number of GPGPU or co-processor deployments at grid resource centres to grow significantly over the next few years.

The purpose of this survey is to gauge how the users currently use, or intend to use, GPGPUs or other accelerated devices in a distributed environment such as grids or clouds, or whether they shall use grid/cloud resources at all for this purpose. In particular, we would like to determine if there is a specific need from European research collaborations and their partners beyond Europe for a more tightly integrated GPGPU capability within the European Grid Infrastructure (www.egi.eu) or its Federated Cloud platform (go.egi.eu/cloud). In addition, we welcome further user comments and feedback on other related topics not covered in the survey.

For further questions or comments, please contact the EGI GPGPU Virtual Team: vt-gpgpu@mailman.egi.eu

More about GPGPU VT at https://wiki.egi.eu/wiki/VT_GPGPU

Information collected through the survey will be aggregated and will be used by the relevant EGI teams to define a technical roadmap.

The closing date for survey responses is Thursday September 13th 2012.

During the survey, you always can go back to the previous questions and only on the last page after you click (done) your answers will be submitted.

Progress bar on the top shows the percentage of the questions being answered.

Questions marked with (*) are mandatory.

1. Do you currently use grid or cloud technologies ?

- Yes
 No

Please provide a comment. Why Yes/No ?

2. Would you be interested in accessing remote GPGPU based resources through a computing infrastructure, such as National Grid Infrastructures or the European Grid Infrastructure (EGI) ?

- Yes
 No

Please provide a comment. Why Yes/No ?

***3. Do you use GPGPU based applications for your scientific computations ?**

- Yes
 No

Please provide a comment. Why Yes/No ?

GPGPU User Community survey

4. GPGPU computation, what is the speed-up achieved compared with not-parallel computing ? (e.g. without message passing, threading, vectorization, not using grid or any other types of distributed computing)

5. How many GPGPU resources do you use for testing activities and/or for full production activities?

GPGPU User Community survey

6. Do you intend to use GPGPU computation within the next 18 months ?

- No
- Yes (Comment a time frame e.g. 3 months)

Please provide a comment.

7. Do you have specific requirements concerning the way remote access of GPGPUs is provided ? (e.g. costing model, security considerations, when and how many GPGPUs would be available, etc)

- Yes
- No

Please provide a comment.

*8. Do you develop or intend to develop any GPGPU based applications ?

- Yes
- No

Please provide a comment. Why Yes/No ?

GPGPU User Community survey

9. Is the code being developed or used inside an international collaboration ?

- No
- Yes (Comment which ones ?)

Please provide a comment. Why Yes/No ?

10. What do you expect from high level programming language abstractions to write parallel programs ? Is there a need for such languages ? (e.g. user-specific knowledge of threads, memory layout, optimized GPU processing, etc.)

11. What particular numerical methods would you like to use on GPGPUs ?

- CUBLAS for linear algebra
- CUSPARSE for sparse matrices
- CURAND for pseudo random number generation
- NPP
- Other

Other (please specify)

12. Do you require double precision floating point arithmetic for your applications ?

- Yes
- No

Please provide a comment. Why Yes/No ?

13. What Application Programming Interface do/will you use ?

- CUDA
- OpenCL
- OpenACC
- Other

Other (please specify)

14. Do you intend to develop code which depends on other application frameworks ? (e.g MPI, BLAST etc.)

- MPI
- BLAST
- Other

Other (please specify)

15. Are there any particular GPGPU applications/solvers/libraries/methods you would like to have on a GPGPU cluster ?

- No
- Yes

Please provide a comment.

16. Is there a need on the market to implement additional applications/solvers/libraries/methods ?

- No
- Yes

Please provide a comment. Why Yes/No ?

***17. Are you performing hybrid computations (CPU + GPU), or CPU is just used for I/O, communication and computation management ?**

Yes

No

Please provide a comment. Why Yes/No ?

GPGPU User Community survey

18. What is the optimal ratio between the number of GPUs and CPU cores for your application ?

19. If you are utilizing all GPU resources on a node, do you expect to be an exclusive user on this node ? (e.g. cost of PCI-Express and RAM bandwidth sharing, etc.)

- No
 Yes

Please provide a comment. Why Yes/No ?

20. Do you expect any particular network topology on the cluster or utilities to exchange the data between the nodes ? (e.g. NVIDIA GPUDirect)

- No
 Yes

Please provide a comment. Why Yes/No ?

GPGPU User Community survey

21. Which scientific disciplines and/or projects of yours is your input to this survey relevant to ?

22. Please provide your name and e-mail (these are not mandatory fields, but we recommend you to fill them so later we may be able to get back to you for further clarification). Your contact details will not be used for any other purposes and will not be forwarded to third parties.

Name

E-mail

23. Please specify here any additional information or feedback relevant to this survey that you wish to provide

Thank You very much for your time!

If you have any other use cases or suggestions or you simply feel that you have more to say - please send us an email: vt-ggpu@egi.eu.

List of currently captured use cases from EGI community: https://wiki.egi.eu/wiki/VT_GPGPU/Use_Case