



20 Jahre TOP500 mit einem Ausblick auf neuere Entwicklungen

Hans Meuer

Prometeus GmbH & Universität Mannheim

hans@meuer.de

ZKI Herbsttagung in Leipzig


11. - 12. September 2012

- Mannheim Supercomputer Statistics 1986 - 1992 & TOP500 Project Start in 1993
- 39 TOP500 Lists / Highlights of the 39th List and Developments
- Competition between Countries, Manufacturers and Sites
- Linpack Pros and Cons & Alternative Supercomputer Rankings
- HPC & Cloud Computing
- Summary


Mannheim Supercomputer Statistics 1986 – 1992

- Presented at ISC'86 – ISC'92/Mannheim Supercomputer Seminars
- Counting of Supercomputers in the World
- 530 systems in the year 1992

Manufacturer Share



Countries Share



Deficits of the Mannheim Supercomputer Statistics


- Not very reliable
- Vector Supercomputers only – „numerical“/
mainly „scientific and engineering“ applications
- Definition of „Supercomputer“ necessary

Top500 Procedure

- ➔ Listing the 500 most powerful computers in the world
- ➔ Yardstick: Rmax of Linpack
 - Solve $Ax=b$, dense problem, matrix is random*
- ➔ Update twice a year:
 - ISC'xy in June in Germany
 - SC'xy in November in the U.S.
- ➔ All information available from the TOP500 webserver at: www.top500.org

Top500 Authors

Project was started in spring 1993 by:



Hans W. Meuer




Erich Strohmaier

Authors since:



06/1993
Jack Dongarra



11/2000
Horst Simon

Dongarra's TOP500 List of World's Fastest Supercomputers Released at Mannheim Conference



„CRPC researcher Jack Dongarra of the University of Tennessee and Oak Ridge National Laboratory is one of three renowned computer scientists who assemble the legendary TOP500 List of the world's fastest supercomputers. Released twice a year since 1993, the list features sites with the most powerful computer systems, determined with information from a questionnaire sent to high-performance computer (HPC) experts, computational scientists, manufacturers, and the Internet community at large. Dongarra, Hans Meuer of the University of Mannheim, and Erich Strohmaier of the University of Tennessee released their June 1999 TOP500 List at the 14th Mannheim Supercomputing Conference and Seminar, held June 10-12 in Mannheim, Germany.“

*Source: <http://www.crpc.rice.edu/WhatsNew/top500.html>
The Center for Research on Parallel Computation at Rice University*

- Mannheim Supercomputer Statistics 1986 - 1992 & TOP500 Project Start in 1993
- ➔ 39 TOP500 Lists / Highlights of the 39th List and Developments
- Competition between Countries, Manufacturers and Sites
- Linpack Pros and Cons & Alternative Supercomputer Rankings
- HPC & Cloud Computing
- Summary


38th List: The TOP10 – Nov 11

Rank	Site	Manufacturer	Computer	Country	Cores	Rmax [Pflops]	Power [MW]
1	RIKEN Advanced Institute for Computational Science	Fujitsu	K Computer SPARC64 VIIIfx 2.0GHz, Tofu Interconnect	Japan	795,024	10.51	12.66
2	National SuperComputer Center in Tianjin	NUDT	Tianhe-1A NUDT TH MPP, Xeon 6C, NVidia, FT-1000 8C	China	186,368	2.566	4.04
3	Oak Ridge National Laboratory	Cray	Jaguar Cray XT5, HC 2.6 GHz	USA	224,162	1.759	6.95
4	National Supercomputing Centre in Shenzhen	Dawning	Nebulae TC3600 Blade, Intel X5650, NVidia Tesla C2050 GPU	China	120,640	1.271	2.58
5	GSIC, Tokyo Institute of Technology	NEC/HP	TSUBAME-2 HP ProLiant, Xeon 6C, NVidia, Linux/Windows	Japan	73,278	1.192	1.40
6	DOE/NNSA/LANL/SNL	Cray	Cielo Cray XE6, 8C 2.4 GHz	USA	142,272	1.110	3.98
7	NASA/Ames Research Center/NAS	SGI	Pleiades SGI Altix ICE 8200EX/8400EX	USA	111,104	1.088	4.10
8	DOE/SC/LBNL/NERSC	Cray	Hopper Cray XE6, 6C 2.1 GHz	USA	153,408	1.054	2.91
9	Commissariat a l'Energie Atomique (CEA)	Bull	Tera 100 Bull bullx super-node S6010/S6030	France	138.368	1.050	4.59
10	DOE/NNSA/LANL	IBM	Roadrunner BladeCenter QS22/LS21	USA	122,400	1.042	2.34

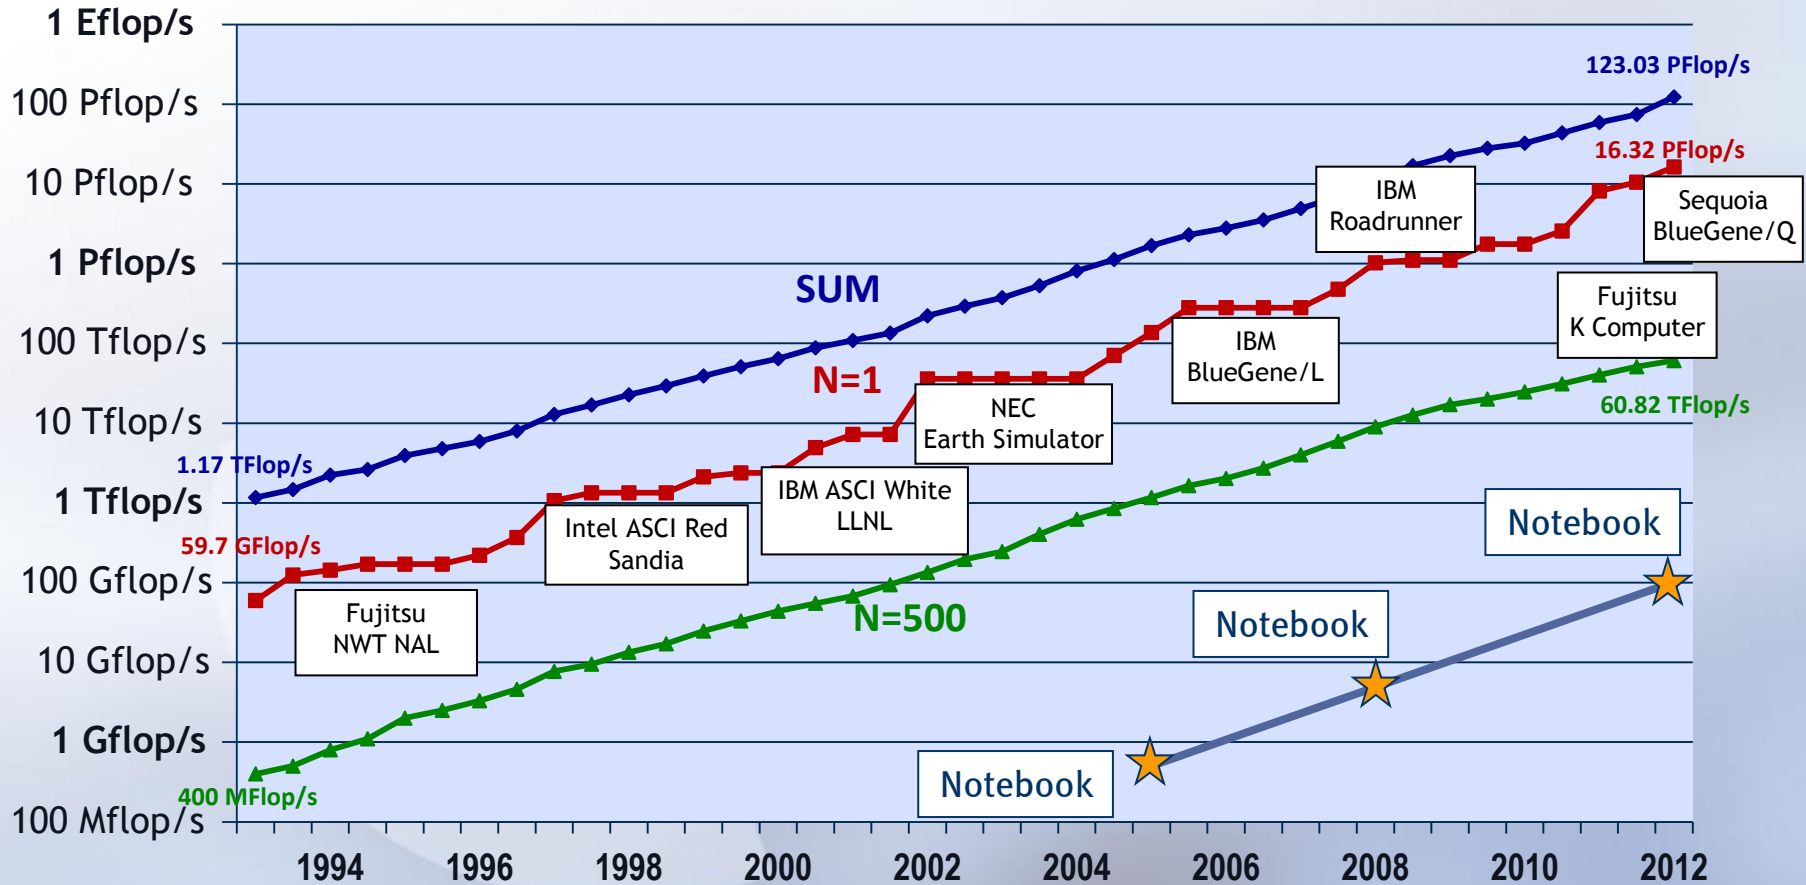
39th List: The TOP10 – Jun 12

#	Site	Manufacturer	Computer	Country	Cores	R _{max} [Pflops]	Power [MW]
1	Lawrence Livermore National Laboratory	IBM	Sequoia BlueGene/Q, Power BQC 16C 1.6GHz, Custom	USA	1,572,864	16.3	7.89
2	RIKEN Advanced Institute for Computational Science	Fujitsu	K Computer SPARC64 VIIIfx 2.0GHz, Tofu Interconnect	Japan	795,024	10.5	12.66
3	Argonne National Laboratory	IBM	Mira BlueGene/Q, Power BQC 16C 1.6GHz, Custom	USA	786,432	8.16	3.95
4	Leibniz Rechenzentrum	IBM	SuperMUC iDataPlex DX360M4, Xeon E5 8C 2.7GHz, Infiniband FDR	Germany	147,456	2.90	3.52
5	National SuperComputer Center in Tianjin	NUDT	Tianhe-1A NUDT TH MPP, Xeon 6C, NVidia, FT-1000 8C	China	186,368	2.57	4.04
6	Oak Ridge National Laboratory	Cray	Jaguar Cray XK6, Opteron 16C 2.2GHz, Gemini, NVIDIA 2090	USA	298,592	1.94	5.14
7	CINECA	IBM	Fermi BlueGene/Q, Power BQC 16C 1.6GHz, Custom	Italy	163,840	1.73	0.82
8	Forschungszentrum Juelich (FZJ)	IBM	JuQUEEN BlueGene/Q, Power BQC 16C 1.6GHz, Custom	Germany	131,072	1.38	0.66
9	Commissariat a l'Energie Atomique CEA/TGCC-GENCI	Bull	Curie thin nodes Bullx B510, Xeon E5 8C 2.7GHz, Infiniband QDR	France	77,184	1.36	2.25
10	National Supercomputing Centre in Shenzhen	Dawning	Nebulae TC3600 Blade, Intel X5650, NVidia Tesla C2050 GPU	China	120,640	1.27	2.58


Replacement Rate



Performance Development




Performance Projection




- Mannheim Supercomputer Statistics 1986 - 1992 & TOP500 Project Start in 1993
- 39 TOP500 Lists / Highlights of the 39th List and Developments
- ➔ Competition between Countries, Manufacturers and Sites
- Linpack Pros and Cons & Alternative Supercomputer Rankings
- HPC & Cloud Computing
- Summary

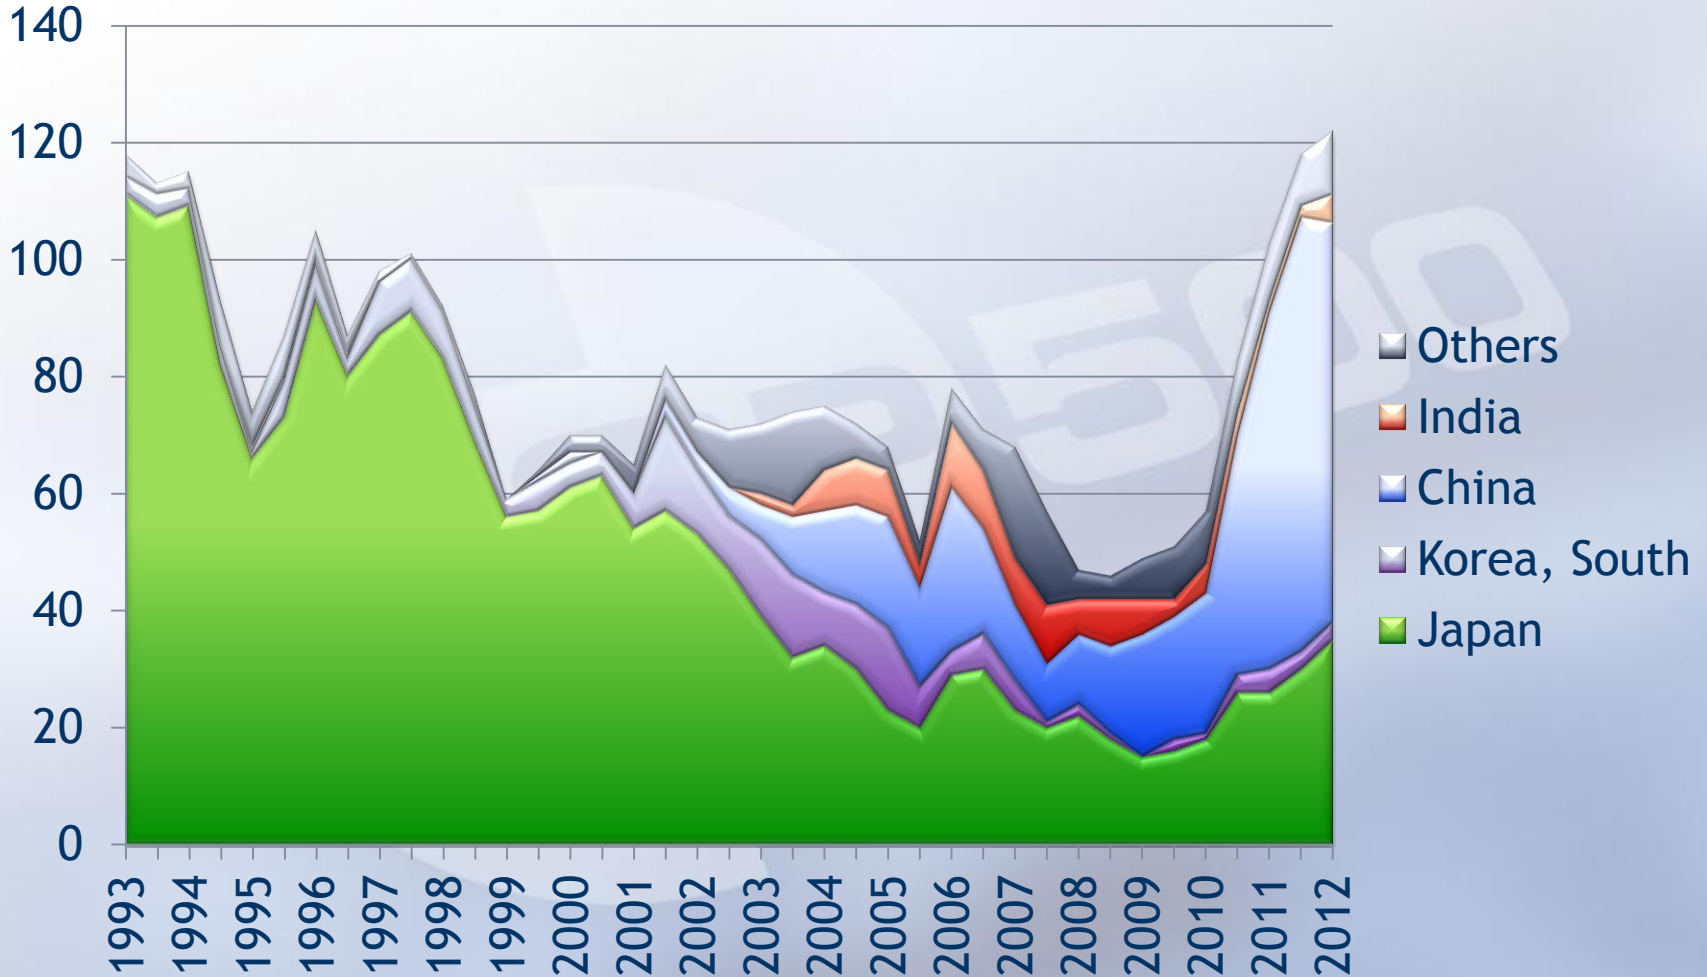
39th List: Countries / System Share




Performance of Countries




Asian Countries




European Countries



39th List: Vendors / System Share



39th List: Vendors (TOP50) / System Share



Top20 Sites since 1993

Rank	Sites	Country	Overall Share (1993-2012)
1	Lawrence Livermore National Laboratory	USA	5.50%
2	Los Alamos National Laboratory	USA	3.30%
3	Sandia National Laboratories	USA	2.61%
4	Government	USA	2.28%
5	DOE/SC/Oak Ridge National Laboratory	USA	2.18%
6	Japan Agency for Marine -Earth Science and Technology	Japan	1.60%
7	NASA/Ames Research Center/NAS	USA	1.47%
8	National Aerospace Laboratory of Japan	Japan	1.29%
9	Forschungszentrum Juelich (FZJ)	Germany	1.25%
10	NCSA	USA	1.07%
11	DOE/SC/LBNL/NERSC	USA	1.07%
12	ECMWF	U.K.	0.96%
13	Naval Oceanographic Office - NAVO DSRC	USA	0.94%
14	RIKEN Advanced Institute for Computational Science (AICS)	Japan	0.94%
15	Semiconductor Company (C)	USA	0.94%
16	University of Tokyo	Japan	0.93%
17	Pittsburgh Supercomputing Center	USA	0.88%
18	ERDC DSRC	USA	0.85%
19	IBM Thomas J. Watson Research Center	USA	0.79%
20	Commissariat a l'Energie Atomique (CEA)	France	0.71%

- Mannheim Supercomputer Statistics 1986 - 1992 & TOP500 Project Start in 1993
- 39 TOP500 Lists / Highlights of the 39th List and Developments
- Competition between Countries, Manufacturers and Sites
- ➔ Linpack Pros and Cons & Alternative Supercomputer Rankings
- HPC & Cloud Computing
- Summary

Pros and Cons of LINPACK

Pros

- Only one figure of merit
- Simple to define and rank
- Allows problem size to change with machine and over time
- Opens competition

Cons

- Emphasizes only ,peak' CPU speed and number of CPUs
- Does not stress local memory bandwidth
- Does not stress the memory system
- Does not stress the internal network
- A single figure of merit cannot reflect the overall performance of an HPC system

GREEN500 - <http://www.green500.org/> Energy is as Important as Performance Status of the GREEN500 Project

- The purpose of the GREEN500 is to provide a ranking of the most energy-efficient supercomputers in the world.
- The emphasis on speed as the ultimate metric has caused other metrics such as reliability, availability, energy efficiency to be largely ignored.
- As a result, there has been an extraordinary increase in the total cost of ownership (TCO) of a supercomputer.
- The inaugural GREEN500 list was announced on November 15, 2008 at SC'08. As a complement to the TOP500, the unveiling of the GREEN500 ushered in a new era where supercomputers can be compared by performance-per-watt.

GREEN500 - <http://www.green500.org/> The Green500 List - June 2012

<i>Green500 Rank</i>	<i>MFLOPS/W</i>	<i>Site*</i>	<i>Computer*</i>	<i>Total Power (kW)</i>
1	2,100.88	DOE/NNSA/LLNL	BlueGene/Q, Power BQC 16C 1.60GHz, Custom	41.10
2	2,100.88	IBM Thomas J. Watson Research Center	BlueGene/Q, Power BQC 16C 1.60GHz, Custom	41.10
3	2,100.86	DOE/SC/Argonne National Laboratory	BlueGene/Q, Power BQC 16C 1.60GHz, Custom	82.20
4	2,100.86	DOE/SC/Argonne National Laboratory	BlueGene/Q, Power BQC 16C 1.60GHz, Custom	82.20
5	2,100.86	Rensselaer Polytechnic Institute	BlueGene/Q, Power BQC 16C 1.60GHz, Custom	82.20
6	2,100.86	University of Rochester	BlueGene/Q, Power BQC 16C 1.60GHz, Custom	82.20
7	2,100.86	IBM Thomas J. Watson Research Center	BlueGene/Q, Power BQC 16C 1.60 GHz, Custom	82.20
8	2,099.56	University of Edinburgh	BlueGene/Q, Power BQC 16C 1.60GHz, Custom	493.10
9	2,099.50	Science and Technology Facilities Council - Daresbury Laboratory	BlueGene/Q, Power BQC 16C 1.60GHz, Custom	575.30
10	2,100.88	Forschungszentrum Juelich (FZJ)	BlueGene/Q, Power BQC 16C 1.60GHz, Custom	657.5

GRAPH500 - <http://www.graph500.org/>

Motivation for GRAPH500

- ➔ Data intensive supercomputer applications are increasingly important for HPC workloads
- ➔ Current benchmarks do not provide useful information on the suitability of supercomputing systems for data intensive applications.
- ➔ A new set of benchmarks is needed in order to guide the design of hardware/ software systems intended to support such applications.
- ➔ Graph algorithms are a core part of many analytics workloads.
- ➔ Backed by a steering committee (50 international experts) GRAPH500 will establish a set of large-scale benchmarks for these applications.

GRAPH500 - <http://www.graph500.org/> Status of the GRAPH500 Project

- Three major application kernels: concurrent search, optimization (single source shortest path), and edge-oriented (maximal independent set).
- Addressing five graph-related business areas: Cyber security, Medical Informatics, Data Enrichment, Social Networks, and Symbolic Networks.
- The GRAPH500 was announced at ISC10 and the first list appeared at SC10 (9 systems ranked).
- Further results have been published at ISC'11 (29 systems) and SC11 (49 systems) and ISC'12 (80 systems)

- Mannheim Supercomputer Statistics 1986 - 1992 & TOP500 Project Start in 1993
- 39 TOP500 Lists / Highlights of the 39th List and Developments
- Competition between Countries, Manufacturers and Sites
- Linpack Pros and Cons & Alternative Supercomputer Rankings
- ➔ **HPC & Cloud Computing**
- Summary

HPC in the Cloud



=



- Cloud is gaining growing interest currently in Industry and Academia.
 - › 70% of CIOs are involved in Cloud - actively resp. observant.
- Scepticism against HPC & Cloud and HPC & Big Data .
- Conferences:
 - ISC Cloud'10, ISC Cloud'11, ISC Cloud'12 in Mannheim

Advantages for HPC in the Cloud



- ➔ HPC Clouds are currently already available: Amazon EC2 , Penguin on Demand, SGI Cyclone, Microsoft Azure, Bull extreme factory, Fujitsu Engineering Cloud, etc
- ➔ No need for own IT resources
- ➔ Saving money: Time, Space , Staff, Cost of ownership
- ➔ Cloud delivers the needed resources: Pay-per-use , Dynamically scale-up and -down
- ➔ In the near future: Applications will sit in the Cloud , Software as a Service

Cloud for Academia



- ➔ High end Supercomputing Applications (Petascale Applications, TOP20 currently in the list) are not suited for the Cloud. They need special tailored HPC Hardware and Software Architectures
- ➔ SME applications with less performance requirements and very often with less parallelization degrees are well suited for the Cloud
- ➔ Ideal solution: Local system in the own computing center for average applications, HPC in the Cloud for overload. National high end supercomputer centers like LRZ Munich, FZ Juelich, HLS Stuttgart if we concentrate on Germany

Challenges of HPC in the Cloud



- ➔ Applications and data have to be transferred over the Internet
- ➔ Loss of control about potential 'crown jewels'
- ➔ Uncertainty about length of transmission time from/to the cloud
- ➔ Trust in cloud service provider ?
- ➔ No information about computers used in the cloud, possible performance deficiency
- ➔ **But the good news are: All these challenges may be solved already today! Please join our forthcoming ISC Cloud'12 event in Mannheim in September 2012**

www.isc-events.com/cloud12



HPC Meets Cloud - The Conference on the Use of Cloud for High Performance Computing (HPC)

Topic: HPC and Big Data in the Cloud

➔ Top Speakers with own hands-on HPC Cloud experiences

➔ Sessions:

- ➔ Keynotes
- ➔ Industry
- ➔ Research and Academia
- ➔ ISV Software in the Cloud
- ➔ Challenges in the Cloud
- ➔ Panels and BoFs

HPC Meets Cloud - The Conference on the Use of Cloud for High Performance Computing (HPC)

Selection of Talks at ISC Cloud'12

- *Rupert Lueck, EMBL: Genome analysis*
- *Wolfgang Lengert, ESA: ESA Earth Observation Flagship Project*
- *Christopher Porter, IBM: The Value of Private Cloud for HPC*
- *Tom-Michael Thamm, Nvidia: Advanced Rendering*
- *Addison Snell, Intersect360 Research: Big Data, HPC, Clouds*
- *Volker Eyrich, Schrödinger Inc.: Virtual Screening*
- *Rolf Sperber, Alcatel-Lucent: Climate Data*
- *Max Lemke, EU: SMEs*
- *Franz-Josef Pfreundt, Fraunhofer ITWM: File Systems*

Selection of Talks at ISC Cloud'12

- *Wolfgang Ziegler, Fraunhofer SCAI: License Management*
- *Giles Hogben, CSA: Cloud Security*
- *Paolo Balboni, ICT Legal Consulting: Legal Aspects*
- *Manish Parashar, Rutgers University: Transforming Science?*
- *David Wallom, Oxford eResearch Centre: Cloud Services*
- *Jens Jensen, Rutherford Appleton Laboratory: Clouds and Security*
- *Frédéric Desprez, INRIA: DIET, a Scalable Platform for Clusters, Grids and Clouds*
- ... and more

- Mannheim Supercomputer Statistics 1986 - 1992 & TOP500 Project Start in 1993
- 39 TOP500 Lists / Highlights of the 39th List and Developments
- Competition between Countries, Manufacturers and Sites
- Linpack Pros and Cons & Alternative Supercomputer Rankings
- HPC & Cloud Computing
- ➔ **Summary**

- Has the TOP500 project been successful and what are the reasons ?

YES. The TOP500 project is the only instrument in the world to look at and evaluate the HPC market. The lists are published two times a year with absolute punctuality, one at ISC in June in Germany and one in November in the US.
- Did the use of Linpack as the yardstick for listing the systems influence the project negatively ?

No. The opposite is true. Without LINPACK the project would have been not so successful.
- Who profits from the TOP500 project and who does not ?

It does not help the users running their application code, of course. But it is very helpful and important for HPC - manufacturers/ vendors, - decision makers/ politicians, - press/media, and – sites.
- Will the TOP500 project continue and how long ?

Definitely YES.
But how long, I ask you to put this question 10 years later to me.

Danke für Ihre Aufmerksamkeit !

Alle Folien dieses Vortrags können hier
runter geladen werden:

<http://www.top500.org/files/meuer-zki2012.pdf>

1.) The TOP500 Project: Looking Back over 15 Years of Supercomputing Experience

Hans Werner Meuer, University of Mannheim & Prometheus GmbH, Germany January 20, 2008

The TOP500 project was launched in 1993 to provide a reliable basis for tracking and detecting trends in high performance computing. Twice a year, a list of the sites operating the world's 500 most powerful computer systems is compiled and released. The best performance on the Linpack benchmark is used as the measurement for ranking the computer systems. The list contains a variety of information including the systems' specifications and major application areas. Information on all 30 TOP500 lists issued to date is available at: www.top500.org

Meuer, Hans Werner (2008), Informatik_Spektrum_31_3_2008, pp 203 - 222, Springer - Verlag 2008

http://www.top500.org/blog/2008/01/20/top500_project_looking_back_over_15_years_supercomputing_experience

2.) The TOP500 Celebrates 20th Anniversary. Will it Survive 20 More?

Tom Tabor, HPCwire, Exclusive Interview with Hans W. Meuer, June 12, 2012

With the upcoming release of the TOP500 next week, the latest rankings are usually a hot topic of discussion this time of year. Over the past 20 years, the list has proven to be a useful and popular compilation of supercomputers for the HPC community. In this exclusive interview, Professor Hans Meuer, considered by many to be the driving force behind the project, offers his thoughts on the TOP500; its past, present, and future.

http://www.hpcwire.com/hpcwire/2012-06-12/the_top500_celebrates_20th_anniversary_will_it_survive_20_more_.html

3.) HPC Lists We'd Like to See

Gary Johnson, Computational Science Solutions, June 15, 2012

Since the release of the first TOP500 list in June of 1993, the HPC community has been motivated by the competition to place high on that list. We're now approaching the twentieth anniversary of the TOP500. In recent years, two additional lists have gained traction: the Green500 and the Graph 500. Would a few more lists be useful? Let's take a look at a some options.

http://www.hpcwire.com/hpcwire/2012-06-15/hpc_lists_we%E2%80%99d_like_to_see.htm

4.) TOP 500 -- 20 Years Later

The ISC Think Tank Series, ISC'12 Hamburg, June 20, 2012

Now in its second year, "The ISC Think Tank Series, sponsored by HPCWire," continues with its dynamic live and uncensored discussion format. These candid, no holds barred panel sessions feature HPC's leading thought leaders from the analyst, industry and end-user communities in discussions on the most controversial and provocative topics facing the industry today. This year's topic: The TOP500 -- Twenty Years Later.

http://www.youtube.com/results?search_query=%22TOP+500+++20+Years+Later%22

5.) Supercomputers – Prestige Objects or Crucial Tools for Science and Industry?

Hans W. Meuer, Horst Gietl, University of Mannheim & Prometheus GmbH, August 3, 2012

This paper - based on an Invited Talk of Hans W. Meuer in the House of Lords on April 18, 2012 - is intended for people interested in High Performance Computing (HPC) in general, in the performance development of HPC systems from the beginning in the 1970s and, above all, in HPC applications in the past, today and tomorrow. Readers do not need to be supercomputer experts.

http://www.top500.org/files/Supercomputers_London_Paper_HWM_HG.pdf

to appear in 'Software Development Practice' in September 2012, UK