

## **GSI+FAIR Computing Meeting am 02.10.2012**

### **Agenda**

Ort: Seminarraum Theorie

Zeit: 10:00 Uhr bis 12:00 Uhr

1. Begrüßung durch Herrn Lindenstruth
2. Annahme des Protokolls vom letzten Meeting am 17.04.2012
3. Frau Miers spricht zu den Themen:
  - Oracle
  - Dark Fiber
  - Dokumentenmanagement
4. Herr Schön spricht zu den Themen:
  - Domainstruktur
  - Neue Massenspeicher
5. Herr Lindenstruth spricht zu den Themen:
  - Planungsstand Green Cube
  - Projizierte Leistungsaufnahme
6. Sonstiges
  - Nächster Termin

# Data Base Strategy

- PostgreSQL is the new standard data base for GSI for all new applications
- a productive PostgreSQL installation will be available in the very near future (managed by HPC)
- existing applications should be migrated if possible (IT will do everything possible, work already started)



# Future of Oracle@GSI

- The GSI Oracle installation will be available for existing applications and products that can not be migrated or linked to another database
- Within the next year the department Controls/Ralph Bär will take over the administration of the GSI Oracle installation (parallel to the accelerator installation)

# Identity Management and others...

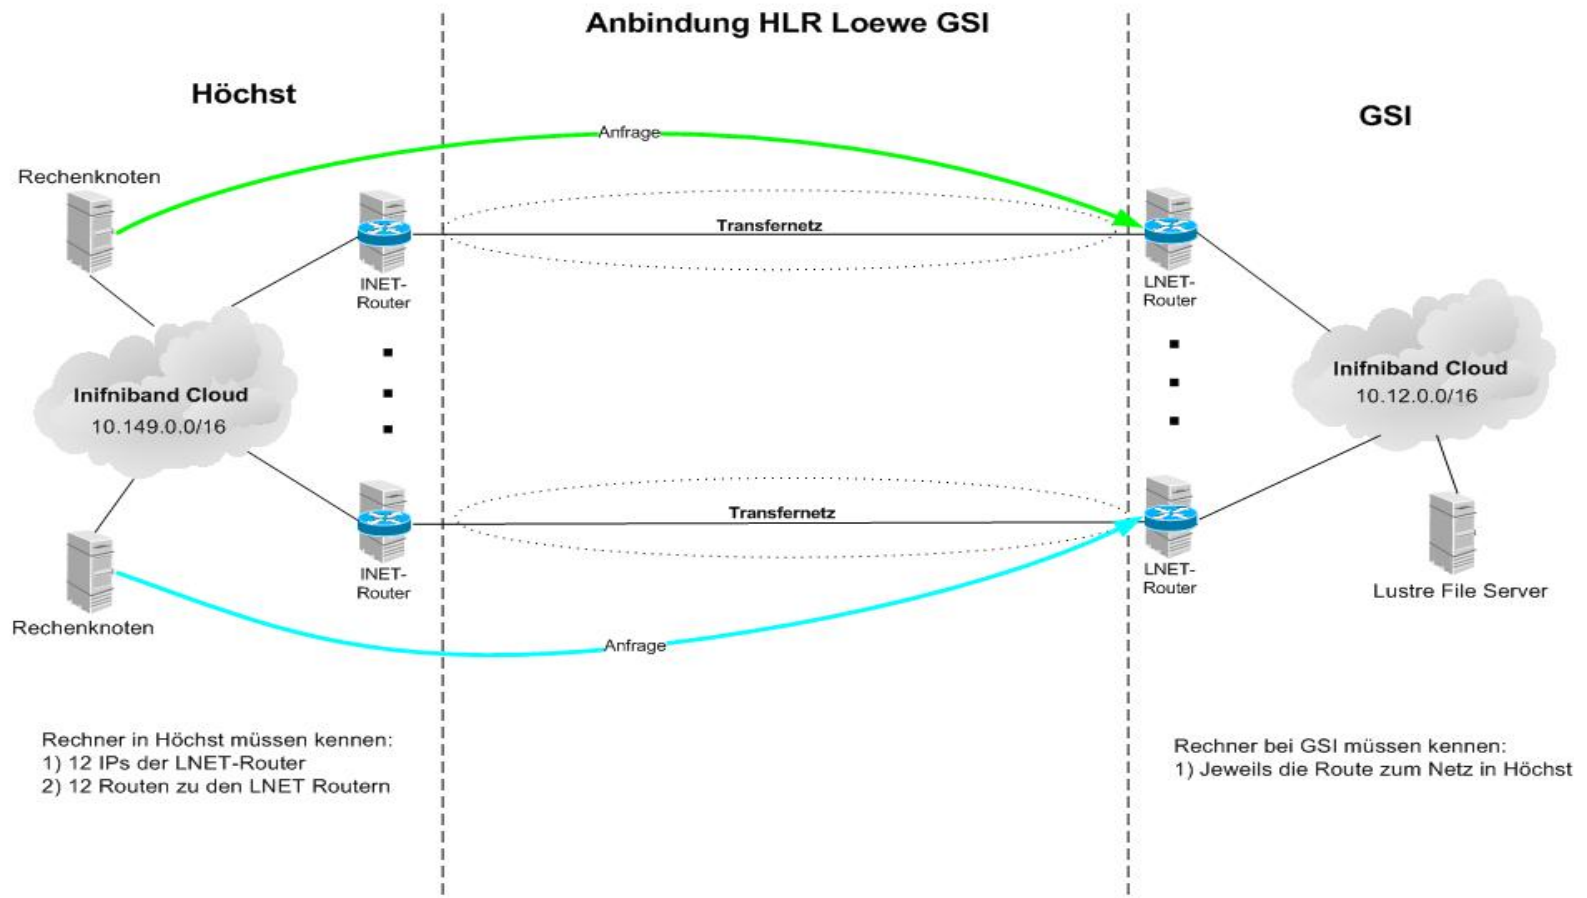
- A result of an evaluation by BWI IT recommends that the administration IT should become more homogenous. This includes that in-house developments should be replaced by standard software
- Same applies to identity management (evaluation of commercial tools started)

# Dark Fiber Connection to Frankfurt

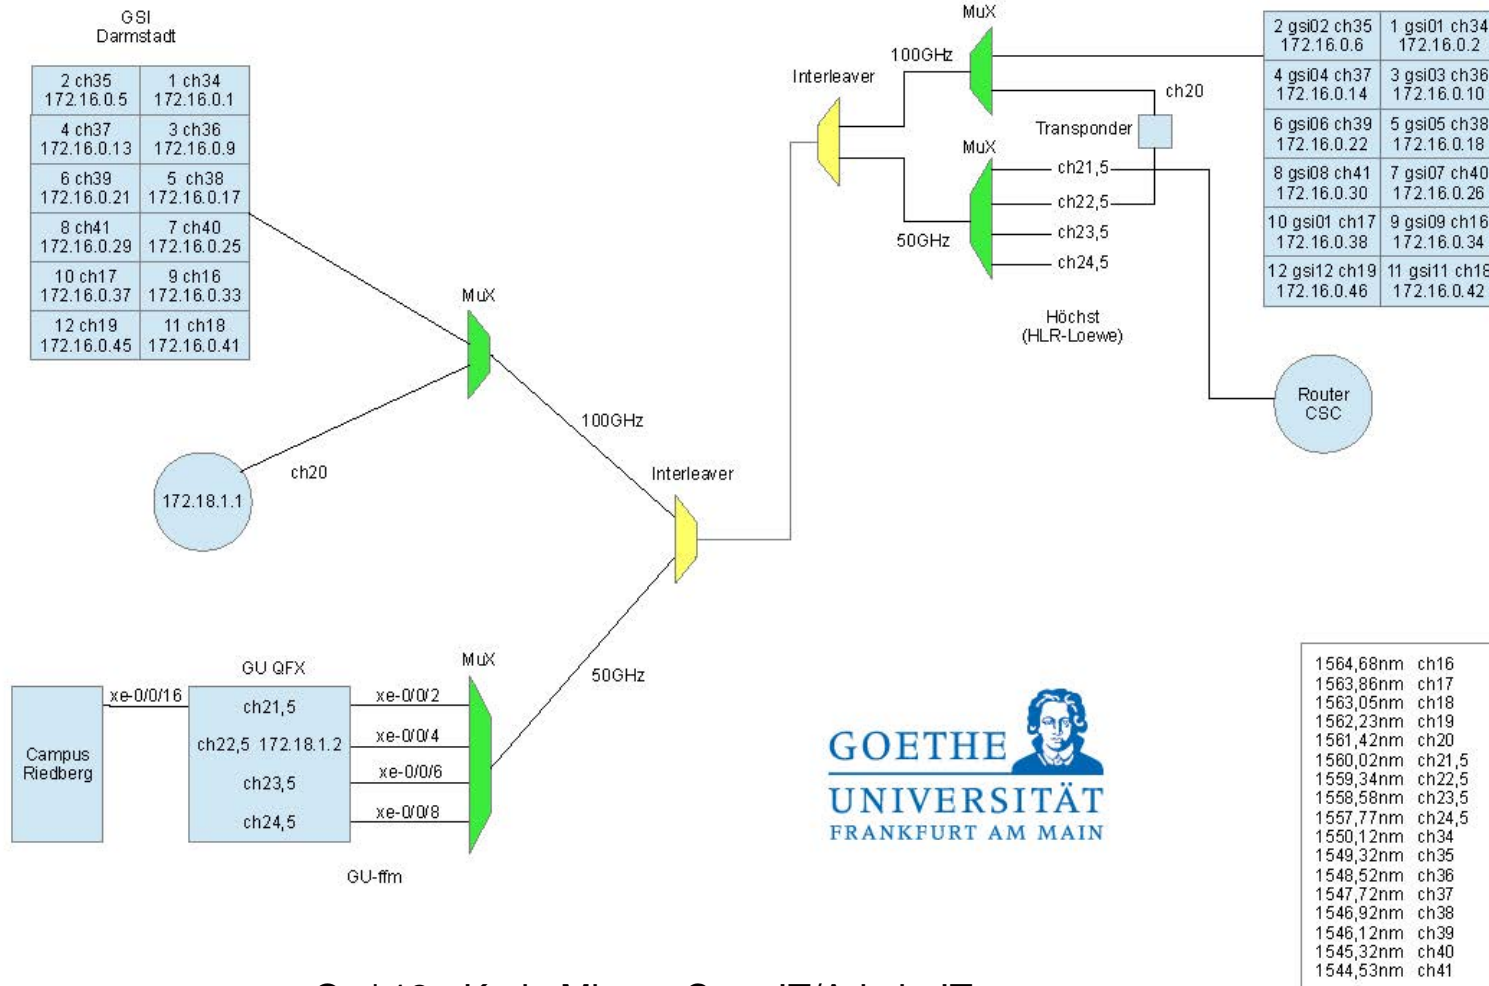
In production since August (bid invitation April, order May/Pan Dacom, delivery July):

- 12x10 Gb LOEWE CSC <-> GSI Lustre
- 10 Gb Uni FFM <-> GSI Lustre
- Using one dark fiber and dense wavelength-division multiplexing technology (DWDM)

# Dark Fiber Connection to Frankfurt



# Dark Fiber Connection to Frankfurt



# Document Management System

- few comments and input
- different presentations at GSI
- very little input and comments from research
- 5 company presentations, one very good, one good, others ~
- first design technical specification for tender, distributed to all interested
- revision of the tender criteria by consultant
- start of bid invitation very soon (within next weeks)



# News from the HPC Department - new people :-)





## - new people :-)

- Victor Penso ( SC -> HPC, Jan 2012)
- Matthias Pausch (TU -> IT Security, April 2012)
- Birgit Alborea (SC -> HPC, 4 months)
- Thomas Stibor (TU Munic → HPC, 1 Month)
- Matteo Dessalvi (HPC) ( U Sardinia → HPC, 1 day )



Thomas Stibor



Matthias Pausch



# News from Prometheus / Hera - IB cluster in the Minicube

- Prometheus: The IB cluster, Hera: The IB lustre
- 2048 cores from Prometheus reserved for lattice tests (HIM)
  - ongoing performance tests with lattice QCD on „bulldozer“ CPUs  
=> needs optimization
  - will be test cluster for „wheezy“ release (improved FP's)
  - Will be test cluster for slurm scheduler
- Beschleuniger Physik: Migrated to Prometheus ( parallel computing)  
=> happy with the performance both commercial code + SIS300Simu
- Hades: First production runs on Hera (successful)
  - happy with the performance
  - thanks to „fair share“, HADES got a lot of computing power for the current DST production  
=> Alice suffered :-(), Beschleuniger Physik suffered :-(), .....
- HPC set „credit decay parameter“ after 1 Month, 50% of the credits are lost ..... as fast technical solution to solve the „problem“



**=> „political meeting“ necessary:**

**Which group gets how much resources  
to avoid conflicts in future**

**..... taking into account, that some groups like  
Beschleuniger Physik or Hades or Theory  
put additional money into the project**

- In addition: Maybe before a „major“ run, a short notice  
Via the HPC-Info list would be a good thing (as I remember,  
Hades announced the run in the computing meeting)**



# News from lustre

- Lustre is full (thats the running gag, but is not fun ;-( )
  - especally /hera, the IB lustre
- Ways out:
  - Deleting files (Alice is on this topic ....)
  - Adding more Fileservers
    - => filling the last (electrified) Rack in Mini Cube with 4 more OSS migrated from Ethernet /lustre
    - In operation next week: adding 120TB user space
  - Adding more (migrated) OSS after Oct 11 (maintenance shutdown of Ethernet lustre (/lustre) necessary to add 14 more OSS to /hera (migrating the USVs to Minicube too)
  - Adding 420 TB more after Oct 15
- => we will add about .5 Petabyte more space for /hera soon
- New Generation of Fileserveres ordered (status: Einkauf)  
estimated time for delivery: >15 Dez
- => will add about 3 Petabyte more Space on /hera  
will be in production early 2013



## News from lustre, II

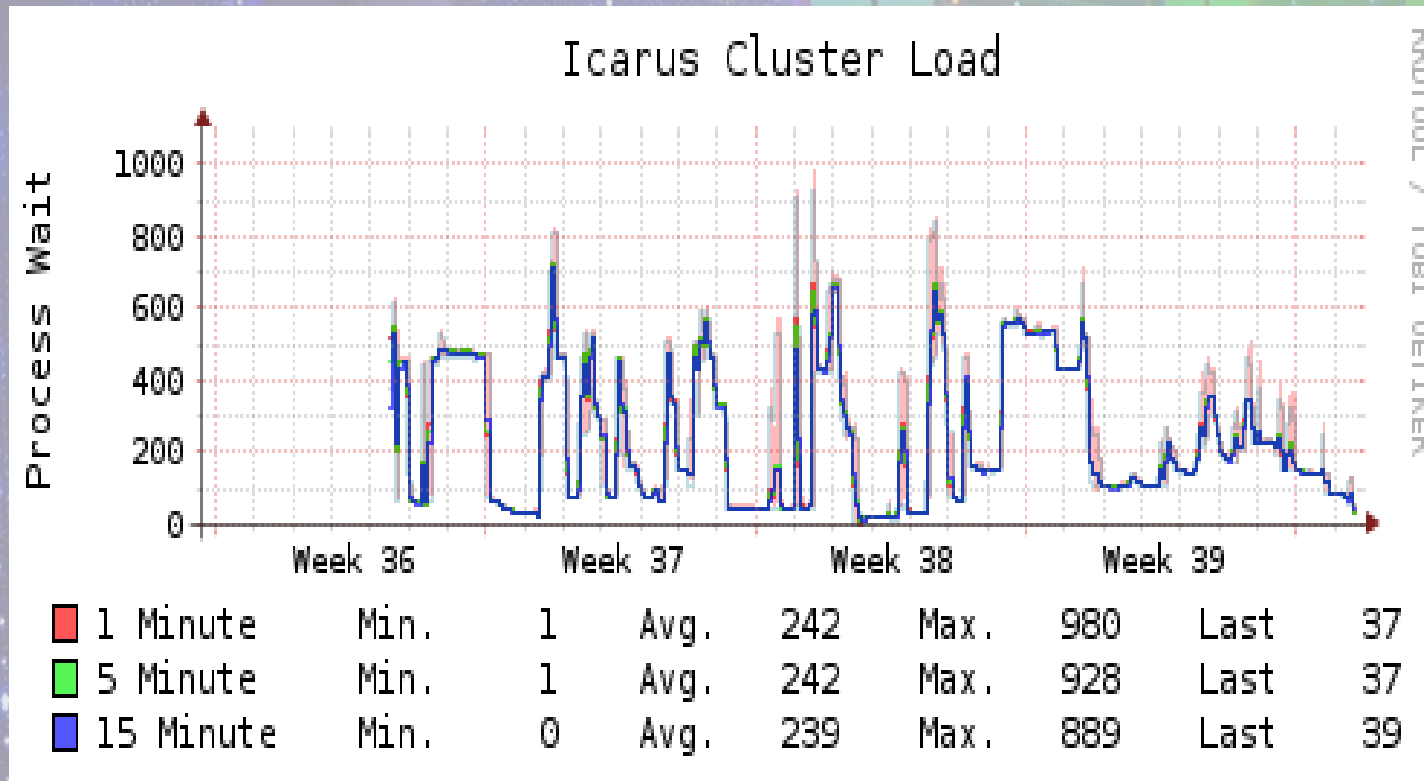
**=> /hera will be growing :-)**

**What is the future of /lustre - the Ethernet lustre?**

- Still 1.7 Petabyte, will loose capacity due to migration of the „new“ OSS (less than 3 years old) to Hera, the IB lustre**
- Will loose capacity due to aged hardware, Data on RAID 5 OSS will be migrated to RAID6 OSS, Then final shutdown ... :-) ..... next months**
- There will be still (reduced) capacity for the Ethernet cluster**

# Ethernet cluster

- Reduced interest on the Icarus cluster

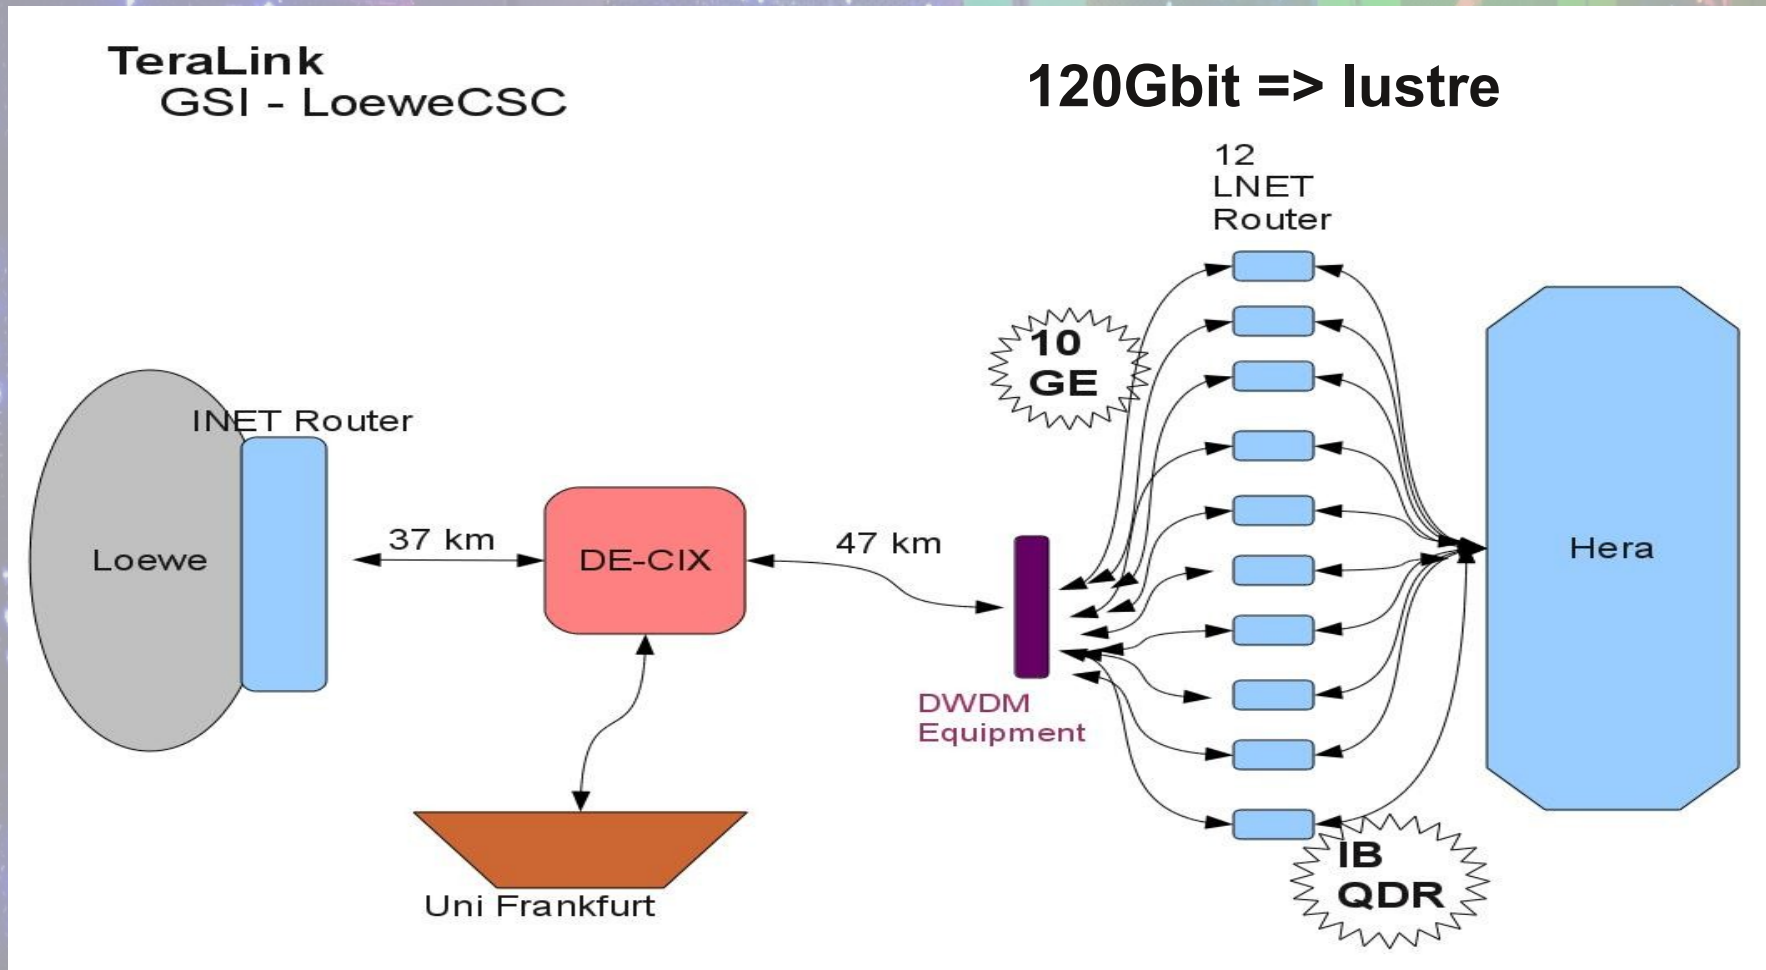


- 2.000 cores, in Peak only 1.000 used
- => we will keep an eye on this cluster, eventually migrate number crunchers to Prometheus/testing



# Connecting Loewe with Hera: First Step for the TeraLink

- Building HPC router cluster (LNET) for the TeraLink  
12 x 10 Gbit Ethernet => lustre IB loadbalancing



# New Developments II

## How to connect sites with Different GID and UID space?

- Mapping of „Alien“ GID/UIDs into a virtual lustre space (kernel)  
=> Thomas (Stibor) merged newest lustre code 2.3xx (GIT) with Kerberos + mapping features, creating Debian packages
  - „Kerberized“ communication between MDS and OSS on the Test system is already working :-)
- => test with clients + performance measurements necessary




# New Domain Structure

**=> launched by IT Security eliminate some „aged“ structure (e.g flat domain structure)**

- **Common project of IT Security, Core IT (Windows,Network), HPC**
- **Hierarchical domain structure**
- **Introducing PKI, strong authentication, smart cards encrypted documents, key recovery**
- **VMs based on KVM (Linux), secured by Linux FWs**

**Status: Working Groups on detailed concept, „pilot“ ready end 2012 (Admin IT)**





2.10.2012

# GSI+FAIR Computing

Volker Lindenstruth  
FIAS, IfI, LOEWE Professur  
Chair of HPC Architecture  
University Frankfurt, Germany  
Phone: +49 69 798 44101  
Fax: +49 69 798 44109  
Email: [voli@compeng.de](mailto:voli@compeng.de)  
WWW: [www.compeng.de](http://www.compeng.de)

# Agenda

## GSI+FAIR Computing Meeting am 02.10.2012

### Agenda

Ort: Seminarraum Theorie

Zeit: 10:00 Uhr bis 12:00 Uhr

1. Begrüßung durch Herrn Lindenstruth
2. Annahme des Protokolls vom letzten Meeting am 17.04.2012
3. Frau Miers spricht zu den Themen:
  - Oracle
  - Dark Fiber
  - Dokumentenmanagement
4. Herr Schön spricht zu den Themen:
  - Domainstruktur
  - Neue Massenspeicher
5. Herr Lindenstruth spricht zu den Themen:
  - Planungsstand Green Cube
  - Projizierte Leistungsaufnahme
6. Sonstiges
  - Nächster Termin

# Neue NIC Forschungsgruppenleiter

- **Prof. Dr. Marcus Bleicher, Direktor HIC for FAIR**
- **Prof. Dr. Owe Philipsen**


















# The FAIR Tier-0 Green Cube









# FAIR Computing Bedarf (reloaded)

Programm	2010			2012		2013			2015		2018		
	CPU # cores	On-line PB	Tape PB/a	CPU # cores	On-line PB	CPU # cores	On-line PB	Tape PB/a	CPU # cores	On-line PB	CPU # cores	On-line PB	Tape PB/a
APPA	100	< 0,1	0,1			200	1,0	2,0			1.000	1,0	2,0
CBM	250	< 0,2	0,2			7.000	15,0	11,0			60.000	15,0	11,0
NUSTAR	100	< 0,1	0,0			300	2,0	2,0			2.000	2,0	2,0
PANDA	250	< 0,2	0,2			8.000	12,0	12,0			66.000	12,0	12,0
ALICE	3.000	1,5	0,0			6.600	3,5	0,0			14.500	8,0	0,0
FOPI	200	0,1	0,2			400	0,2	0,3					
HADES	1.000	1,0	1,5			1.000	2,0	3,0					
Gitter-QCD	2.500	< 0,3	0,3			20.000	0,5	0,5			100.000	1,0	2,0
Theorie	1.000	< 0,2	0,2			5.000	0,5	0,5			20.000	1,0	2,0
rest	500	< 0,3	0,3			2.000	1,5	1,5			10.000	3,0	3,0
summe	8.900	3	3	13000 +500TF	7	50.500	38	33	?	?	273.500	43	34

**8 M€ in 2015 nach Fertigstellung des Green Cubes verfügbar**

# Rack Verteilung Green Cube

Resource	Racks 2010	Racks 2012	Racks 2018
generic IT	28	27	50
ALICE Tier-2	20	20	30
Server	17	35	50
Lustre mass storage	24	67	100
Theory, LQCD	2	20	150
CBM FLES			100
PANDA			110
other			10
<b>total</b>	<b>91</b>	<b>169</b>	<b>600</b>

# Power

Resource	Racks 2012	Leistung 2012 kW	Racks 2018	Leistung 2018 kW
generic IT	27	160	50	300
ALICE Tier-2	20	100	30	300
Server	35	110	50	500
Lustre mass storage	67	160	100	500
Theory, LQCD	20	250	150	1500
CBM FLES			100	1000
PANDA			110	1100
other			10	
<b>total</b>	<b>169</b>	<b>780</b>	<b>600</b>	<b>5200</b>



# FAIR Ereignis Rekonstruktions Workshop

**Save the date – doodle follows**

- **Mo, 29.10.**
- **Fr, 2. 11.**
- **Mo, 5.11.**
- **Mo, 12.11.**
- **Fr, 7.12.**

