LMU LUDWIG-MAXIMILIANS-UNIVERSITÄT MÜNCHEN

MNM TEAM
Munich Network Management Team

lrz

UNISDR Global Assessment Report - Current and Emerging Data and Compute Challenges

Matti Heikkurinen*, Dieter Kranzlmüller

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

Contents

Environmental computing case study

1. Leibniz Supercomputing Centre, MNM-Team, UNISDR
2. What is environmental computing?
3. UNISDR collaboration in detail

Leibniz Supercomputing Centre
of the Bavarian Academy of Sciences and Humanities

With approx. 230 employees
for more than 100.000 students and
for more than 30.000 employees
including 8.500 scientists

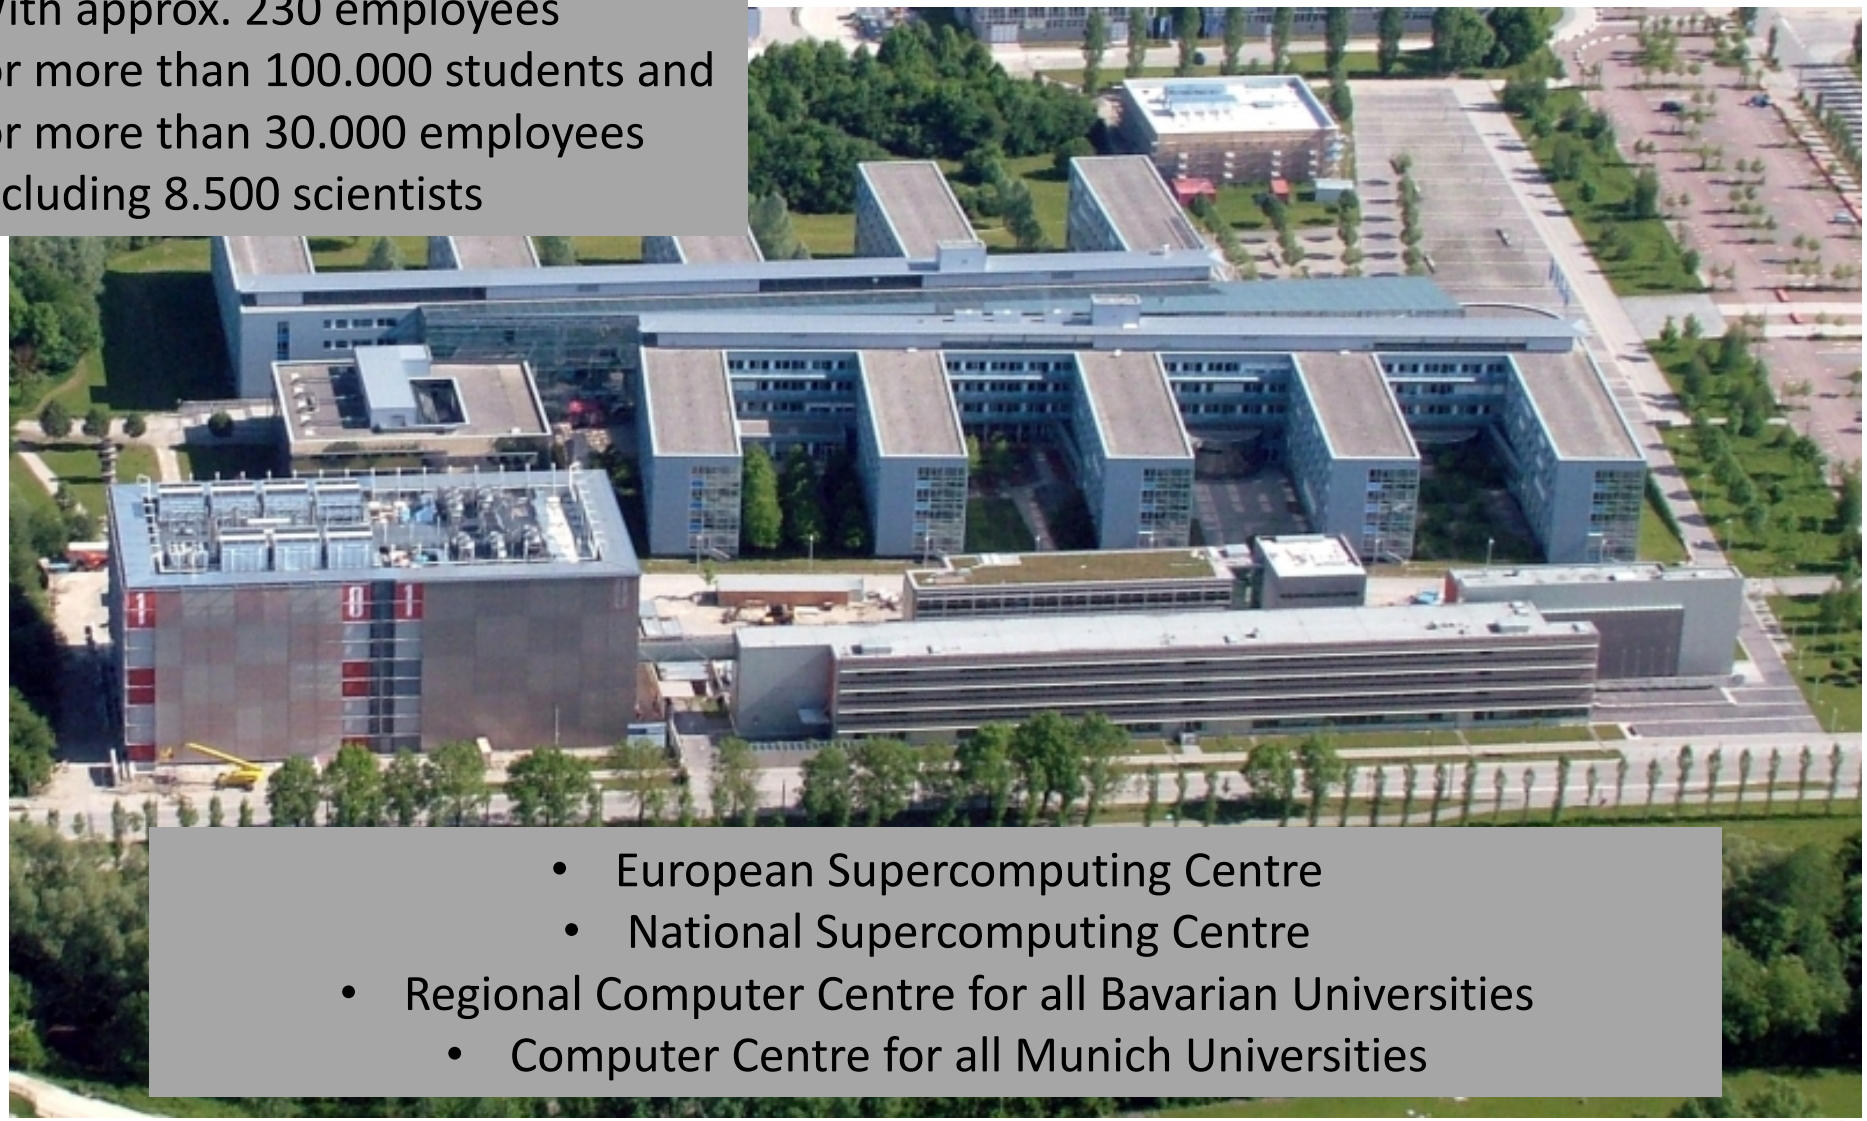

- European Supercomputing Centre
- National Supercomputing Centre
- Regional Computer Centre for all Bavarian Universities
- Computer Centre for all Munich Universities

Photo: Ernst Graf

SuperMUC Phase 1 + 2

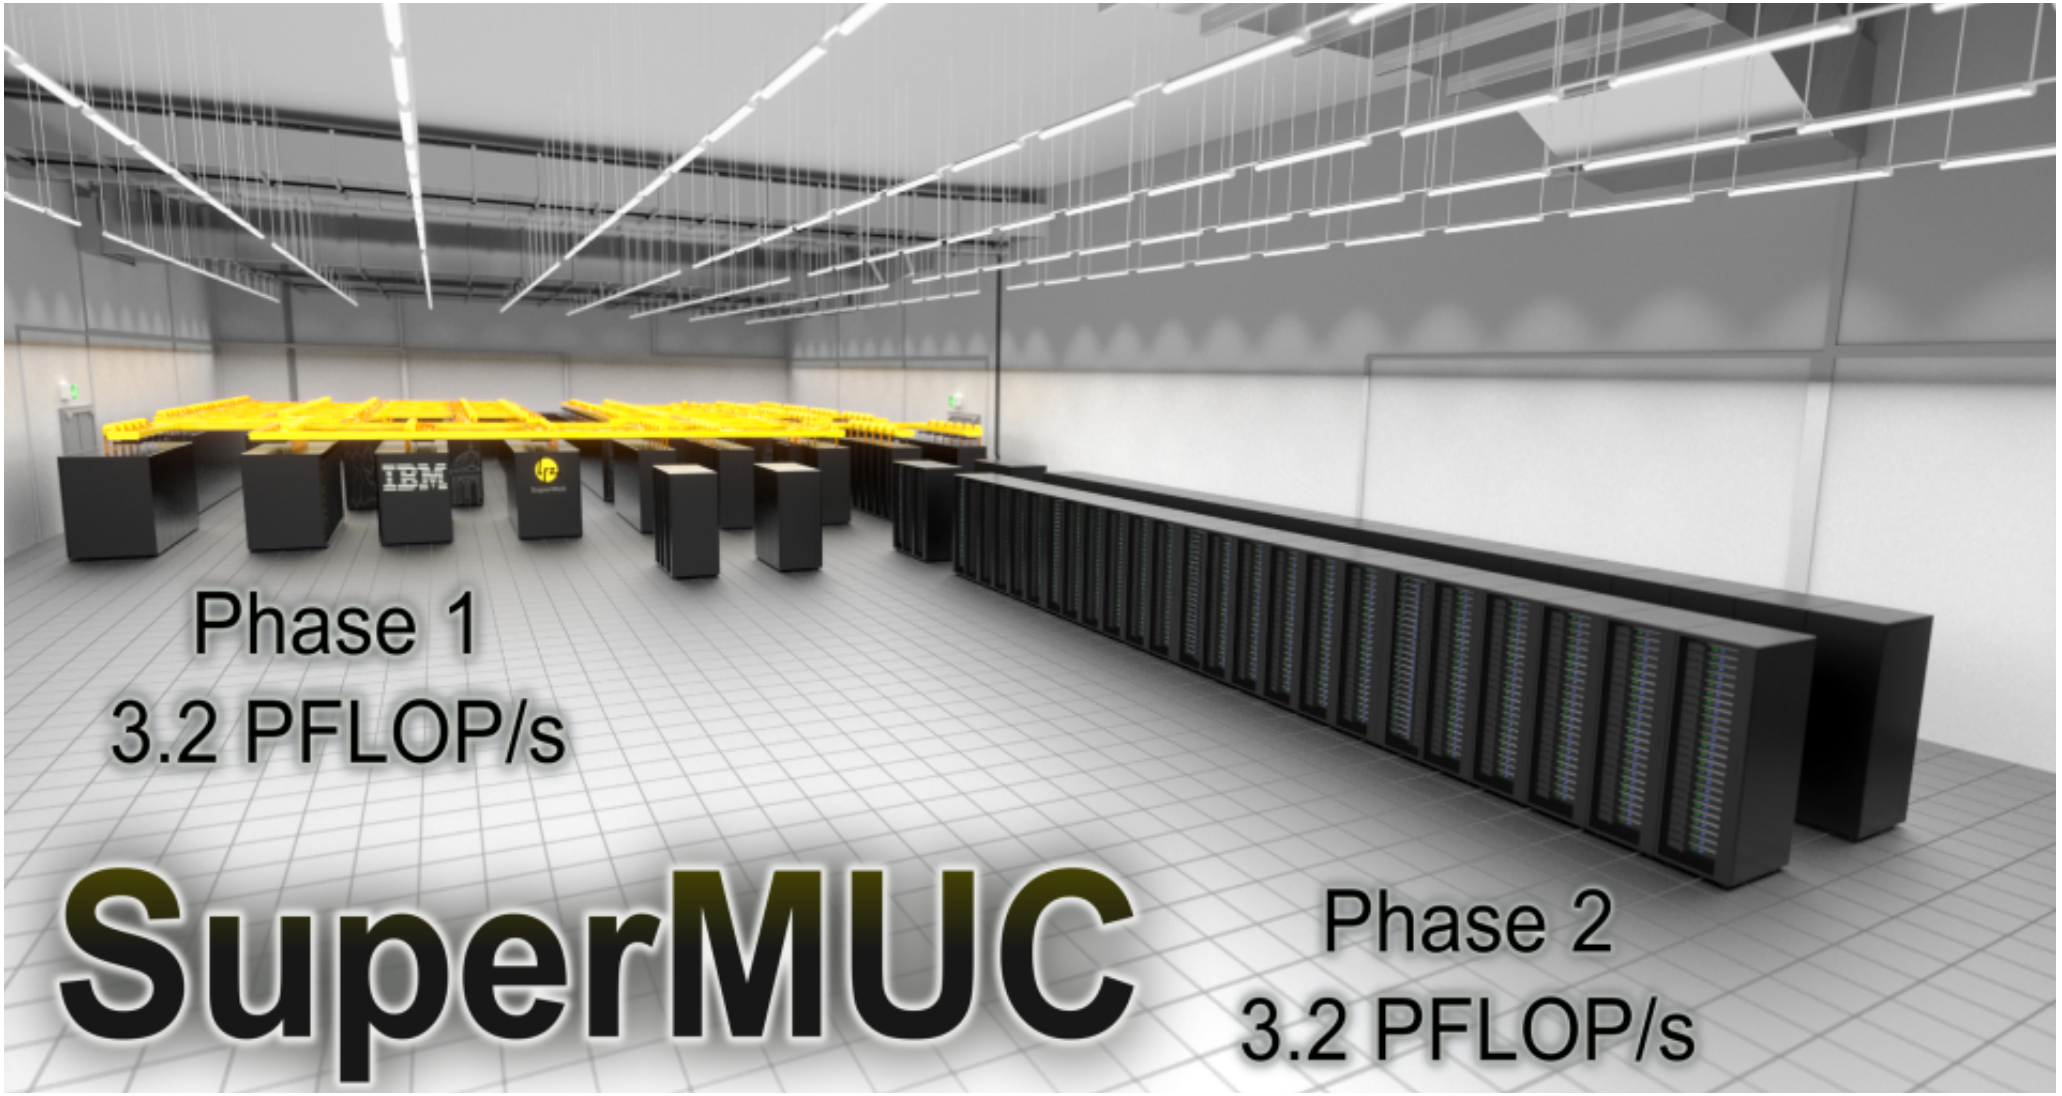

Top 500 Supercomputer List (June 2012)

Rank	Site	Computer/Year Vendor	Cores	R_{max}	R_{peak}	Power
1	DOE/NNSA/LLNL United States	**Sequoia** - BlueGene/Q, Power BQC 16C 1.60 GHz, Custom / 2011 IBM	1572864	16324.75	20132.66	7890.0
2	RIKEN Advanced Institute for Computational Science (AICS) Japan	K computer, SPARC64 VIIIfx 2.0GHz, Tofu interconnect / 2011 Fujitsu	705024	10510.00	11280.38	12659.9
3	DOE/SC/Argonne National Laboratory United States	**Mira** - BlueGene/Q, Power BQC 16C 1.60GHz, Custom / 2012 IBM	786432	8162.38	10066.33	3945.0
4	Leibniz Rechenzentrum Germany	**SuperMUC** - iDataPlex DX360M4, Xeon E5-2680 8C 2.70GHz, Infiniband FDR / 2012 IBM	147456	2897.00	3185.05	3422.7
5	National Supercomputing Center in Tianjin China	**Tianhe-1A** - NUDT YH MPP, Xeon X5670 6C 2.93 GHz, NVIDIA 2050 / 2010 NUDT	186368	2566.00	4701.00	4040.0
6	DOE/SC/Oak Ridge National Laboratory United States	**Jaguar** - Cray XK6, Opteron 6274 16C 2.200GHz, Cray Gemini interconnect, NVIDIA 2090 / 2009 Cray Inc.	298592	1941.00	2627.61	5142.0
7	CINECA Italy	**Fermi** - BlueGene/Q, Power BQC 16C 1.60GHz, Custom / 2012 IBM	163840	1725.49	2097.15	821.9
8	Forschungszentrum Juelich (FZJ) Germany	**JuQUEEN** - BlueGene/Q, Power BQC 16C 1.60GHz, Custom / 2012 IBM	131072	1380.39	1677.72	657.5
9	CEA/TGCC-GENCI France	**Curie thin nodes** - Bullx B510, Xeon E5-2680 8C 2.700GHz, Infiniband QDR / 2012 Bull	77184	1359.00	1667.17	2251.0
10	National Supercomputing Centre in Shenzhen (NSCS) China	**Nebulae** - Dawning TC3600 Blade System, Xeon X5650 6C 2.66GHz, Infiniband QDR, NVIDIA 2050 / 2010 Dawning	120640	1271.00	2984.30	2580.0

www.top500.org

SuperMUC System @ LRZ

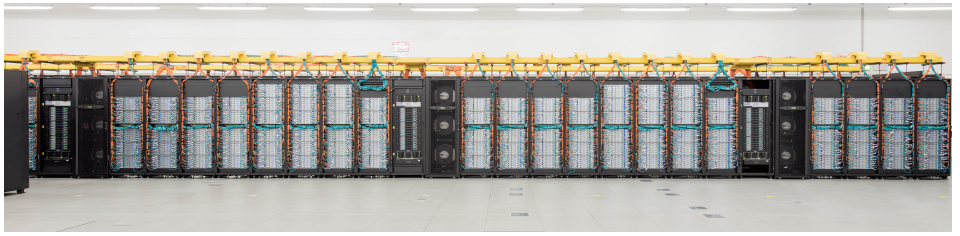

Phase 1 (IBM System x iDataPlex):

- 3.2 PFlops peak performance
- 9216 IBM iDataPlex dx360M4 nodes in 18 compute node islands
- 2 Intel Xeon E5-2680 processors and 32 GB of memory per compute node
- 147,456 compute cores
- Network Infiniband FDR10 (fat tree)

Phase 2 (Lenovo NeXtScale WCT):

- 3.6 PFlops peak performance
- 3072 Lenovo NeXtScale nx360M5 WCT nodes in 6 compute node islands
- 2 Intel Xeon E5-2697v3 processors and 64 GB of memory per compute node
- 86,016 compute cores
- Network Infiniband FDR14 (fat tree)

Common GPFS file systems with 10 PB and 5 PB usable storage size respectively
Common programming environment
Direct warm-water cooled system technology

LRZ Application Mix

- Computational Fluid Dynamics: **Optimisation of turbines** and wings, noise reduction, air conditioning in trains
- Fusion: Plasma in a future fusion reactor (ITER)
- Astrophysics: Origin and evolution of stars and galaxies
- Solid State Physics: Superconductivity, surface properties
- **Geophysics: Earth quake scenarios**
- Material Science: Semiconductors
- Chemistry: Catalytic reactions
- Medicine and Medical Engineering: Blood flow, aneurysms, air conditioning of operating theatres
- **Biophysics:** Properties of viruses, genome analysis
- **Climate research: Currents in oceans, hydrometeorology**

Link with Computer Science and management: MNM-Team

- MNM-Team history
  - Established 25 years ago, LMU, TUM, LRZ
  - One of the first groups to address IT management
  - People processes behind 80 percent of mission critical IT service downtime
- Research interests
  - Manageability of networked systems: concepts, tools, processes
  - From basic IT research to providing research IT services (code to consulting)
- Ongoing activities
  - PiCS partnership: redefining the interface between computational scientist and supercomputing
  - Environmental computing: supporting interdisciplinary production of „actionable knowledge“

UNISDR – The United Nations Office for Disaster Risk Reduction

https://www.unisdr.org/

GAR – Global Assessment Report on Disaster Risk Reduktion 2015

Global Assessment Report on Disaster Risk Reduction 2015

Making development sustainable: The future of disaster risk management

Home | Pocket GAR | **GAR 2015 Main Report** | Documents | Data | Download | Press

Foreword | At a glance | Preface | Introduction | Part I | Part II | Part III | Backmatter

Most disasters that could happen have not happened yet.

Economic losses from disasters such as earthquakes, tsunamis, cyclones and flooding are now reaching an average of US$250 billion to US$300 billion each year. **Future losses** (expected annual losses) are now estimated at US$314 billion in the built environment alone. **This is the amount that countries should set aside each year to cover future disaster losses.** (→ Chapter 3)

http://www.preventionweb.net/english/hyogo/gar/2015/en/home/GAR_2015/GAR_2015_6.html

Number of Disasters per Region

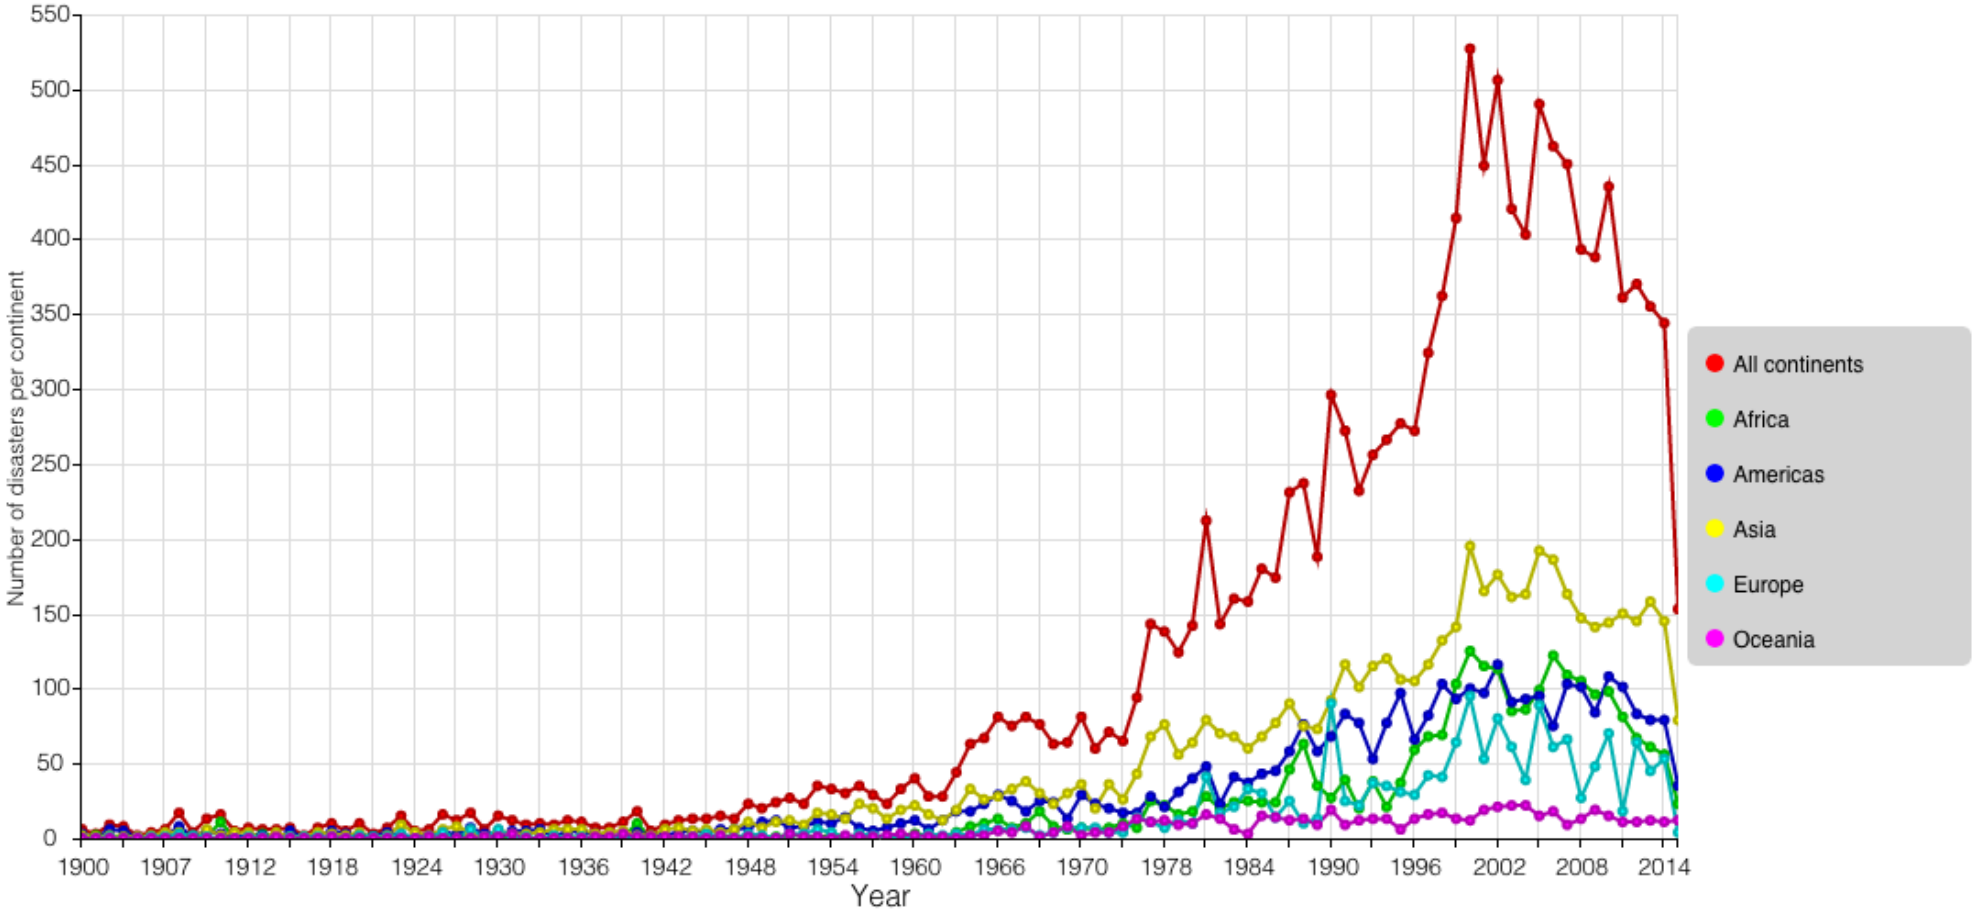

http://www.emdat.be/disaster_trends/index.html

Munich Re – Loss Events Worldwide 2014

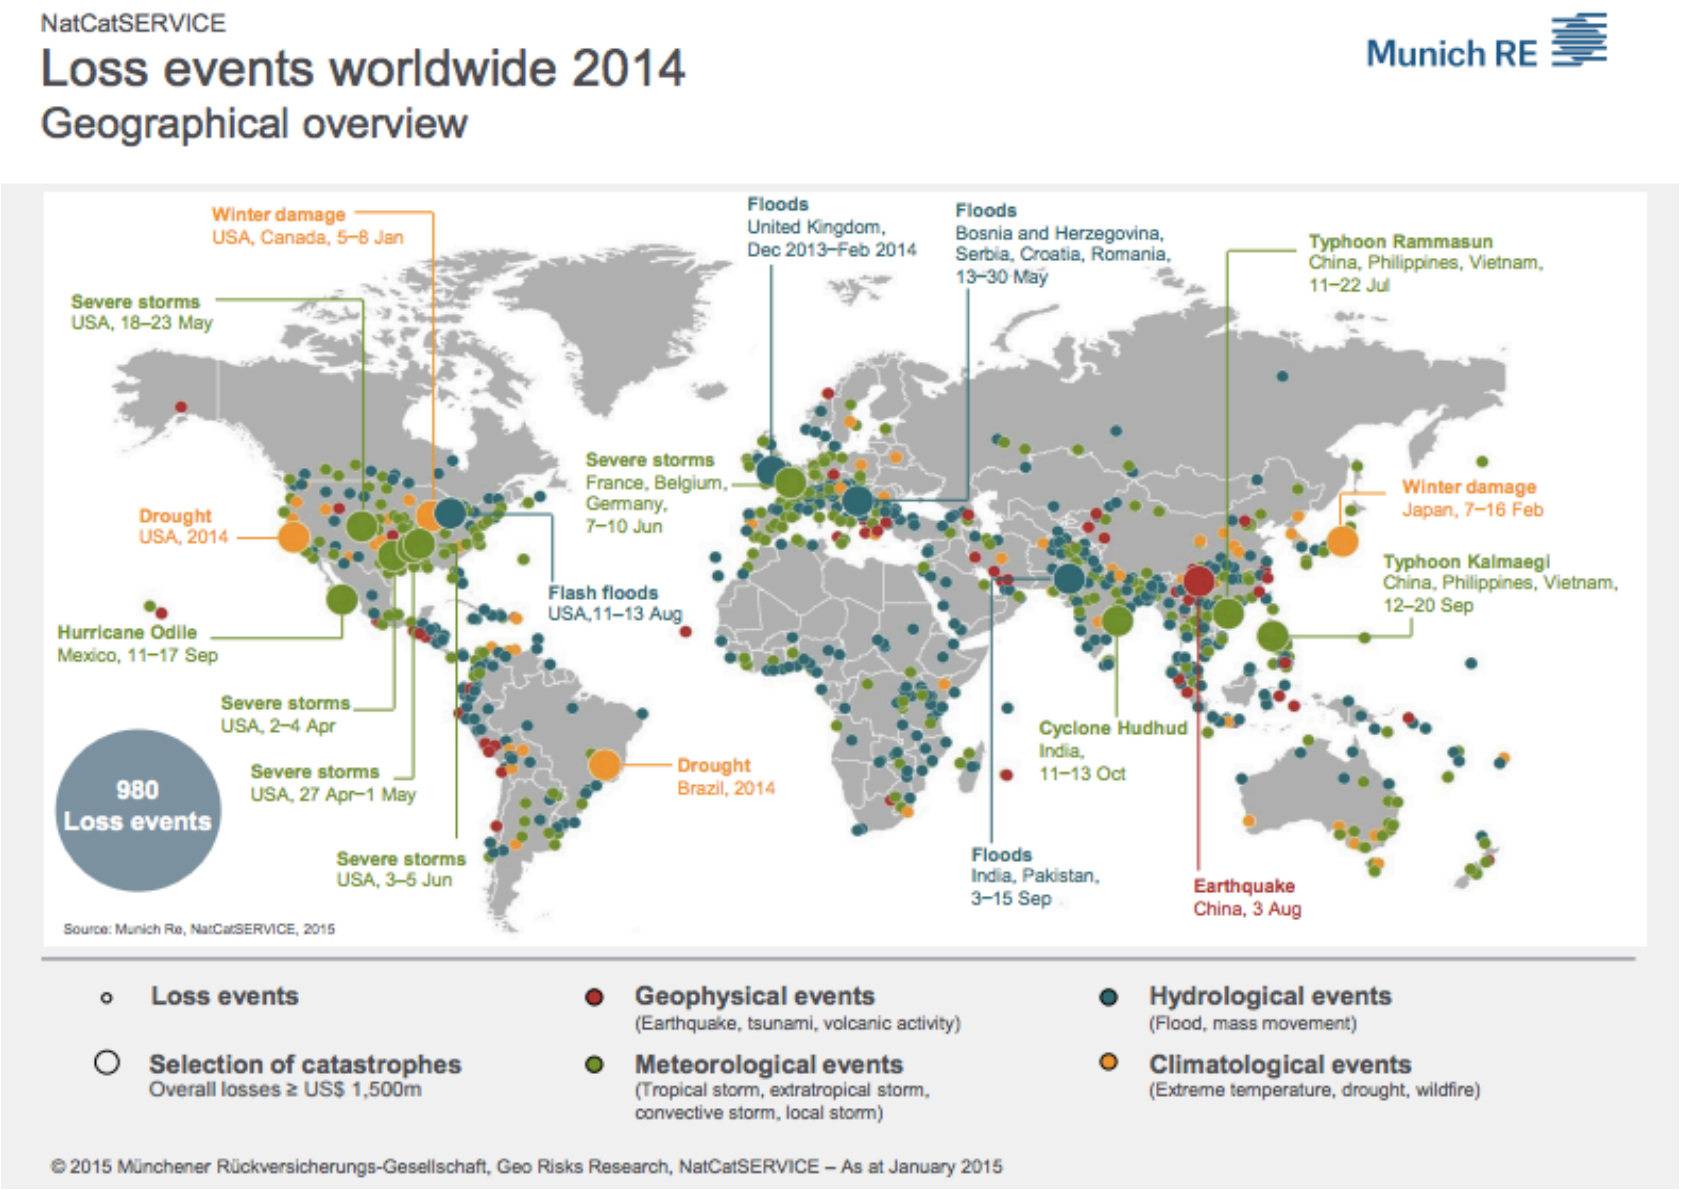

http://www.preventionweb.net/files/41773_munichreworldmapnaturalcatastrophes.pdf

What is environmental computing?

- …and why LRZ & MNM-Team are interested?

What if we could predict flash flood?

- From no warning to even a few hour's warning
  - Prevent loss of life
  - Reduce material damages considerably
- What was needed?
  - Computing capacity
  - Model chains
- Cans of worms opened:
  - Every model requires different environment (OS, libraries,…)
  - Every model describes these requirements differently
  - Data standards tend to be different
  - …and described in bespoke way
  - Optimal hardware varies

Case DRIHM

- Workflow linking rainfall, discharge and water level and flow
- ”Plug and play” framework for models and data
- Case studies demonstrating capability for better advance warnings

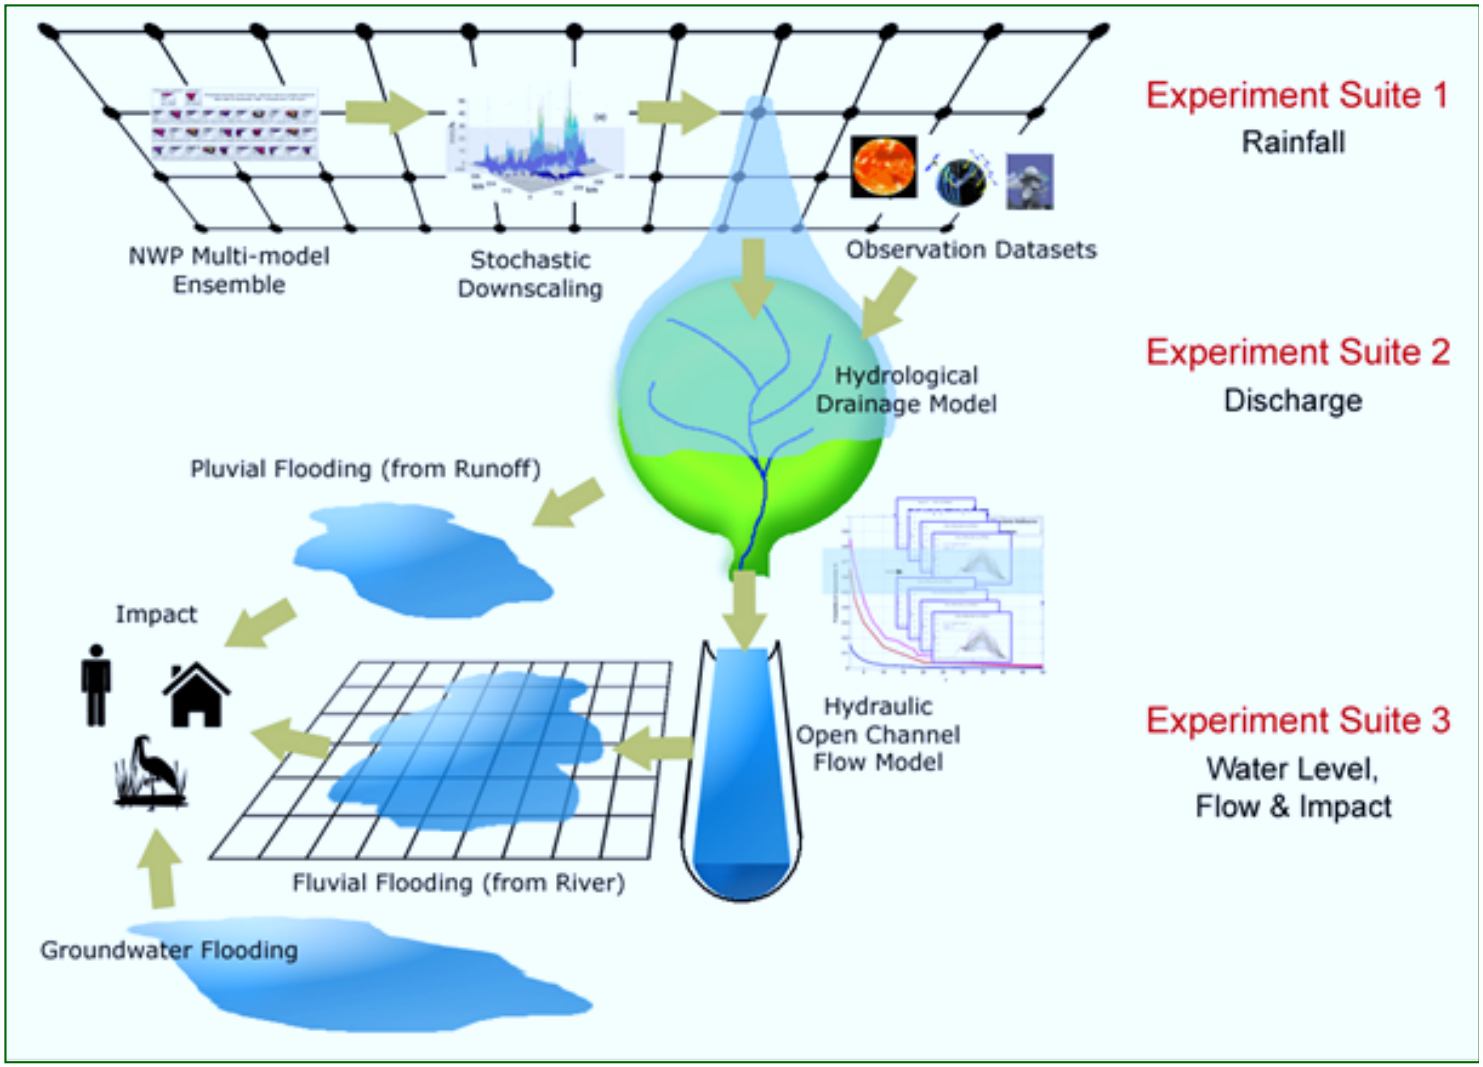

The origin of the idea – déjà vu

- Most of the inter- or transdisciplinary projects merging their solutions into new large-scale services go through the similar process in understanding
  - Relationships between the components
  - Relationship with the components of the IT infrastructure
  - Understanding who is the “customer”
  - Consensus about the high-level service description
  - Semantics of the “glue” linking components together
- The exact solutions tend to be different, but already shared awareness of the issue helps
- Community building as a way to catalyse (eventual) standardisation

Working definition of environmental computing?

- Producing **actionable knowledge** related to environmental phenomena using advanced modelling approaches – e.g. **multi-model, multi-scale, multi-data**
  - Using non-trivial computing resources
  - Service orientation: reusable solution, “robust” in different contexts

Why start a „new discipline“

- Having a name
  - Brings people together
    - Including people in funding agencies and industry
  - Reduces academic career risk of doing something inter-disciplinary
    - “Fringe Dwellers” seen as essential by practitioners
    - Formal rules: underperforming or non-relevant
  - Sparks the development of common body of knowledge

Success story based on similar transdisciplinary initiative

- History of medical informatics
  - Roots in the 1950's - Computerised EKG analysis, electronic patient records
  - ACM SIG in Biomedical Computing 1960
  - Name of the field discussed throughout the 70's
- Structures
  - First associations in 80's
  - Curricula recommendations in 90's
- Situation today:
  - Market size estimates between 6,5 and 12,5 b$ (2012, 2015)
  - Thousands of registered members in professional associations
  - Recognised specialty in recruitment

Environmental computing conclusions

- Supporting advanced environmental modelling requires new approaches
  - To manage **both** the interface to IT services **and** between different specialisations developing models
- Networking initiative rather than formal definition
  - We expect definition to emerge gradually through shared experiences and cross-pollination
- Development of the IT platforms bring opportunities and challenges
  - New IT architectures require adaptation of software
  - Adapt LRZ approaches to extreme scaling (workshops, PiCS parnership model)

- The UNISDR data and compute challenge

UNISDR GAR process

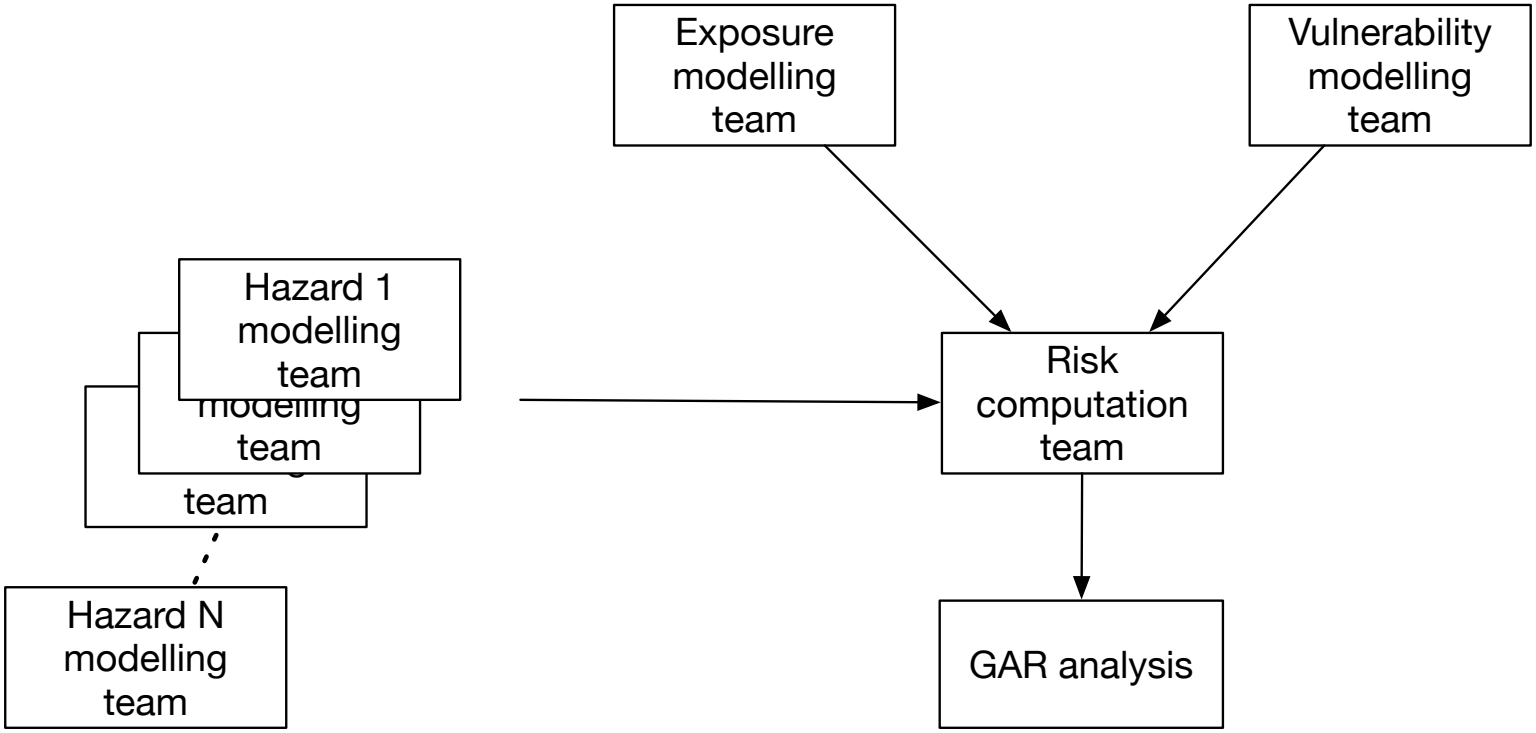

History of the UNISDR collaboration

- Informal discussions in 2014, contact through DRIHM project
- Discussion re, speeding up the loss calculation
  - "China calculation takes 5 weeks, can you lend us a supercomputer"
- Informal collaboration, calculation time to few days
  - First parallel version into operational use

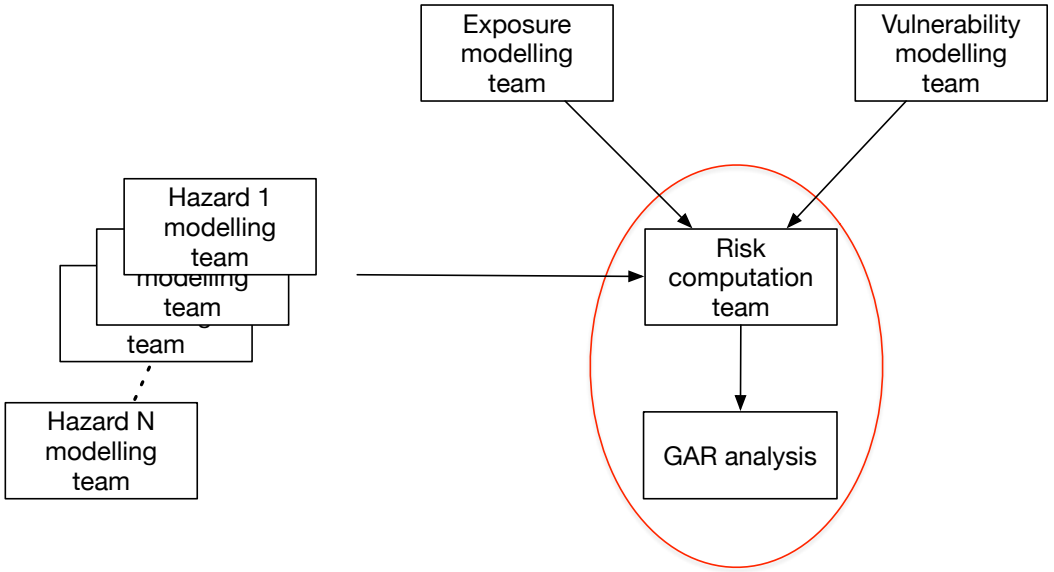

Next challenge: open GAR data

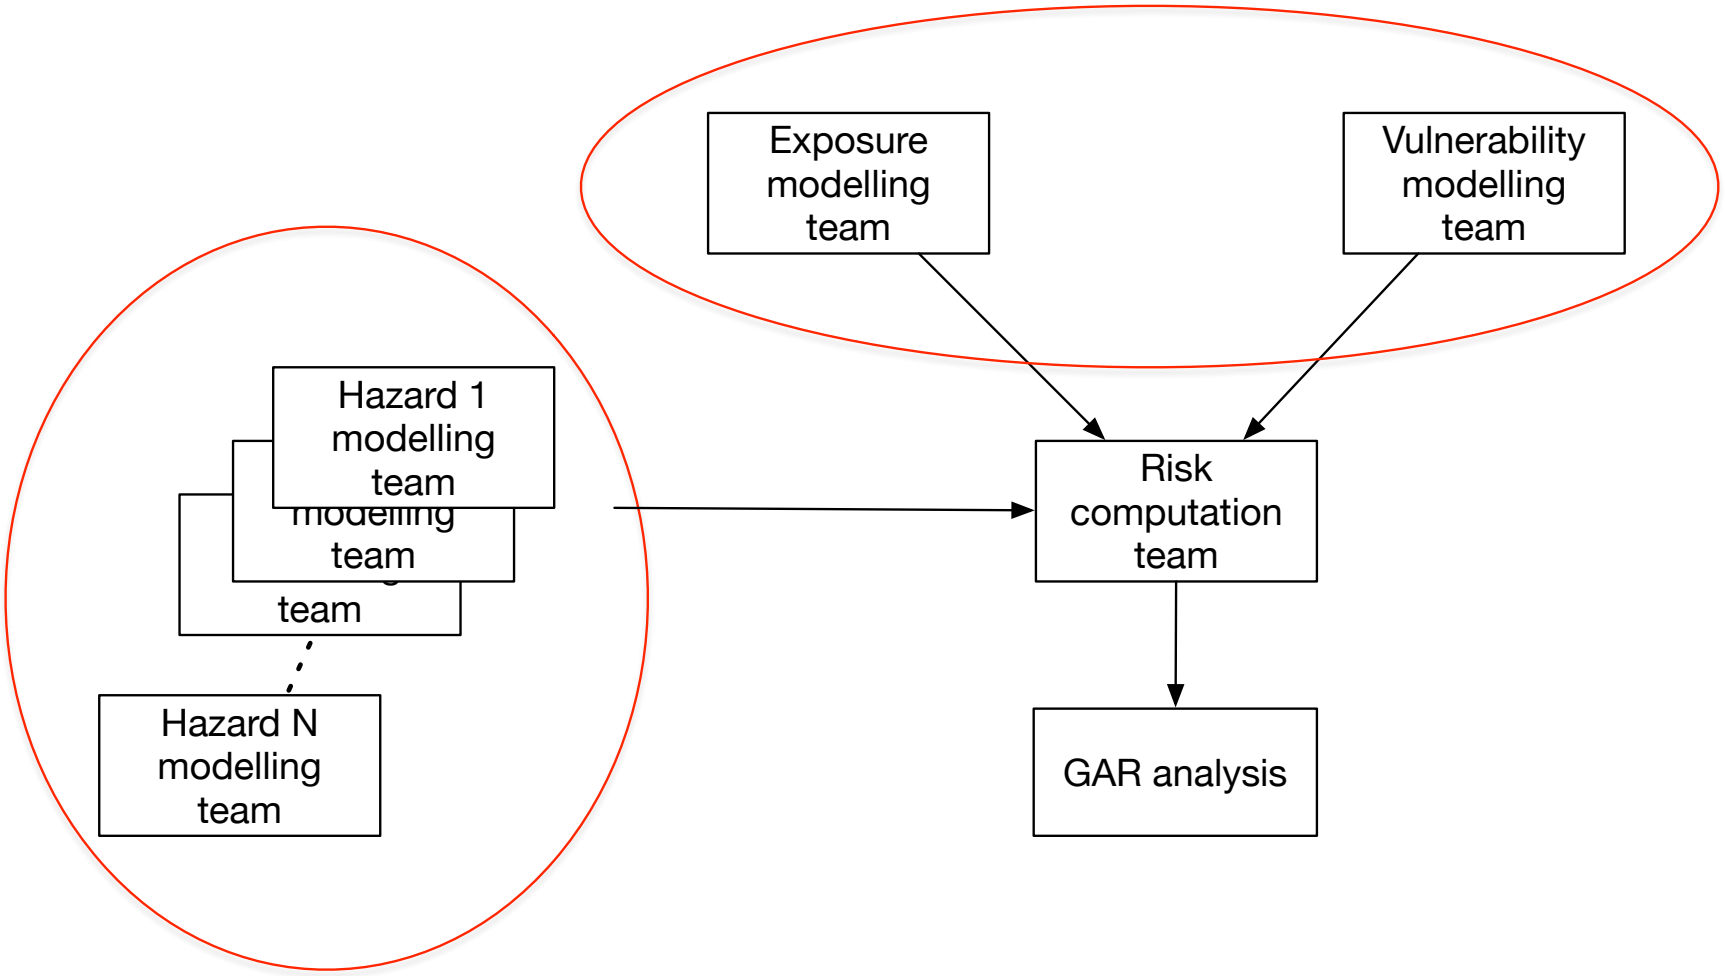

GAR open data product

- Existing demand from the community
  - Currently shared on request
  - Open data to lower the threshold of reuse
  - Main issue: sufficient infrastructure
- Current production process driven by the GAR cycle
  - New document every two years
  - Simple versioning, implicit metadata
- Future challenges
  - On-demand process, multiple versions?
  - More diverse uses, more opportunities for misunderstandings

GAR open data solution history

- Version 0
  - Simply mount the data disk on a web server managed by MNM-Team
- Version 1
  - ”Shell” around the directory hierarchy
  - New look and feel
  - Additional functionality (directory and file descriptions, download directory contents)
- Version 2
  - Based on advanced research data management systems
  - Automatic generation of metadata, workflows, versioning,…
  - Main issue: no support for current implicit metadata (directory path)
- Version 3
  - Refinement of version 1(!)

GAR data portal - beta

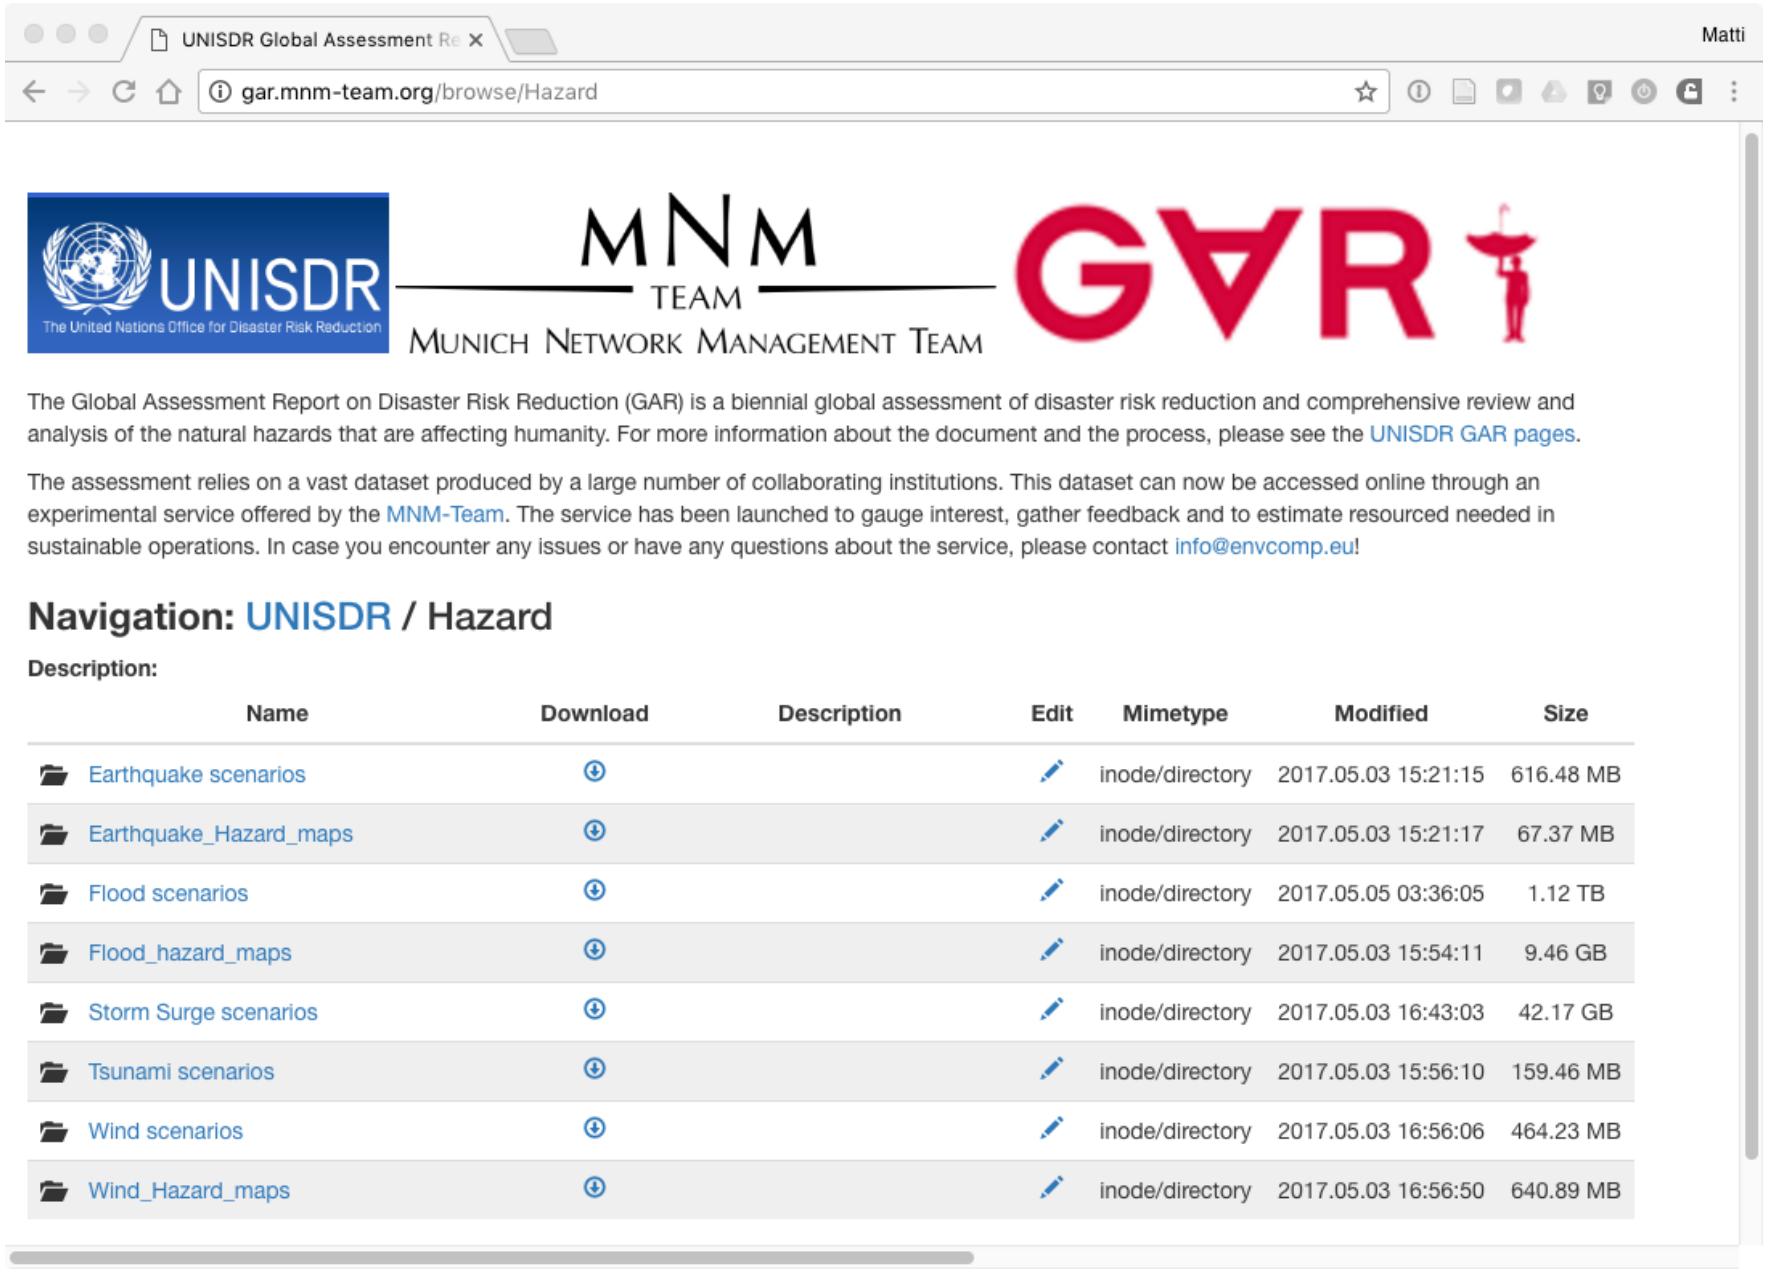

Next steps

- Observe the use of data portal
  - Who will use it?
  - How?
  - Are the new use cases in addition to expected ones?
- Adapt to the new GAR process
  - To be determined in Cancun Global Platform event
- Compute and data solution for the whole GAR lifecycle process?

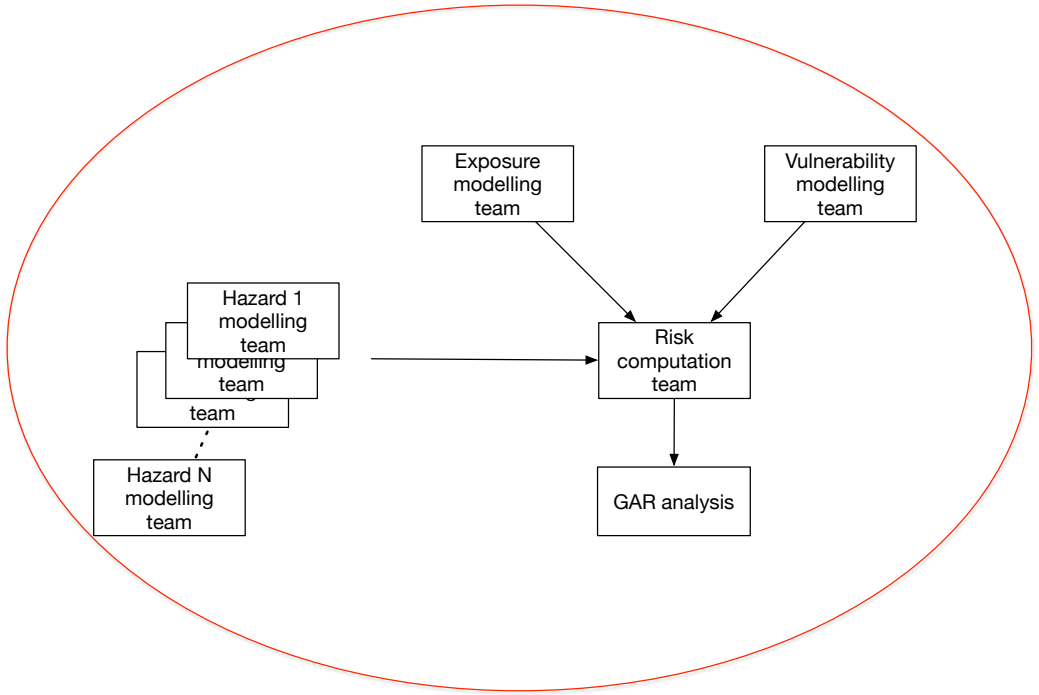

- Conclusions

Observations related to LRZ and MNM-Team role

- We understand the potential impact of advanced environmental modelling
  - And have a mandate to pursue research in the domain
- The gap between the Big Data/Supercomputing practices and the reality of most of the practitioners
  - State of the art solutions may have high up-front organisational or technical investments
  - Need to have the initial success to give time for reflection!
- LRZ and MNM-Team are promoting environmental computing as a method to fill this gap
  - ”Branding” to make interest visible

Data and compute challenge

- Between everyday data and Big Data exists the “Awkward data” domain
  - TB range
  - Typically collaborative efforts with emerging standards and practices etc).
  - UNISDR a perfect example
- Computing challenge similar
  - The road from a single threaded or shared memory parallelism to efficient cluster/supercomputing approaches is challenging
  - Effort/return ratio is very unlinear!
- There is probably no single technical solution to this challenge, but organisation/process the key
  - “PiCS approach to environmental computing”
  - Workshops
  - Networking

Upcoming events, contact

- UNISDR Global Platform
  - 22-26 May, Cancun, Mexico
  - http://www.unisdr.org/conferences/2017/globalplatform
- Upcoming environmental computing workshops
  - ICCS, June 14th 2017 (http://www.envcomp.eu/ICCS17)
  - Enviroinfo 2017, Luxemburg 13-15 September
    - Workshop "Applied Environmental Modelling – Operation and Impact" (http://www.enviroinfo2017.org/)
  - eScience 2017, 24-27 October, Aucland, New Zealand (http://escience2017.org.nz/)
- Contact and more information
  - info@envcomp.eu
  - Envcomp website: www.envcomp.eu