

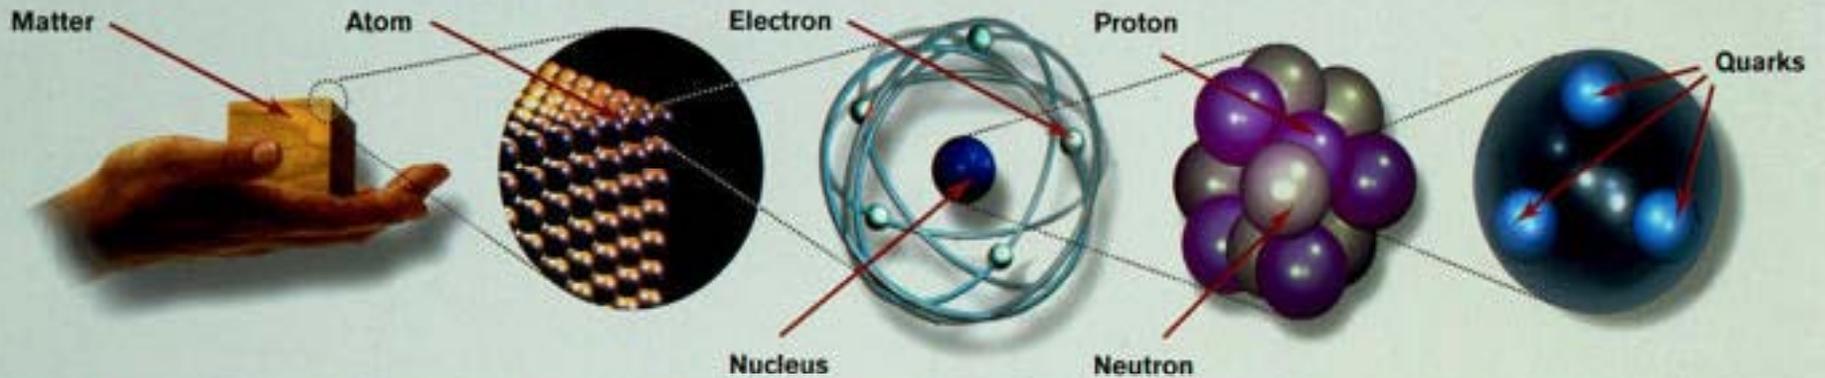
CERN

European Organization for Nuclear Research  
Organisation Européenne pour la Recherche Nucléaire

# Data and Database Challenges at CERN

Tony Cass  
Leader, Database Services Group  
Information Technology Department

29<sup>th</sup> May 2010



**Matter particles**  
All ordinary particles belong to this group

LEPTONS		
<b>FIRST FAMILY</b>	<b>Electron</b> Responsible for electricity and chemical reactions; it has a charge of -1	<b>Electron neutrino</b> Particle with no electric charge, and possibly no mass; billions fly through your body every second
<b>SECOND FAMILY</b>	<b>Muon</b> A heavier relative of the electron; it lives for two-millionths of a second	<b>Muon neutrino</b> Created along with muons when some particles decay
<b>THIRD FAMILY</b>	<b>Tau</b> Heavier still, it is extremely unstable. It was discovered in 1975	<b>Tau neutrino</b> Not yet discovered but believed to exist

QUARKS		
<b>Up</b> Has an electric charge of plus two-thirds; protons contain two, neutrons contain one		<b>Down</b> Has an electric charge of minus one-third; protons contain one, neutrons contain two
<b>Charm</b> A heavier relative of the up; found in 1974		<b>Strange</b> A heavier relative of the down; found in 1964
<b>Top</b> Heavier still		<b>Bottom</b> Heavier still; measuring bottom quarks is an important test of electroweak theory

**Force particles**  
These particles transmit the four fundamental forces of nature although gravitons have so far not been discovered

**Gluons**  
Carriers of the **strong force** between quarks

Felt by quarks

The explosive release of nuclear energy is the result of the **strong force**

**Photons**  
Particles that make up light; they carry the **electromagnetic force**

Felt by quarks and charged leptons

Electricity, magnetism and chemistry are all the results of **electro-magnetic force**

**Intermediate vector bosons**  
Carriers of the **weak force**

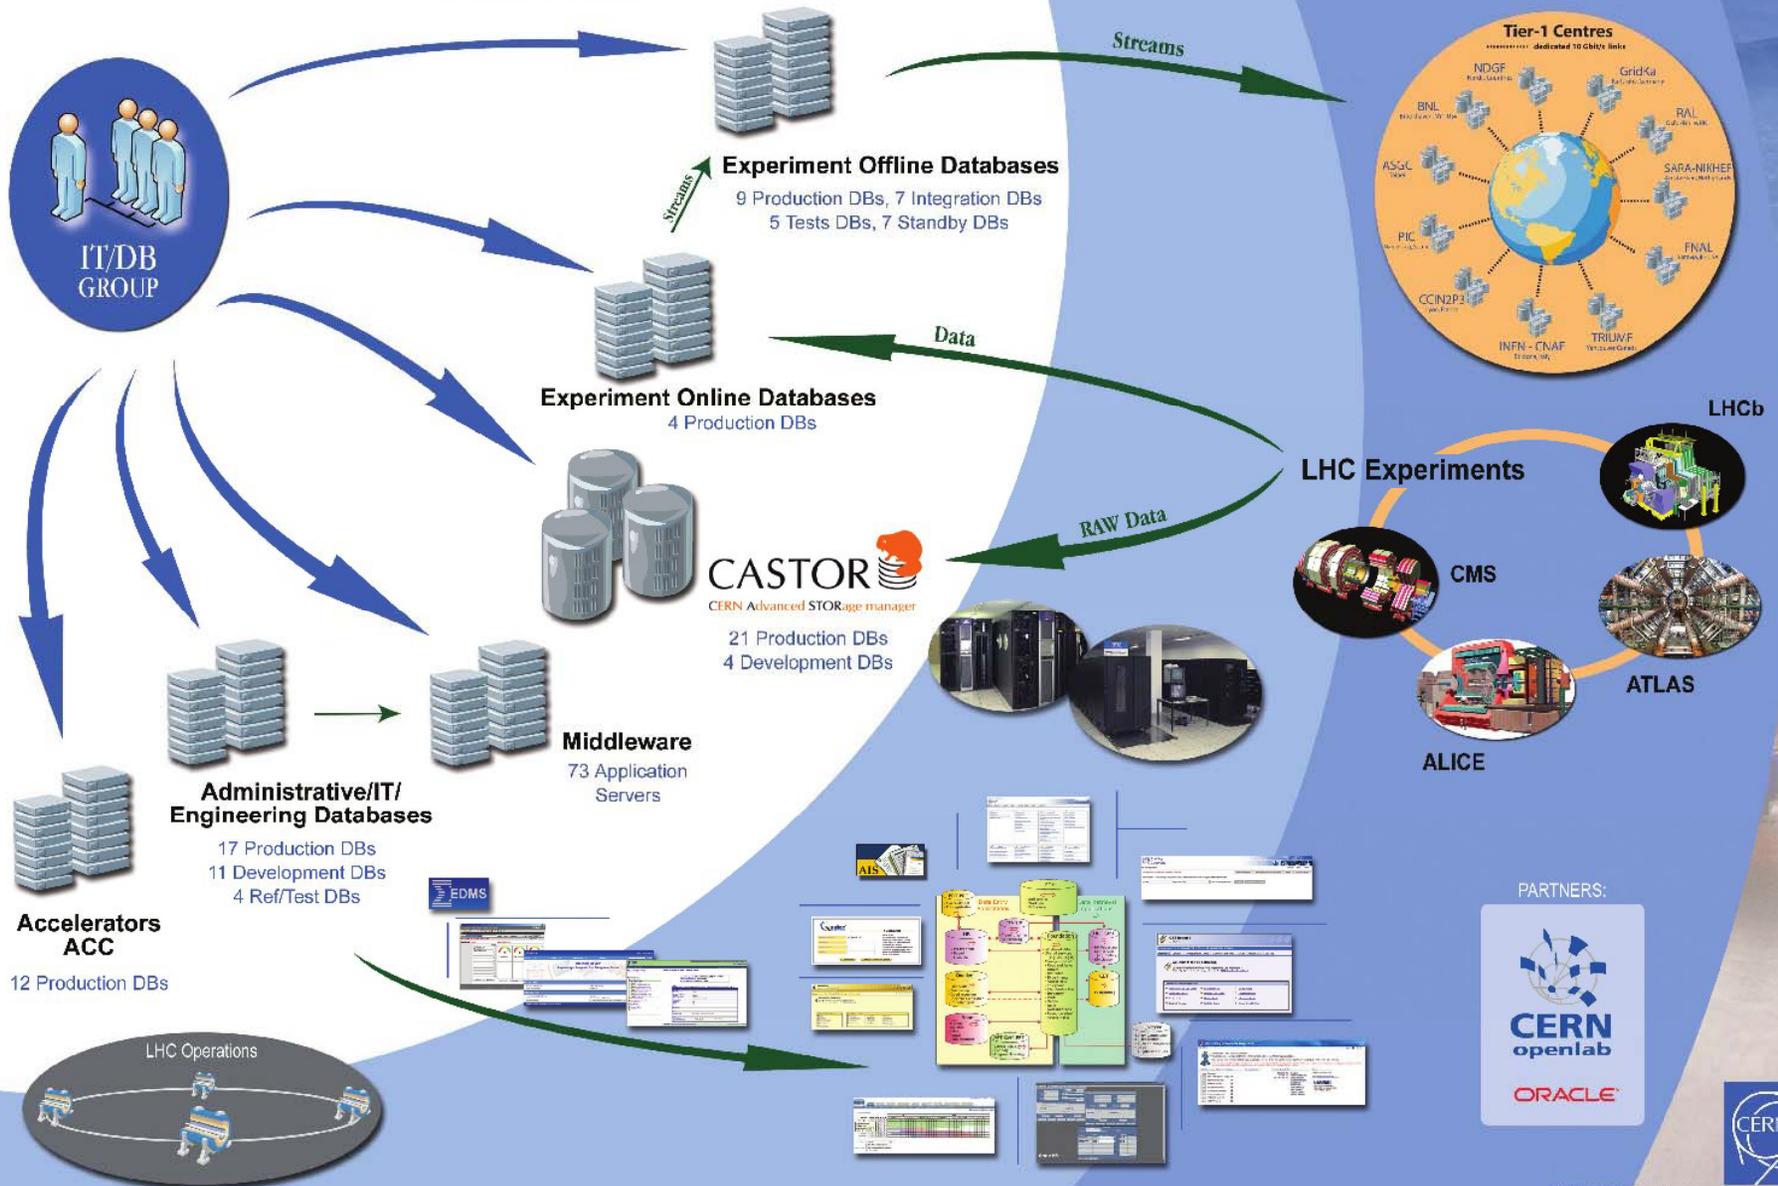
Felt by quarks and leptons

Some forms of radio-activity are the result of the **weak force**

**Gravitons**  
Carriers of **gravity**

Felt by all particles with mass

All the weight we experience is the result of the **gravitational force**





# Outline

- Introduction to CERN, Experiments & Data
- Challenges
- Summary/Conclusion



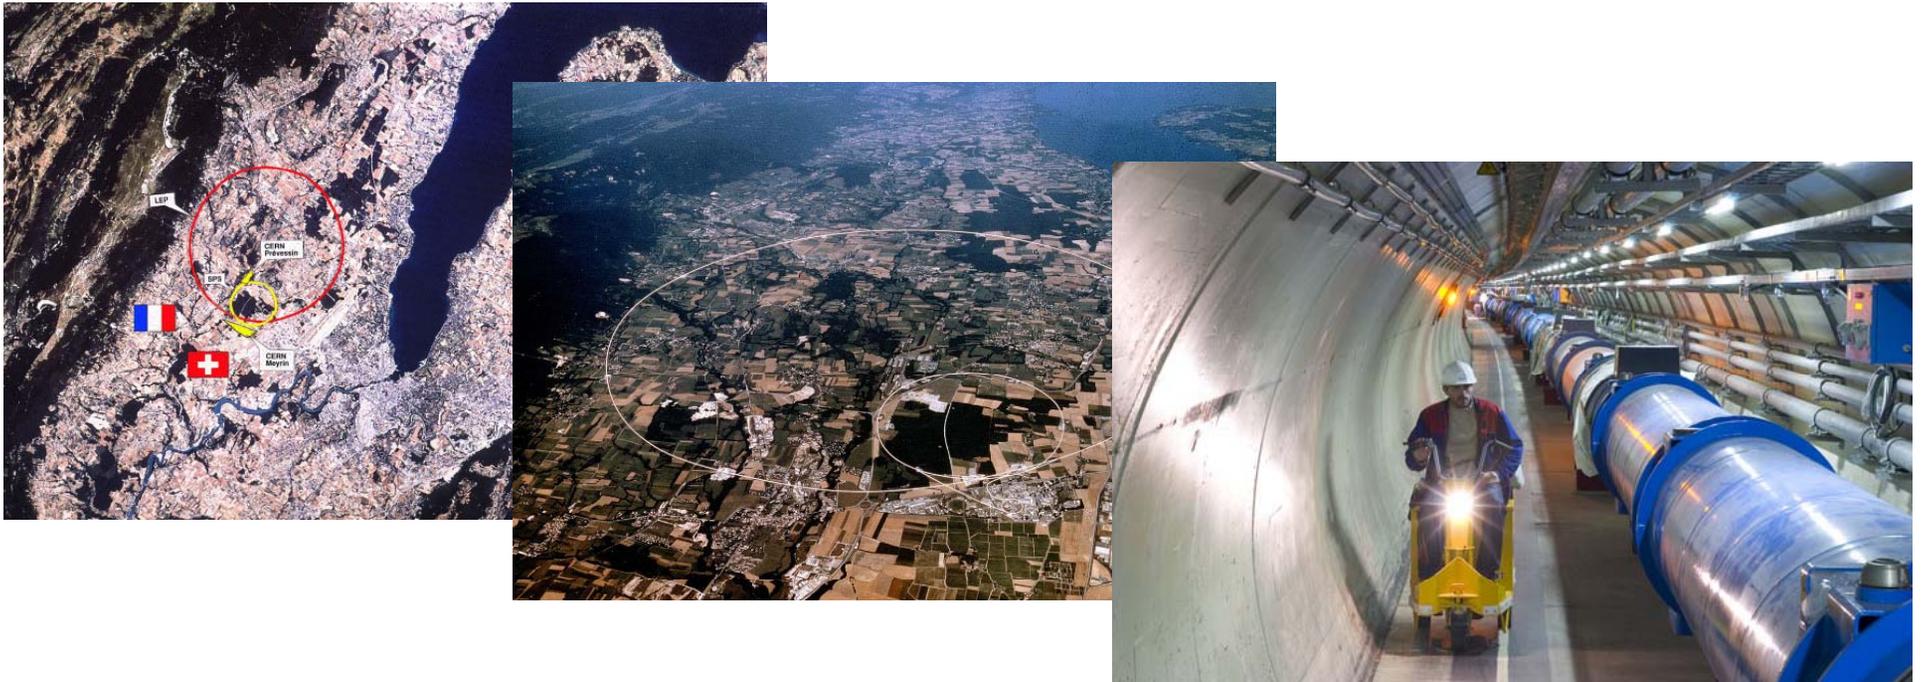
# Outline

- Introduction to CERN, Experiments & Data
- Challenges
- Summary/Conclusion



# CERN

The fastest racetrack on the planet...

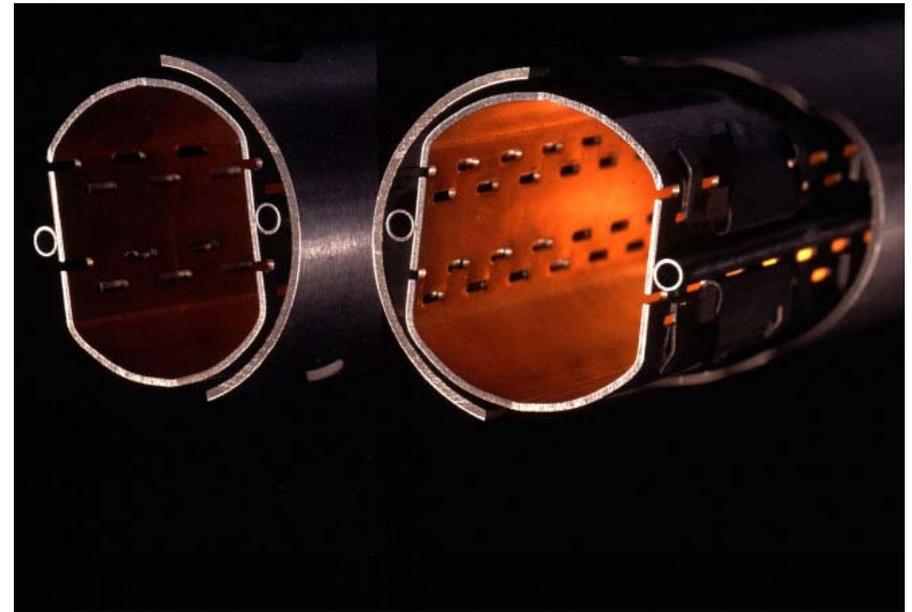
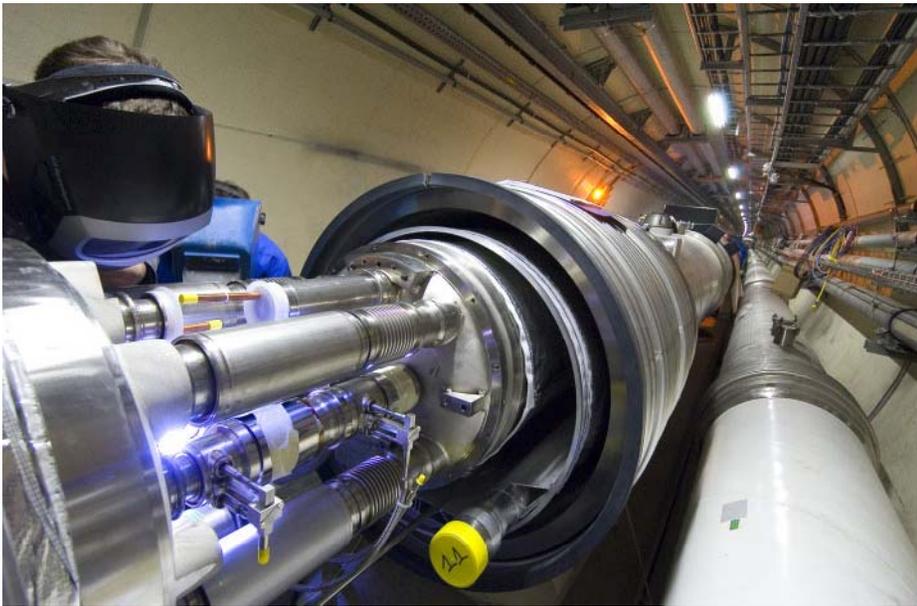


Trillions of protons will race around the 27km ring in opposite directions over 11,000 times a second, travelling at 99.999999991 per cent the speed of light.



# CERN

The emptiest space in the solar system...



To accelerate protons to almost the speed of light requires a vacuum as empty as interplanetary space. There is 10 times more atmosphere on the moon than there will be in the LHC.



# CERN

One of the coldest places in the universe...

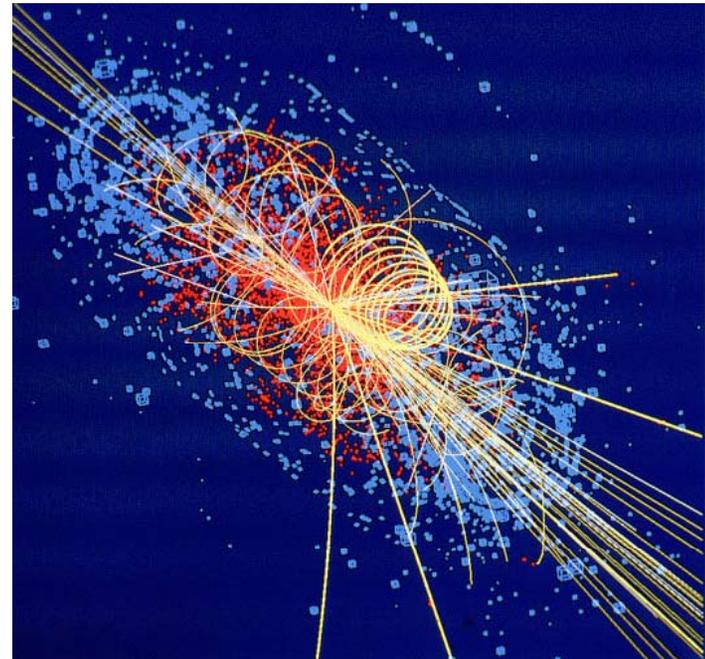
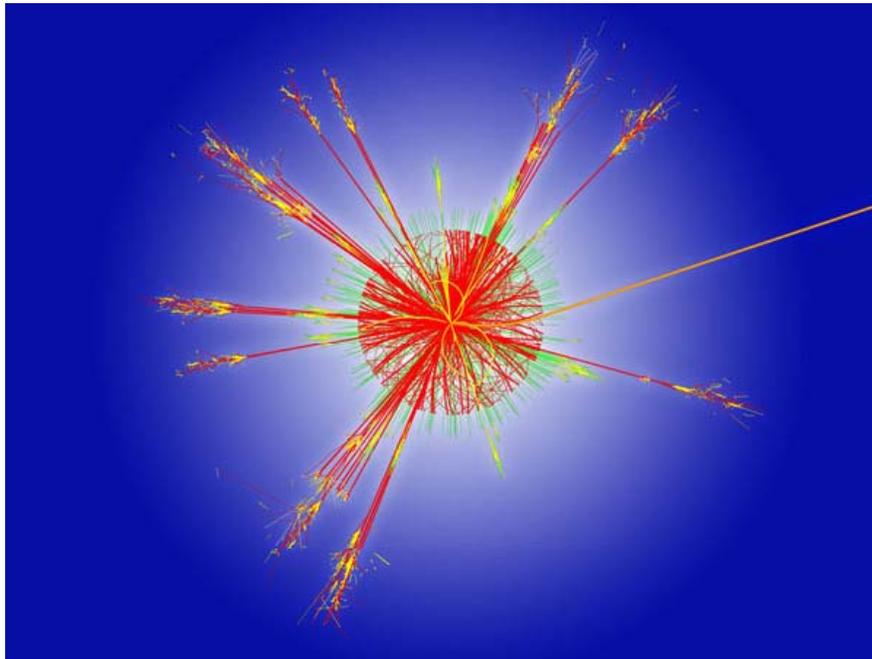


With an operating temperature of about -271 degrees Celsius, just 1.9 degrees above absolute zero, the LHC is colder than outer space.



# CERN

The hottest spots in the galaxy...

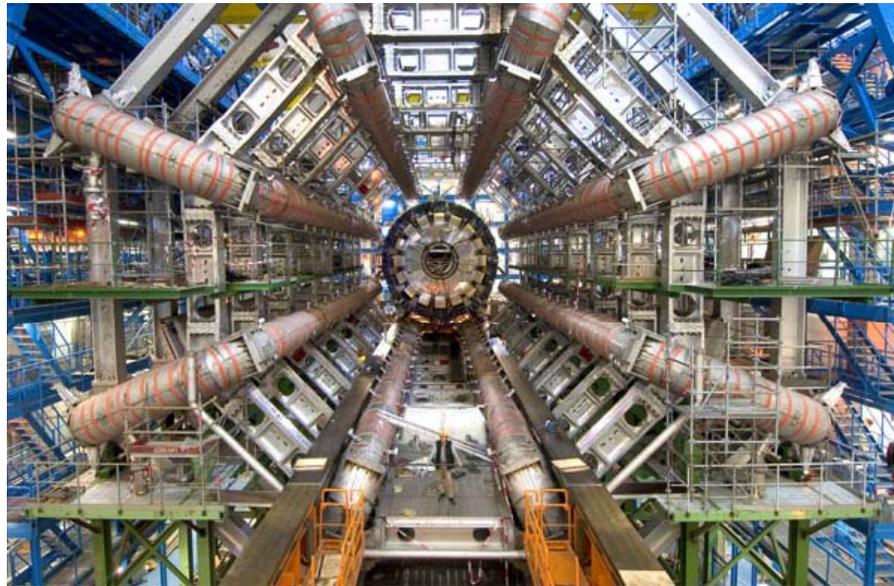


When two beams of protons collide, they will generate temperatures 1000 million times hotter than the heart of the sun, but in a minuscule space.



# CERN

The biggest most sophisticated detectors ever built...



To sample and record the debris from up to 600 million proton collisions per second, scientists are building gargantuan devices that measure particles with micron precision.



# CERN

The most extensive computer system in the world...



To analyse the data, tens of thousands of computers around the world are being harnessed in the Grid. The laboratory that gave the world the web, is now taking distributed computing a big step further.



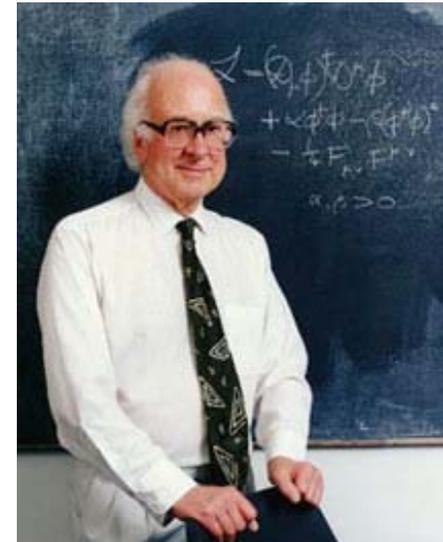
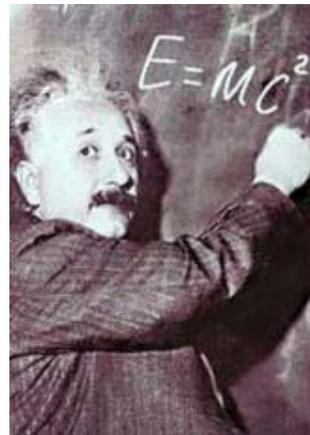
# CERN

Why?



# CERN

To push back the frontiers of knowledge...

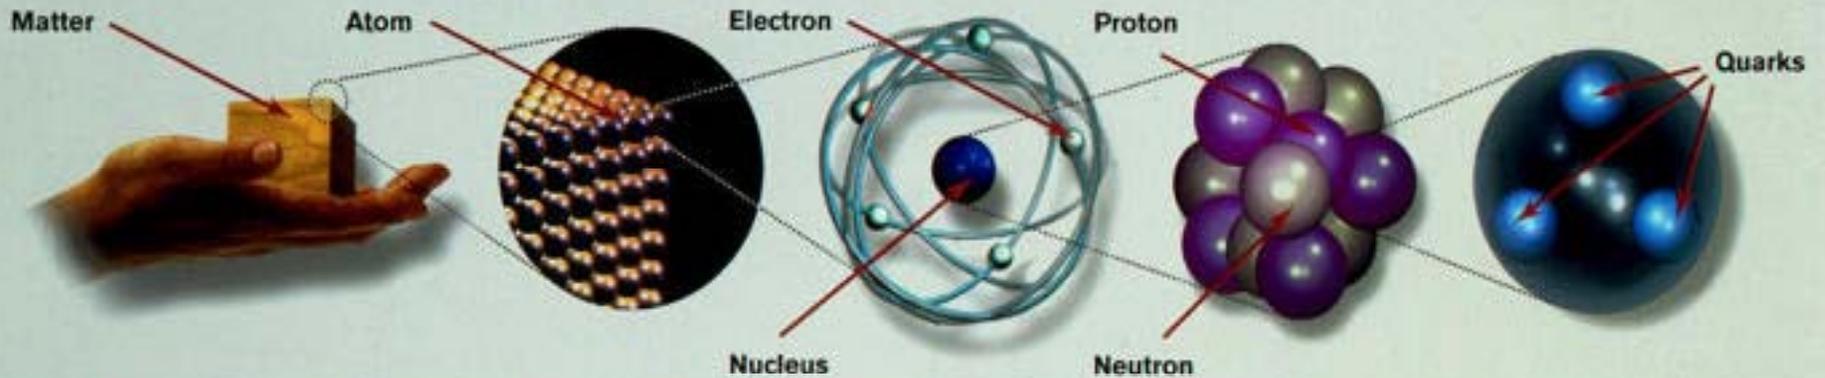


Newton's unfinished business... what is mass?

Science's little embarrassment... what is 96% of the Universe made of?

Nature's favouritism... why is there no more antimatter?

The secrets of the Big Bang... what was matter like within the first second of the Universe's life?



**Matter particles**  
All ordinary particles belong to this group

LEPTONS		
<b>FIRST FAMILY</b>	<b>Electron</b> Responsible for electricity and chemical reactions; it has a charge of -1	<b>Electron neutrino</b> Particle with no electric charge, and possibly no mass; billions fly through your body every second
<b>SECOND FAMILY</b>	<b>Muon</b> A heavier relative of the electron; it lives for two-millionths of a second	<b>Muon neutrino</b> Created along with muons when some particles decay
<b>THIRD FAMILY</b>	<b>Tau</b> Heavier still, it is extremely unstable. It was discovered in 1975	<b>Tau neutrino</b> Not yet discovered but believed to exist

QUARKS		
<b>Up</b> Has an electric charge of plus two-thirds; protons contain two, neutrons contain one		<b>Down</b> Has an electric charge of minus one-third; protons contain one, neutrons contain two
<b>Charm</b> A heavier relative of the up; found in 1974		<b>Strange</b> A heavier relative of the down; found in 1964
<b>Top</b> Heavier still		<b>Bottom</b> Heavier still; measuring bottom quarks is an important test of electroweak theory

**Force particles**  
These particles transmit the four fundamental forces of nature although gravitons have so far not been discovered

**Gluons**  
Carriers of the **strong force** between quarks

Felt by quarks

The explosive release of nuclear energy is the result of the **strong force**

**Photons**  
Particles that make up light; they carry the **electromagnetic force**

Felt by quarks and charged leptons

Electricity, magnetism and chemistry are all the results of **electro-magnetic force**

**Intermediate vector bosons**  
Carriers of the **weak force**

Felt by quarks and leptons

Some forms of radio-activity are the result of the **weak force**

**Gravitons**  
Carriers of **gravity**

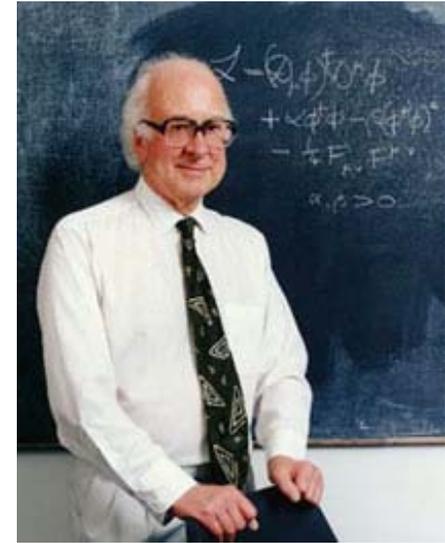
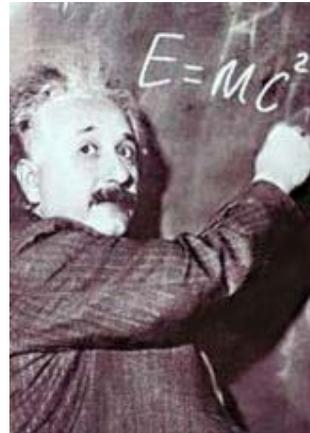
Felt by all particles with mass

All the weight we experience is the result of the **gravitational force**



# CERN

To push back the frontiers of knowledge...



Newton's unfinished business... what is mass?

Science's little embarrassment... what is 96% of the Universe made of?

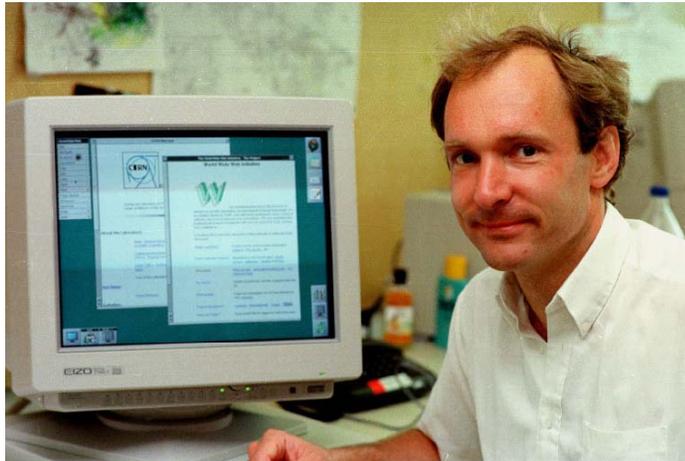
Nature's favoritism... why is there no more antimatter?

The secrets of the Big Bang... what was matter like within the first second of the Universe's life?



# CERN

To develop new technologies...



Information technology - the Web and the Grid

Medicine - diagnosis and therapy

Security - scanning technologies for harbours and airports

Vacuum - new techniques for flat screen displays or solar energy devices



# CERN

To unite people from different countries and cultures...



20 Member states

38 Countries with cooperation agreements

111 Nationalities

10000 People



# CERN

To train the scientists and engineers of tomorrow...

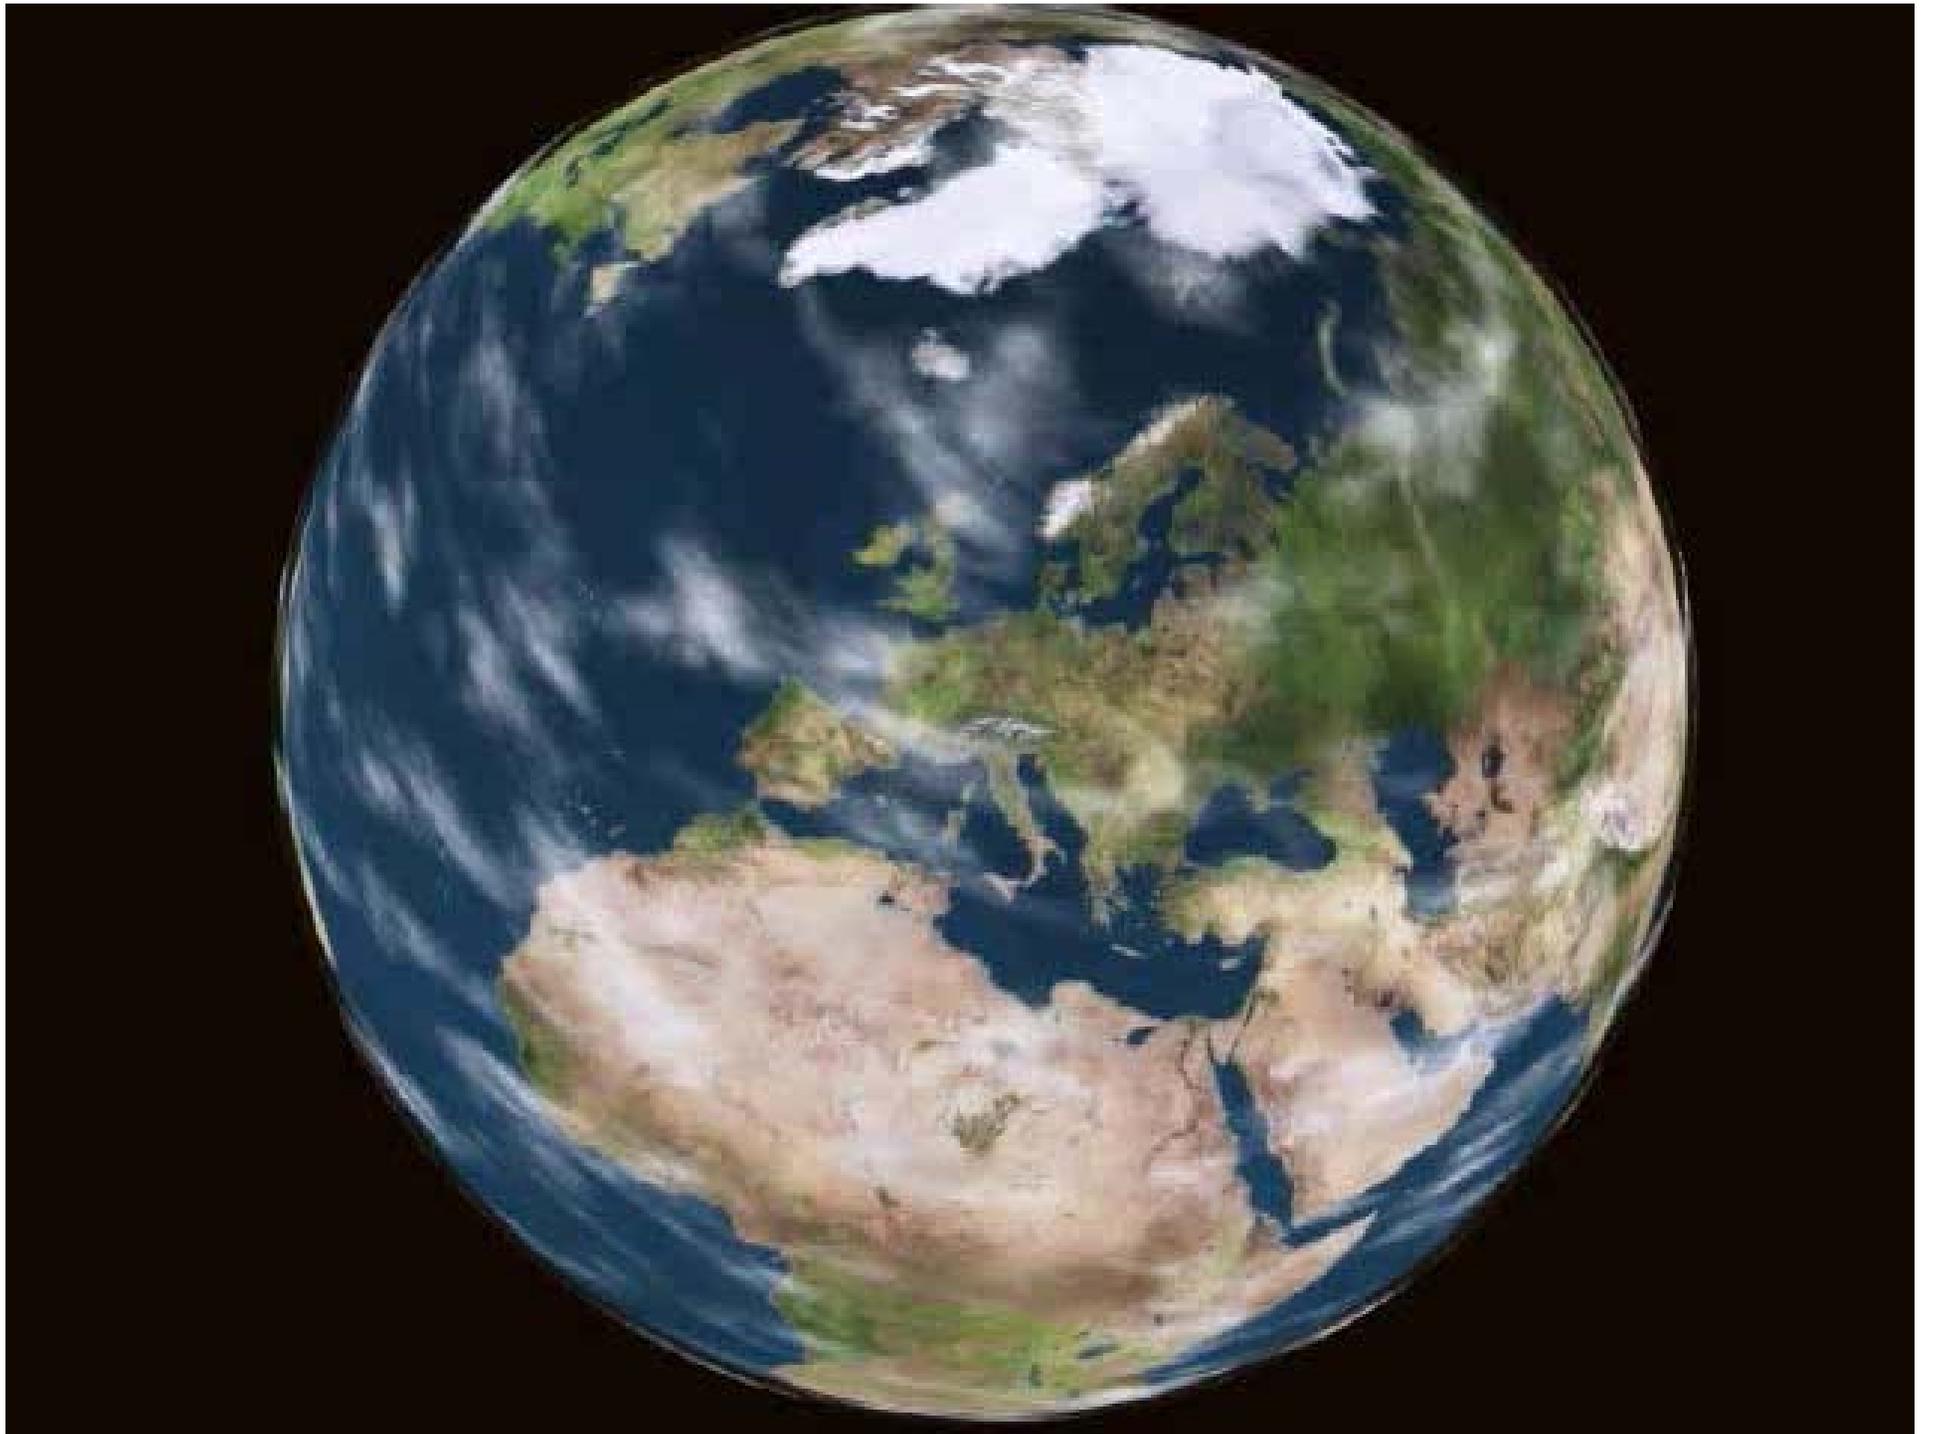


From mini-Einstein workshops for five to sixes, through to professional schools in physics, accelerator science and IT, CERN plays a valuable role in building enthusiasm for science and providing formal training..



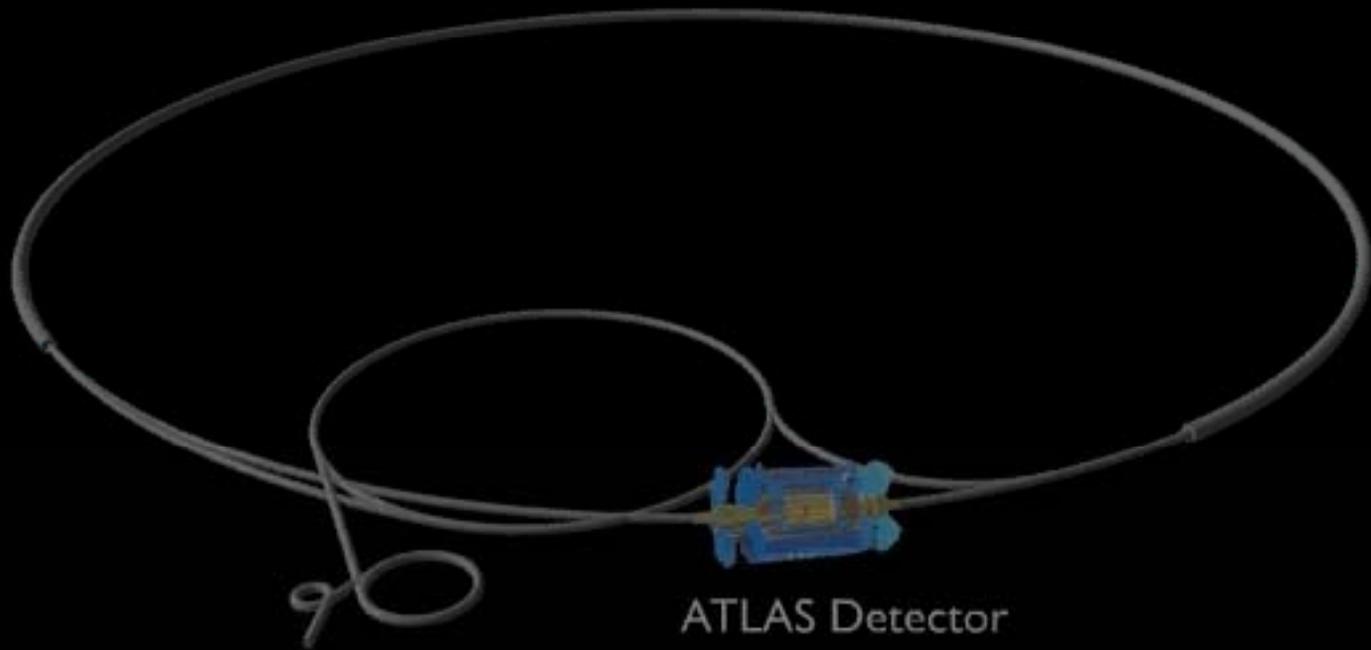
# "Compact" Detectors!





PLAY ▶

Large Hadron Collider



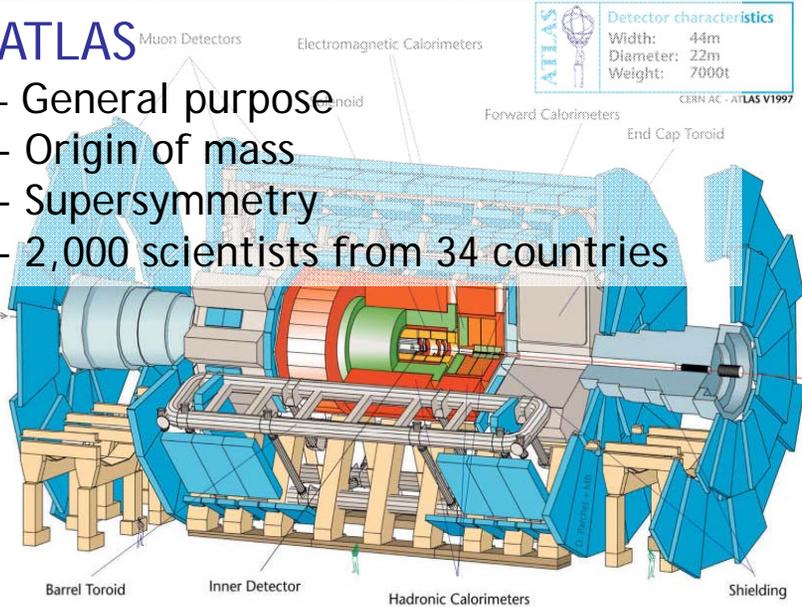
ATLAS Detector



# The Four LHC Experiments...

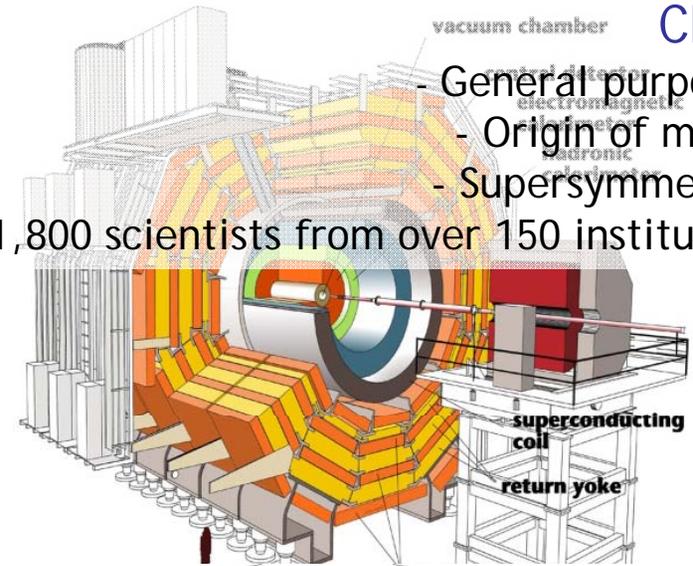
## ATLAS

- General purpose
- Origin of mass
- Supersymmetry
- 2,000 scientists from 34 countries



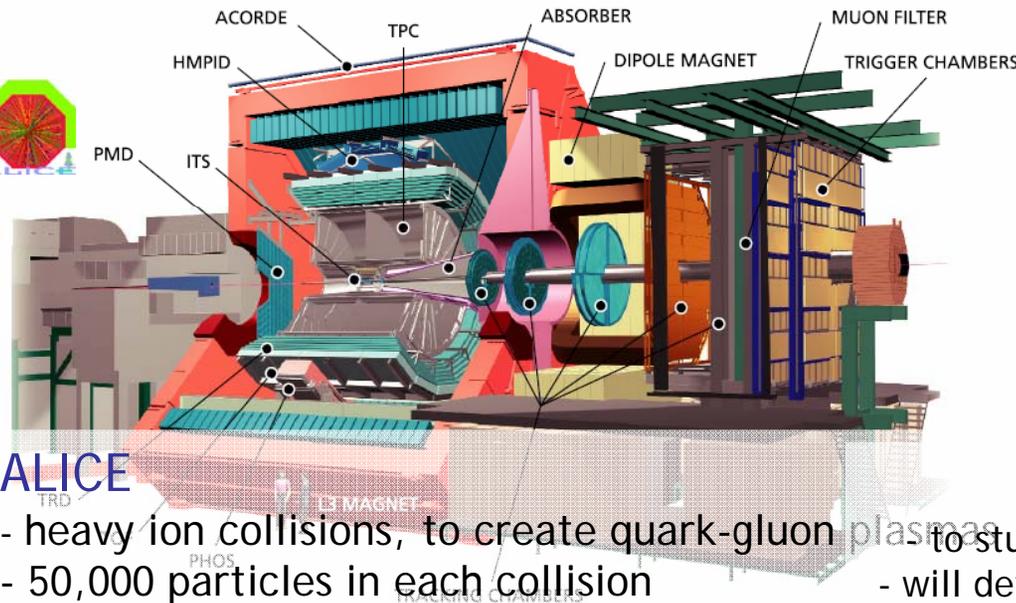
## CMS

- General purpose
- Origin of mass
- Supersymmetry
- 1,800 scientists from over 150 institutes



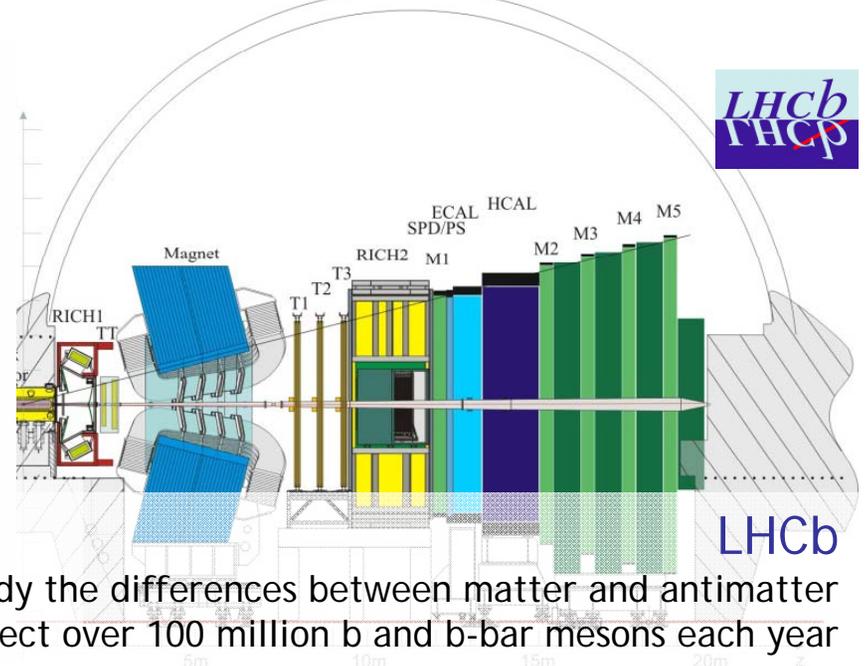
## ALICE

- heavy ion collisions, to create quark-gluon plasmas
- 50,000 particles in each collision



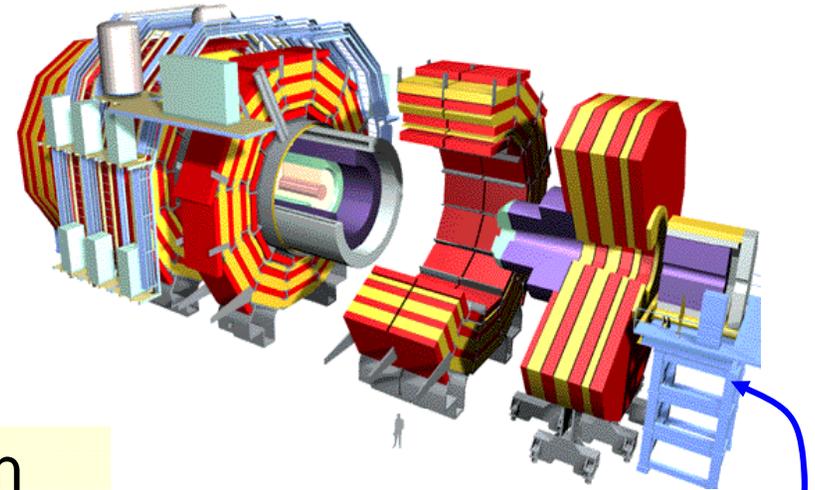
## LHCb

- will detect over 100 million b and b-bar mesons each year





... generate lots of data ...



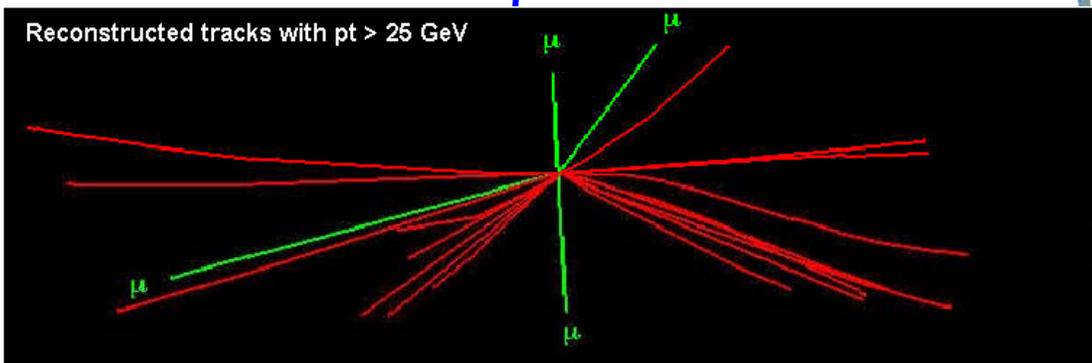
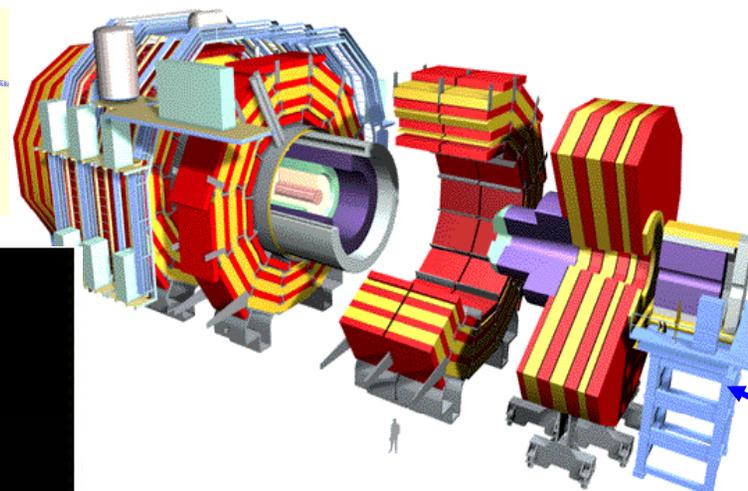
The accelerator generates 40 million particle collisions (events) every second at the centre of each of the four experiments' detectors





# ... generate lots of data ...

reduced by online computers to a few hundred "good" events per second.



Which are recorded on disk and magnetic tape  
at 100-1,000 MegaBytes/sec → **~15 PetaBytes** per year  
for all four experiments





# which is distributed worldwide

## Tier-0 (CERN):

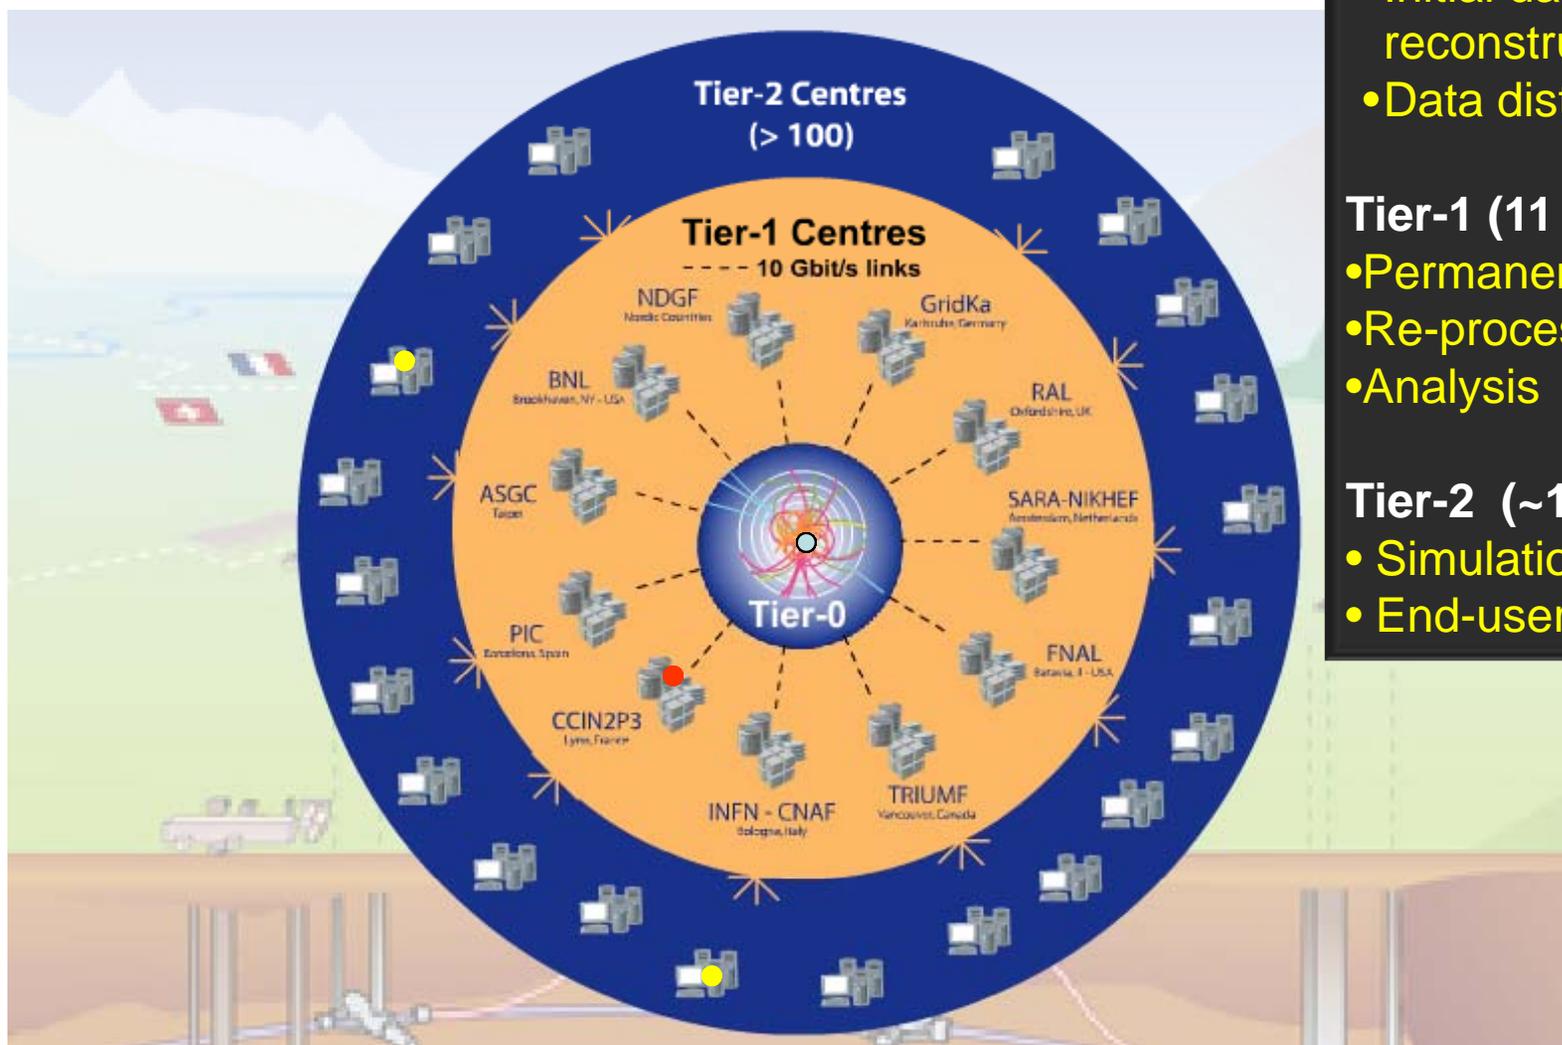
- Data recording
- Initial data reconstruction
- Data distribution

## Tier-1 (11 centres):

- Permanent storage
- Re-processing
- Analysis

## Tier-2 (~130 centres):

- Simulation
- End-user analysis





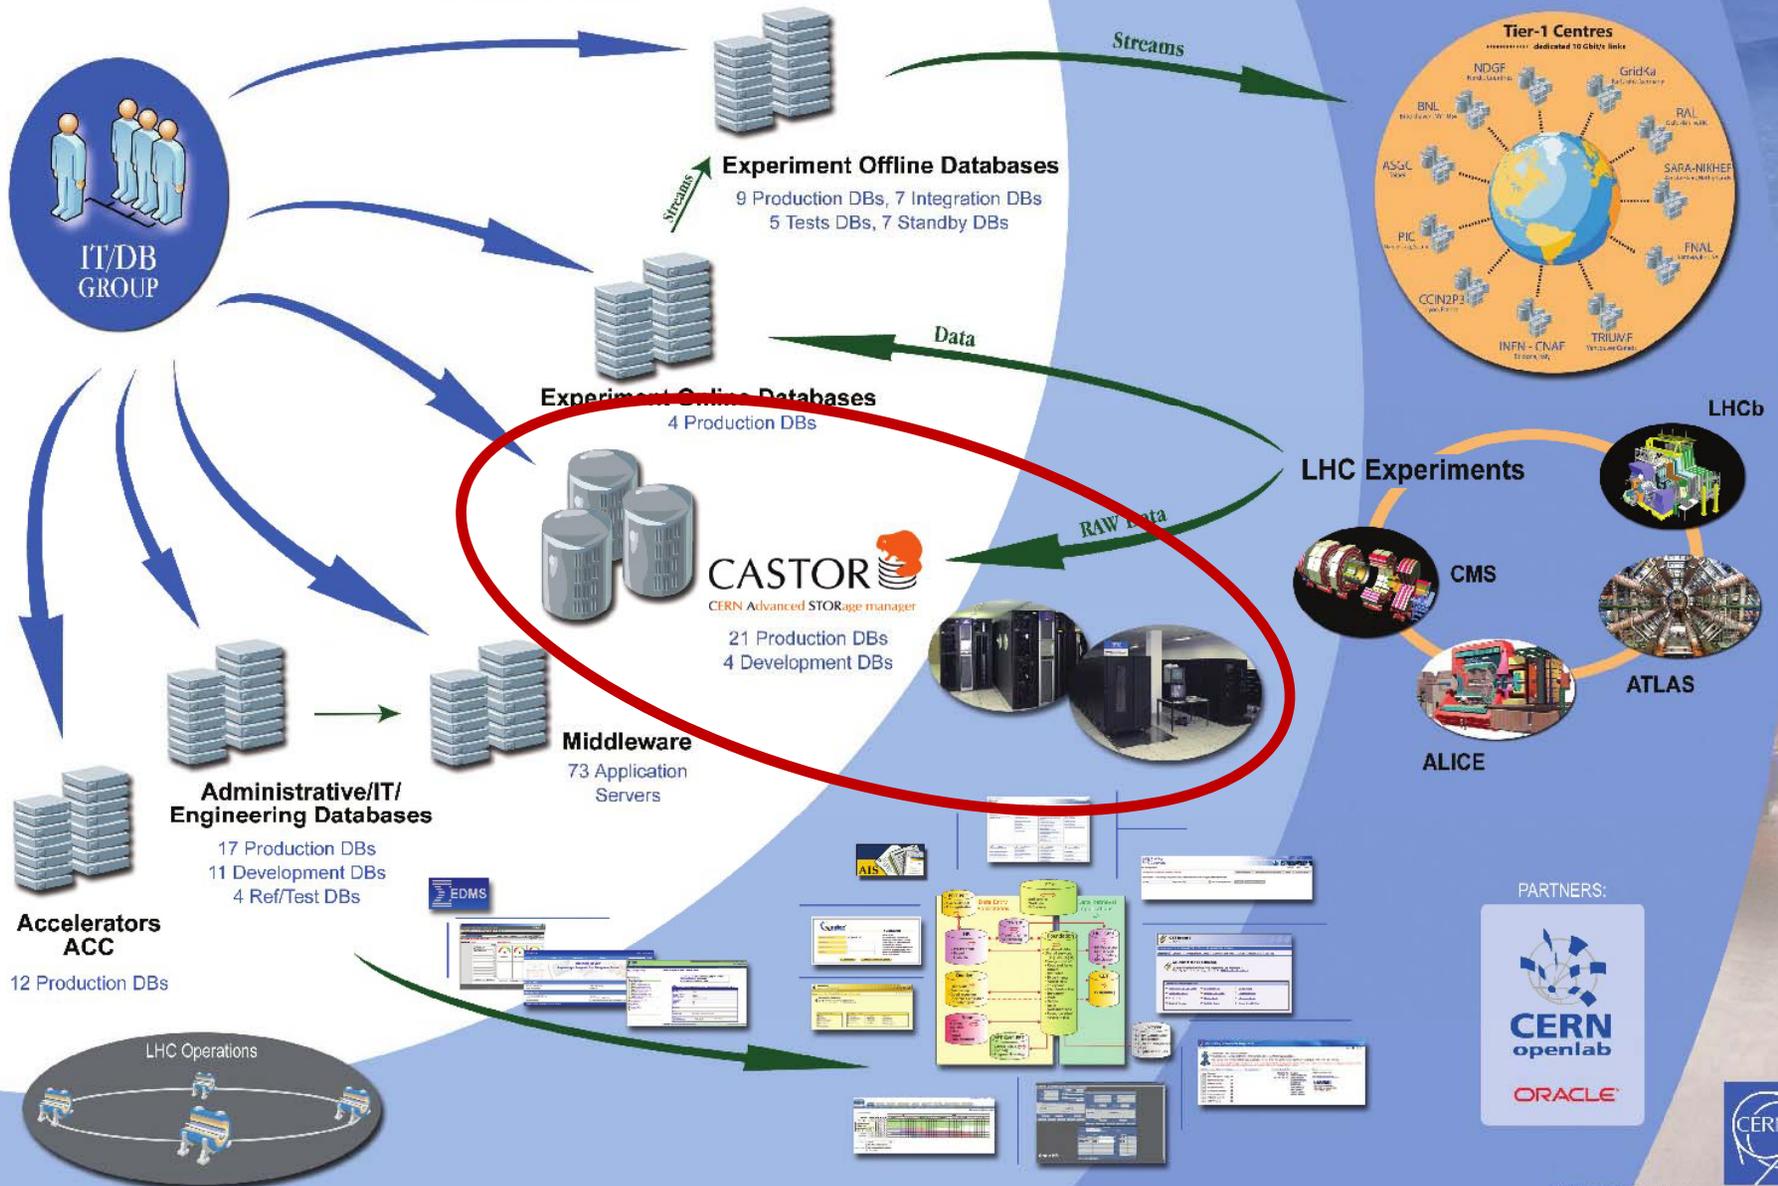
# Outline

- Introduction to CERN, Experiments & Data
- **Challenges**
- Summary/Conclusion



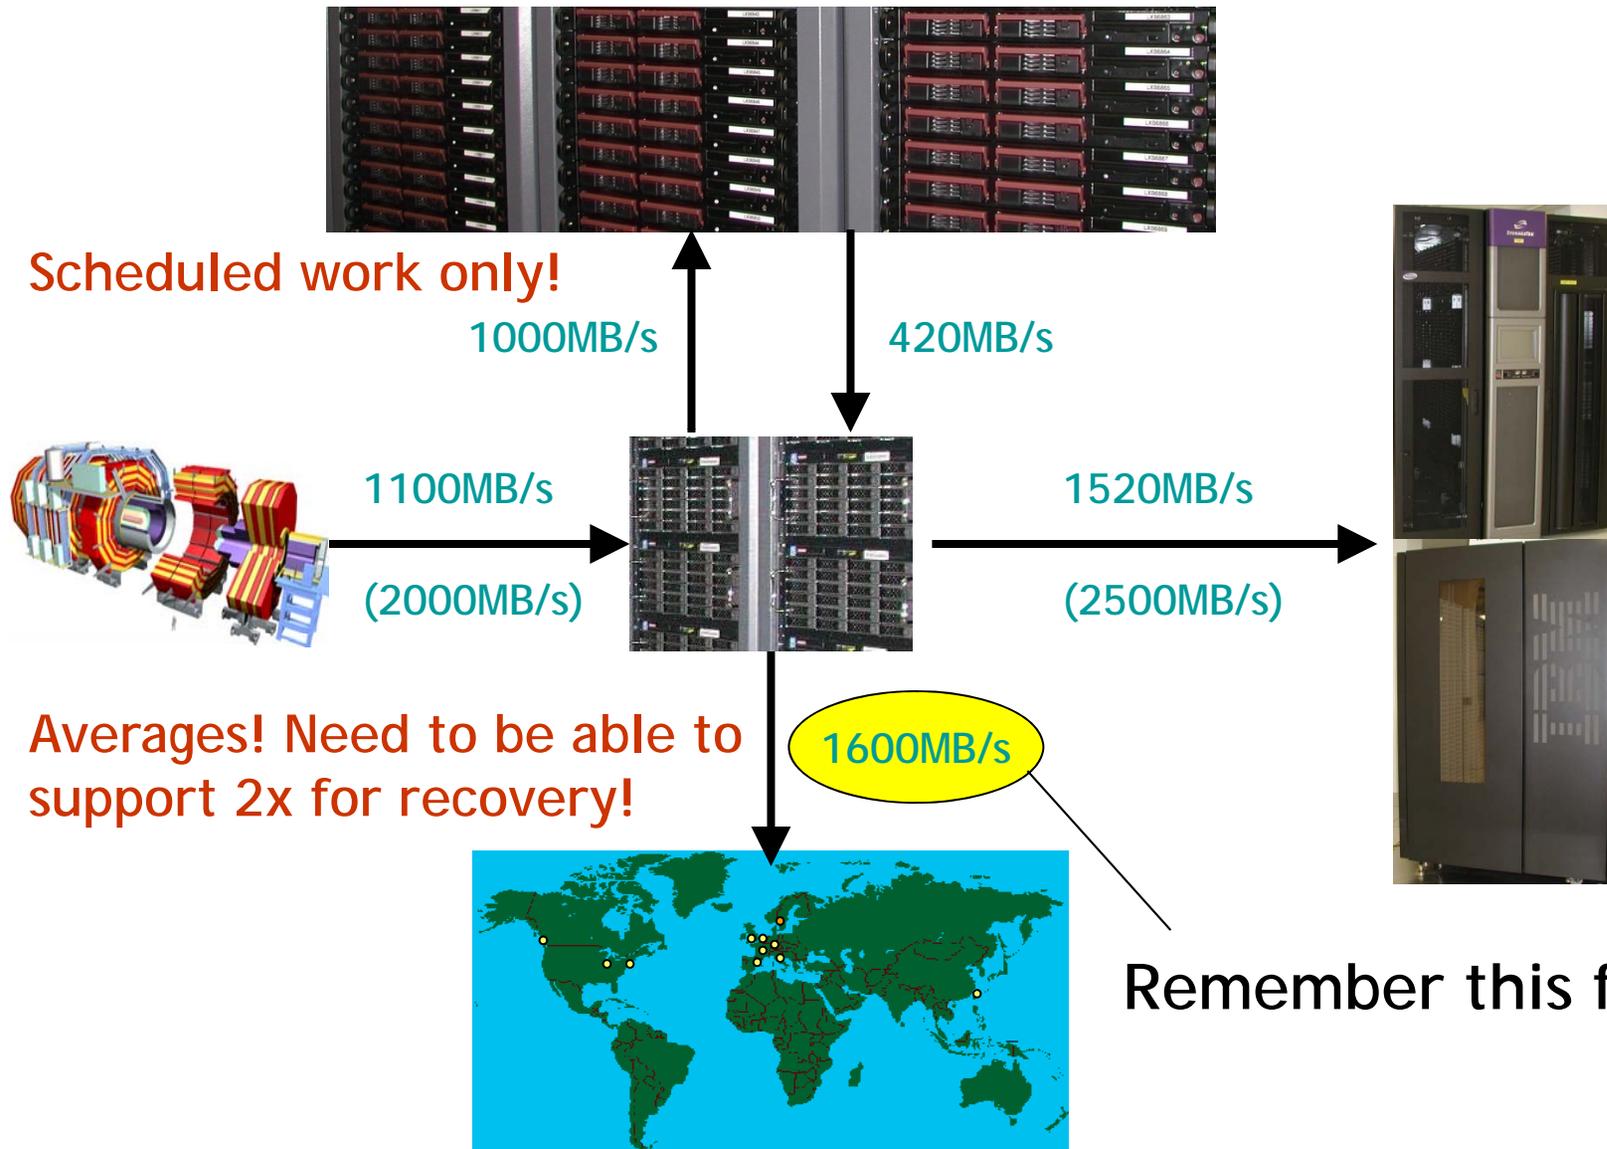
# Outline

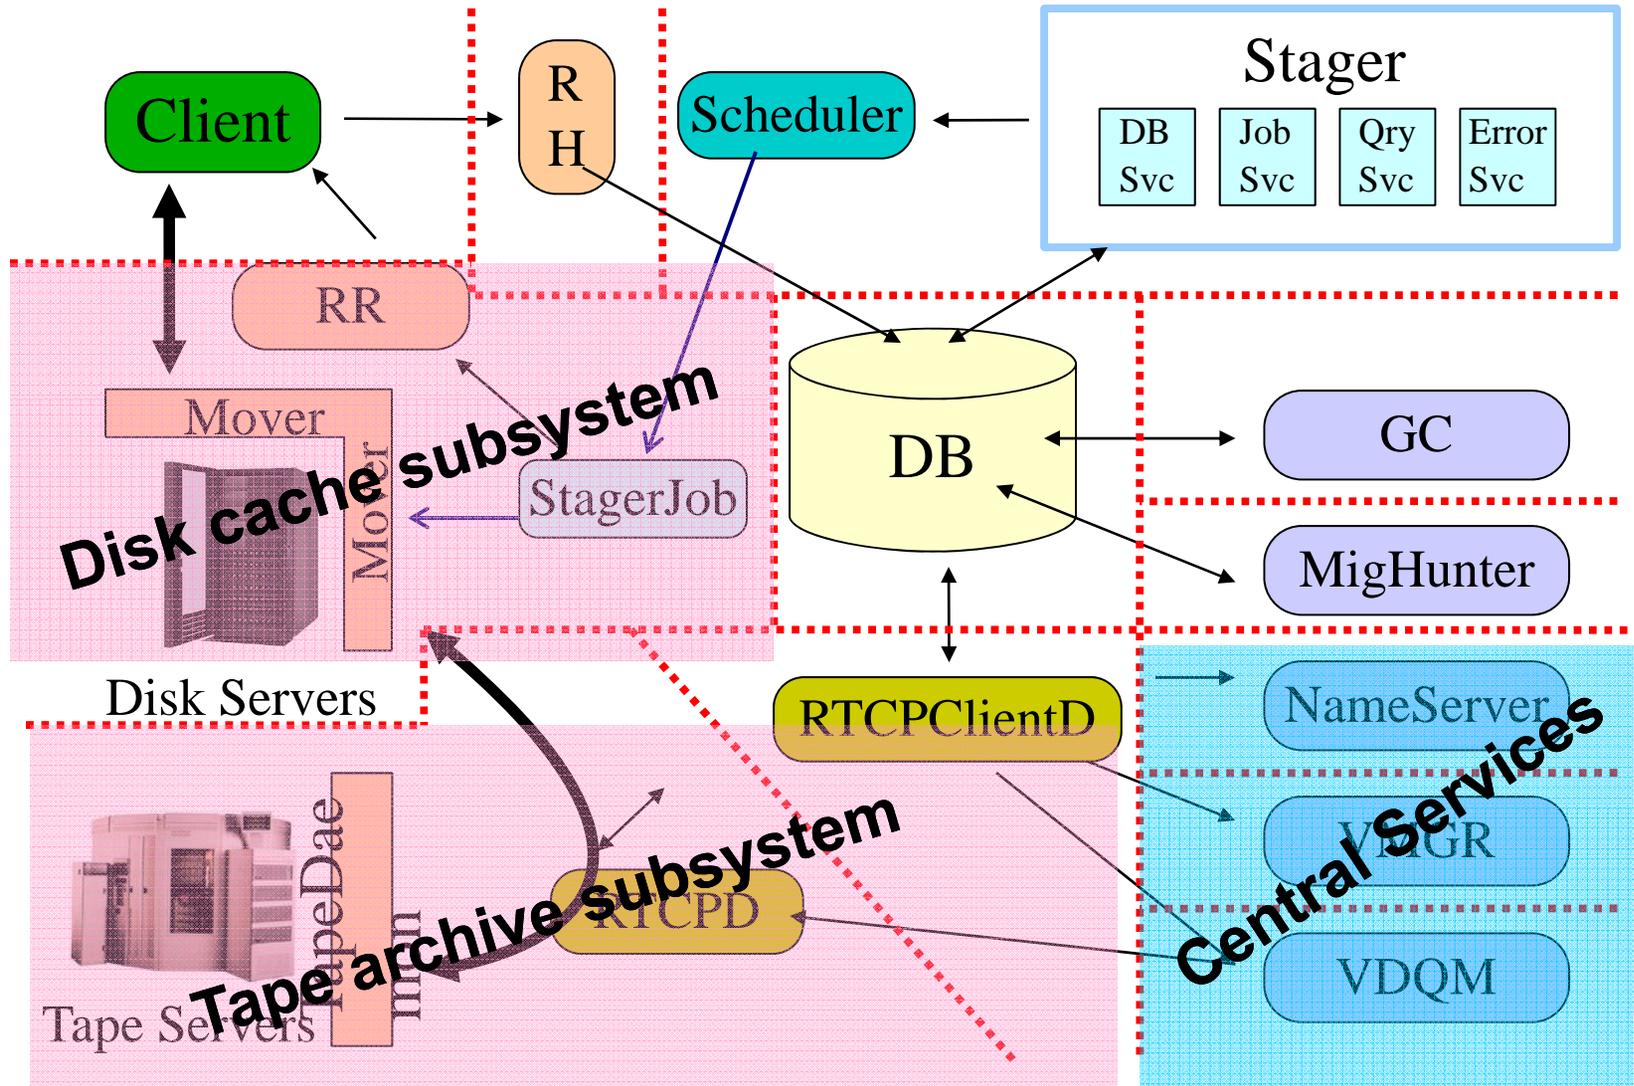
- Introduction to CERN, Experiments & Data
- **Challenges**
  - Data Storage & Distribution
- Summary/Conclusion





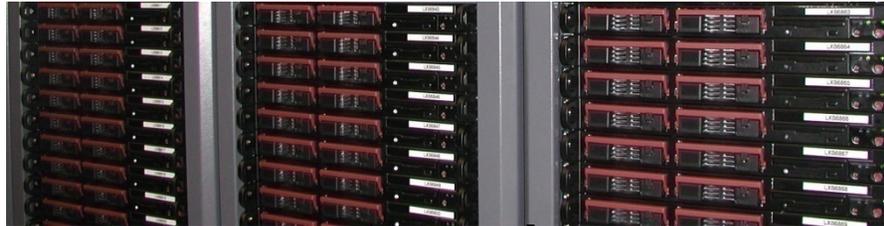
# Dataflows and rates



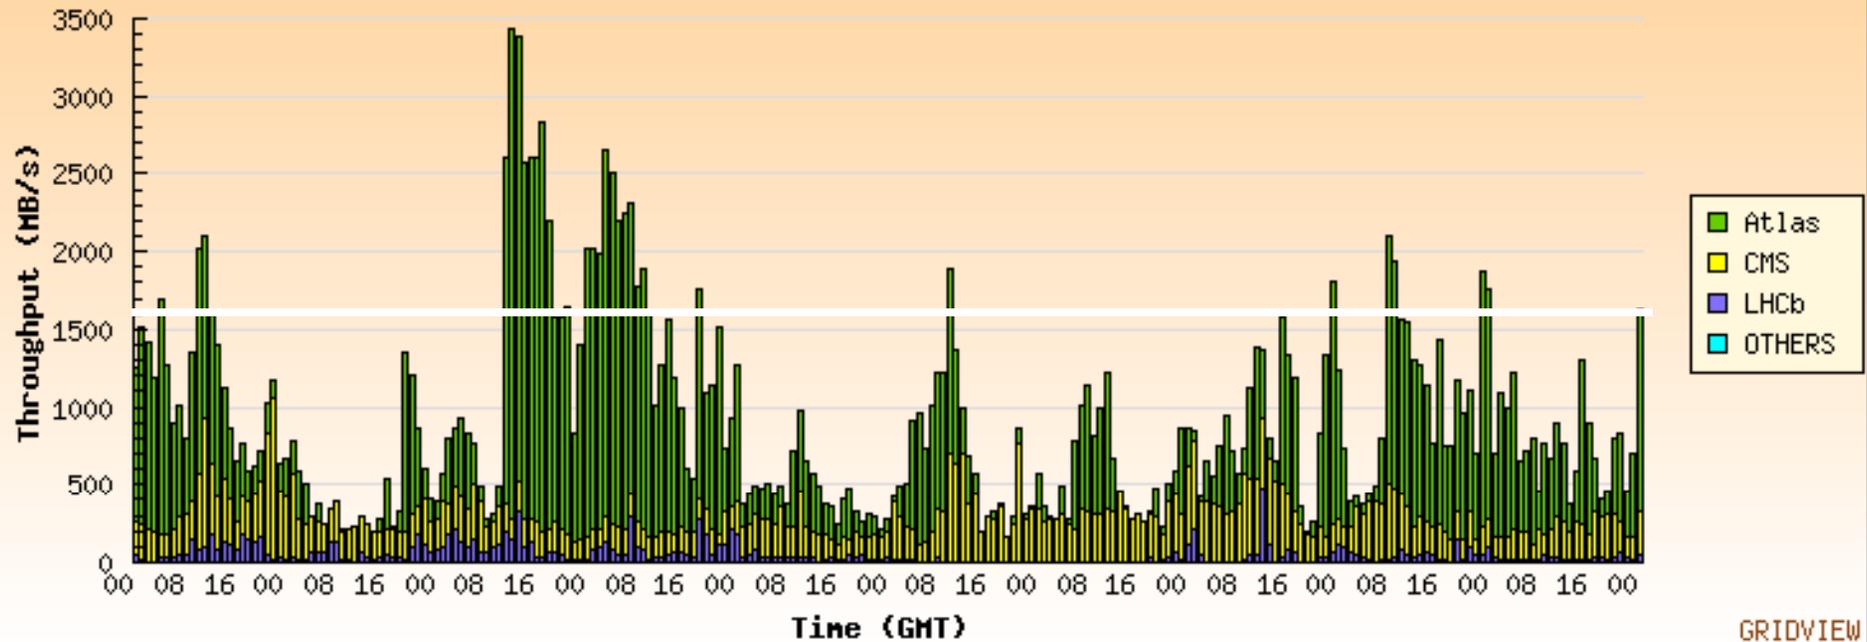




# Dataflows and rates



Averaged Throughput from 00 Hrs on 17/05/10 to 00 Hrs on 27/05/10  
VO-wise Data Transfer From CERN-PROD To All T1 Sites

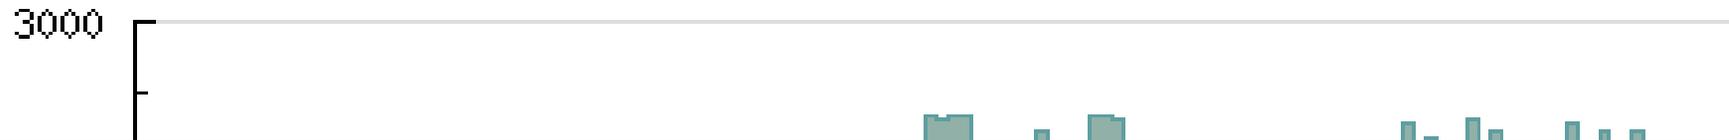




# The real challenge? Users!

## TOEXPRESS Service Class

Active transfers - last 24 hours



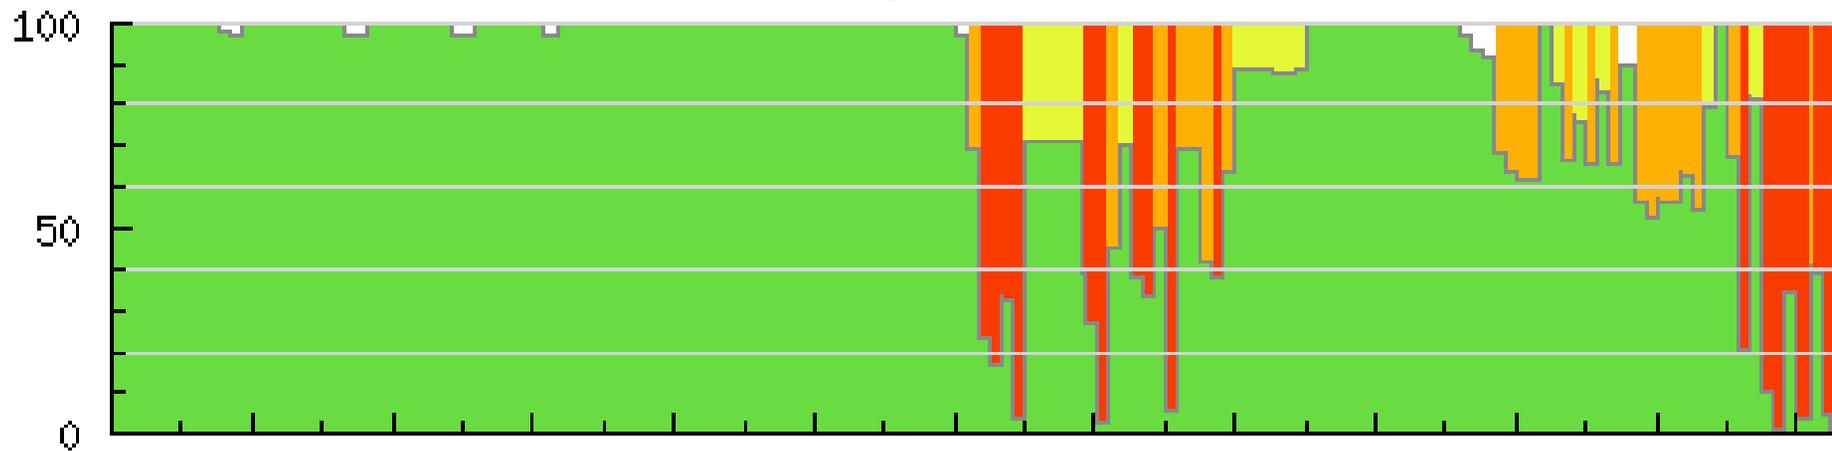
## TOEXPRESS Service Class

Queued transfers - last 24 hours



## TOEXPRESS Service Class

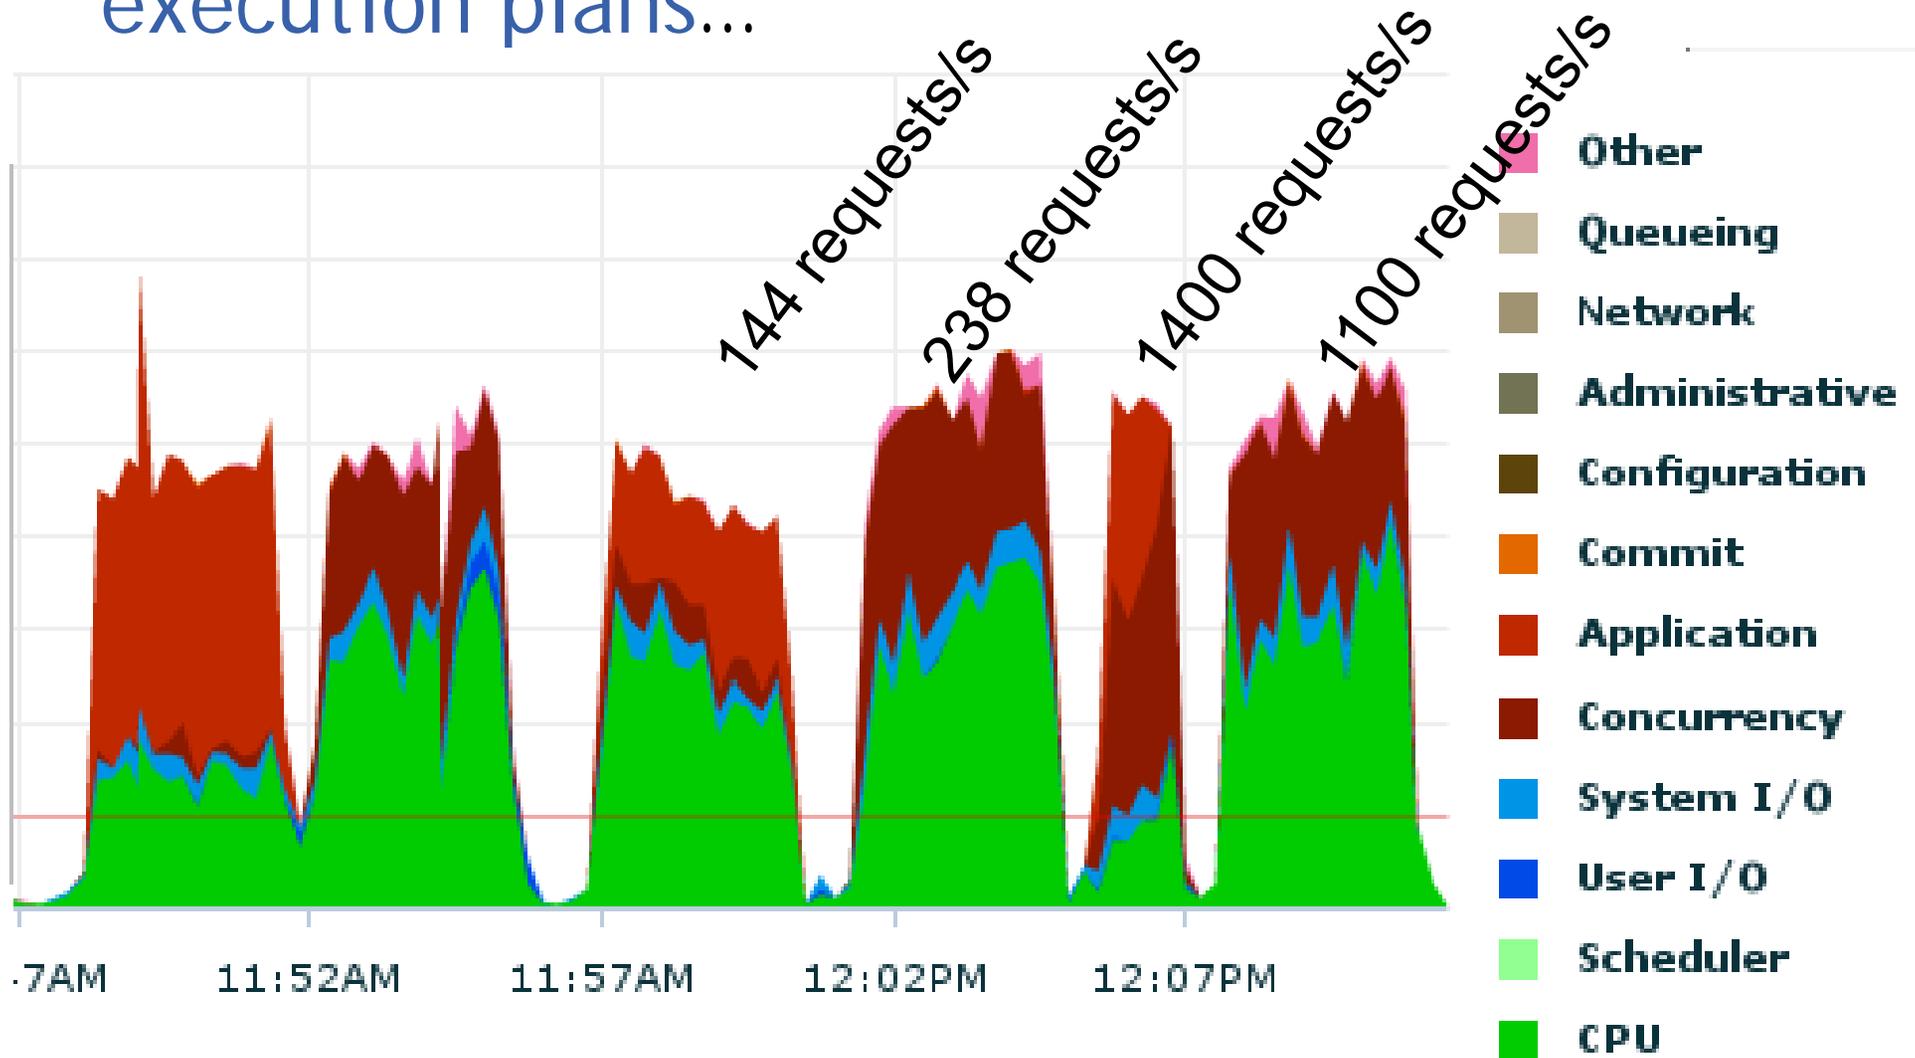
availability - last 24 hours





# The real challenge? Users!

- Such user behaviour can affect Oracle execution plans...

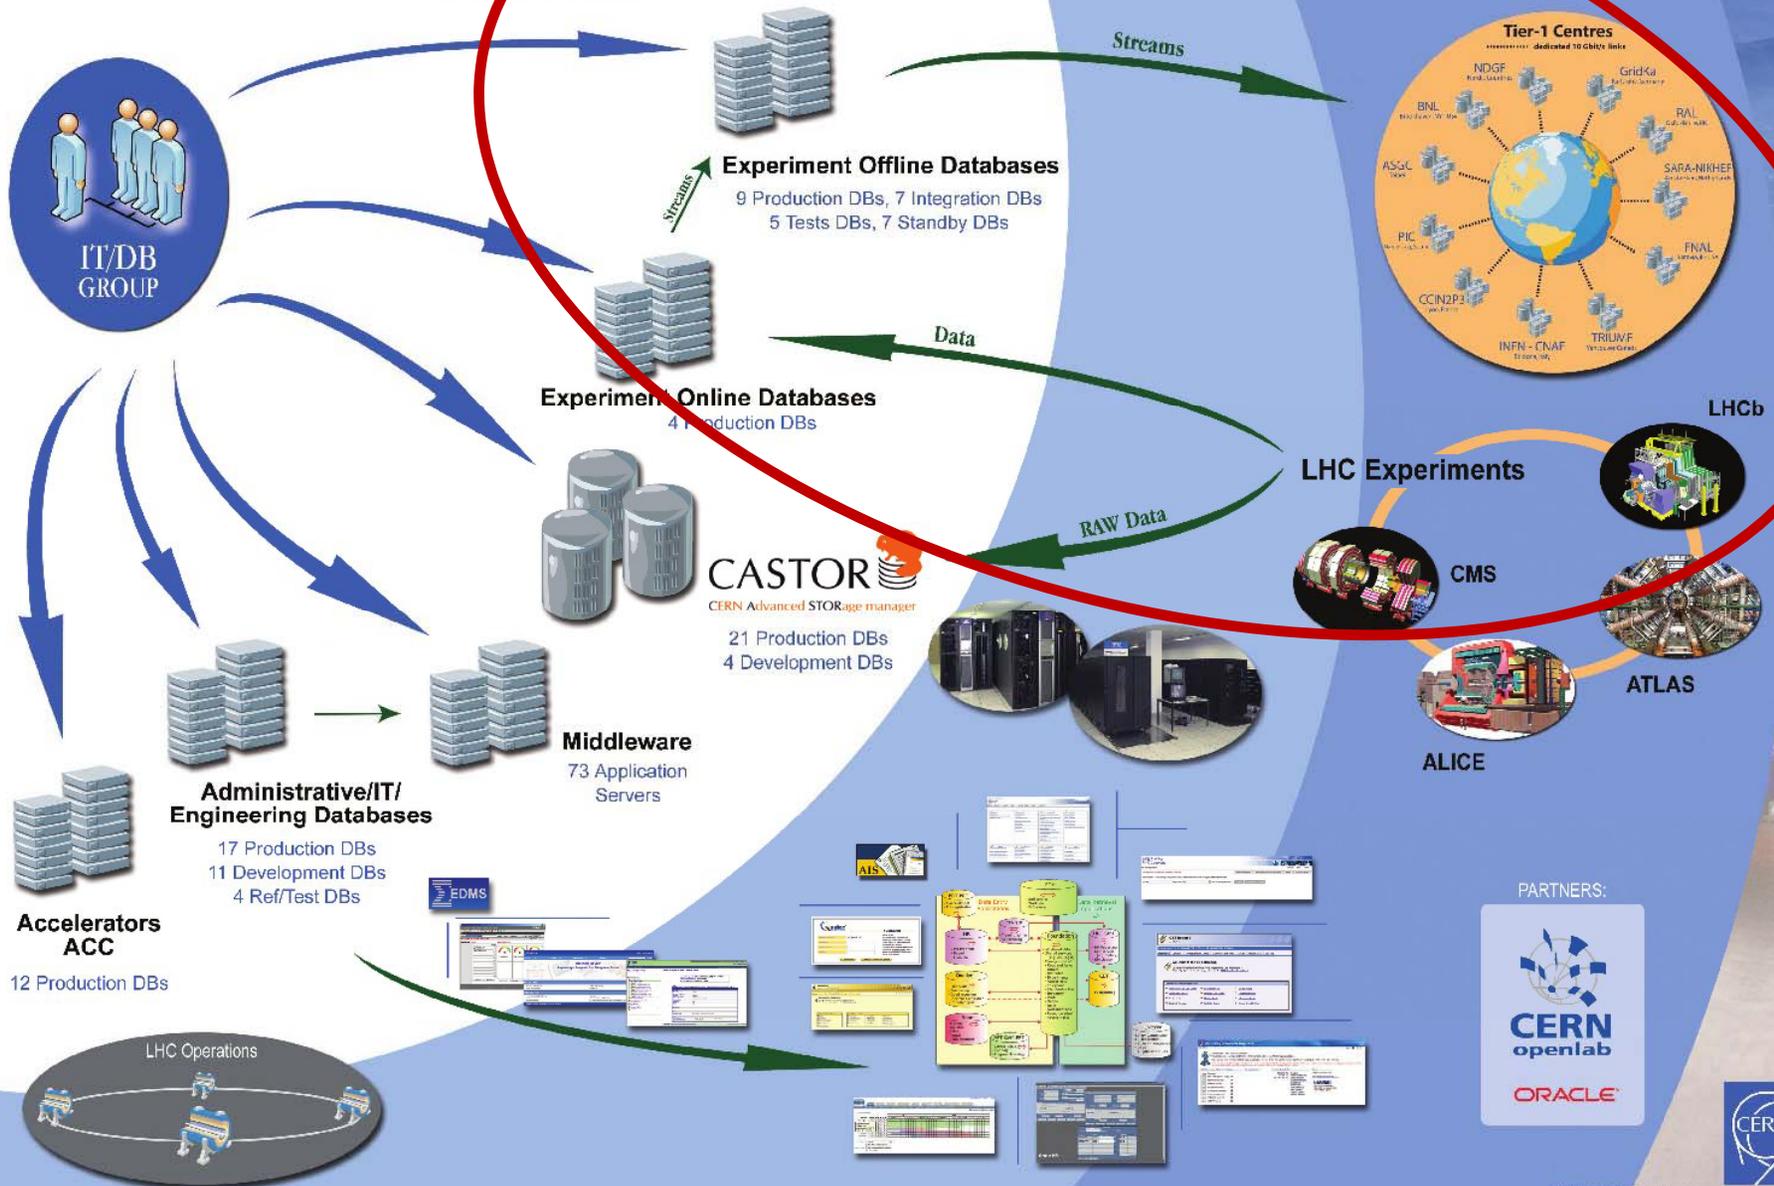




# Outline

- Introduction to CERN, Experiments & Data
- **Challenges**
  - Metadata distribution
- Summary/Conclusion

<http://cern.ch/it-dep/db/>



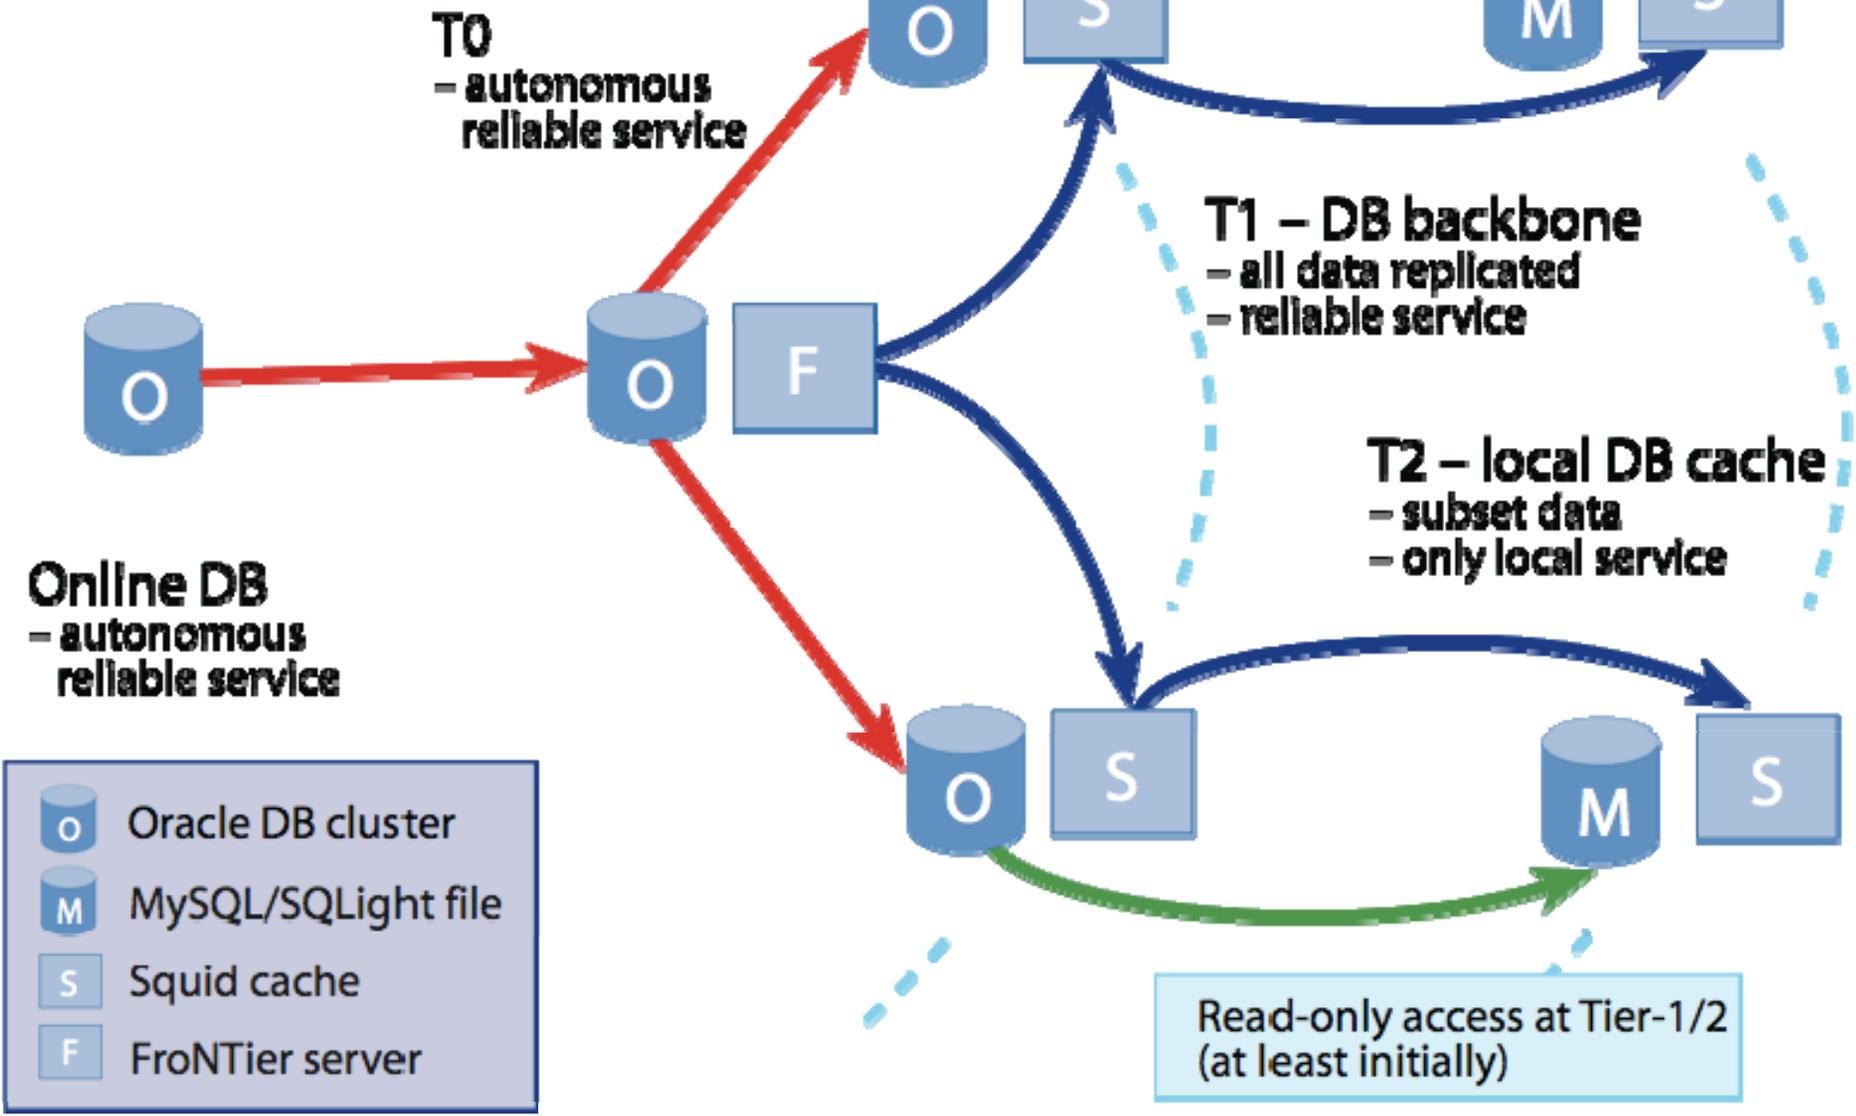
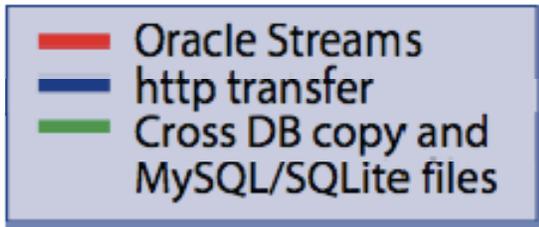
PARTNERS:





# Metadata Distribution

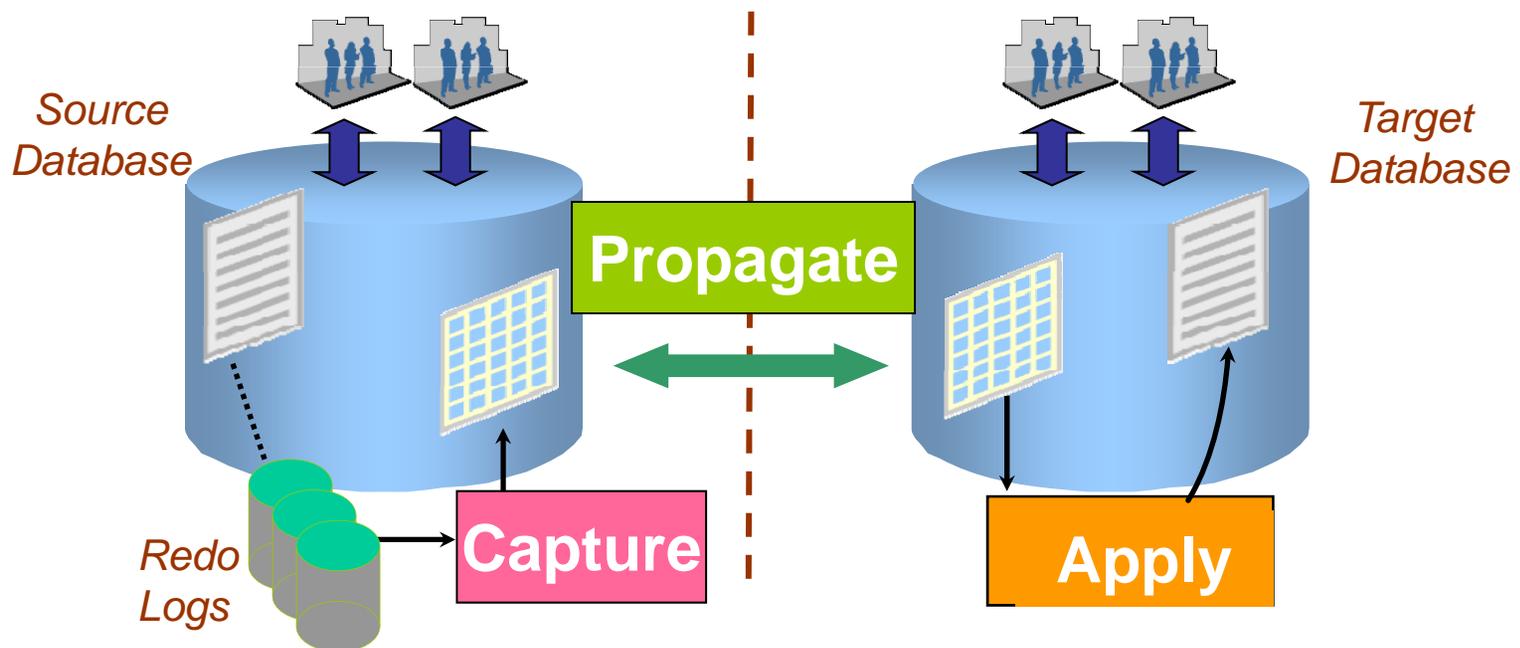
- To make sense of the **raw data** generated by the detectors, physicists need data about any conditions at the time it was taken that affect the **detector calibration**.
- This **conditions data** is stored in a relational database and needs **distributing to Tier1 centres** to enable future reprocessing of the raw data.
- **Oracle Streams** enables this distribution...





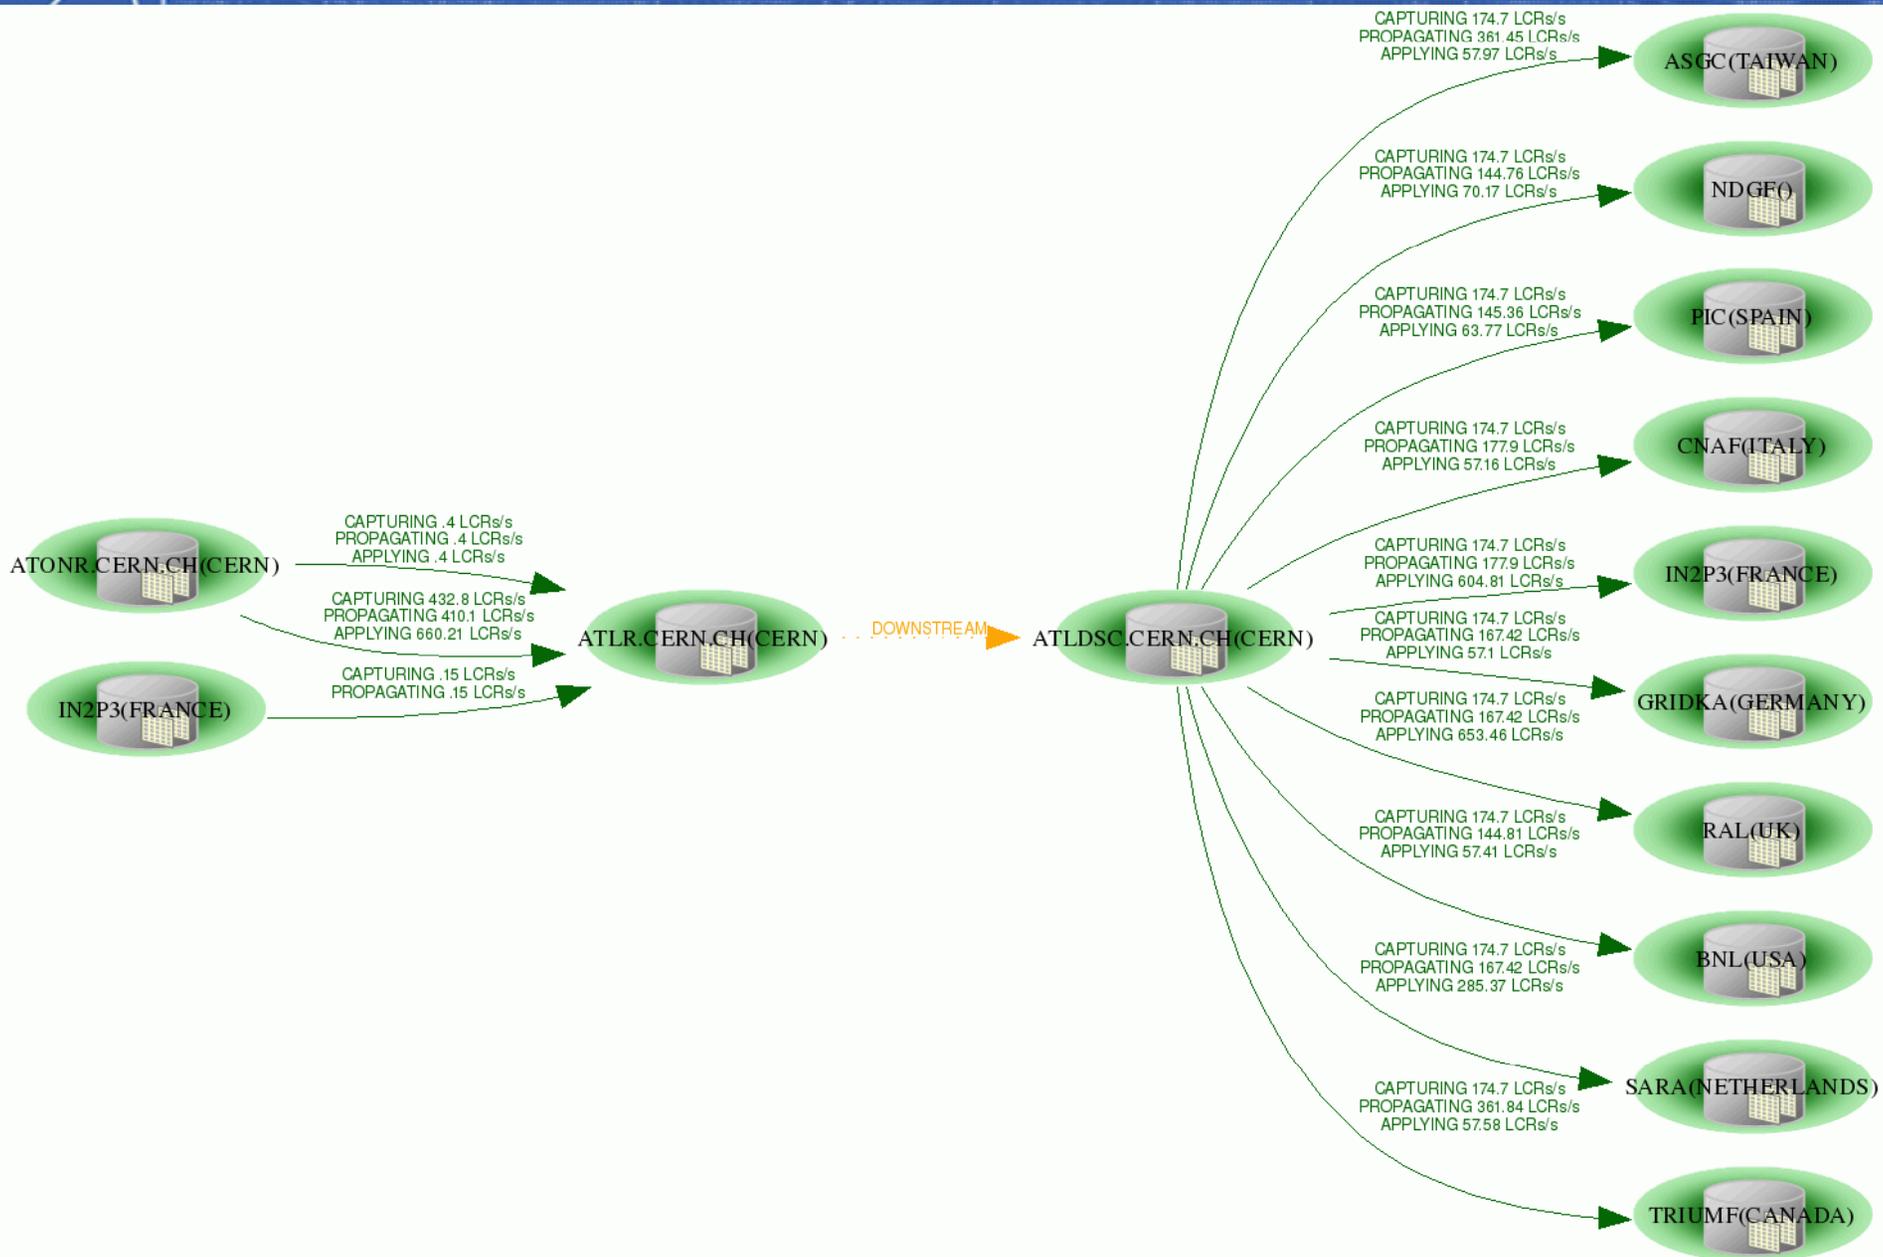
# Oracle Streams Replication

- Technology for **sharing information** between databases
- Database changes captured from the redo-log and propagated asynchronously as Logical Change Records (LCRs)





# Streams setup for ATLAS





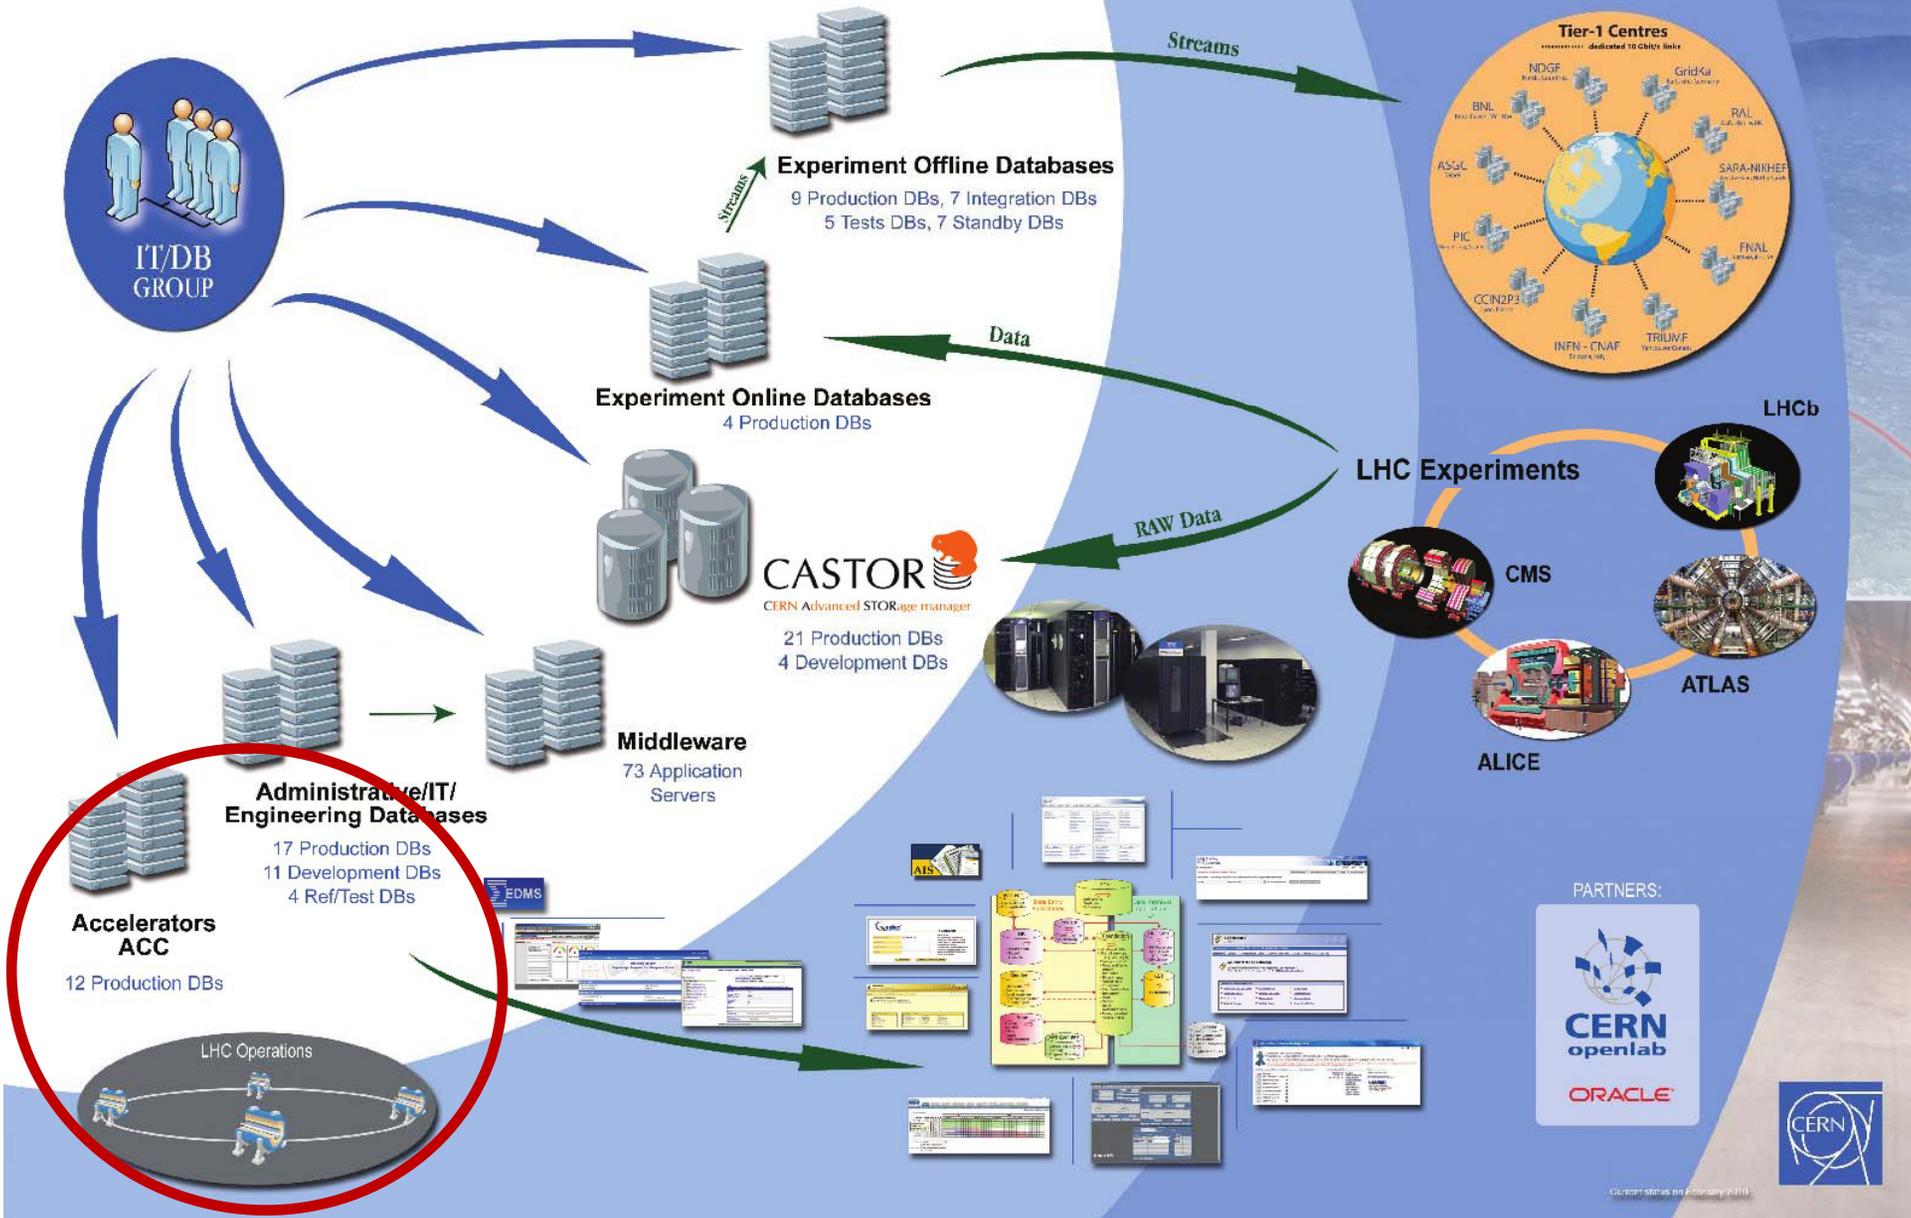
# Metadata Distribution

- To make sense of the **raw data** generated by the detectors, physicists need data about any conditions at the time it was taken that affect the **detector calibration**.
- This **conditions data** is stored in a relational database and needs **distributing to Tier1 centres** to enable future reprocessing of the raw data.
- **Oracle Streams** enables this distribution...
- We will be evaluating **GoldenGate** in the near future as this offers greater flexibility.



# Outline

- Introduction to CERN, Experiments & Data
- **Challenges**
  - The users again...
- Summary/Conclusion





# Databases for Accelerators

ORACLE Enterprise Manager 10g  
Grid Control

Home Targets Deployments Alerts Comp

Hosts | Databases | Middleware | Web Applications | Services | Systems | Groups | All Targets

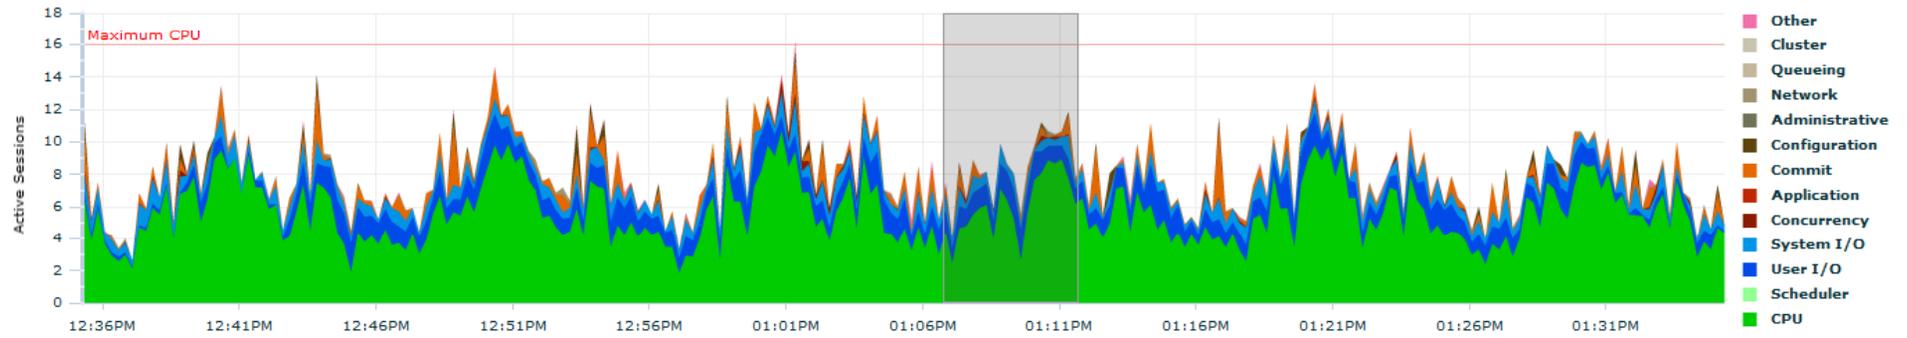
Cluster: accmeas\_cluster > Cluster Database: ACCMEAS.cern.ch > Database Instance: ACCMEAS.cern.ch ACCMEAS3 >

## Top Activity

Switch Database Instance ACCMEAS

Drag the shaded box to change the time period for the detail section below.

View Data Re



## Detail for Selected 5 Minute Interval

Start Time 30-Mar-2010 13:06:47 o'clock MEST

### Top SQL

Actions Schedule SQL Tuning Advisor

Select All | Select None

Select Activity (%)	SQL ID	SQL Type
<input type="checkbox"/>	3guzpbs8t9r6q	SELECT
<input type="checkbox"/>	3pkrm4ccbknata	INSERT
<input type="checkbox"/>	3k1gkn3n5sr2b	189
<input type="checkbox"/>	2a6aduxkh4apt	SELECT
<input type="checkbox"/>	2bwzcbnwk1awz	SELECT
<input type="checkbox"/>	c8krkrfxcnmg3	SELECT
<input type="checkbox"/>	7pd3cywtvut4b	SELECT
<input type="checkbox"/>	qjugsph43w51v	SELECT
<input type="checkbox"/>	dwnh1mnszqxs	SELECT
<input type="checkbox"/>	7ibtz6rnqdv4z	PL/SQL EXECUTE

Actions Schedule SQL Tuning Advisor

Total Sample Count: 2,285

### Top Sessions

View Top Sessions

Activity (%)	Session ID	User Name	Program
9.56	508	MEASDB	
7.68	528	MEASDB_PUB	python@cs-ccr-abbi2 (TNS V1-V3)
7.18	503	MEASDB	
5.30	441	MEASDB	
5.07	491	MEASDB	
4.88	281	MEASDB	oracle@dbsrva251.cern.ch (J001)
3.95	279	MEASDB	
3.69	398	MEASDB	
3.65	33	MEASDB	
2.99	298	SYS	oracle@dbsrva251.cern.ch (LGWR)

Total Sample Count: 2,605

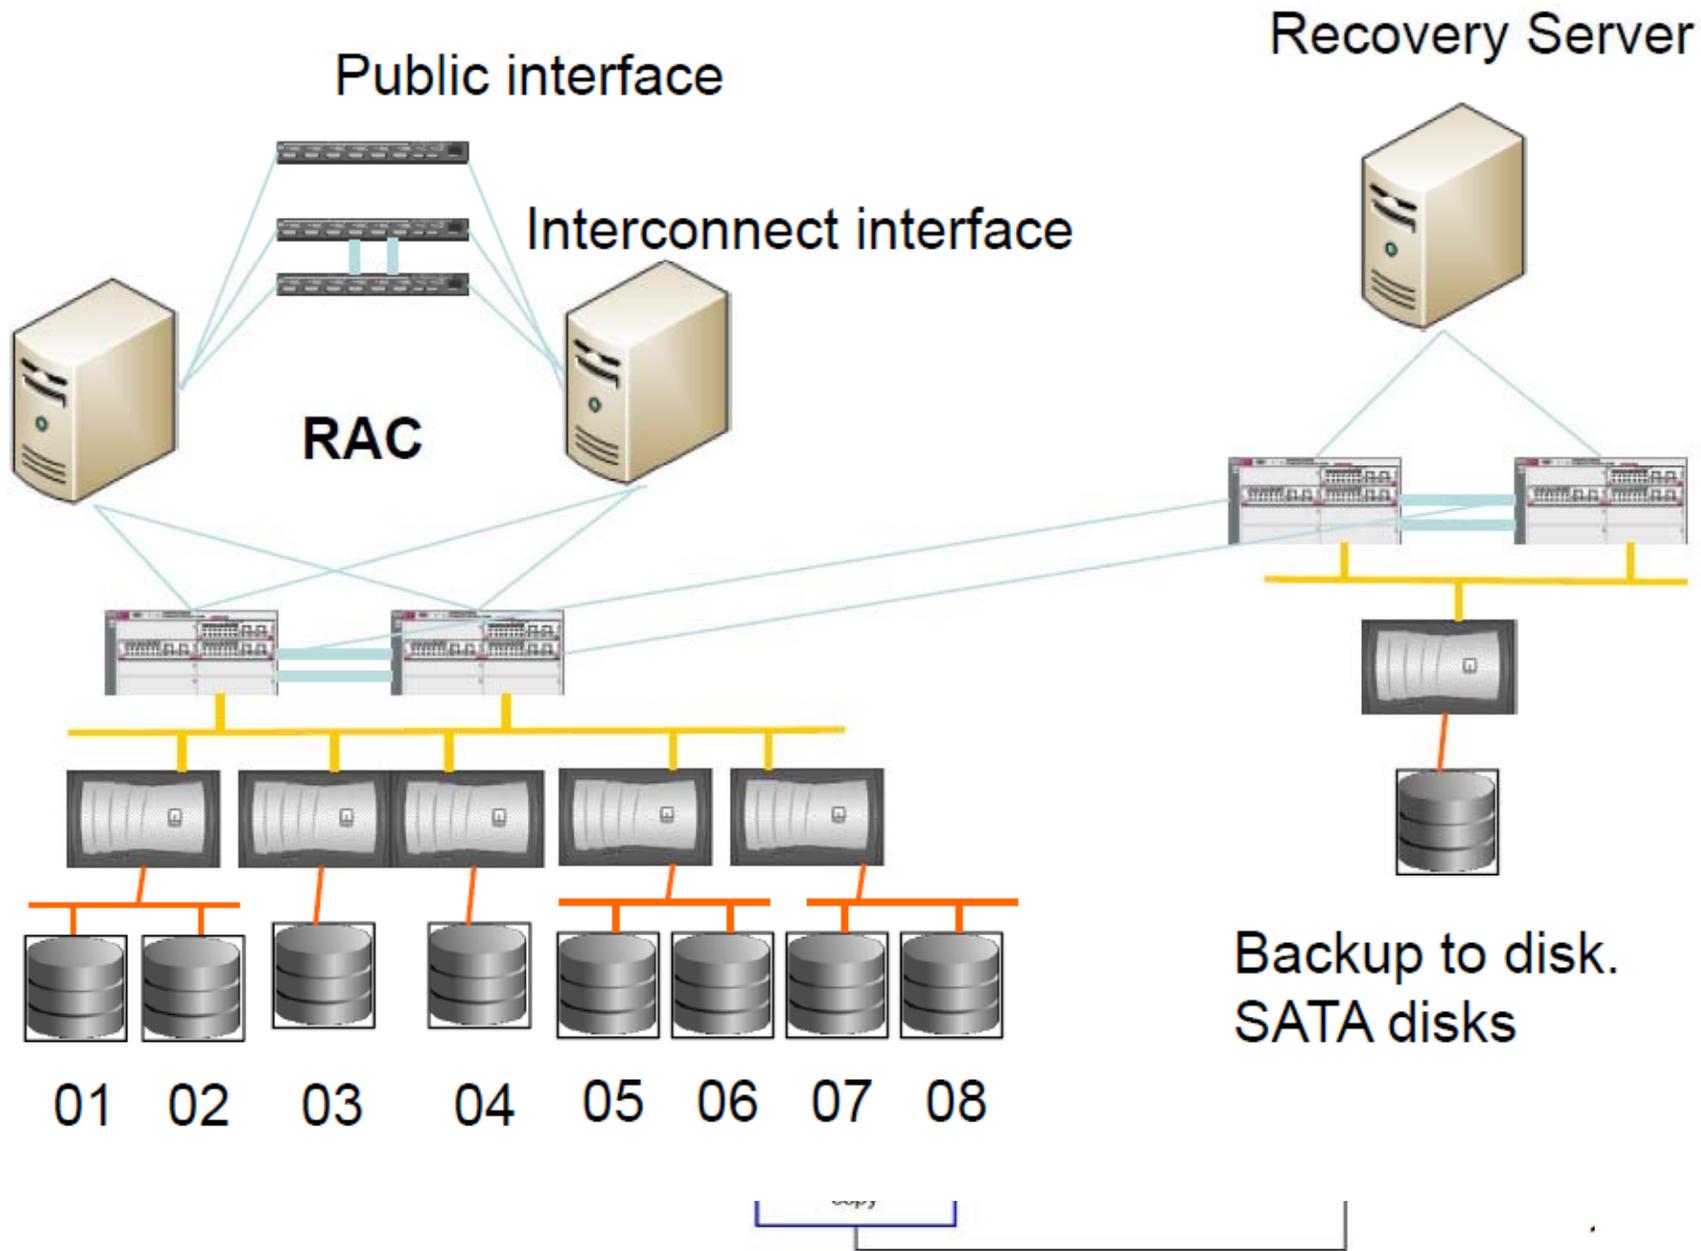
0



# Backups: A digression

- Backing up databases is a **chore**, not a **challenge**...
- ...the challenge is to **recover** a database...
  - ... and this had better be rare, not a chore!
- To help reduce the challenge of, and increase confidence in, recovery procedures, we have developed an **automated recovery test tool**.

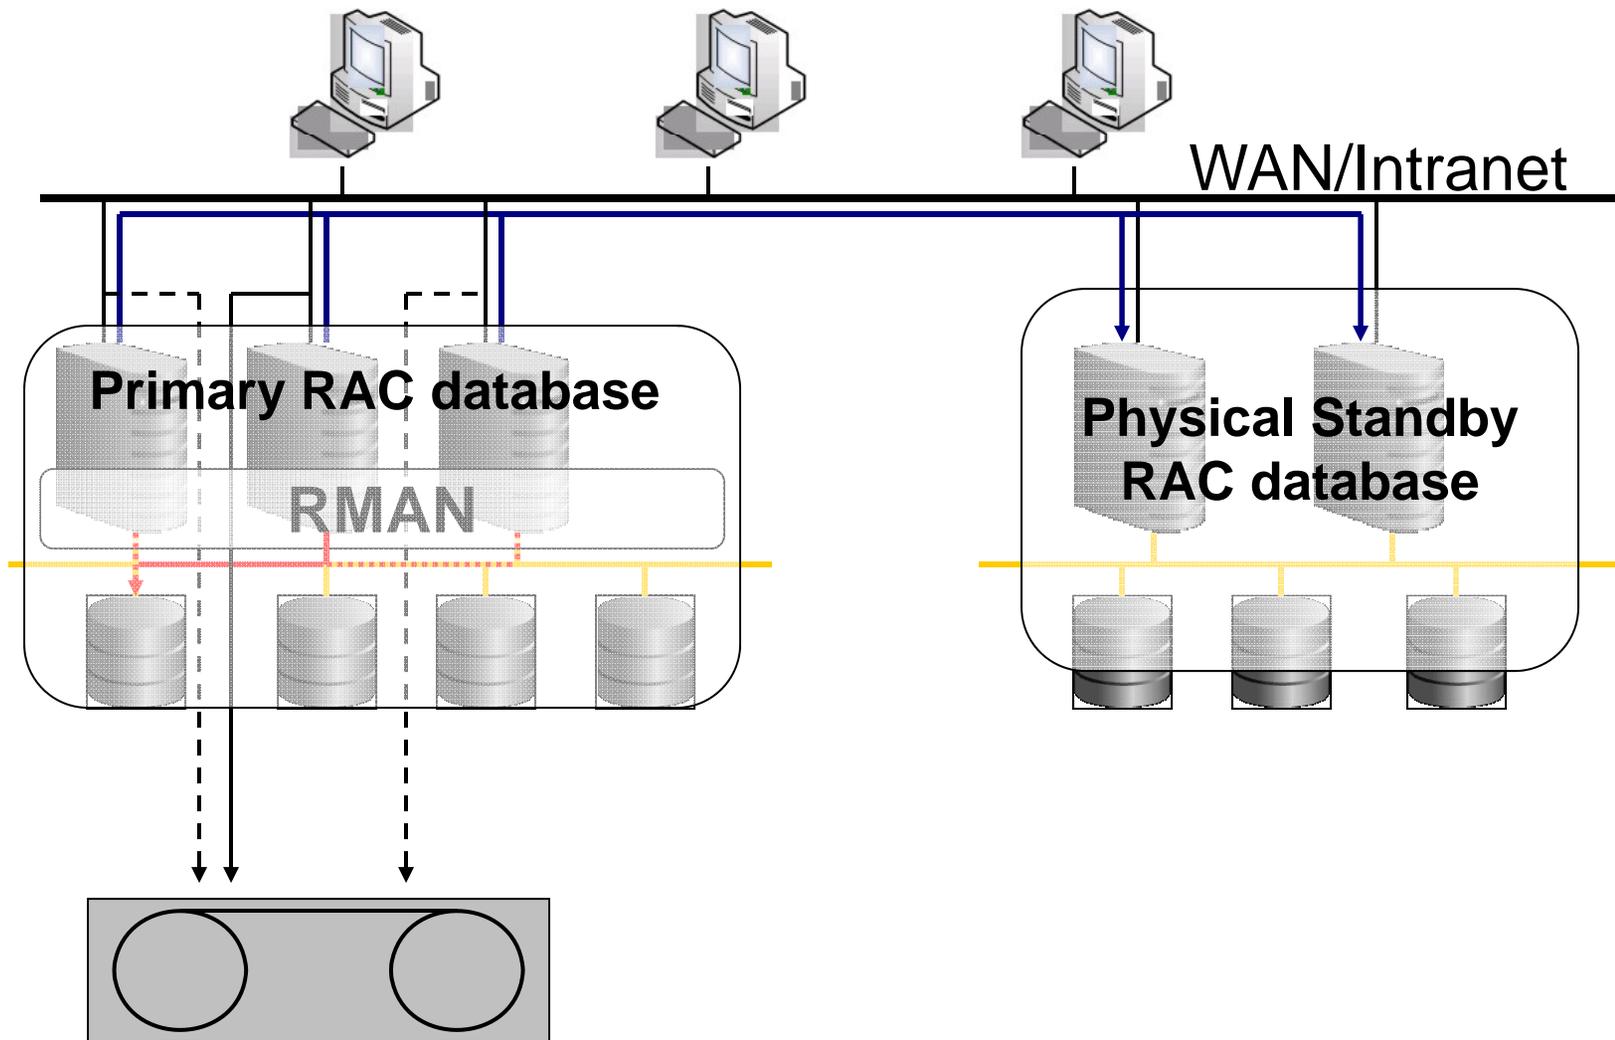
Initial set-up:  
32 or 64 bits





# Backups: A digression

- Backing up databases is a **chore**, not a **challenge**...
- ...the challenge is to **recover** a database...
  - ... and this had better be rare, not a chore!
- To help reduce the challenge of, and increase confidence in, recovery procedures, we have developed an **automated recovery test tool**.
- Problem: **days to recover some databases**
  - At least 2.5 days for *partial restore* of the critical database for accelerator operations...
  - ➔ use Data Guard to create standby databases.

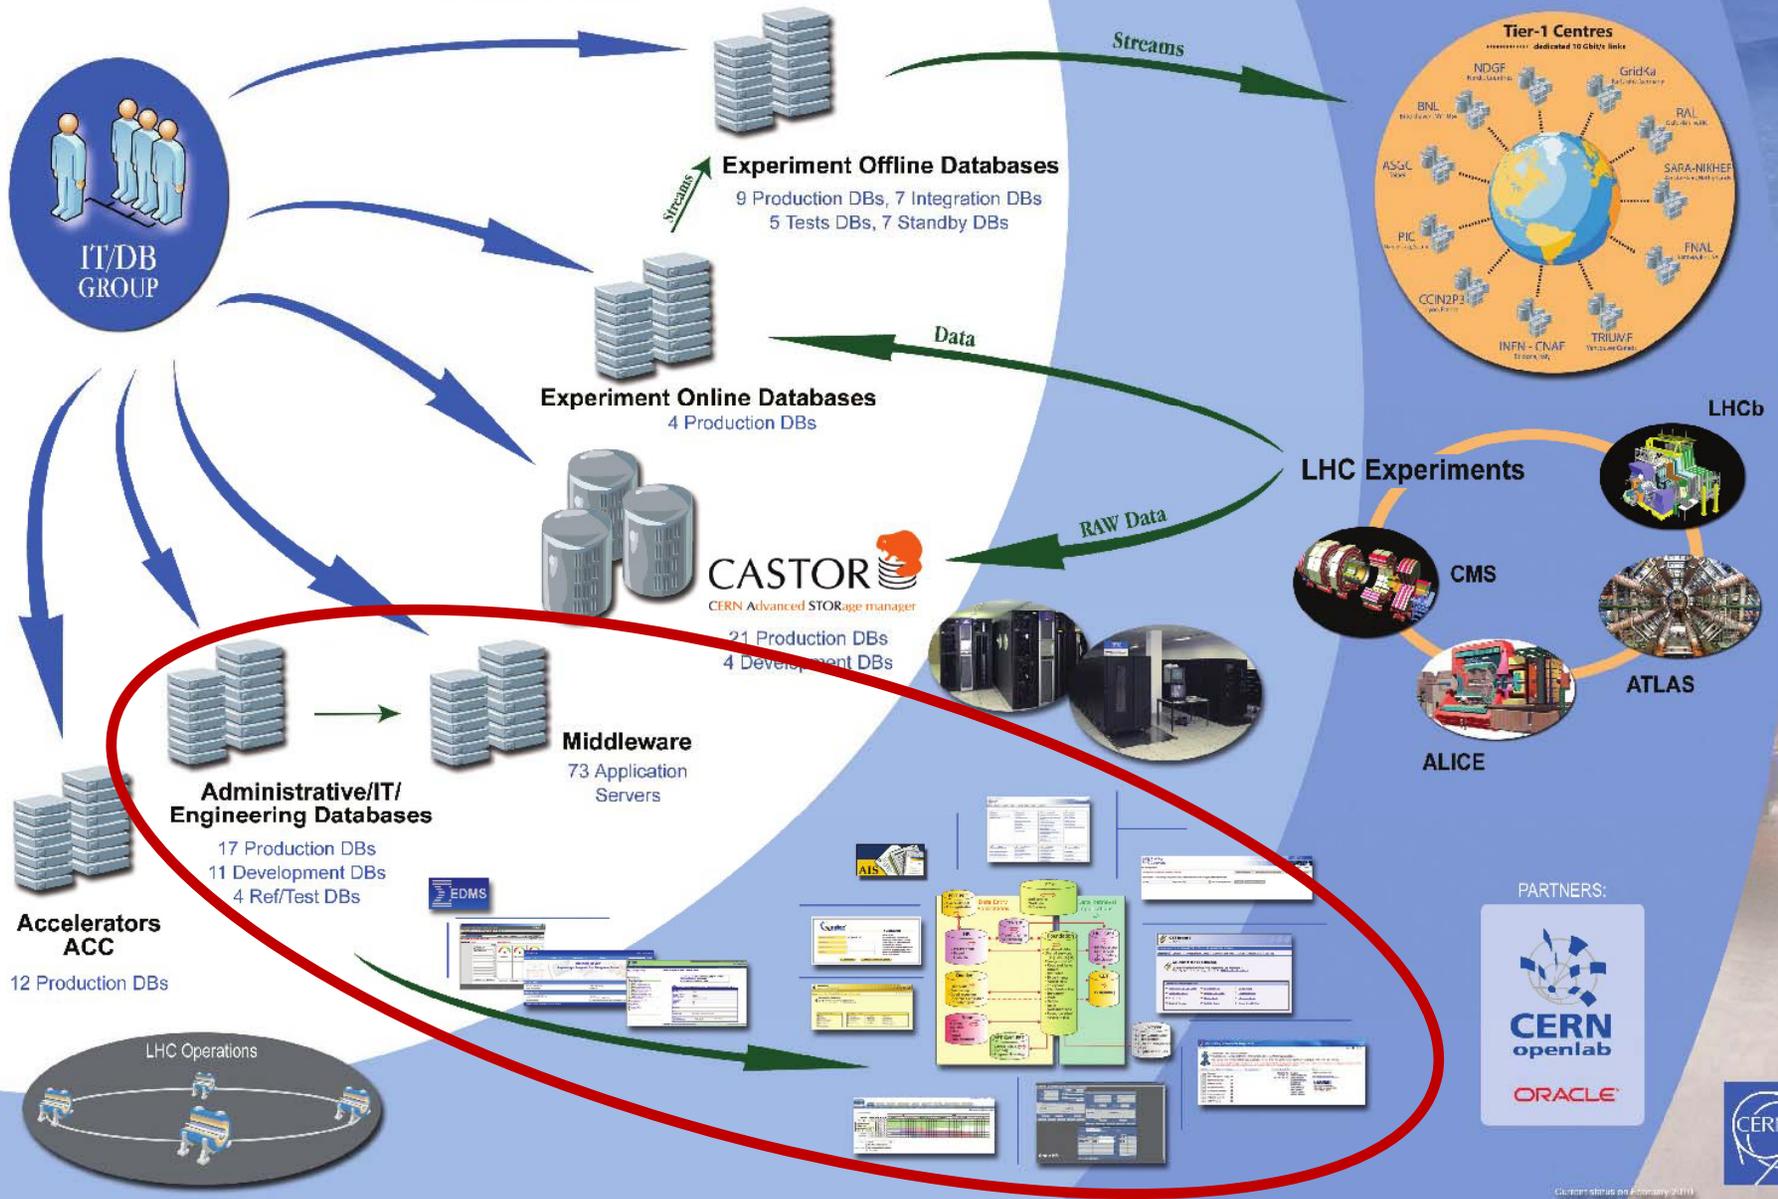


**We are extremely keen to migrate to 11g and exploit Active Data Guard!**



# Outline

- Introduction to CERN, Experiments & Data
- **Challenges**
  - Virtualisation & Fabric Management
- Summary/Conclusion



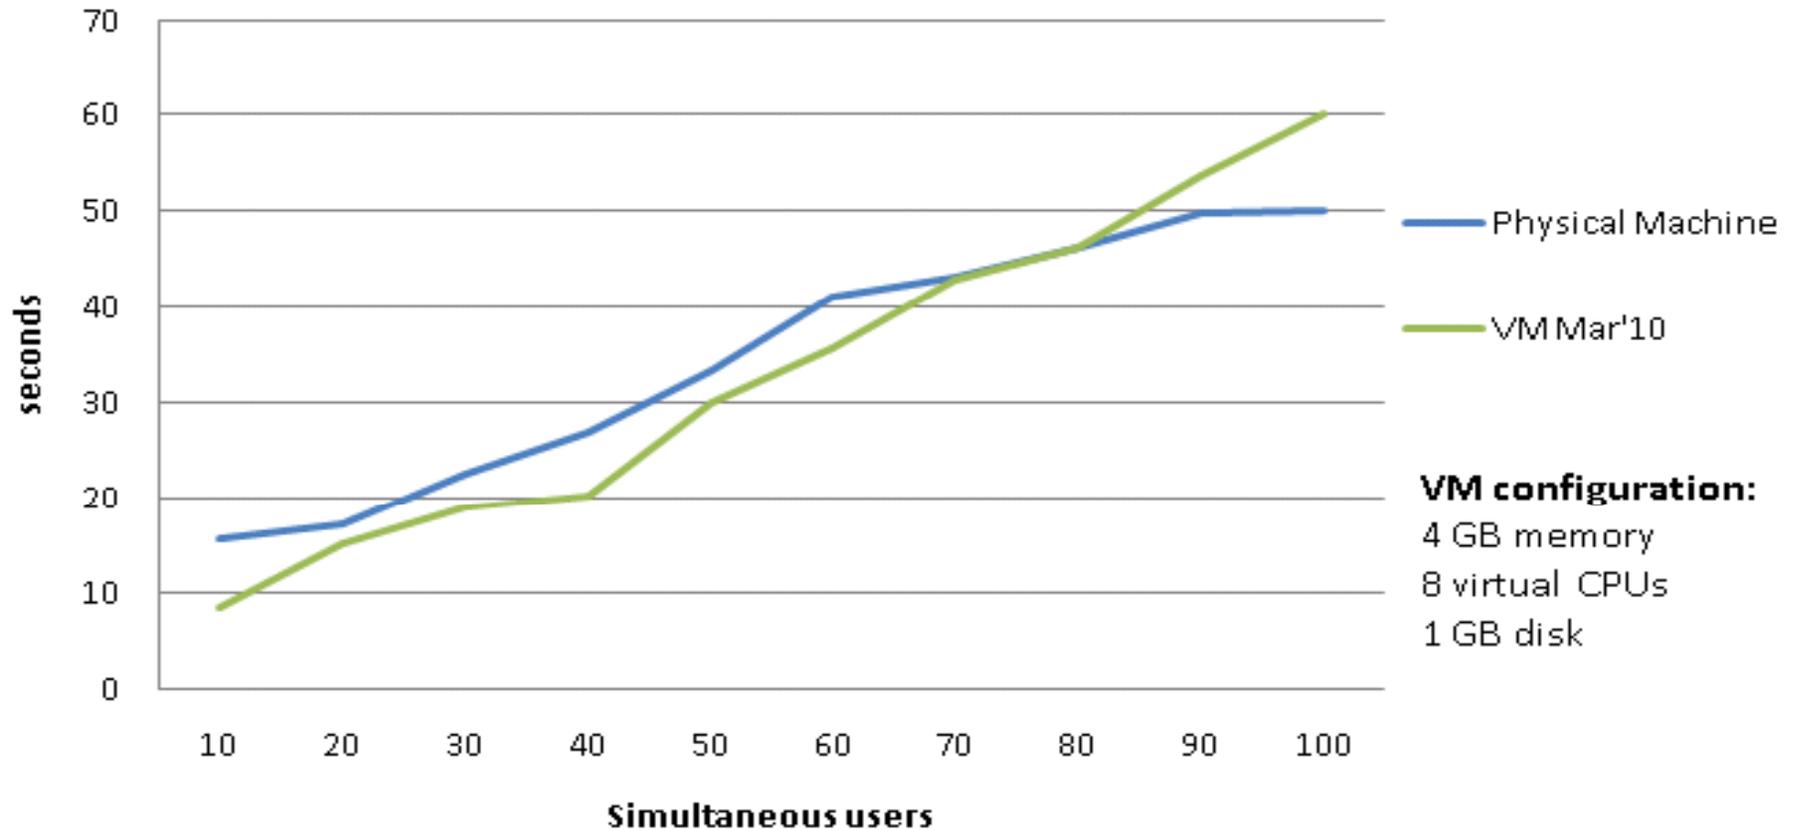


## Size isn't all...

- As well as being **vital for accelerator and experiment operations**, Oracle databases also **underpin CERN's administrative applications**.
- Individual databases are **small**, but we have **many servers**
- Average load is **very low**, so **clear opportunity for server consolidation** with **virtualisation**.
- Excellent results from tests with Oracle VM and Web Logic Server virtual edition.

# Physical vs Virtual

## Physical vs VM'10



- Comparison of a Physical Machine (4GB memory, 8 cores) running WebLogic Server versus a Virtual Machine (4GB memory, 8 Virtual CPUs) running WebLogic Server – Virtual Edition
- An unreleased version of the kernel of JRockit has been used for this tests.



# Fabric Management

- As most large organisations, CERN has **standards and procedures** for network and system configuration.
- “Easy to install” applications can be good, but they should be **configurable** to work correctly in a tightly managed environment.
- **Good collaboration around Oracle VM**
  - e.g. IP/MAC address binding, assumptions about NFS as shared file system.



# Outline

- Introduction to CERN, Experiments & Data
- **Challenges**
  - Monitoring
- Summary/Conclusion

<http://cern.ch/it-dep/db/>

Service Level Status overview

Home Search KPIs Tags Admin Documentation Help

Services provided by IT/DB group 28 May 2010 Fri 14:24:46

- IT/DB services**
- availability:  (more)
- percentage: 99%
- status: **available**
- this service consists of:
- Accelerator databases
  - AIS databases
  - EDMS databases
  - General purpose databases
  - IT infrastructure databases
  - CASTOR and SRM databases
  - ALICE databases
  - ATLAS databases
  - CMS databases
  - LHCb databases
  - Other experiments databases
  - WLCG database
  - Physics DB Streams
  - J2EE Public Service
  - EDMS Service
  - RMAN Jobs

**Additional information**

full name: **Services provided by IT/DB group**

short name: IT/DB services

group: IT-DB

email: [Helpdesk@cern.ch](mailto:Helpdesk@cern.ch)

service Eric Grancher

managers: Nilo Segura Chinchilla  
Eva Dafonte Perez

**Availability update**

last update: 14:16:53, 28 May 2010 (8 minutes ago)

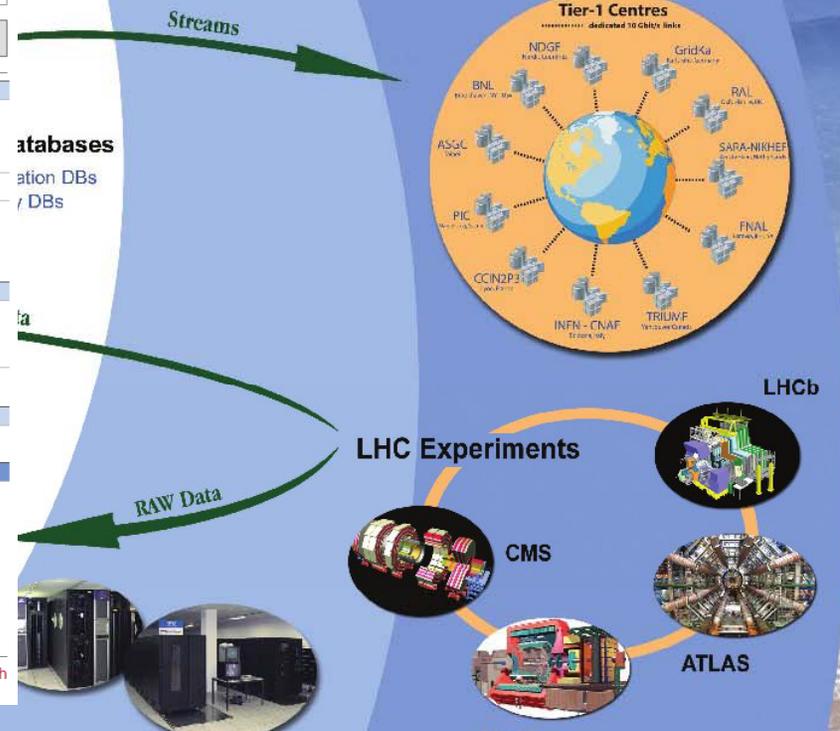
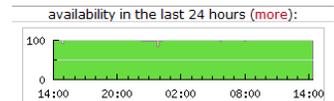
expires after: 57 minutes

[rss feed with status changes](#)

Part of (subservice of): none / not declared

**Admin**

[admin tools](#)



SLS by CERN IT/CF

[SLS.Support@cern.ch](mailto:SLS.Support@cern.ch)

**Administrative/IT/ Engineering Databases**

- 17 Production DBs
- 11 Development DBs
- 4 Ref/Test DBs

**Middleware**

73 Application Servers

**Accelerators ACC**

12 Production DBs



PARTNERS:

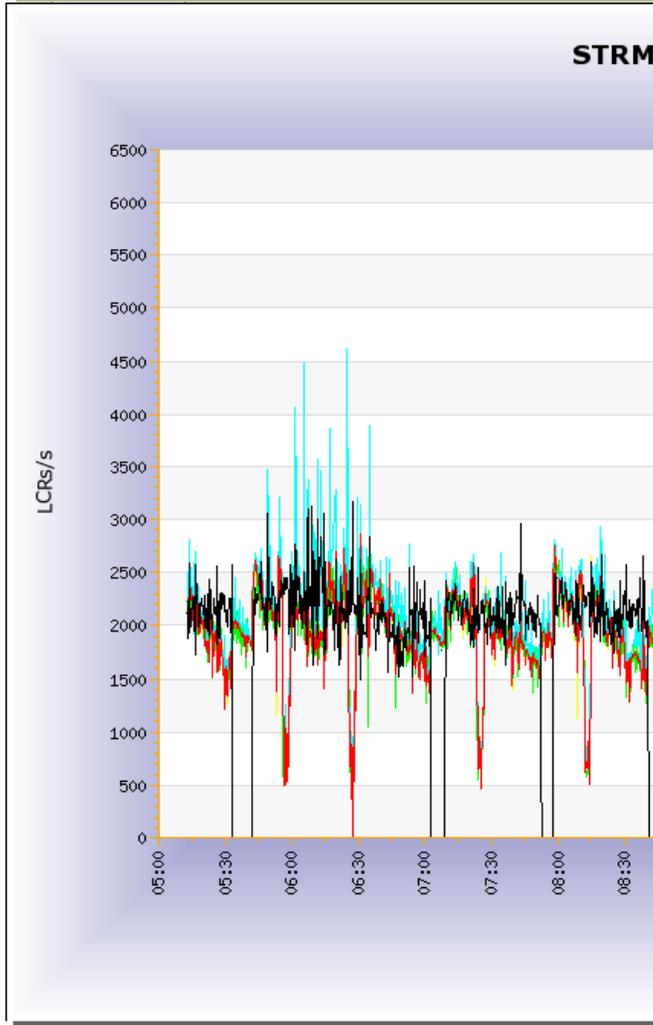




# Up and seen to be up

- Understanding the current state of our many database servers and applications is a **major challenge**.
- **Collaboration around Oracle Enterprise Manager** has been a really fruitful aspect of the CERN openlab partnership with Oracle.
- Many tools developed at CERN for monitoring Streams have led to features in OEM 10.2 and 11.1.

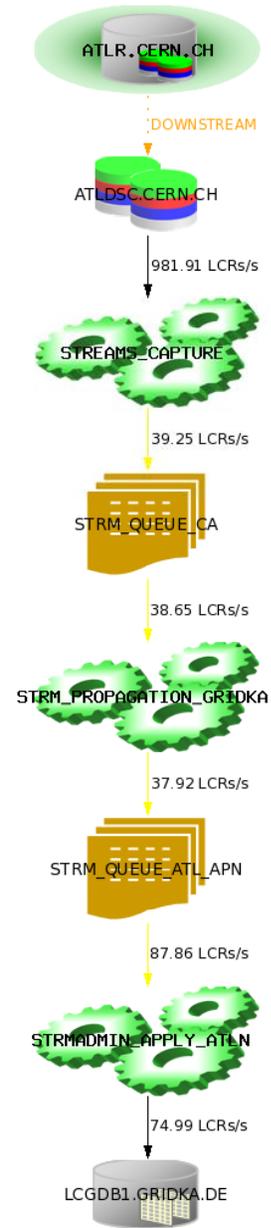
Streams



Remove Streams Configuration

STREAM DETAILS

SOURCE		ATLDSC.CERN.CH
<b>CAPTURE</b>		
Name	STREAMS_CAPTURE	
SCN read	6110514671773	
LCRs captured	135508158 (981.91/s)	
SCN captured	6110514662689	
LCRs enqueued	4098569 (39.25/s)	
SCN enqueued	6110514662572	
Capture Latency	33 sec	
Info	ATLDSC.CERN.CH	
State	CAPTURING CHANGES	
Error Time		
Error Msg		
<b>QUEUE</b>		
Name	STRM_QUEUE_CA	
Id	42947	
Outstanding Msg	24	
Cumulative Msg	6223281 (38.69/s)	
Cumulative Msg Spilled	2693626 (0/s)	
<b>PROPAGATION</b>		
Name	STRM_PROPAGATION_GRIDKA	
LCRs Propagated	3783492 (38.65/s)	
Bytes Propagated	0 (0/s)	
State	ENABLED	
Error Time		
Error Msg		
<b>DESTINATION</b>		
		LCGDB1.GRIDKA.DE
<b>QUEUE</b>		
Name	STRM_QUEUE_ATL_APN	
Id	198958	
Outstanding Msg	0	
Cumulative Msg	240899548 (37.92/s)	
Cumulative Msg Spilled	53380 (0/s)	
<b>APPLY</b>		
Name	STRMADMIN_APPLY_ATLN	
LCRs Dequeued	240899548 (87.86/s)	
Total LCRs Applied	235085002 (74.99/s)	
SCN Dequeued	6110514660055	
Dequeue Latency	27 sec	
Transaction Received	4435350	
Transaction Assigned	4435352	
Transaction Applied	4435350	
State	IDLE	
HWM SCN	6110514639800	
HWM Apply Latency	33 sec	
Last applied LCR	2010-05-27 16:18:04 (11 sec ago)	
Data Latency	33 sec	
Error Time		
Error Msg		





# Up and seen to be up

- Understanding the current state of our many database servers and applications is a **major challenge**.
- **Collaboration around Oracle Enterprise Manager** has been a really fruitful aspect of the CERN openlab partnership with Oracle.
- Many tools developed at CERN for monitoring Streams have led to features in OEM 10.2 and 11.1.
- As in the fabric management area, integration with other tools is desirable. A future challenge?

User-Defined Numeric Metric: Script SESSIONS\_LIMIT\_IDLE\_2H: Last 24 hours

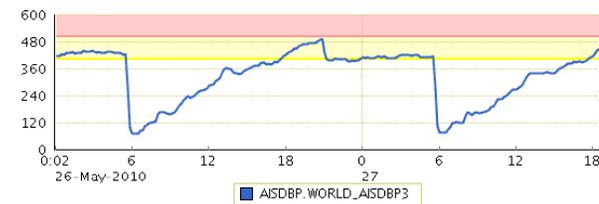
Last Updated 27-May-2010 18:32:57 MEST  
View Data Last 24 hours

Script SESSIONS\_LIMIT\_IDLE\_2H

Statistics

Last Known Value 448  
Average Value 339.33  
High Value 493  
Low Value 79  
Warning Threshold 400  
Critical Threshold 500  
Occurrences Before Alert 1  
Corrective Action None

Metric Value



Alert History

Comment for Most Recent Alert

Add Comment

Severity	Timestamp	Message	Last Comment	Details
Warning	27-May-2010 17:42:57	UDM alert (SESSIONS_LIMIT): key = %Key_value = 406		-
Success	27-May-2010 05:42:57	CLEARED - UDM alert (SESSIONS_LIMIT): key = %Key_value = 283		-
Warning	26-May-2010 23:52:57	UDM alert (SESSIONS_LIMIT): key = %Key_value = 407		-
Success	26-May-2010 22:52:57	CLEARED - UDM alert (SESSIONS_LIMIT): key = %Key_value = 395		-
Warning	26-May-2010 21:52:57	UDM alert (SESSIONS_LIMIT): key = %Key_value = 406		-
Success	26-May-2010 21:22:57	CLEARED - UDM alert (SESSIONS_LIMIT): key = %Key_value = 397		-
Warning	26-May-2010 17:42:57	UDM alert (SESSIONS_LIMIT): key = %Key_value = 408		-

Related Links

Acknowledge

Compare Objects Script

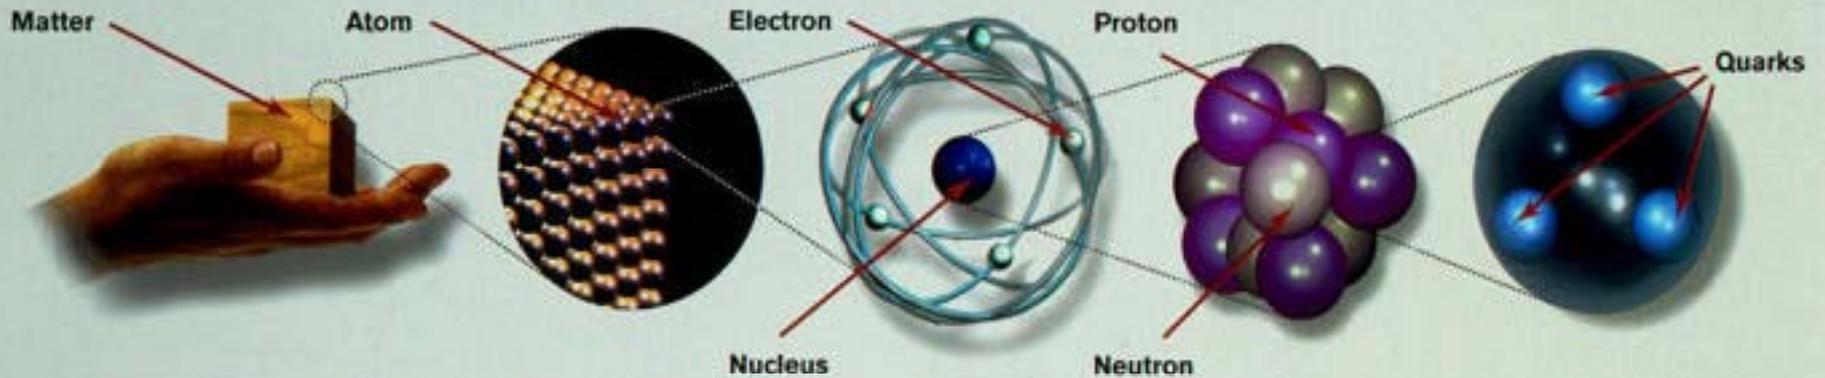
Compare Targets

Metric and Policy Settings



# Outline

- Introduction to CERN, Experiments & Data
- Challenges
- **Summary/Conclusion**



**Matter particles**  
All ordinary particles belong to this group

LEPTONS		
<b>FIRST FAMILY</b>	<b>Electron</b> Responsible for electricity and chemical reactions; it has a charge of -1	<b>Electron neutrino</b> Particle with no electric charge, and possibly no mass; billions fly through your body every second
<b>SECOND FAMILY</b>	<b>Muon</b> A heavier relative of the electron; it lives for two-millionths of a second	<b>Muon neutrino</b> Created along with muons when some particles decay
<b>THIRD FAMILY</b>	<b>Tau</b> Heavier still, it is extremely unstable. It was discovered in 1975	<b>Tau neutrino</b> Not yet discovered but believed to exist

QUARKS		
<b>Up</b> Has an electric charge of plus two-thirds; protons contain two, neutrons contain one		<b>Down</b> Has an electric charge of minus one-third; protons contain one, neutrons contain two
<b>Charm</b> A heavier relative of the up; found in 1974		<b>Strange</b> A heavier relative of the down; found in 1964
<b>Top</b> Heavier still		<b>Bottom</b> Heavier still; measuring bottom quarks is an important test of electroweak theory

**Force particles**  
These particles transmit the four fundamental forces of nature although gravitons have so far not been discovered

**Gluons**  
Carriers of the **strong force** between quarks

Felt by quarks

The explosive release of nuclear energy is the result of the **strong force**

**Photons**  
Particles that make up light; they carry the **electromagnetic force**

Felt by quarks and charged leptons

Electricity, magnetism and chemistry are all the results of **electro-magnetic force**

**Intermediate vector bosons**  
Carriers of the **weak force**

Felt by quarks and leptons

Some forms of radio-activity are the result of the **weak force**

**Gravitons**  
Carriers of **gravity**

Felt by all particles with mass

All the weight we experience is the result of the **gravitational force**

# The Twenty Member States of CERN



## Member States (Dates of Accession)

 AUSTRIA (1959)	 DENMARK (1953)	 GREECE (1953)	 NORWAY (1953)	 SPAIN (1/1961-12/1968-1/1983)
 BELGIUM (1953)	 FINLAND (1991)	 HUNGARY (1992)	 POLAND (1991)	 SWEDEN (1953)
 BULGARIA (1999)	 FRANCE (1953)	 ITALY (1953)	 PORTUGAL (1986)	 SWITZERLAND (1953)
 CZECH FR (1993)	 GERMANY (1953)	 NETHERLANDS (1953)	 SLOVAK FR (1993)	 UNITED KINGDOM (1953)

