



# **Computer Centre Visit**

## **Austrian Teacher's Programme 2015**

**Thursday 19<sup>th</sup> November 2015**




**Dr. Andreas Wagner**

**Deputy Group Leader - Operating Systems and Infrastructure Services**

**CERN IT Department**



# The IT Department & The Worldwide LHC Computing Grid (WLCG)



**Dr. Andreas Wagner**  
**Operating Systems and Infrastructure Services**  
**Deputy Group Leader, IT Department**

- **Mission**

- Physics Computing
- General Purpose and Administrative Computing



- **Projects**

- With Physics community (WLCG)
- With organisations, and industry

**More?**

<http://cern.ch/information-technology/about>



- **Significant Numbers**

- 7500 PCs in offices,  
2500 Macs


- **Mail:**  
~20000 mailboxes,  
6000 mailing lists,  
3 Million messages/day

- **Web:**  
Number of Web Sites ~15000  
Hits per month ~1 billion

- These services are widely used by people at CERN  
and physicist from their institutes!



Web traffic per month



# The LHC Data Challenge (I)




CERN IT Department  
CH-1211 Genève 23  
Switzerland  
[www.cern.ch/it](http://www.cern.ch/it)

Andreas Wagner, CERN, IT Department

# View of the ATLAS detector (during construction)

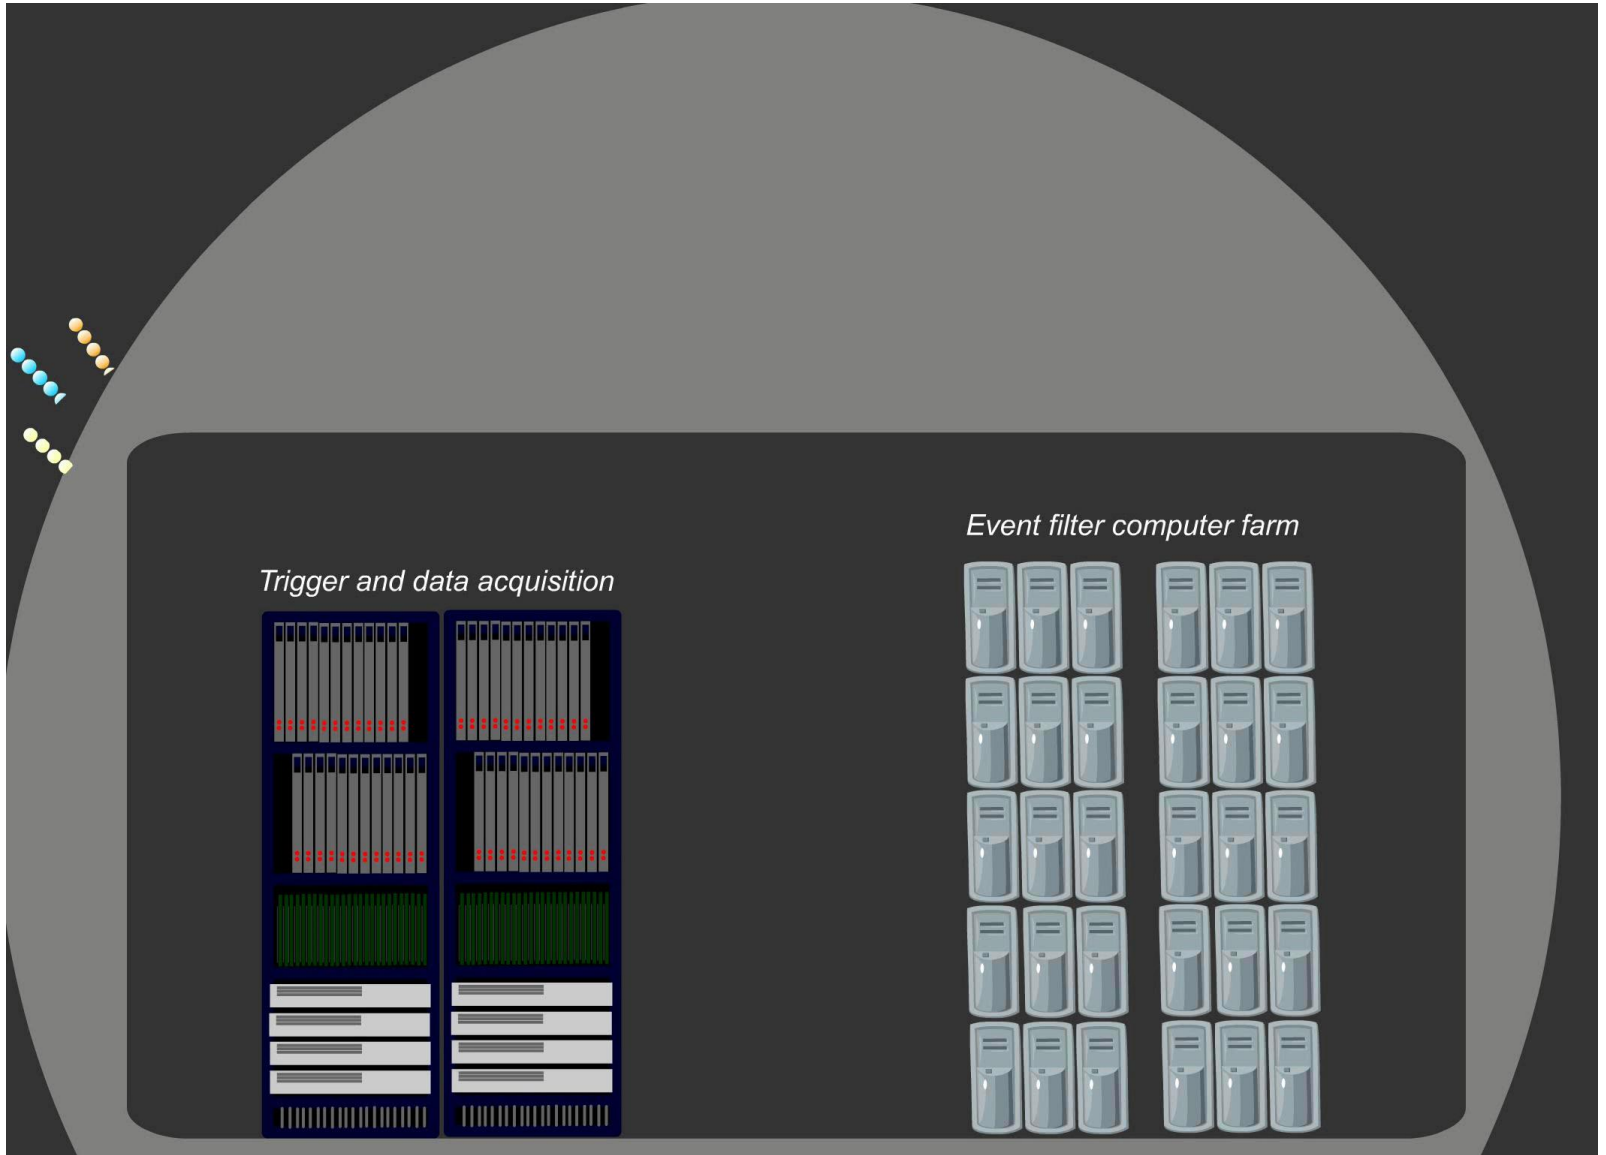
> 150 million sensors deliver data ...  
... 40 million times per second

# The LHC Data Challenge (II)






# The LHC Data Challenge (III)




# Tier 0 at CERN: Acquisition, First pass reconstruction, Storage & Distribution



# The LHC Data Challenge (III)

- The accelerator will run for 20 years
- Experiments will produce **>25 Million Gigabytes** of data each year (about 5 million DVDs, stacked would be **tower of 6 km!**)
- Total data stored at CERN 100 Pbyte (~ 700 years of HD quality movies)
- More than **480 million experiment files** stored in data centre
- LHC data analysis requires a computing power equivalent to **~250,000 of today's CPU cores**
- Requires many cooperating computer centres (>160 data centres), as CERN can **only** provide **~15% of the resources**




# Solution: the Grid

- Use the Grid to unite computing resources of particle physics institutes around the world

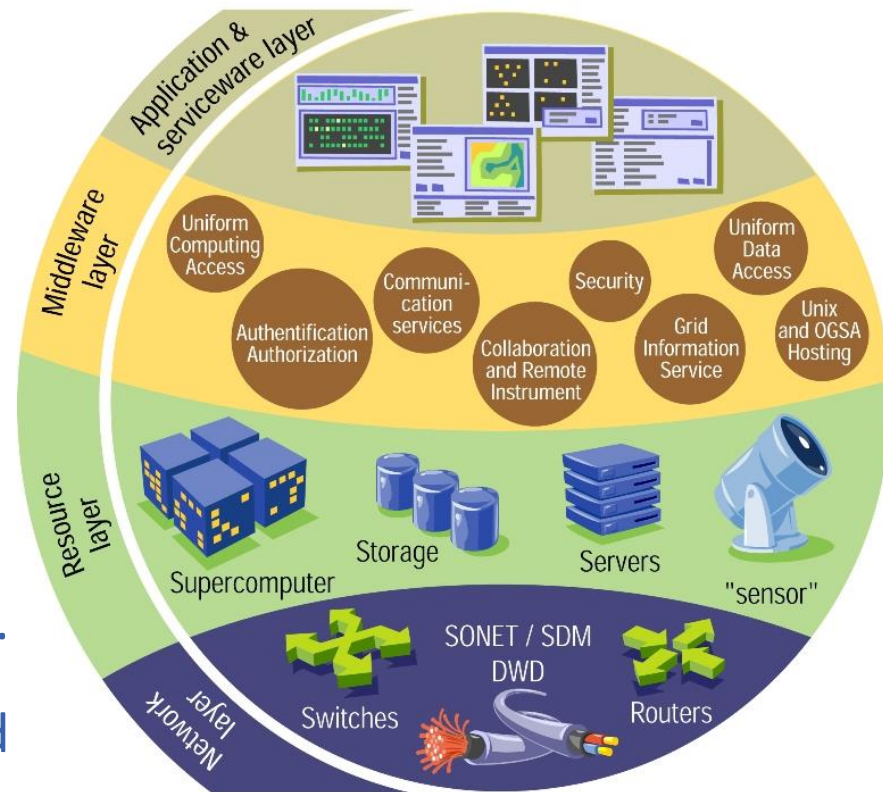
The **World Wide Web** provides seamless access to information that is stored in many millions of different geographical locations


The **Grid** is an infrastructure that provides seamless access to computing power and data storage capacity distributed over the globe



# How does the Grid work?

- It makes multiple computer centres look like a single system to the end-user
- Advanced software, called middleware, automatically finds the data the scientist needs, and the computing power to analyse it.
- Middleware balances the load on different resources. It also handles security, accounting, monitoring and much more.






- Tier-0 (CERN):**
  - Data recording
  - Initial data reconstruction
  - Data distribution
- Tier-1 (11 centres):**
  - Permanent storage
  - Re-processing
  - Analysis
- Tier-2 (~130 centres):**
  - Simulation
  - End-user analysis

**Austrian Federated Tier-2:**

- Wien: HEPHY, Austrian Academy of Science, ... (for CMS)
- Innsbruck: University of Innsbruck (for ATLAS)

# Grid in Österreich




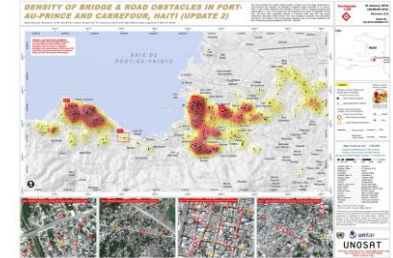
# Broader Impact of the LHC Computing Grid

- WLCG has been leveraged on both sides of the Atlantic, to benefit the wider scientific community
  - Europe:
    - Enabling Grids for E-science (EGEE) 2004-2010
    - European Grid Infrastructure (EGI) since 2010
  - USA:
    - Open Science Grid (OSG) since 2004
- Many scientific applications →

**Archeology**  
**Astronomy**  
**Astrophysics**  
**Civil Protection**  
**Comp. Chemistry**  
**Earth Sciences**  
**Finance**  
**Fusion**  
**Geophysics**  
**High Energy**  
**Physics**  
**Life Sciences**  
**Multimedia**  
**Material Sciences**



- **Collaboration with Institutions: UNOSAT**
  - Satellite image analysis for crisis response
- **Collaboration with Industry: CERN openlab**
  - Evaluates state-of-the-art technologies in a very complex environment and improves them;
  - Test in a research environment today what will be used in industry tomorrow
- **Citizen Cyberscience Centre**
  - **Computing for Clean Water** optimizing nanotube based water filters by large scale simulation on volunteer PCs
  - **AfricaMap** volunteer thinking to generate maps of regions of Africa from satellite images, with UNOSAT
  - **LHC@home** volunteer project for public participation in LHC collision simulations, using VM technology
- **Health e-Child**
  - Study biomedical information for clinical practice, medical research, and personalised healthcare for the citizens of the EU.



www.cern.ch/openlab

PARTNERS



ORACLE

SIEMENS

# The CERN Tier-0 in Numbers

- Data Centre Operations (Tier 0)
  - 24x7 operator support and System Administration services to support 24x7 operation of all IT services.
  - Hardware installation & retirement
    - ~7,000 hardware movements/year; ~1800 disk failures/year
  - Management and Automation framework for large scale Linux clusters

## Overview: Data Centre

### MEYRIN DATA CENTRE

● Number of Cores in Meyrin	119,093
● Number of Drives in Meyrin	78,130
● Number of Memory Modules in Meyrin	79,124
● Number of 10G NIC in Meyrin	4,280
● Number of 1G NIC in Meyrin	22,205
● Number of Processors in Meyrin	21,311
● Number of Servers in Meyrin	11,557
● Total Disk Space in Meyrin (TB)	143,956
● Total Memory Capacity in Meyrin (TB)	443


### WIGNER DATA CENTRE

● Number of Cores in Wigner	20,576
● Number of Drives in Wigner	15,075
● Number of Memory Modules in Wigner	10,263
● Number of 10G NIC in Wigner	1,211
● Number of 1G NIC in Wigner	2,224
● Number of Processors in Wigner	2,574
● Number of Servers in Wigner	1,290
● Total Disk Space in Wigner (TB)	47,694
● Total Memory Capacity in Wigner (TB)	83

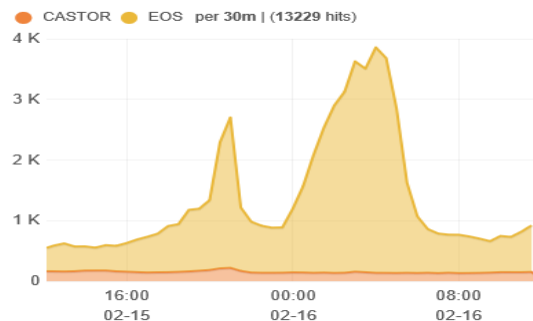
### NETWORK AND STORAGE

● Tape Drives	115
● Tape Cartridges	23,186
● Data Volume on Tape (TB)	105,023
● Free Space on Tape (TB)	30,689
● Routers (GPN)	134
● Routers (TN)	26
● Routers (Others)	99
● Star Points	633
● Switches	3,496


### # BATCH JOBS



### # ACTIVE DATA TRANSFERS




### # VM CREATED



# Scaling CERN Data Center(s) to anticipated Physics needs

CERN Data Center dates back to the 70's

- Upgraded in 2005 to support LHC (2.9 MW)
- Still optimizing the current facility (cooling automation, temperatures, infrastructure)



Renovation of the “barn” for accommodating 450 KW of “critical” IT loads (increasing 513 total to 3.5 MW)

Exploitation of 100 KW of remote facility down town

- Understanding costs, remote dynamic management, ensure business continuity



Exploitation of a remote Data center in Hungary

- Max. 2.7 MW (N+1 redundancy)
- Business continuity
- 100 Gbps connections

## CC Visit Point

- Visualisations
- Computer Museum

## Visitors Gallery


- View into upper machine room

## Foyer - Computer Museum

- GridView

## Walk Through the Computer Centre Machine Rooms

- Ground floor – Physics clusters
- Basement – Tape robots



For more information about the Grid:



[www.cern.ch/lcg](http://www.cern.ch/lcg)



[www.eu-egee.org](http://www.eu-egee.org)



[www.eu-egi.org/](http://www.eu-egi.org/)



[www.gridcafe.org](http://www.gridcafe.org)



Thank you for your kind attention!

# Thank you!



[www.cern.ch](http://www.cern.ch)