HPC.NRW

HPC.NRW in a Nutshell: Application for computing time

How can I get access to supercomputers?

Tim Cramer

GREAT COMPUTING COMES WITH GREAT SUPPORT.

Introduction

HPC.NRW

- HPC resources are expensive
  - Up to 15 Mio € per Tier-2 system (expected life time: ~5 years)
  - Local staff for administration, maintenance, support, review processes, procurements, etc.
  - ~ 1 Mio € power consumption per year / system (depending on system size)

- Funding agencies (DFG, NRW, Bund, etc.)
  - Usage only for scientific purpose (e.g. crypto mining strictly forbidden)
  - HPC operators have to ensure scientific usage
    - Compute time application & review
    - Project monitoring
    - Project reports

HPC Pyramid

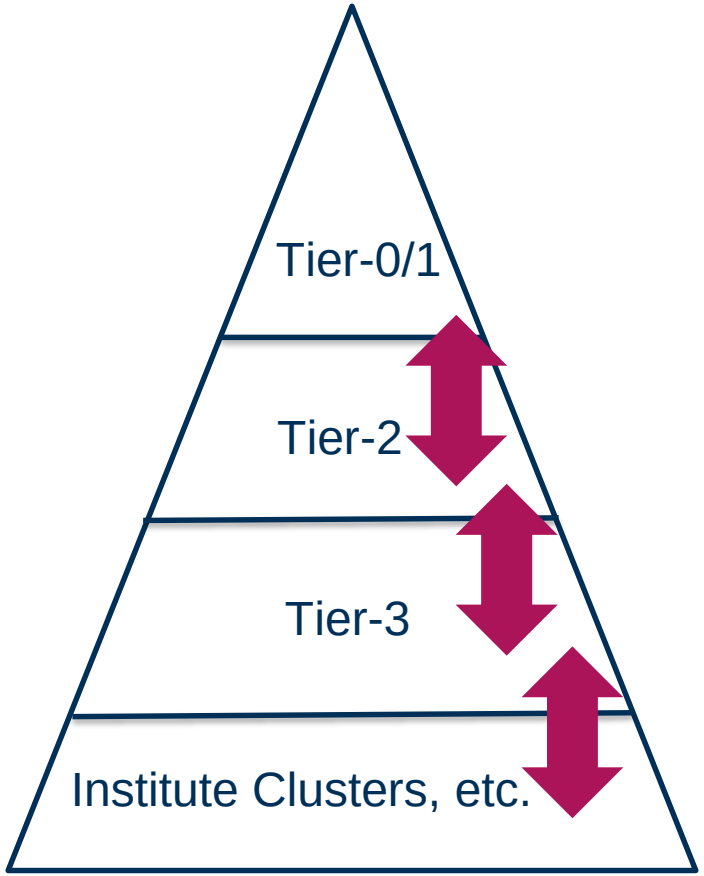

Tier 0: European Level

- Partnership for Advanced Computing in Europe (PRACE)
- https://prace-ri.eu/hpc-access/calls-for-proposals

Tier 1: National Level (large scale)

- Gauss Centre for Supercomputing (GCS)
- Jülich (JSC), Munich (LRZ), Stuttgart (HLRS)
- https://www.gauss-centre.eu/for-users/hpc-access

Tier 2: Regional-National Level

- Gauss Allianz (GA): https://gauss-allianz.de
- Aachen, Cologne, Paderborn (and others outside NRW)
- Nationales Hochleistungsrechnen (NHR): https://www.nhr-gs.de

Tier 3: Regional Level

- E.g. local universities

Quick Reference Cards

– Information as quick reference cards: https://hpc.dh.nrw/de/quick-reference-cards

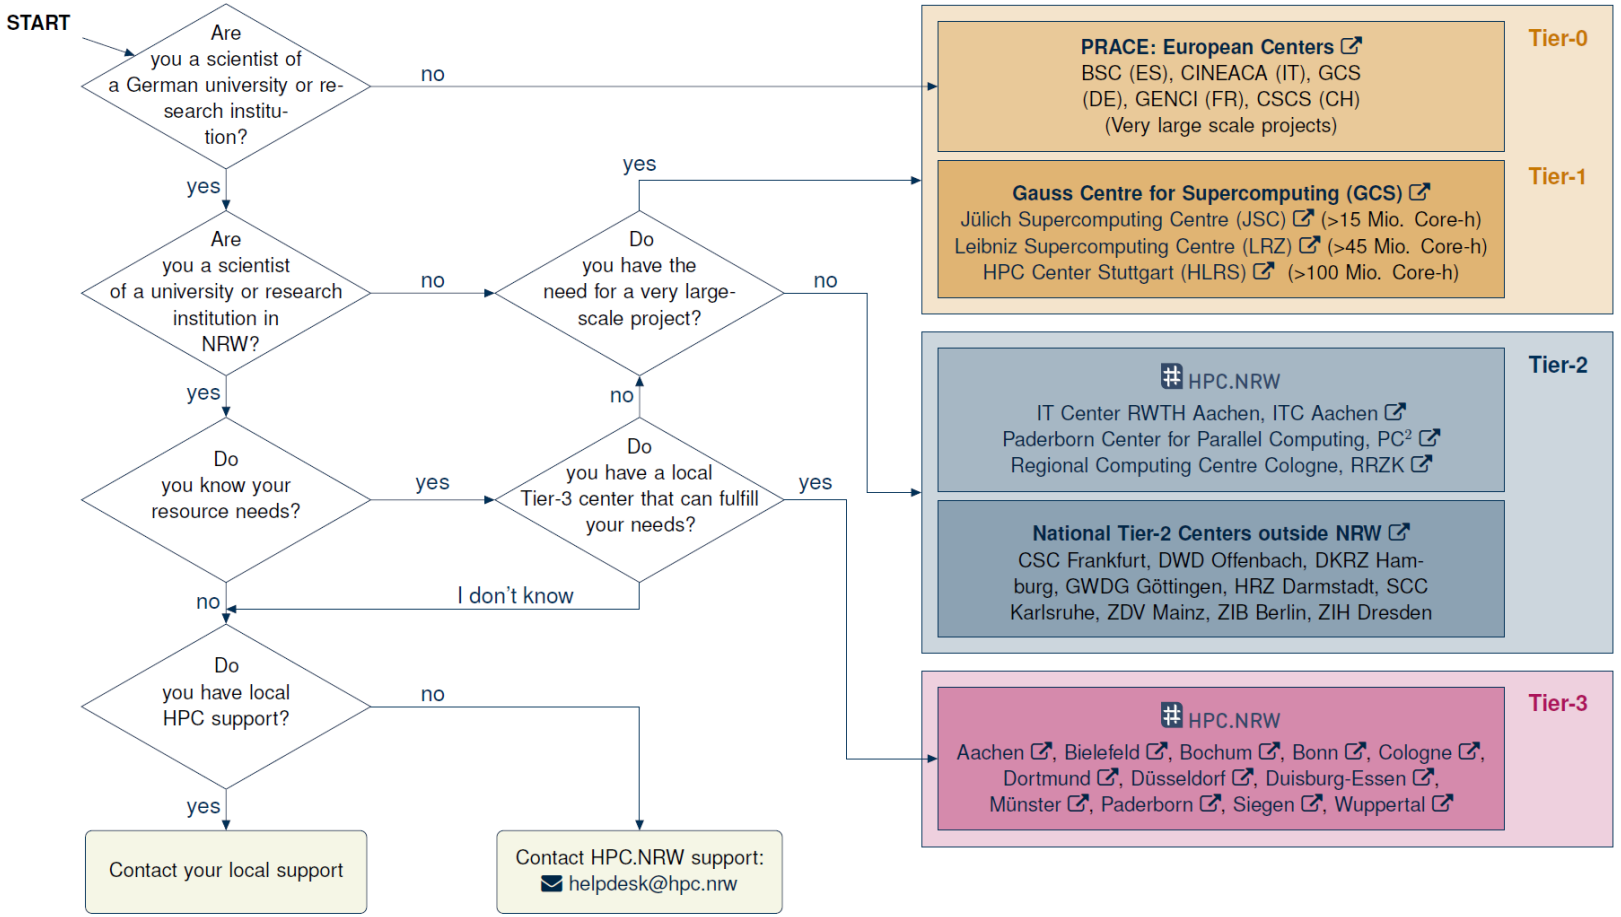

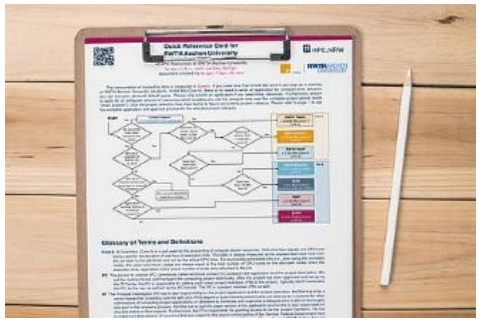

Who can use the computing resources in NRW?

HPC.NRW

- Tier-3 Centers
  - Mainly persons from the local university
  - Exceptions might exist, e.g.:
    - Members of FH Dortmund, Bochum can use the local university resources
    - Limited use of CPUs in Siegen possible

- Tier-2 Centers
  - Aachen, Cologne, Paderborn
  - Scientific staff of FHs only (no students), Principle Investigator (PI) required (Professor or Ph.D)
  - Scientific driven application & review process

- Citizens of countries that are subject to the export control policy of the German Federal Government may need additional authorization from the German Federal Office for Economic Affairs and Export Control (BAFA) before they are allowed to use HPC resources

How to apply for computing resources?

Most effort for you here, we will take care about the other steps

Really? So complicated? I just want to do research!

→ Don’t be afraid! It is not that bad ;-)

Project Preparation 1/3

- Effort for proposal depends on amount of required resources
- Concrete guidelines differ, but are quite similar
  - Check QRCs for details: https://hpc.dh.nrw/de/quick-reference-cards
- Resource estimation (allocation in Mio-Core-h)
  - Core-h := Usage / reservation of one core for one hour
  - Example Core-h: Using one compute node with 48 cores for one year (24/7):
    48 cores * 24 h * 365 days = 0.42 Mio Core-h
  - Example Memory: Many HPC systems are equipped with 2-4 GB per core.
- Trial accounts / test projects might be available

Project Preparation 2/3

- Select a HPC center
- Identify a fitting project category

[1] RWTH and FZJ only
[2] NRW universities without own HPC centers and Paderborn only
[3] Depending on purpose also: UzK WGGC, UzK secure, UzK acceleration (GPU) projects
[4] Future Tier-2 categories, might change

Project category	Aachen	Paderborn	Cologne[4]
PREP / Test project	Max. Core-h depends on purpose	Max. Core-h depends on purpose	< 1000 Core-h
Small project	< 0.24 Mio Core-h [1]	< 4 Mio Core-h [2]	< 1000 Core-h
Medium / Normal project	< 8 Mio Core-h	< 12 Mio Core-h	< 5 Mio Core-h
Large project	8 – 35 Mio Core-h	12 – 60 Mio Core-h	> 5 Mio Core-h [3]

- For larger projects you need to provide scaling information
  - Tutorial: https://hpc-wiki.info/hpc/Scaling_tutorial
- Prepare a project description (templates can typically be found on the local websites)
  - Most important part, will be scientific reviewed

Project Preparation 3/3

HPC.NRW

- NHR “periodical procedure” (German: “getaktetes Verfahren”)
  - Paderborn and Aachen
  - **Only** the Resource Allocation Board (RAB) decides about computing resources

Call	Submission Deadline	TecRev Deadline	SciRev Deadline	RAB Meeting	User Board Meeting	Project Start
1	**01.10.**	15.10.	20.11.	**End Nov**	Mid Dec	**01.01.**
2	**01.01.**	15.01.	20.02.	**End Feb**	Mid Mar	**01.04.**
3	**01.04.**	15.04.	20.05.	**End May**	Mid Jun	**01.07.**
4	**01.07.**	15.07.	20.08.	**End Aug**	Mid Sept	**01.10.**

→ Even for NHR normal: Formal resource allocation might be after the meeting of the RAB

Proposal Submission

HPC.NRW

1. Use the local submission system
   - Aachen (online): https://www.itc.rwth-aachen.de/hpc-project
   - Cologne (PDF): https://rrzk.uni-koeln.de/en/hpc-projects/hpc/access-and-use-instructions
   - Paderborn (online): https://pc2.uni-paderborn.de/go/access

2. Principal Investigator (PI) has to sign the application

3. Send signed and scanned proposal

Formal Evaluation & Technical Review

Formal Evaluation

- Formal aspects of a project are verified by the HPC center (e.g., Is the PI a professor or owns an Ph.D?)
- PI (or contact person) will be contacted if questions/problems show up
- Duration: usually some work days

Technical Review

- HPC experts will check your proposal for technical feasibility (e.g., availability of requested resources, software, etc.)
- Duration: up to two week

Scientific Review

- Not done for test (or smaller) projects
- We select 1-3 reviewer (=experts in the corresponding scientific domain)
- Depending on project size: Internal or external experts
- Single-blind review
- Reviewers check the scientific soundness
- Duration:
  - Rolling calls: usually 2 weeks to 3 month (depending on reviewer performance)
  - Fixed date calls: > 3 month (depending on deadlines)
- Note: If your project proposal is successful, you might be requested as review for other projects in future

Resource Allocation and Monitoring 1/2

HPC.NRW

- The resource allocation board (“Vergabegremium”) decides about the resources for the project
- In case of success:
  - generate an account (if not done already)
  - add members to the approved compute project
  - prepare and submit job scripts for the project
  - obtain the project account information (quota, usage, etc.)
- Typical time for a project: One year
  - Uniform resource consumption on monthly-base expected
  - You might borrow Core-h from next month or use from last month

Resource Allocation and Monitoring 2/2

HPC.NRW

- Technically every project on the cluster has a certain budget
- Example Aachen (command line tool):
  - Sliding Window (3 months)
    - 1000 (remainder from previous month)
    + 50000 (for the current month)
    + 50000 (for next month)
    - 59000 (consumed this month)
    = 40000 Core-h left over to be consumed this month at most!

    OR:

    Core-h left over for this month*:
    50000 * -20% = -10000

```
$ r_wlm_usage -p bund1234 -q
Account:                                  bund1234
Type:                                        bund
Start of Accounting Period:             01.11.2020
End of Accounting Period:               31.10.2021
State of project:                           active
-------------------------------------------------
Quota monthly (core-h):                      50000
Remaining  core-h of prev. month:            -1000
Consumed   core-h current month:             59000
Consumed   core-h last 4 weeks:              65000
Consumable core-h (%):                         -20
Consumable core-h:                           40000
-------------------------------------------------
Total quota (core-h):                   0.600 Mio
Total consumed core-h so far:           0.500 Mio
-------------------------------------------------
Default partition:                           c18m
Allowed partitions:                   c18m,c18g
Max. allowed wallclocktime:           24.0 hours
Max. allowed cores per job:                  384
```

* 200%: No core-hours were used during the previous and the current month
-101%: The usage for the current and the previous month is > three months' quota

Project Reports

- After the project you have to provide a report about the scientific results
- Acknowledgments in related publications required

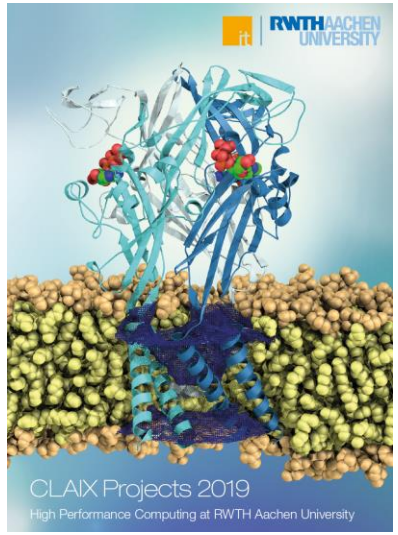

→ Both helps the HPC centers to argue for new future HPC resources

Conclusion

HPC.NRW

Project Preparation | Proposal Submission | Formal Evaluation | Technical Review | Scientific Review | Resource Allocation and Monitoring

Questions?

All information as quick reference card available:




Backup slides

Resource Allocation and Monitoring Backup

HPC.NRW

NHR Normal Aachen

New application
(< 8 Mio. Core-h)

Technical Review
(< 2 weeks)

Formal reject: Resubmission of improved application recommended

Application technical approved

Significant technical issues

Provisioning of start contingent
(0.24 Mio Core-h p. a. / 0.02 Mio Core-h p. m.)

User will be informed (w.r.t the project start), start contingent is part of the project

Scientific Review (< 5 weeks)

2 scientific reviews >0

Only 1 sci review after 7 weeks >0

At least one reject
OR
No answer from reviewer

Conditional Allocation
(Minimum of reviews)

Conditional Allocation
(Minimum of reviews)

User information

RAB decision

Final resource allocation

User will be informed about final contingent