

# In Need of Partnerships

## An Essay about the Collaboration between Computational Sciences and IT Services

Dieter Kranzlmüller

Munich Network Management Team  
Ludwig-Maximilians-Universität München (LMU) &  
Leibniz Supercomputing Centre (LRZ)  
of the Bavarian Academy of Sciences and Humanities






4

Leibniz Rechenzentrum  
Germany

SuperMUC - iDataPlex DX360M4, Xeon  
E5-2680 8C 2.70GHz, Infiniband FDR /  
2012  
IBM


147456 2897.00 3185.05 3422.7

www.top500.org, June 2012




Video: SuperMUC rendered on SuperMUC by LRZ

<http://youtu.be/OIAS6iiqWrQ>




next to come (2014):  
SuperMUC Phase II  
6.4 PFlop/s



# SuperMUC and its predecessors



# SuperMUC and its predecessors



Picture: Horst-Dieter Steinhöfer

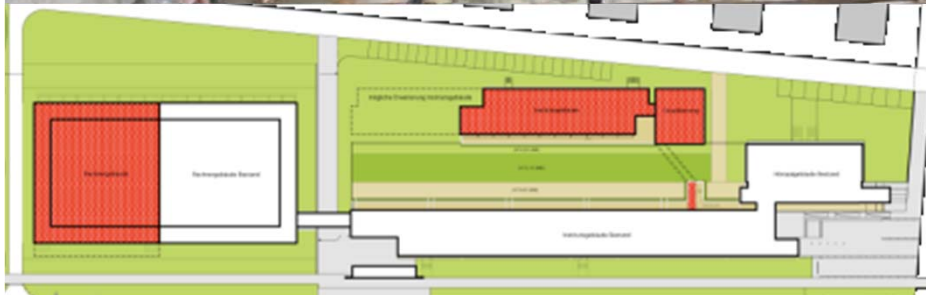


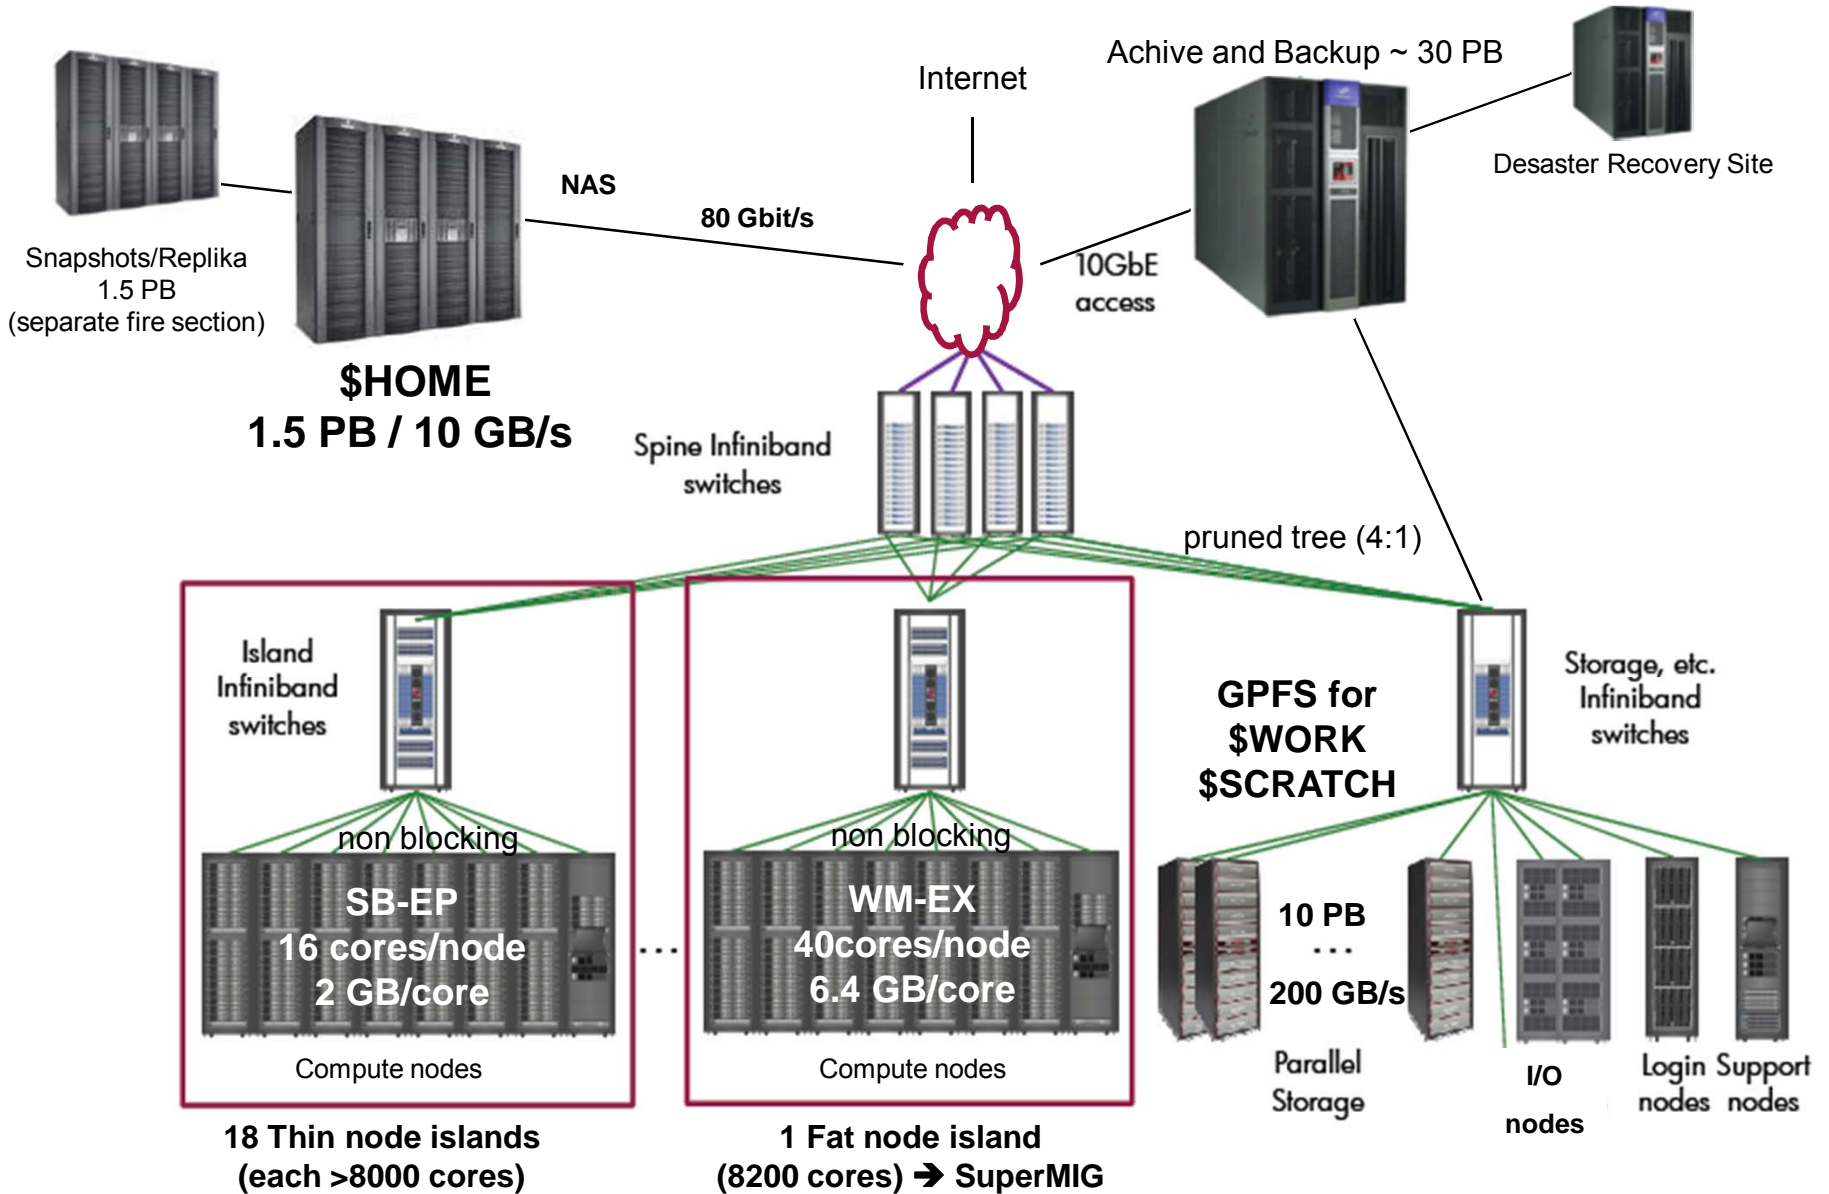
Figure: Herzog+Partner für StBAM2 (staatl. Hochbauamt München 2)



Picture: Ernst A. Graf



Date	System	Flop/s	Cores
2000	HLRB-I	2 Tflop/s	1512
2006	HLRB-II	62 Tflop/s	9728
2012	SuperMUC	3200 Tflop/s	155656
2014	SuperMUC Phase II	3.2 + 3.2 Pflop/s	229960



- How to use today's supercomputers?
- How to cope with the complexity?
- How to use these machines efficiently?
- How to scale applications?
- How to do I/O?
- How about resilience?
- ...

- July 2013:

  - First SuperMUC Extreme Scale Workshop**

- Participants:

  - 15 international projects

- Prerequisites:

  - Successful run on 4 islands (32768 cores)

- Participating Groups (Software packages):

  - LAMMPS, VERTEX, GADGET, WaLBerla, BQCD, Gromacs, APES, SeisSol, CIAO

- Successful results (> 64000 Cores):

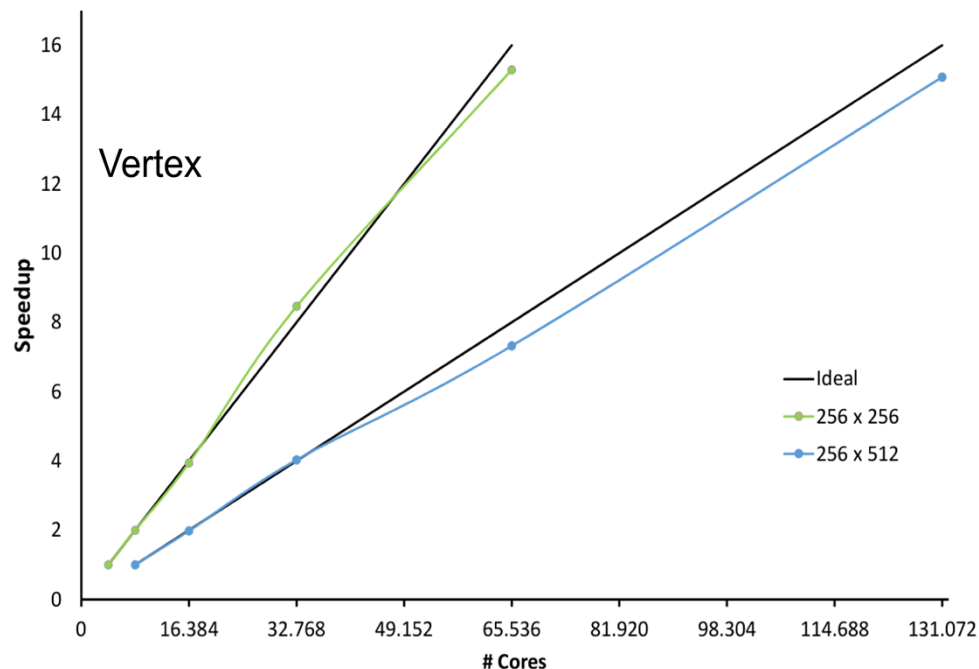
  - Invited to participate in PARCO Conference (Sept. 2013) including a publication of their approach

- Regular SuperMUC operation
  - 4 Islands maximum
  - Batch scheduling system
- Entire SuperMUC reserved 2,5 days for challenge:
  - 0,5 Days for testing
  - 2 Days for executing
  - 16 (of 19) Islands available
- Consumed computing time for all groups:
  - 1 hour of runtime = 130.000 CPU hours
  - 1 year in total

Name	MPI	# cores	Description	TFlop/s/island	TFlop/s max
Linpac	IBM	★ 128000	TOP500	161	2560
Vertex	IBM	★ 128000	Plasma Physics	15	245
GROMACS	IBM, Intel	☆ 64000	Molecular Modelling	40	110
Seissol	IBM	☆ 64000	Geophysics	31	95
waLBerla	IBM	★ 128000	Lattice Boltzmann	5.6	90
LAMMPS	IBM	★ 128000	Molecular Modelling	5.6	90
APES	IBM	☆ 64000	CFD	6	47
BQCD	Intel	★ 128000	Quantum Physics	10	27

- 5 Software packages were running on max 16 islands:
  - LAMMPS
  - VERTEX
  - GADGET
  - WaLBerla
  - BQCD


- VERTEX reached 245 TFlop/s on 16 islands (A. Marek)



- Hybrid (MPI+OpenMP) on SuperMUC still slower than pure MPI (e.g. GROMACS), but applications scale to larger core counts (e.g. VERTEX)
- Core pinning needs a lot of experience by the programmer
- Parallel IO still remains a challenge for many applications, both with regard to stability and speed.
- Several stability issues with GPFS were observed for very large jobs due to writing thousands of files in a single directory. This will be improved in the upcoming versions of the application codes.



- LRZ Extreme Scale Benchmark Suite (LESS) will be available in two versions: public and internal
- All teams will have the opportunity to run performance benchmarks after upcoming SuperMUC maintenances
- Second LRZ Extreme Scaling Workshop → 2-5 June 2014
- Initiation of the LRZ Partnership Initiative  $\pi$ CS



Slices through the three-dimensional gas density (top panels) and vorticity (bottom panels) for fully developed, highly compressible, supersonic turbulence, generated by solenoidal driving (left-hand column) and compressive driving (right-hand column), and a grid resolution of  $4096^3$  cells.

Federrath C MNRAS 2013;mnras.stt1644

MONTHLY NOTICES  
of the Royal Astronomical Society

Dr. Christian Pelties, Department of Earth and Environmental Sciences (LMU)  
Prof. Michael Bader, Department of Informatics (TUM)

1,42 Petaflop/s on 147.456 Cores of SuperMUC  
(44,5 % of Peak Performance)

[http://www.uni-muenchen.de/informationen\\_fuer/presse/presseinformationen/2014/pelties\\_seisol.html](http://www.uni-muenchen.de/informationen_fuer/presse/presseinformationen/2014/pelties_seisol.html)

Picture: Alex Breuer (TUM) / Christian Pelties (LMU)

- Effective usage of High Performance Computing infrastructures requires substantial amount of knowledge and expertise
- Collaboration between Computational Sciences and IT Services leads to new research results
- A partnership between CS and IT is preferred over a provider-user relationship
- Incentives are needed to ensure fruitful collaboration
- LRZ has established the **Partnership Initiative Computational Sciences  $\pi$ CS**

rendered on SuperMUC by LRZ



# In Need of Partnerships

An Essay about the Collaboration  
between Computational Sciences and IT Services

Dieter Kranzlmüller  
[kranzlmueeller@lrz.de](mailto:kranzlmueeller@lrz.de)