Peter Chochula
CERN/ALICE

THE CYBERSECURITY IN ALICE - AS SEEN FROM USER'S PERSPECTIVE

The ALICE experiment at CERN LHC

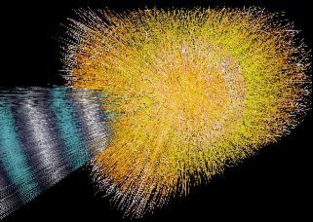

Basic task – study of heavy ion collisions at LHC
The design of ALICE allows for p-p studies as well

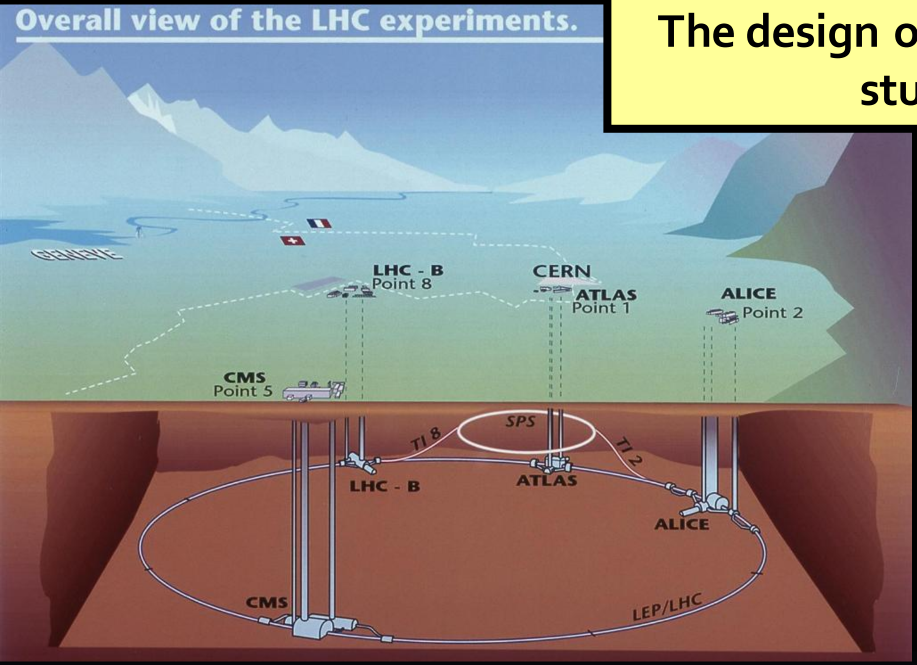

Experiment
Size: 16 x 26 metres (some detectors placed >100m from the interaction point)
Mass: 10,000,000 kg
Detectors: 20
Magnets: 2

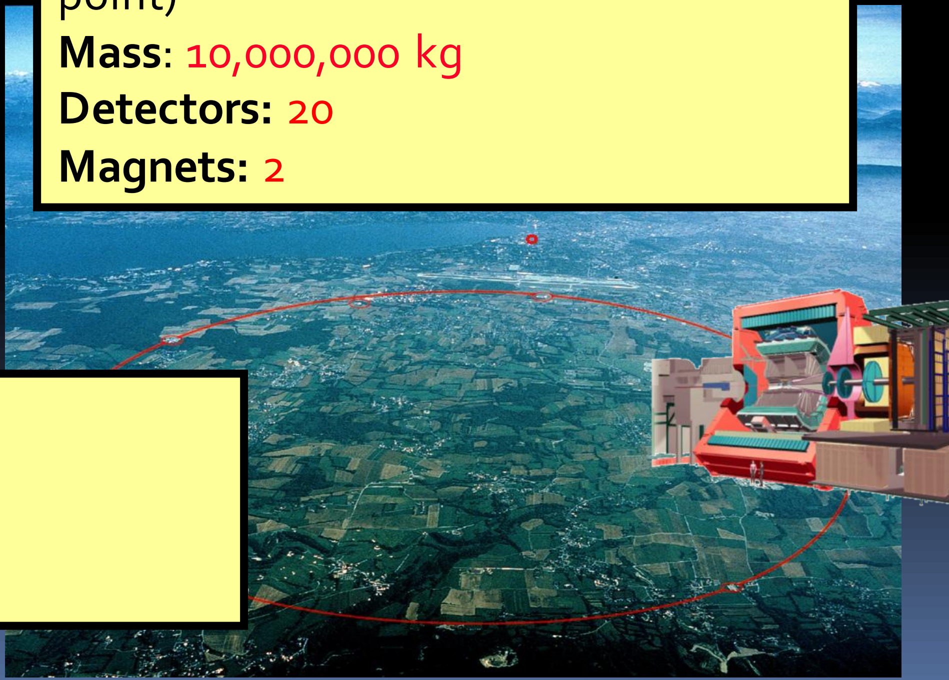

The ALICE Collaboration:
Members: 1300
Institutes: 116
Countries: 33

- ALICE - a very visible object, designed to detect the invisible…


CICERON

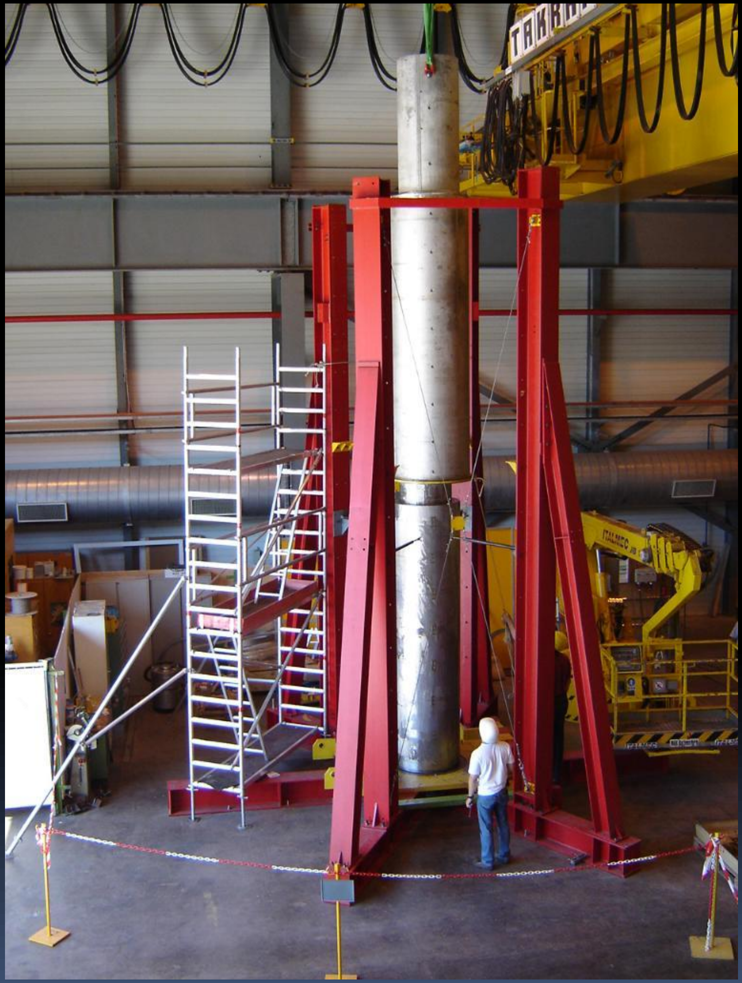

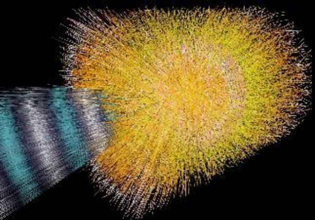

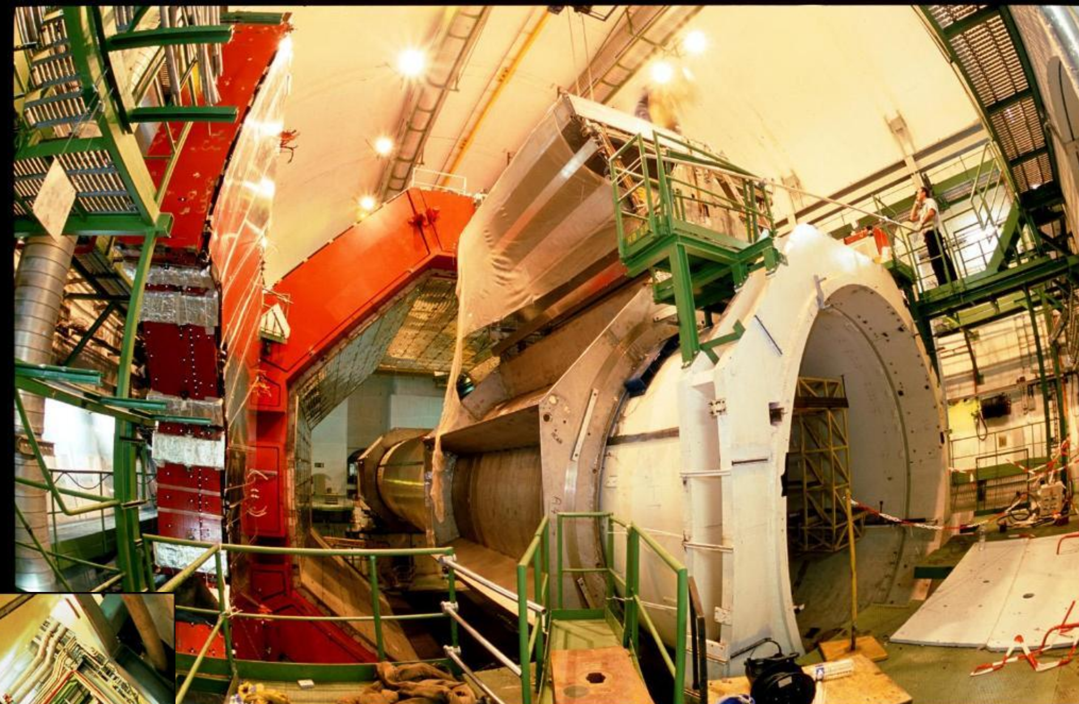


ALICE

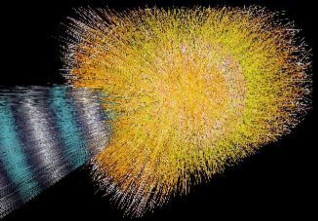


CICERON

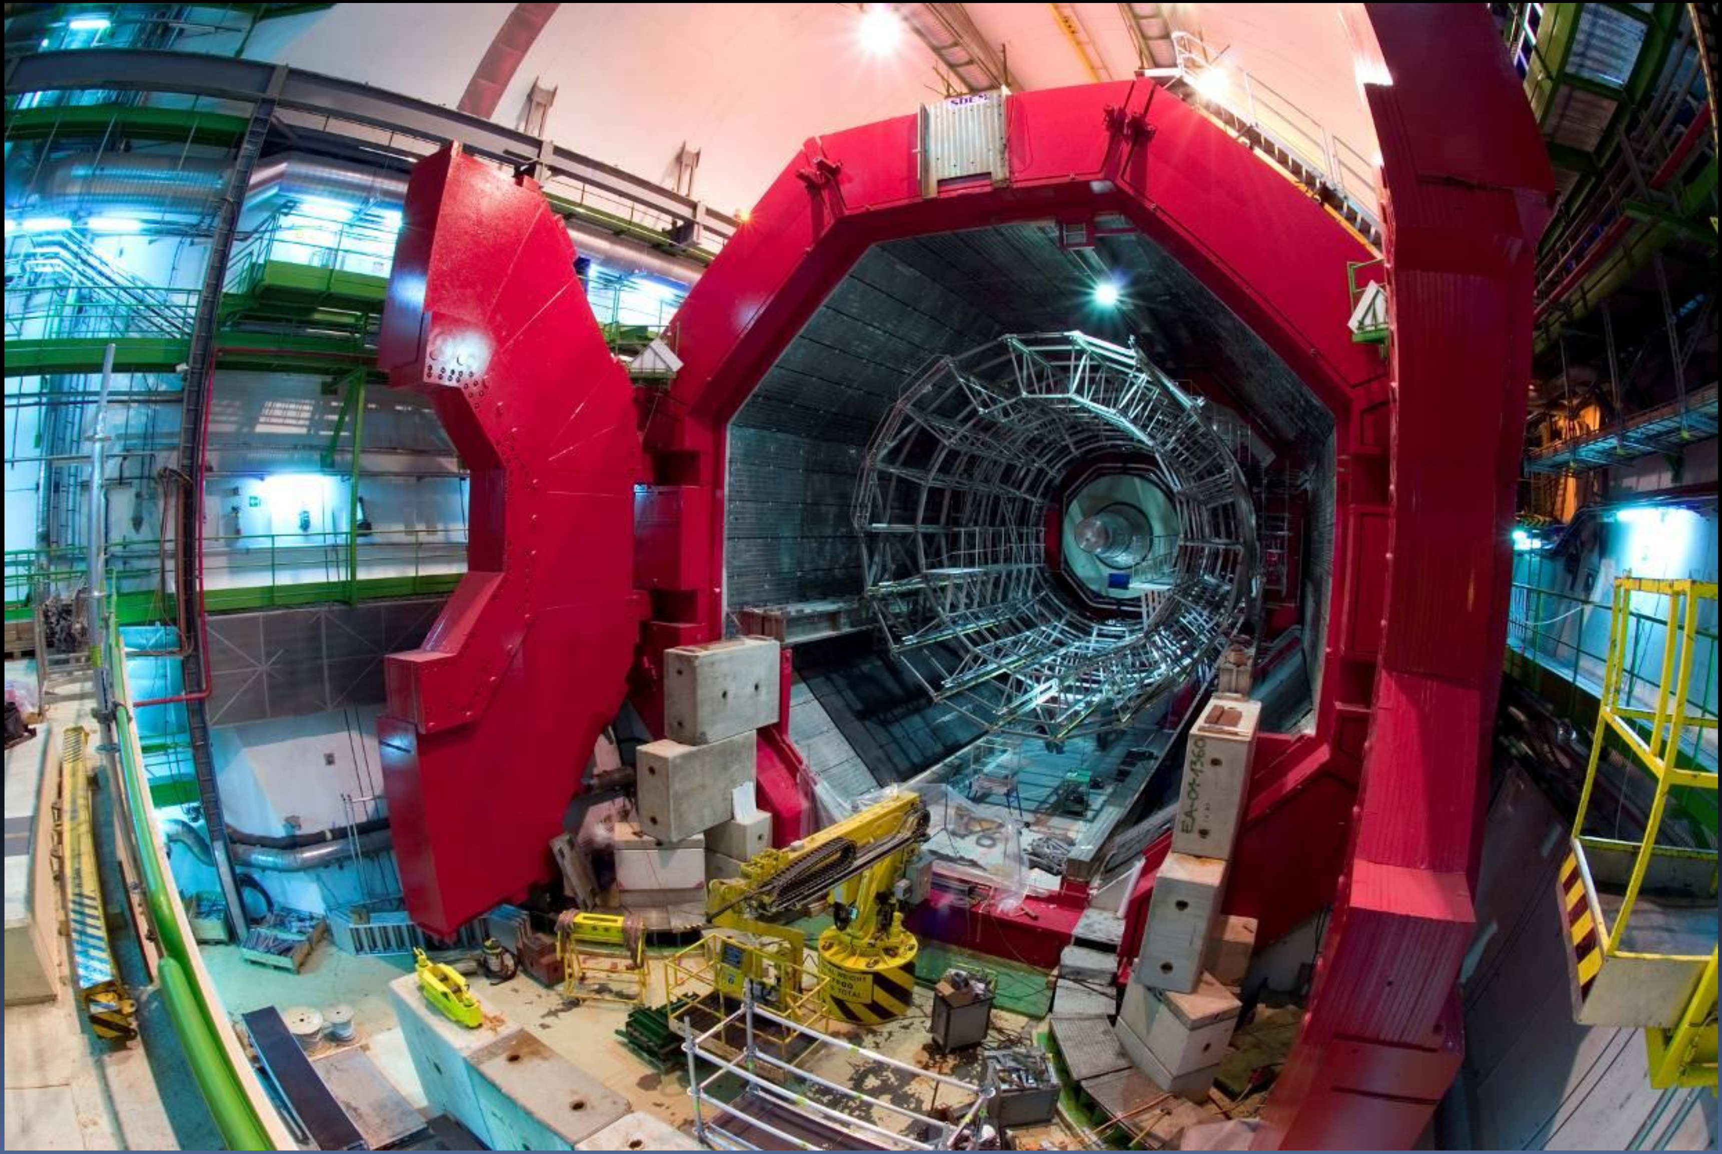

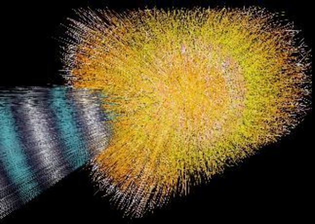

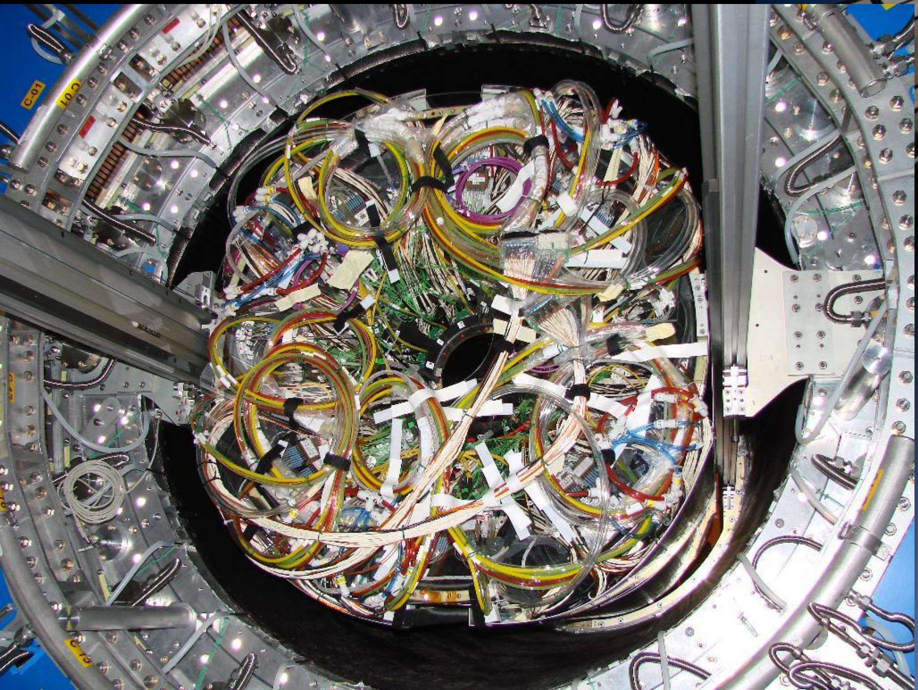

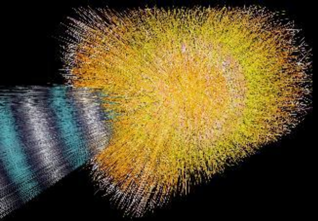

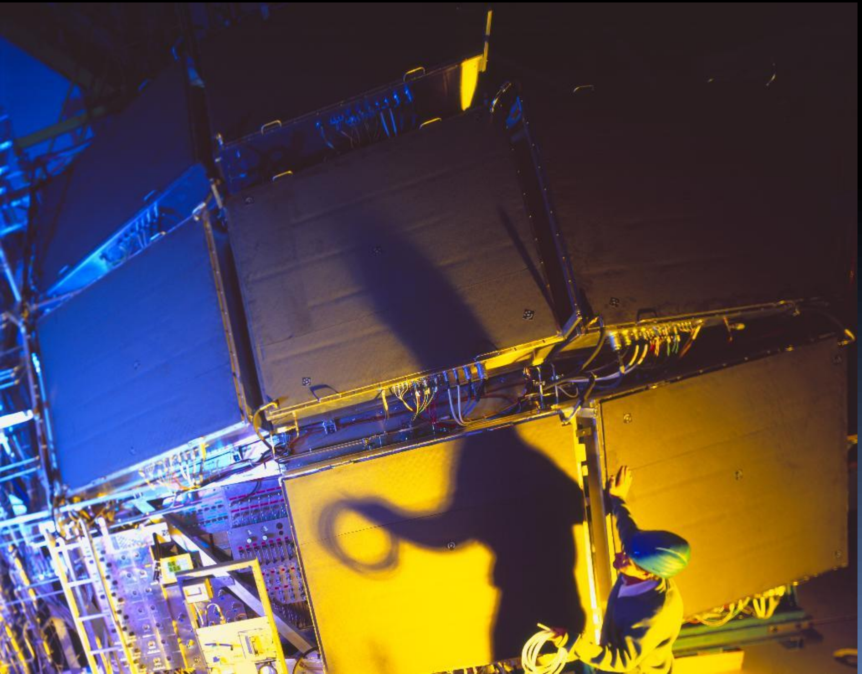

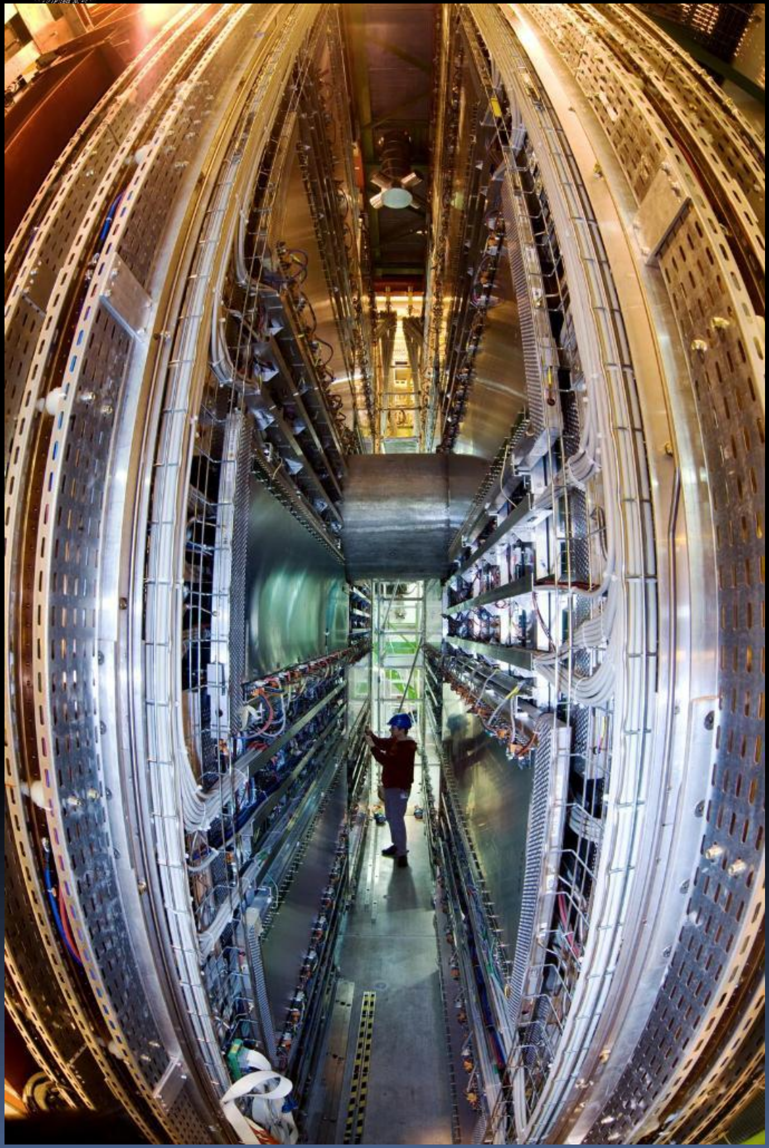

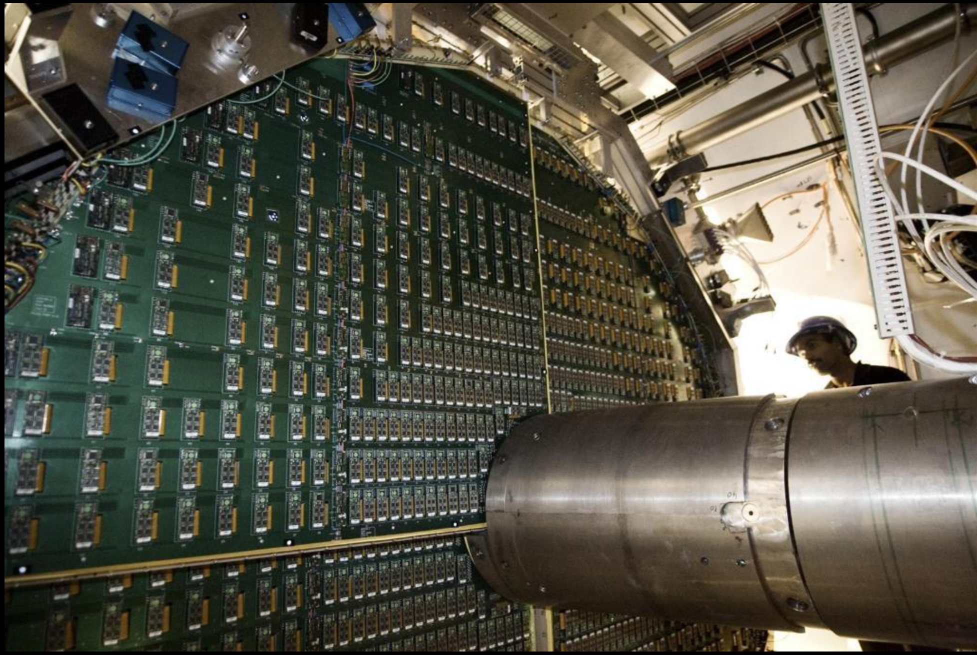

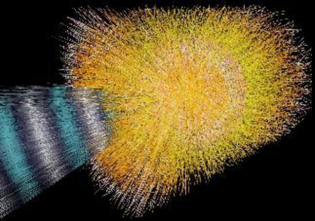

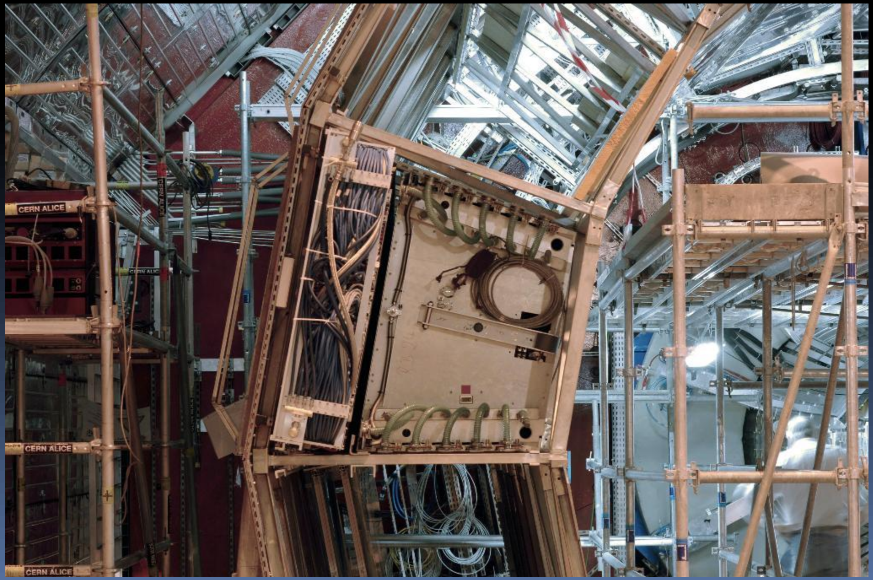
CERN ALICE

Operational since the very beginning

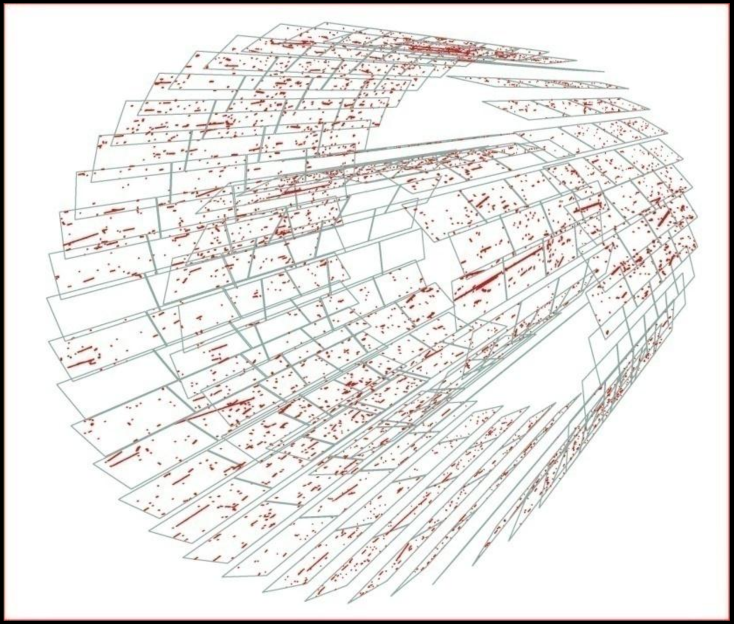

Historically first particles in LHC were detected by ALICE pixel detector
Injector tests, June 15 2008

First proton collisions

Luminosity monitor (V0)

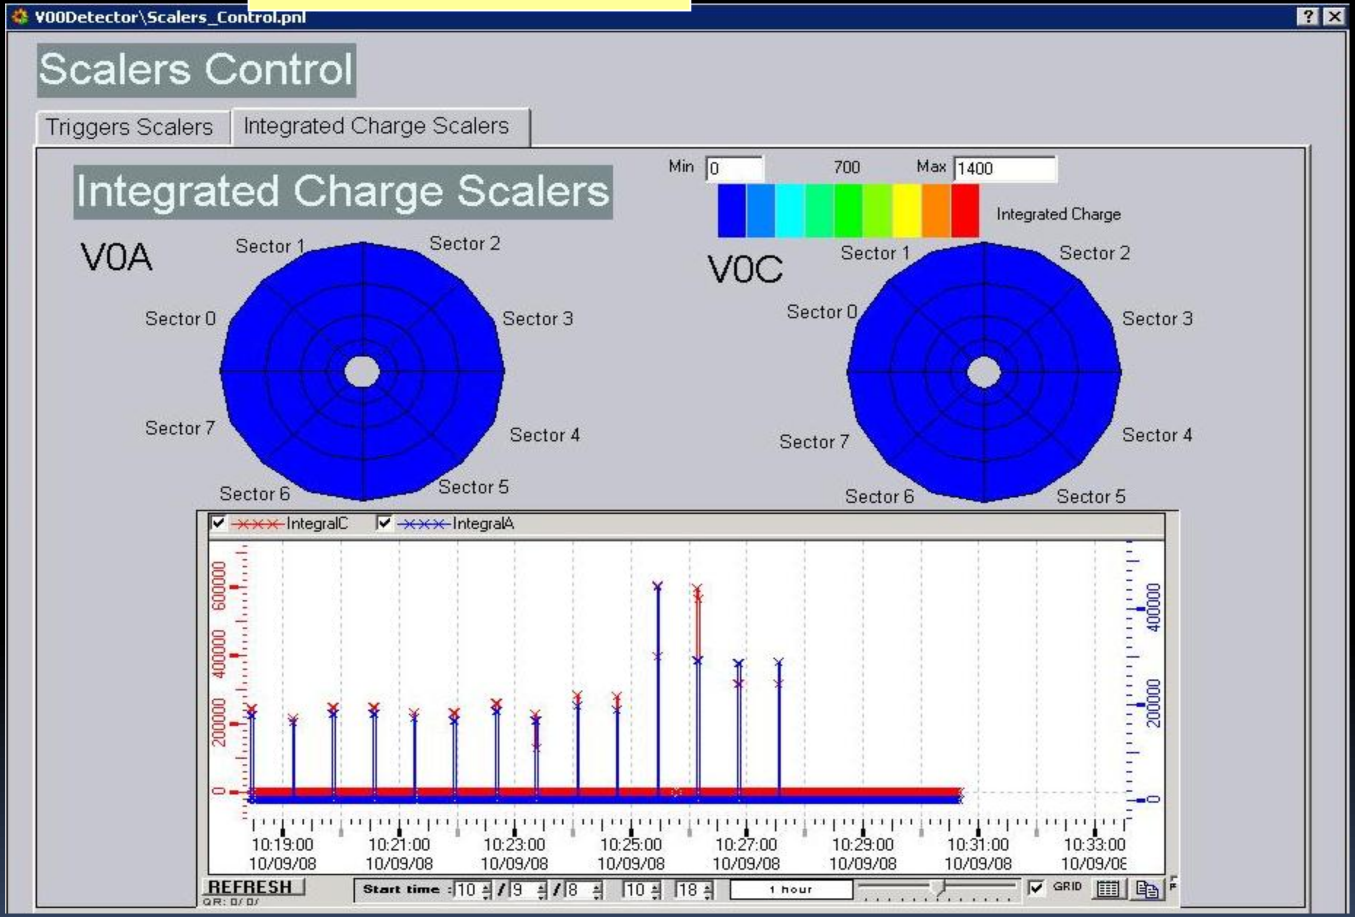

First ion collisions

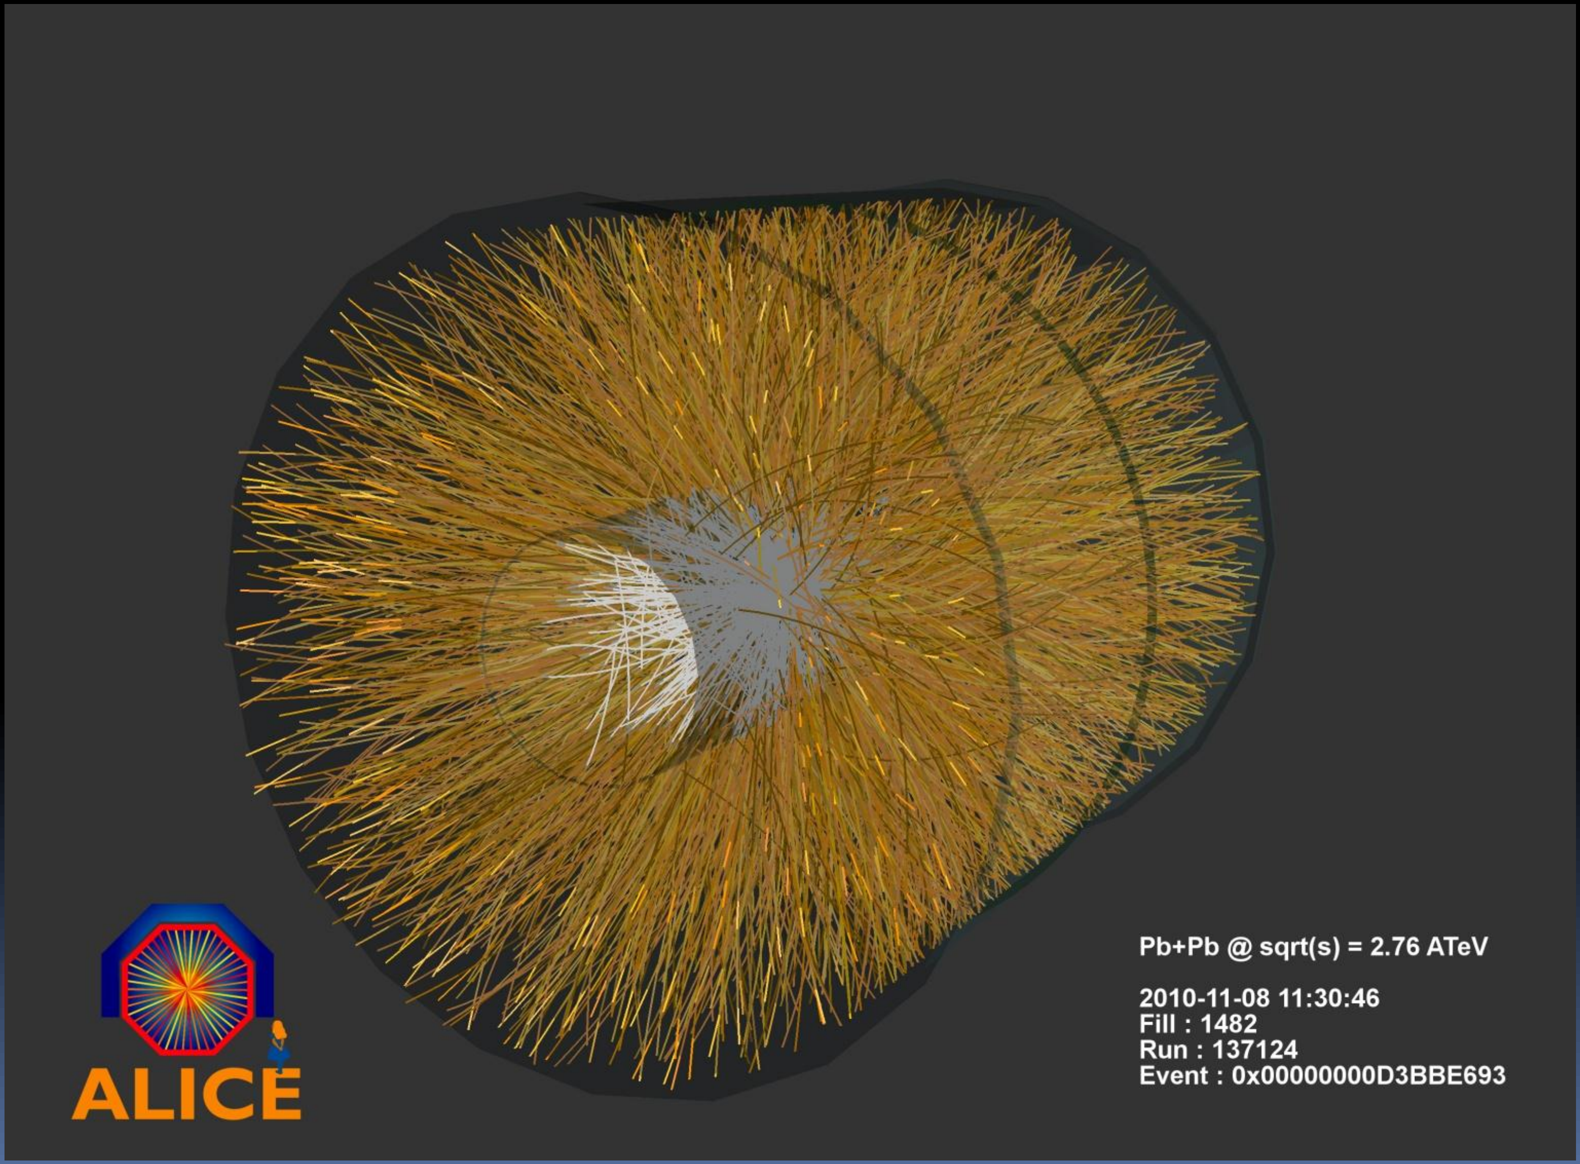

What do we do?

The particle collisions allow us to recreate conditions which existed very short (μs) after the Big Bang

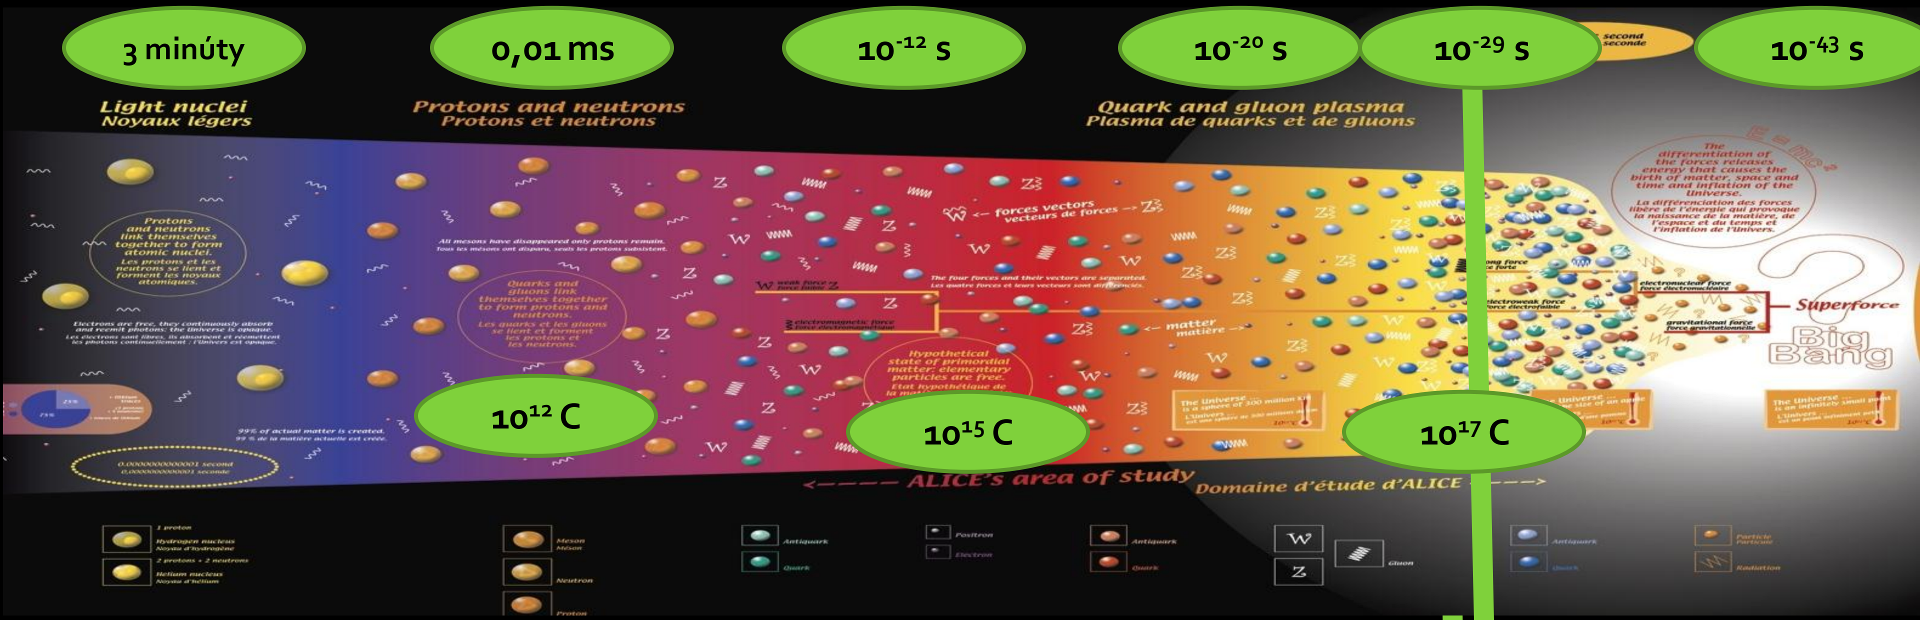

- CERN is trying to answer many questions
- Where does the mass come from?
  - How does the matter behave at temperatures higher than in the middle of the Sun?
    - Why is the mass of protons and neutrons 100times higher than the mass of contained quarks?
      - Can we release quarks from protons and neutrons?
        - Why is the mass of the Universe much higher that we can explain?
          - Where did all the antimatter go?
            - ......??

- ALICE is primary interested in ion collisions
  - Focus on last weeks of LHC operation in 2011 (Pb-Pb collisions)

- During the year ALICE is being improved
- In parallel, ALICE participates in p-p programme
- So far, in 2011 ALICE delivered:
  - 1000 hours of stable physics data taking
  - 2.0 10^9 events collected
  - 2.1 PB of data
  - 5300 hours of stable cosmics datataking, calibration and technical runs
  - 1.7 10^{10} events
  - 3.5 PB of data
    - IONS STILL TO COME IN 2011!

- Where is the link to cyber security?

- The same people who built and exploit ALICE are also in charge of its operation
  - In this talk we focus only at part of the story, the Detector Control System (DCS)

- The ALICE Detector Control System (DCS)

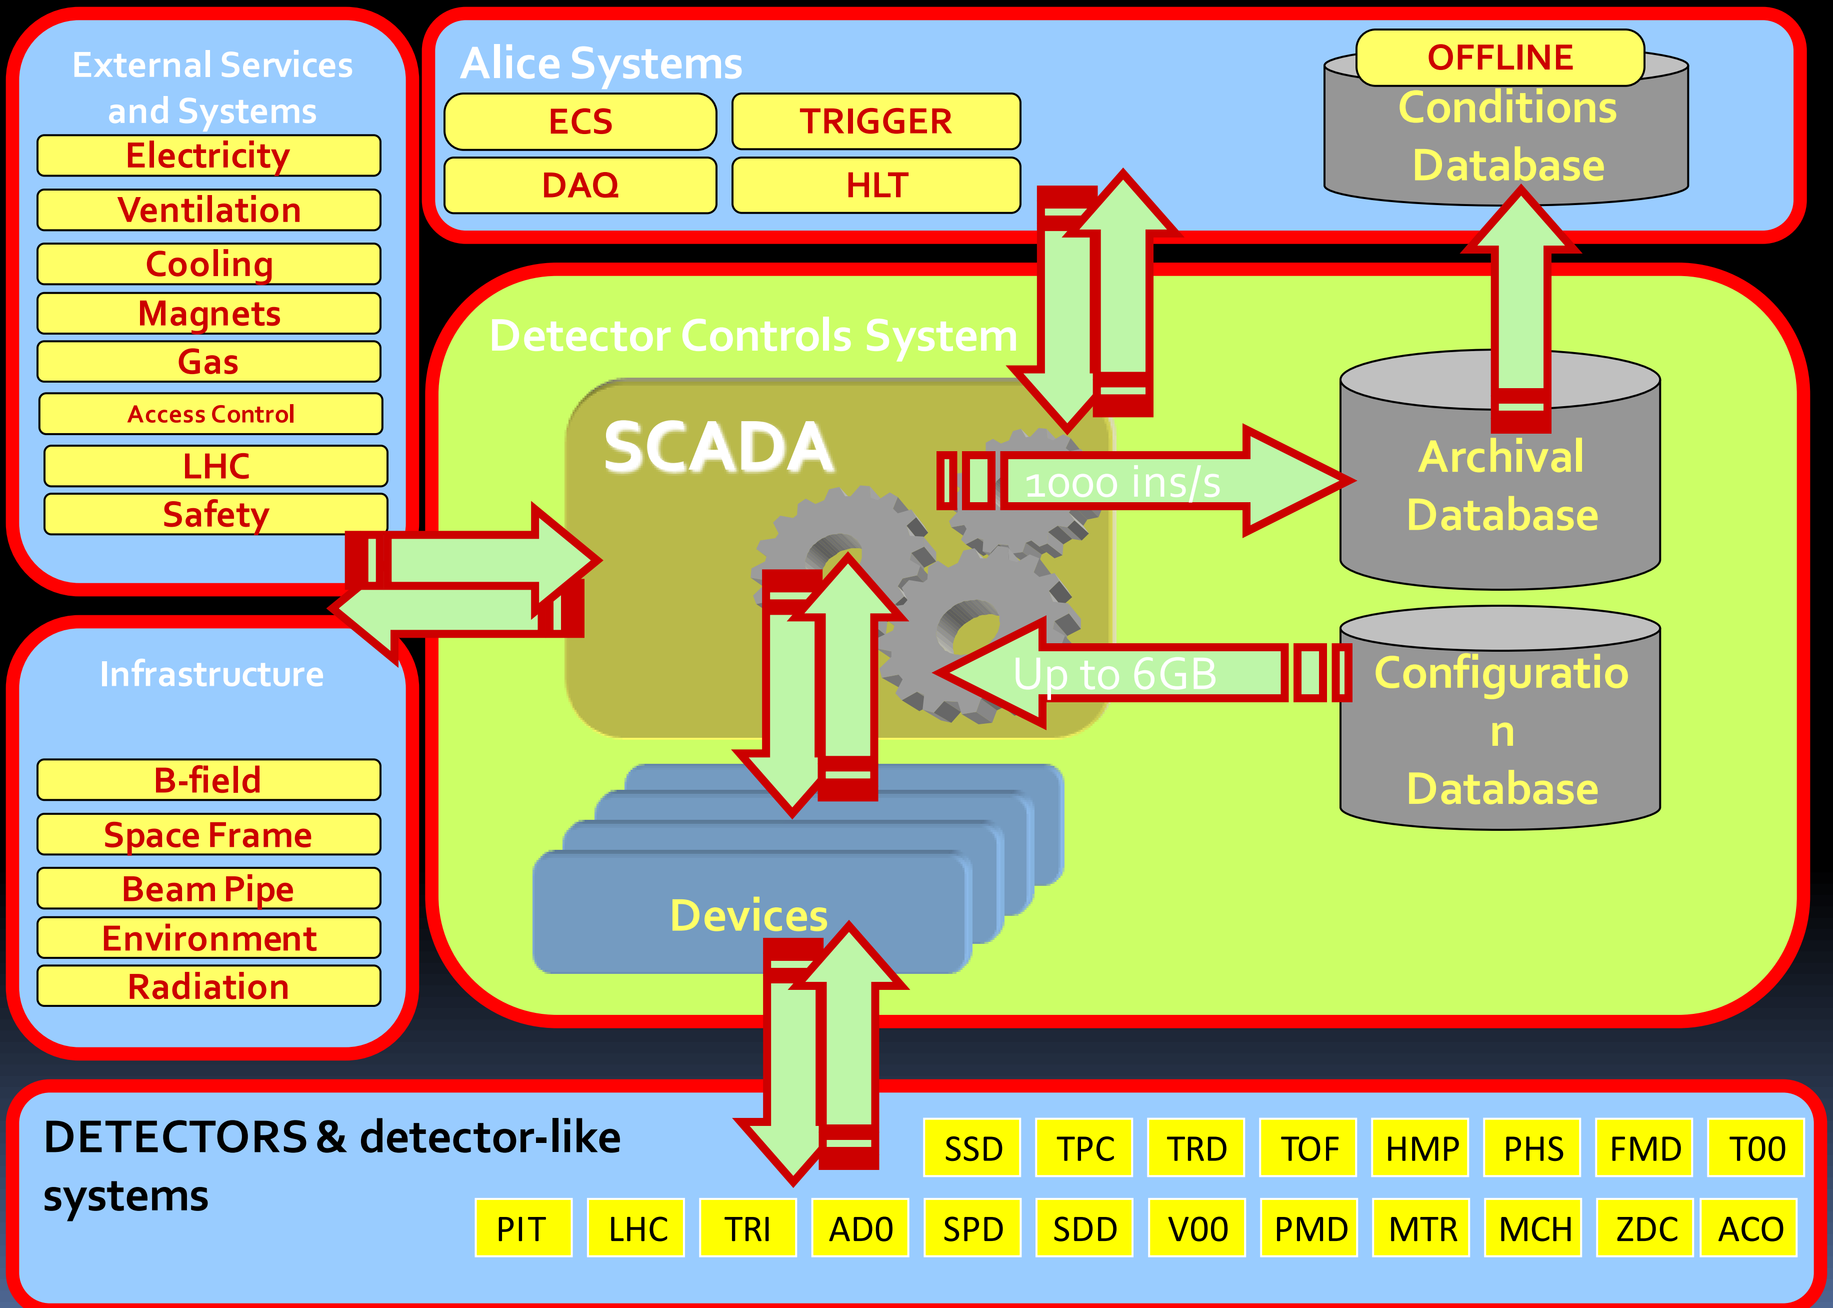
External Services and Systems
Electricity
Ventilation
Cooling
Magnets
Gas
Access Control
LHC
Safety
Infrastructure
B-field
Space Frame
Beam Pipe
Environment
Radiation
Alice Systems
ECS
TRIGGER
DAQ
HLT
OFFLINE
Conditions Database
Detector Controls System
SCADA
1000 ins/s
Archival Database
Up to 6GB
Configuratio n Database
Devices
DETECTORS & detector-like systems
SSD
TPC
TRD
TOF
HMP
PHS
FMD
T00
PIT
LHC
TRI
AD0
SPD
SDD
V00
PMD
MTR
MCH
ZDC
ACO

- Dataflow in ALICE DCS
  - 6GB of data is needed to fully configure ALICE for operation
  - Several stages of filtering applied to acquired data

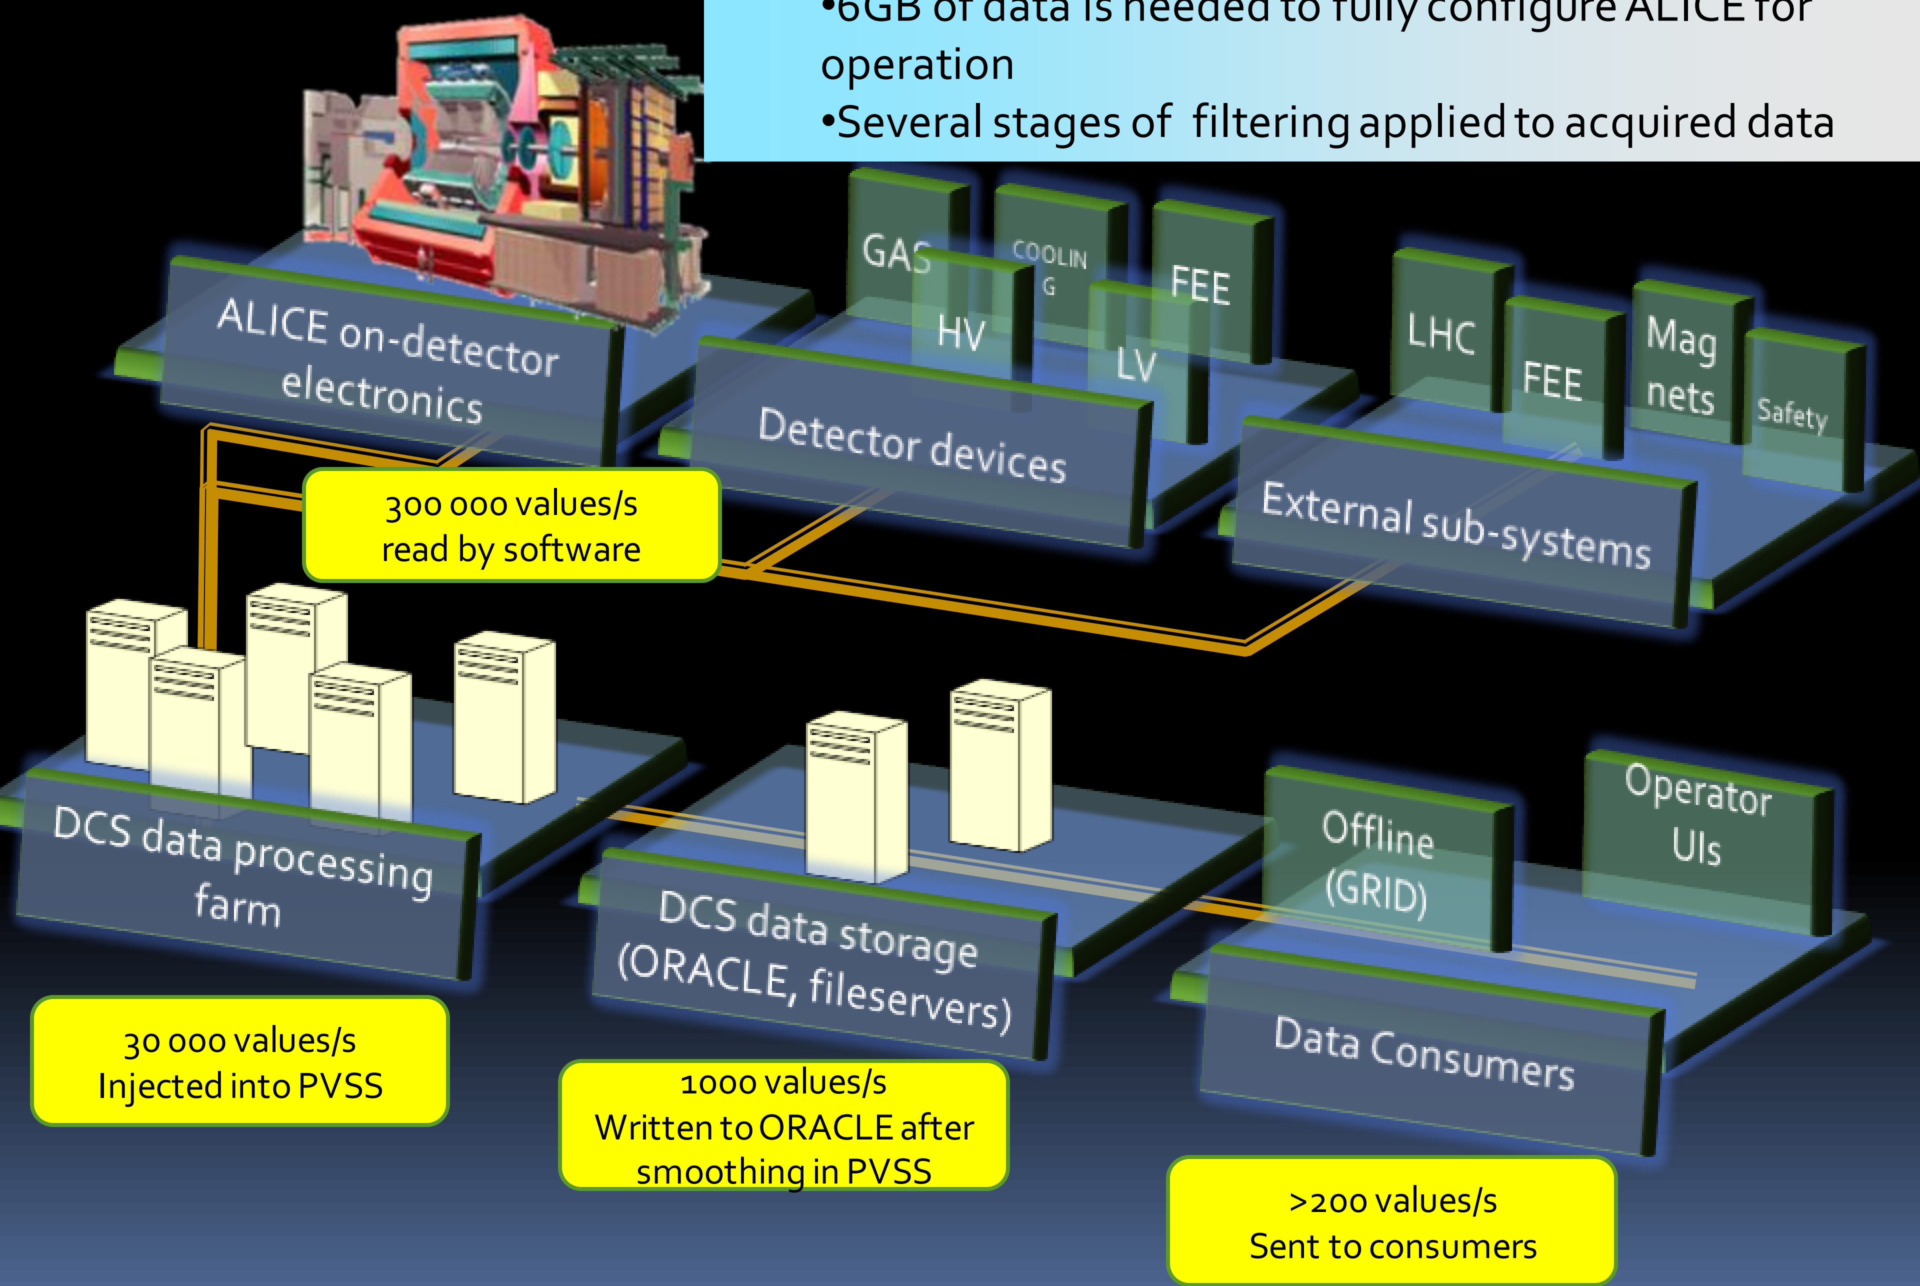

Building blocks of ALICE DCS

1200 network-attached devices
270 crates (VME and power supplies)
4 000 controlled voltage channels

- 18 detectors with different requirements
  - Effort to device standardization
    - Still large diversity mainly in FEE part
      - Large number of busses (CANbus, JTAG, Profibus, RS232, Ethernet, custom links…)

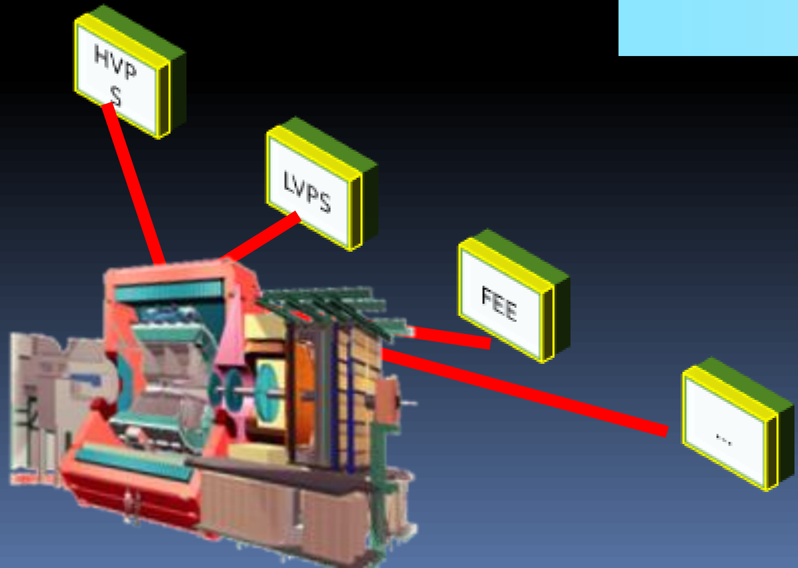

180 000 OPC items
100 000 Front-End (FED) services
1 000 000 parameters supervised by the DCS
Monitored at typical rate of 1Hz

- Hardware diversity is managed through standard interfaces
  - OPC servers for commercial devices
  - FED servers for custom hardware
    - Provides hardware abstraction, uses CERN DIM (TCP/IP based) protocol for communication

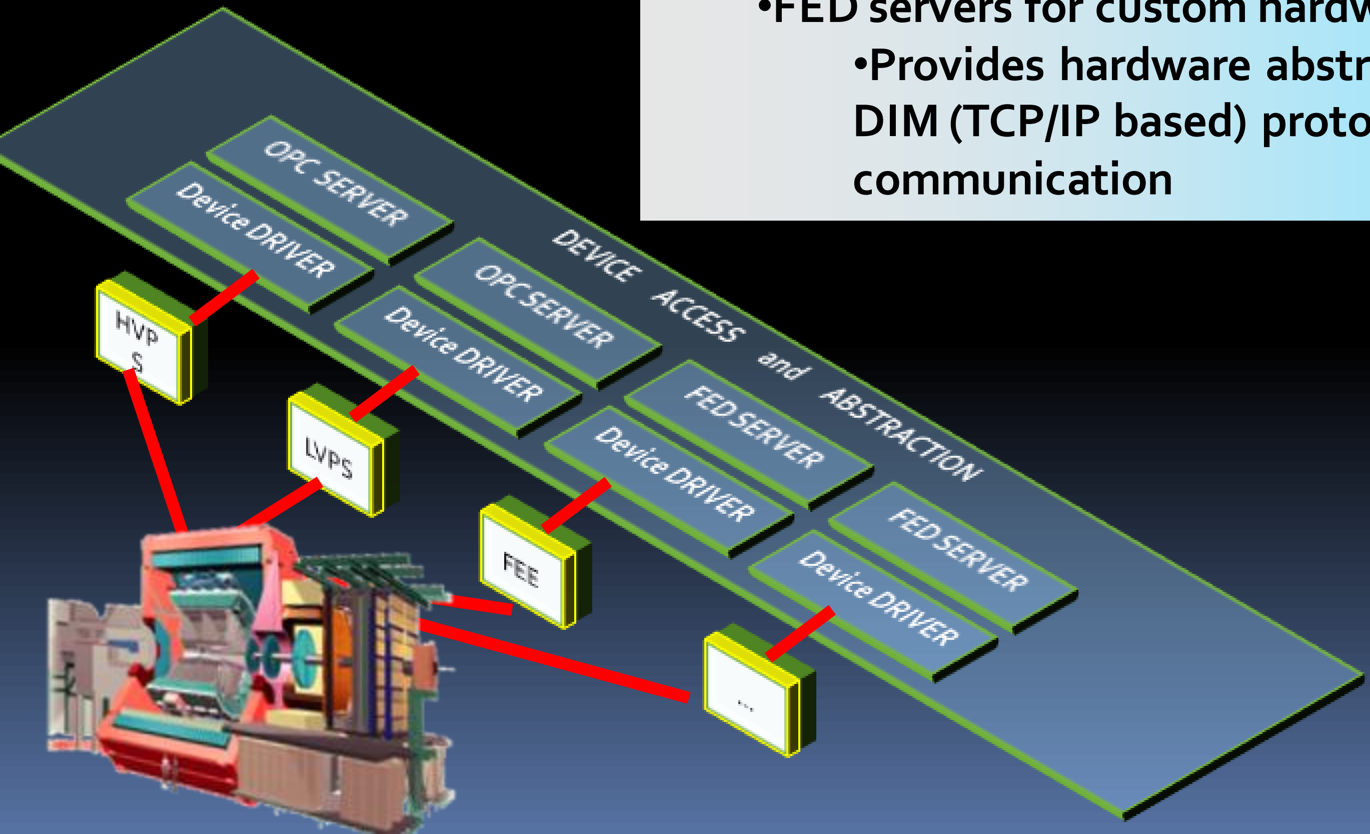

110 detector computers
60 backend servers
DCS Oracle RAC (able to process up to 150 000 inserts/s)

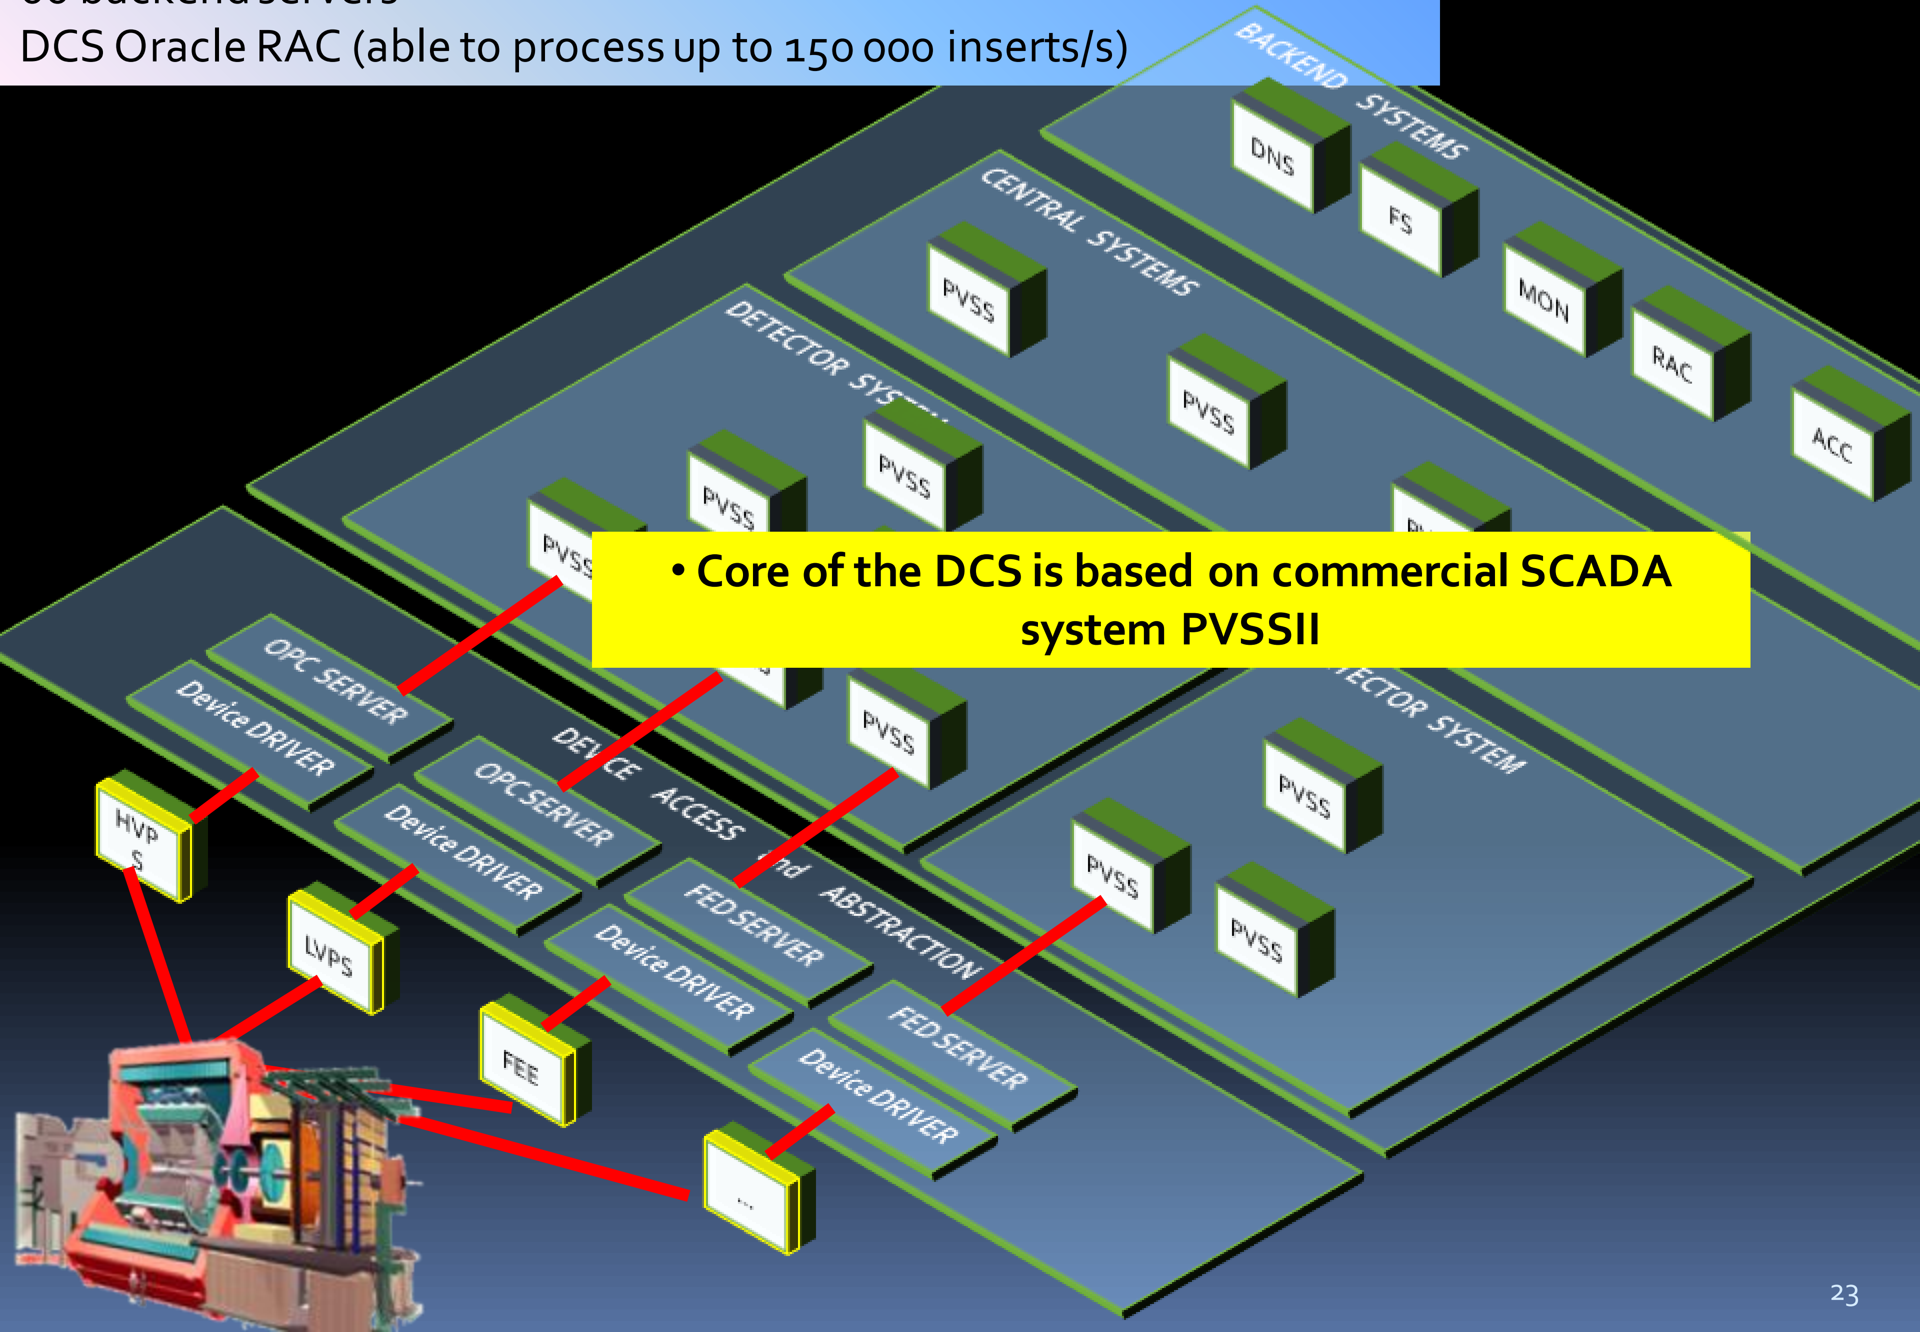

- **Core of the DCS is based on commercial SCADA system PVSSII**

PVSSII Architecture

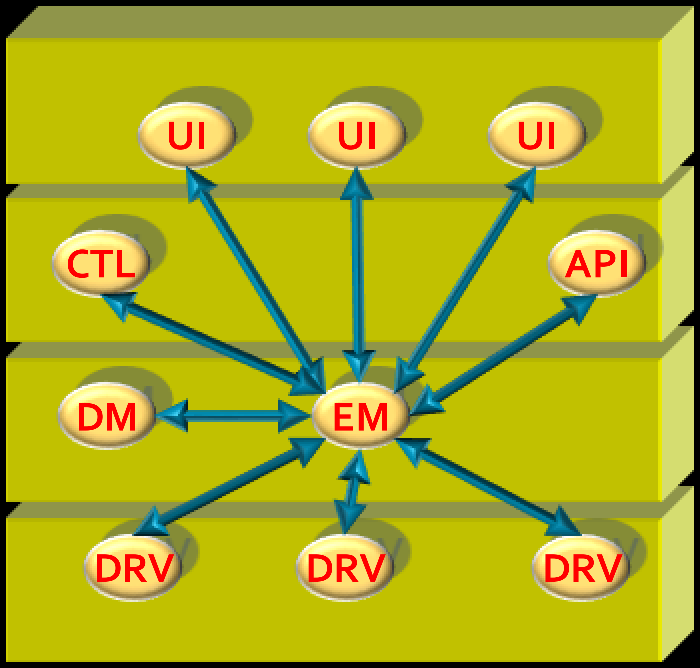

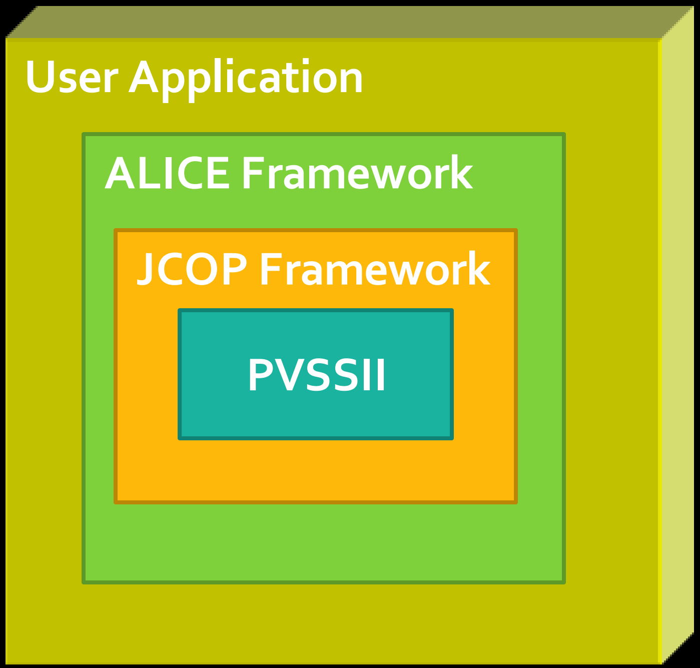

- PVSSII system is composed of specialized program modules (managers)
- Managers communicate via TCP/IP
- ALICE DCS is built from 100 PVSS systems composed of 900 managers
- PVSSII is extended by JCOP and ALICE frameworks on top of which User applications are built

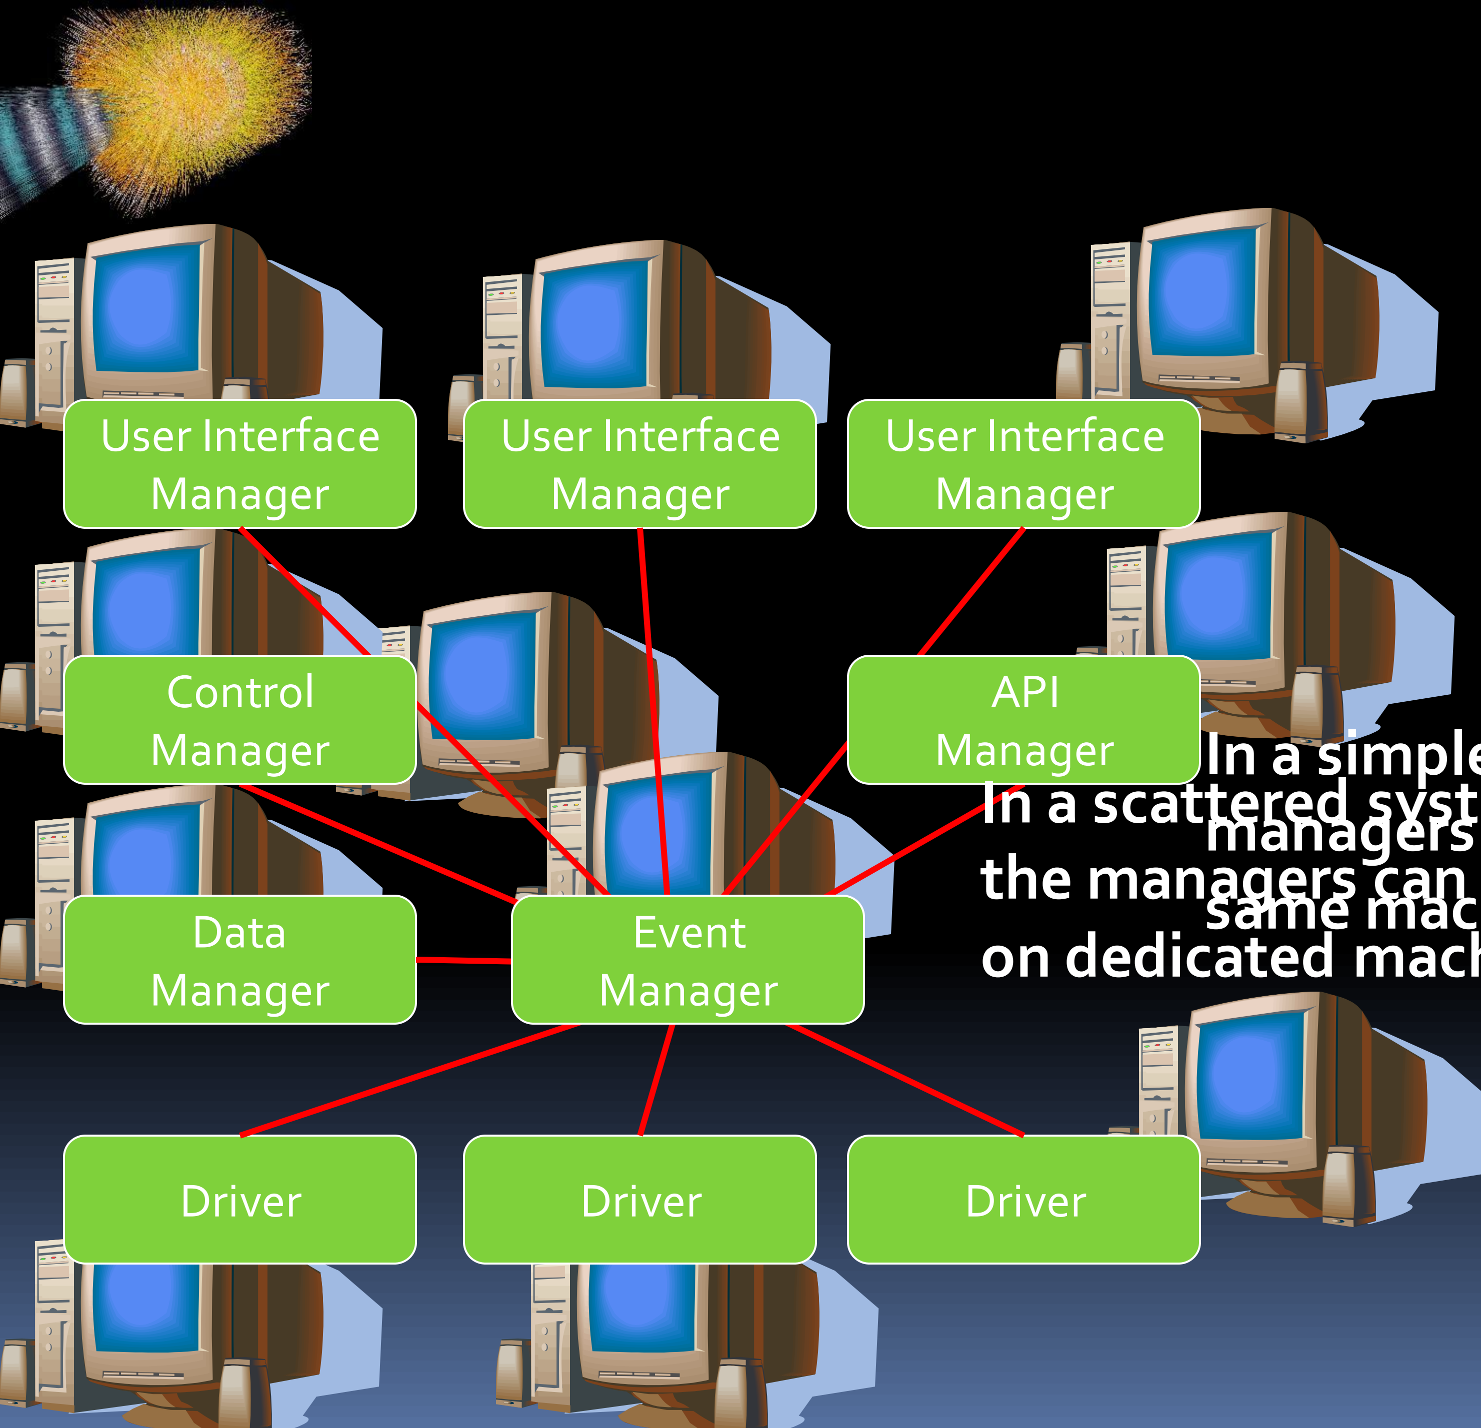

In a simple system all managers run on the same machine

In a scattered system, the managers can run on dedicated machines

In a distributed system several PVSSII systems (simple or scatered) are interconnected

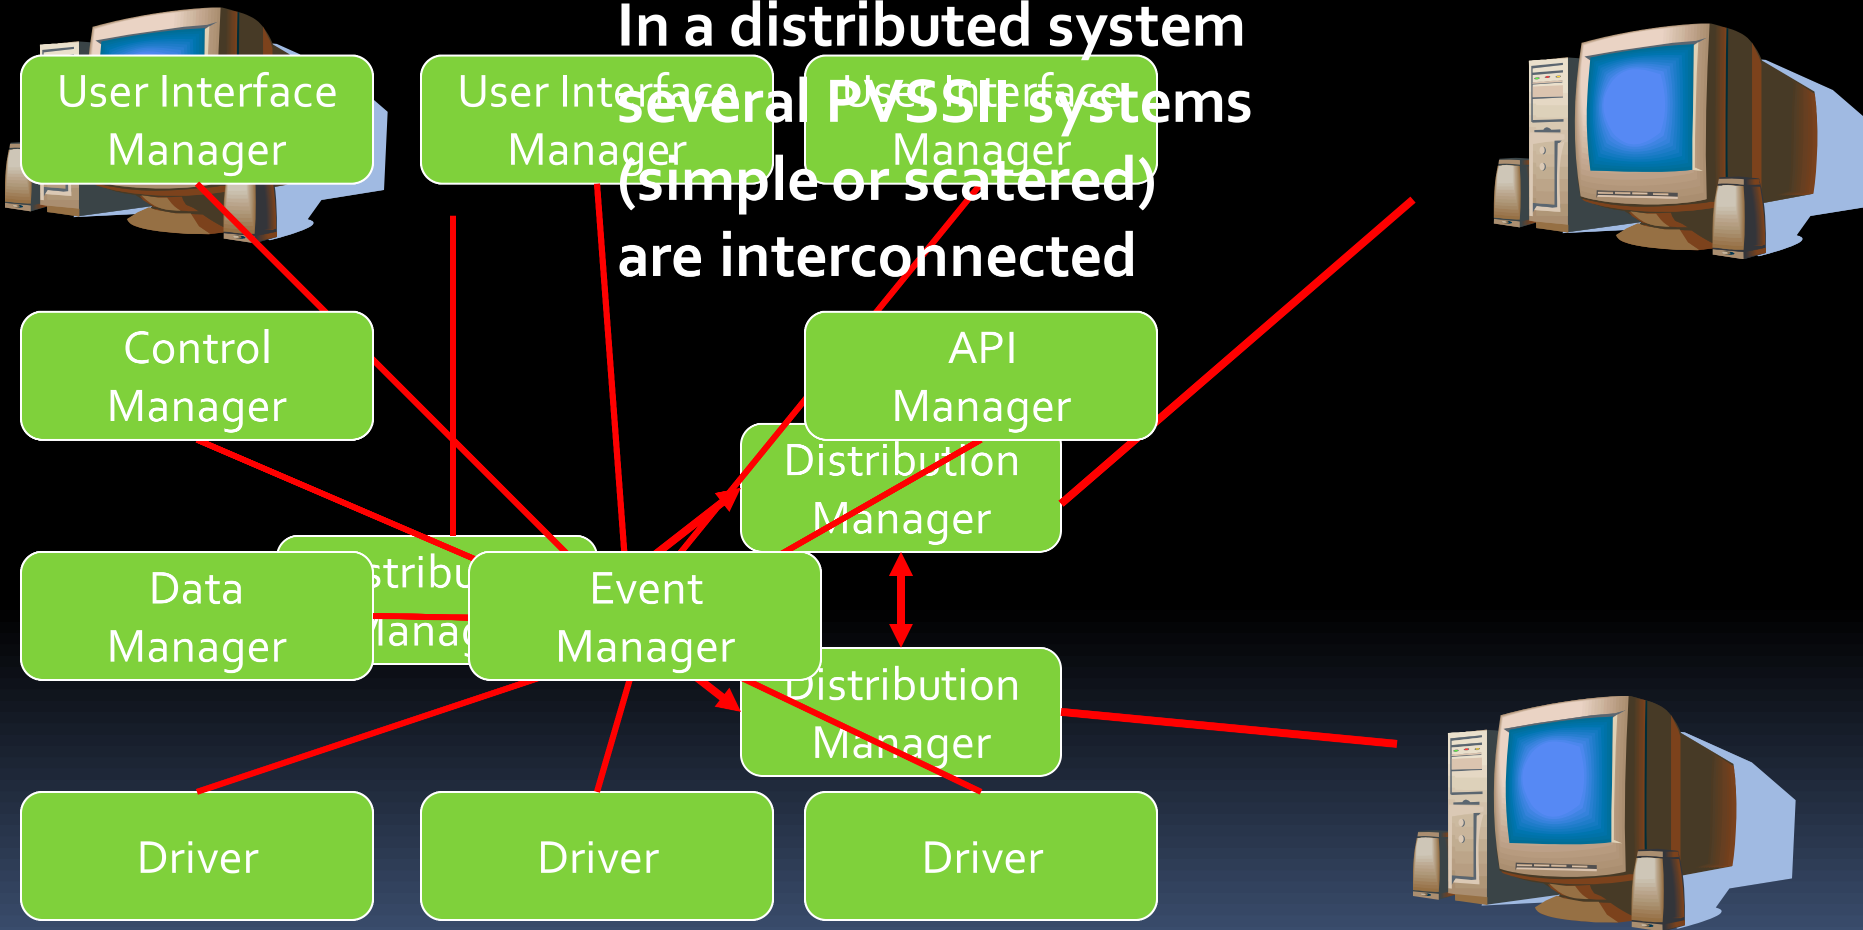

Each detector DCS is built as a distributed PVSSII system

- **Mesh, no hierarchical topology**
- **Detector specific**

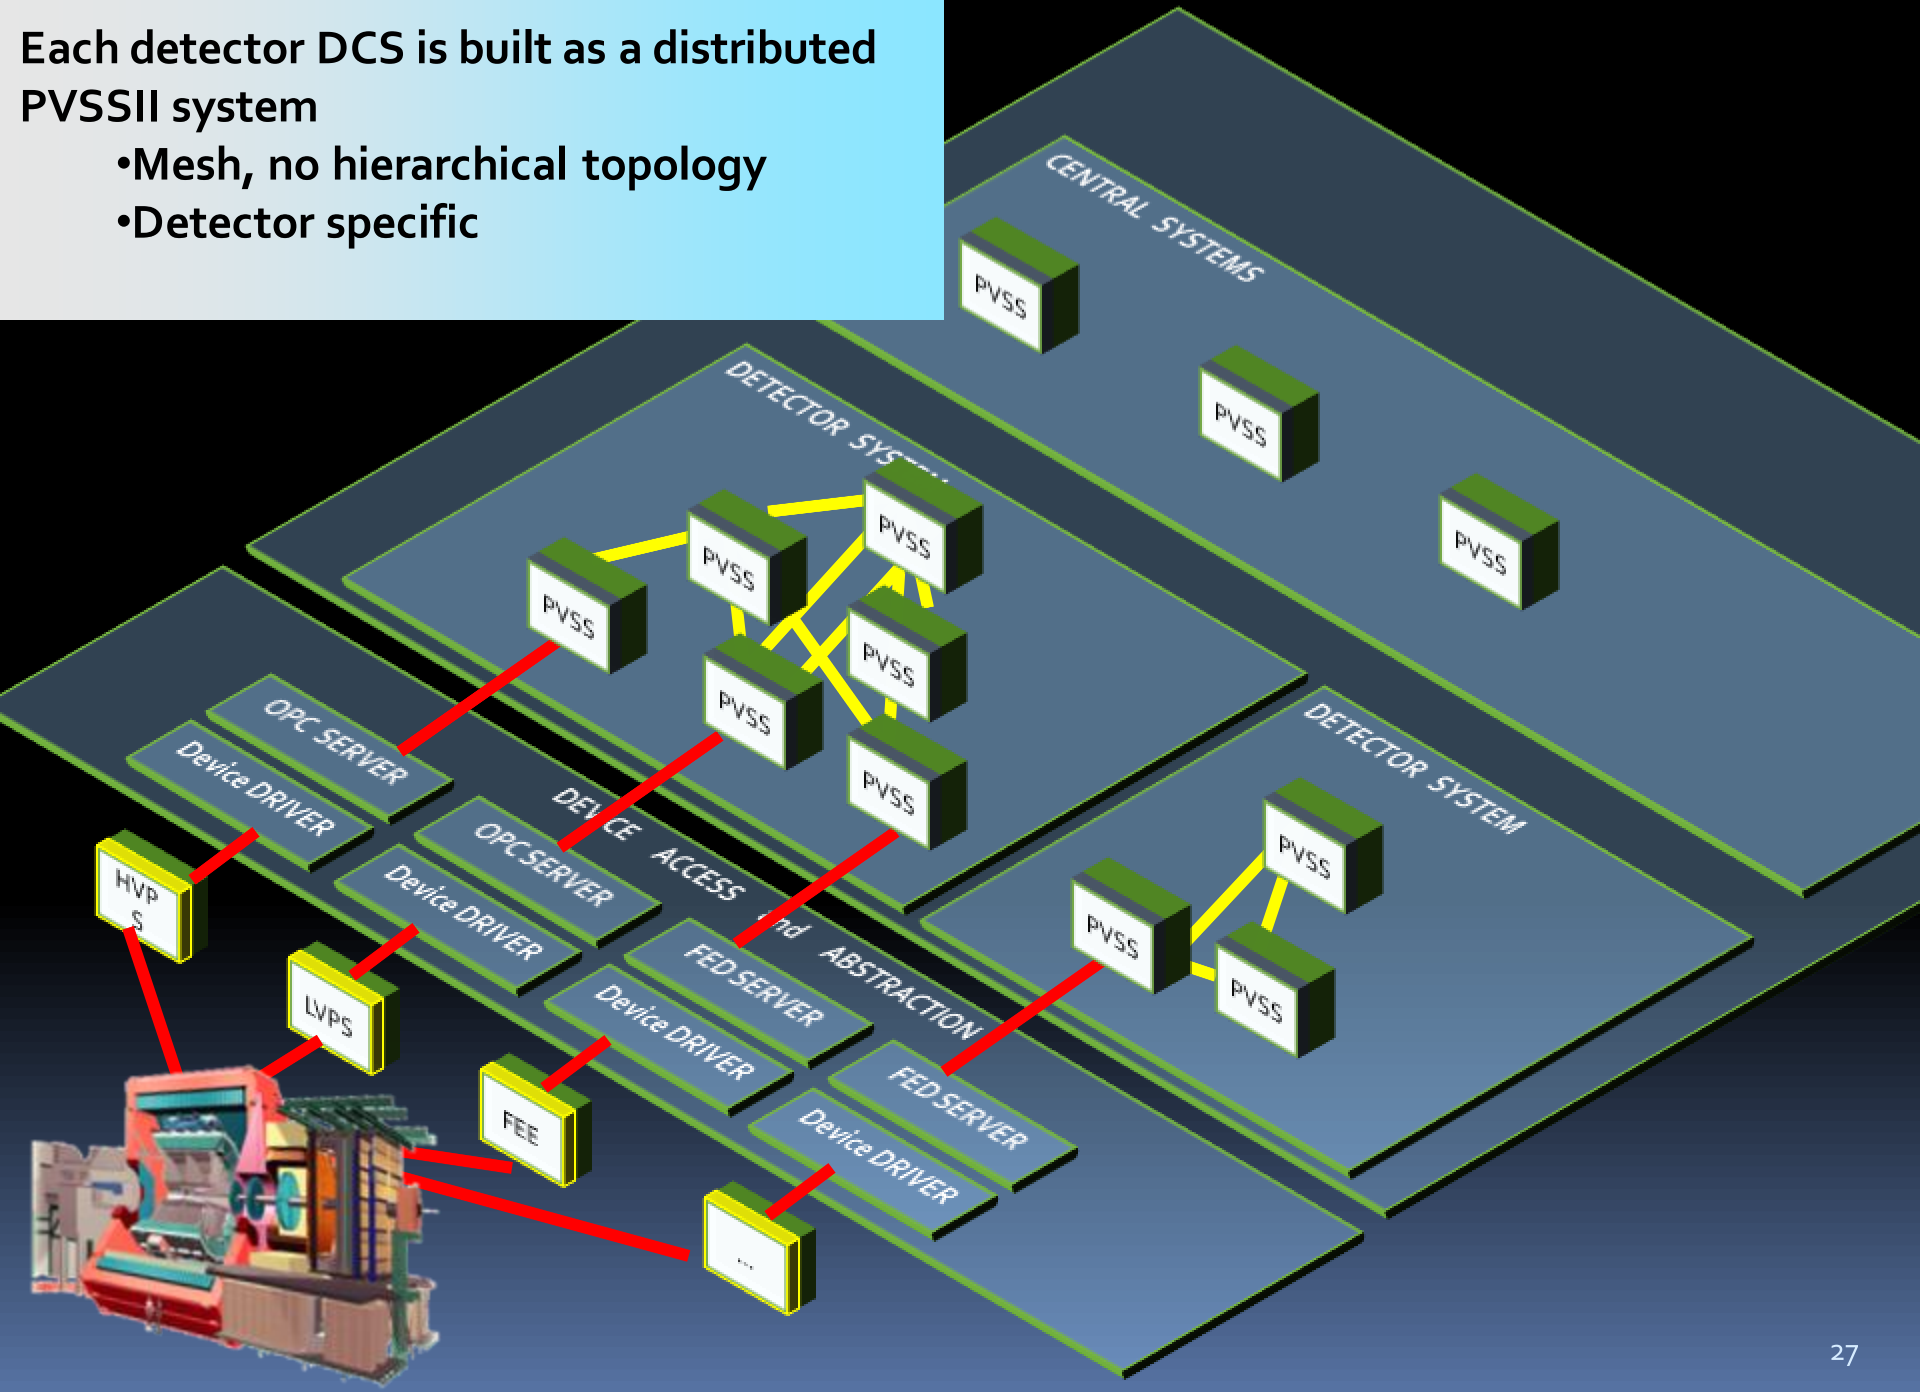

ALICE DCS is built as a distributed system of detector systems

Central servers connect to ALL detector systems

- global data exchange
- synchronization
- Monitoring…

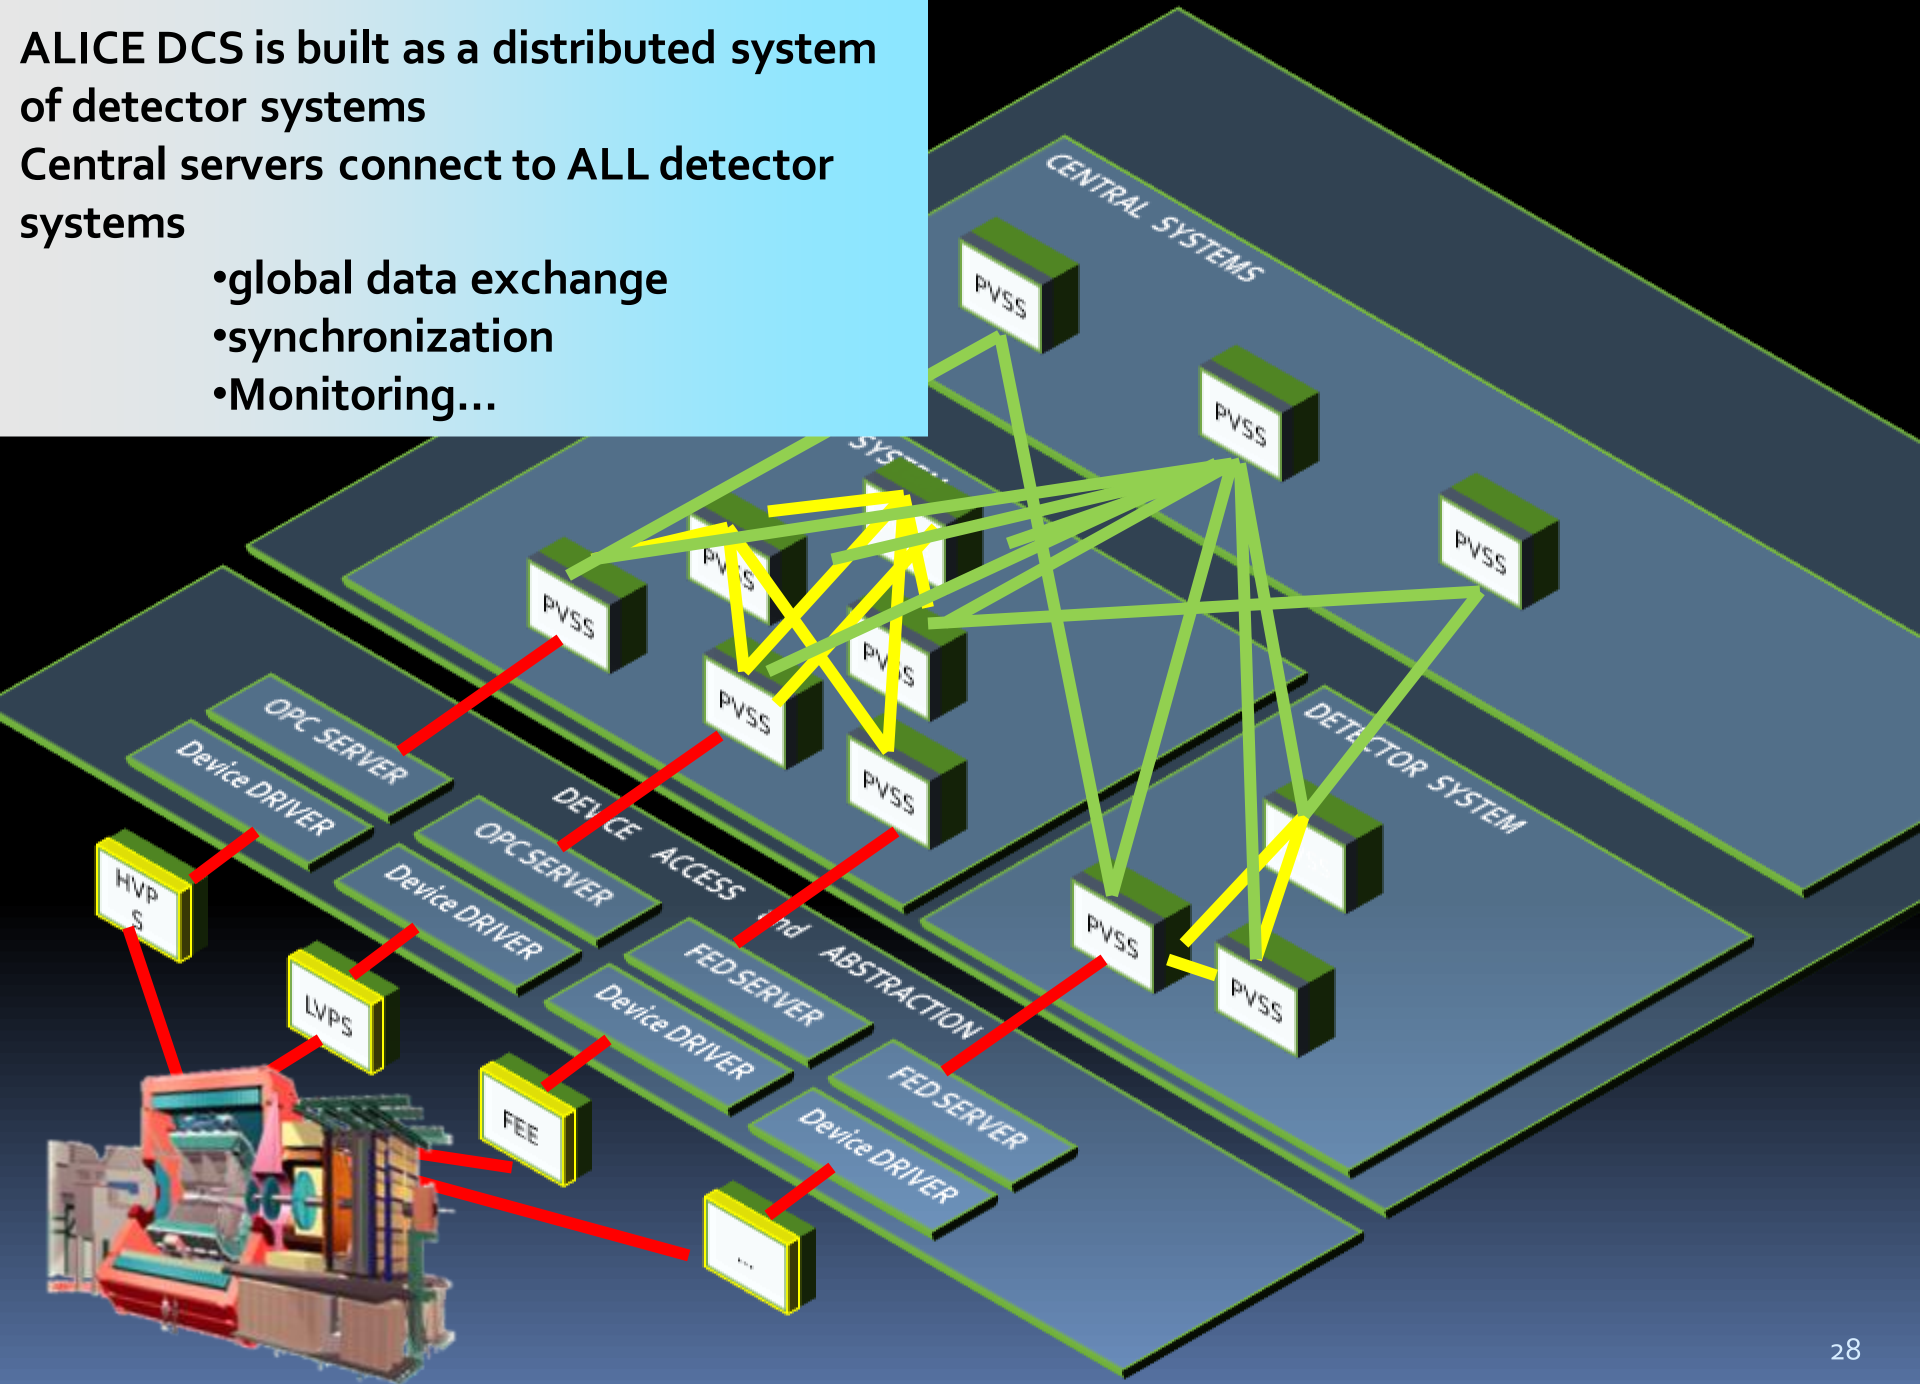

• PVSSII distributed system is not a natural system representation for the operator
•ALICE DCS Is modeled as a FSM using CERN SMI++ tools
  - •Hide experiment complexity
  - •Focus on operational aspect

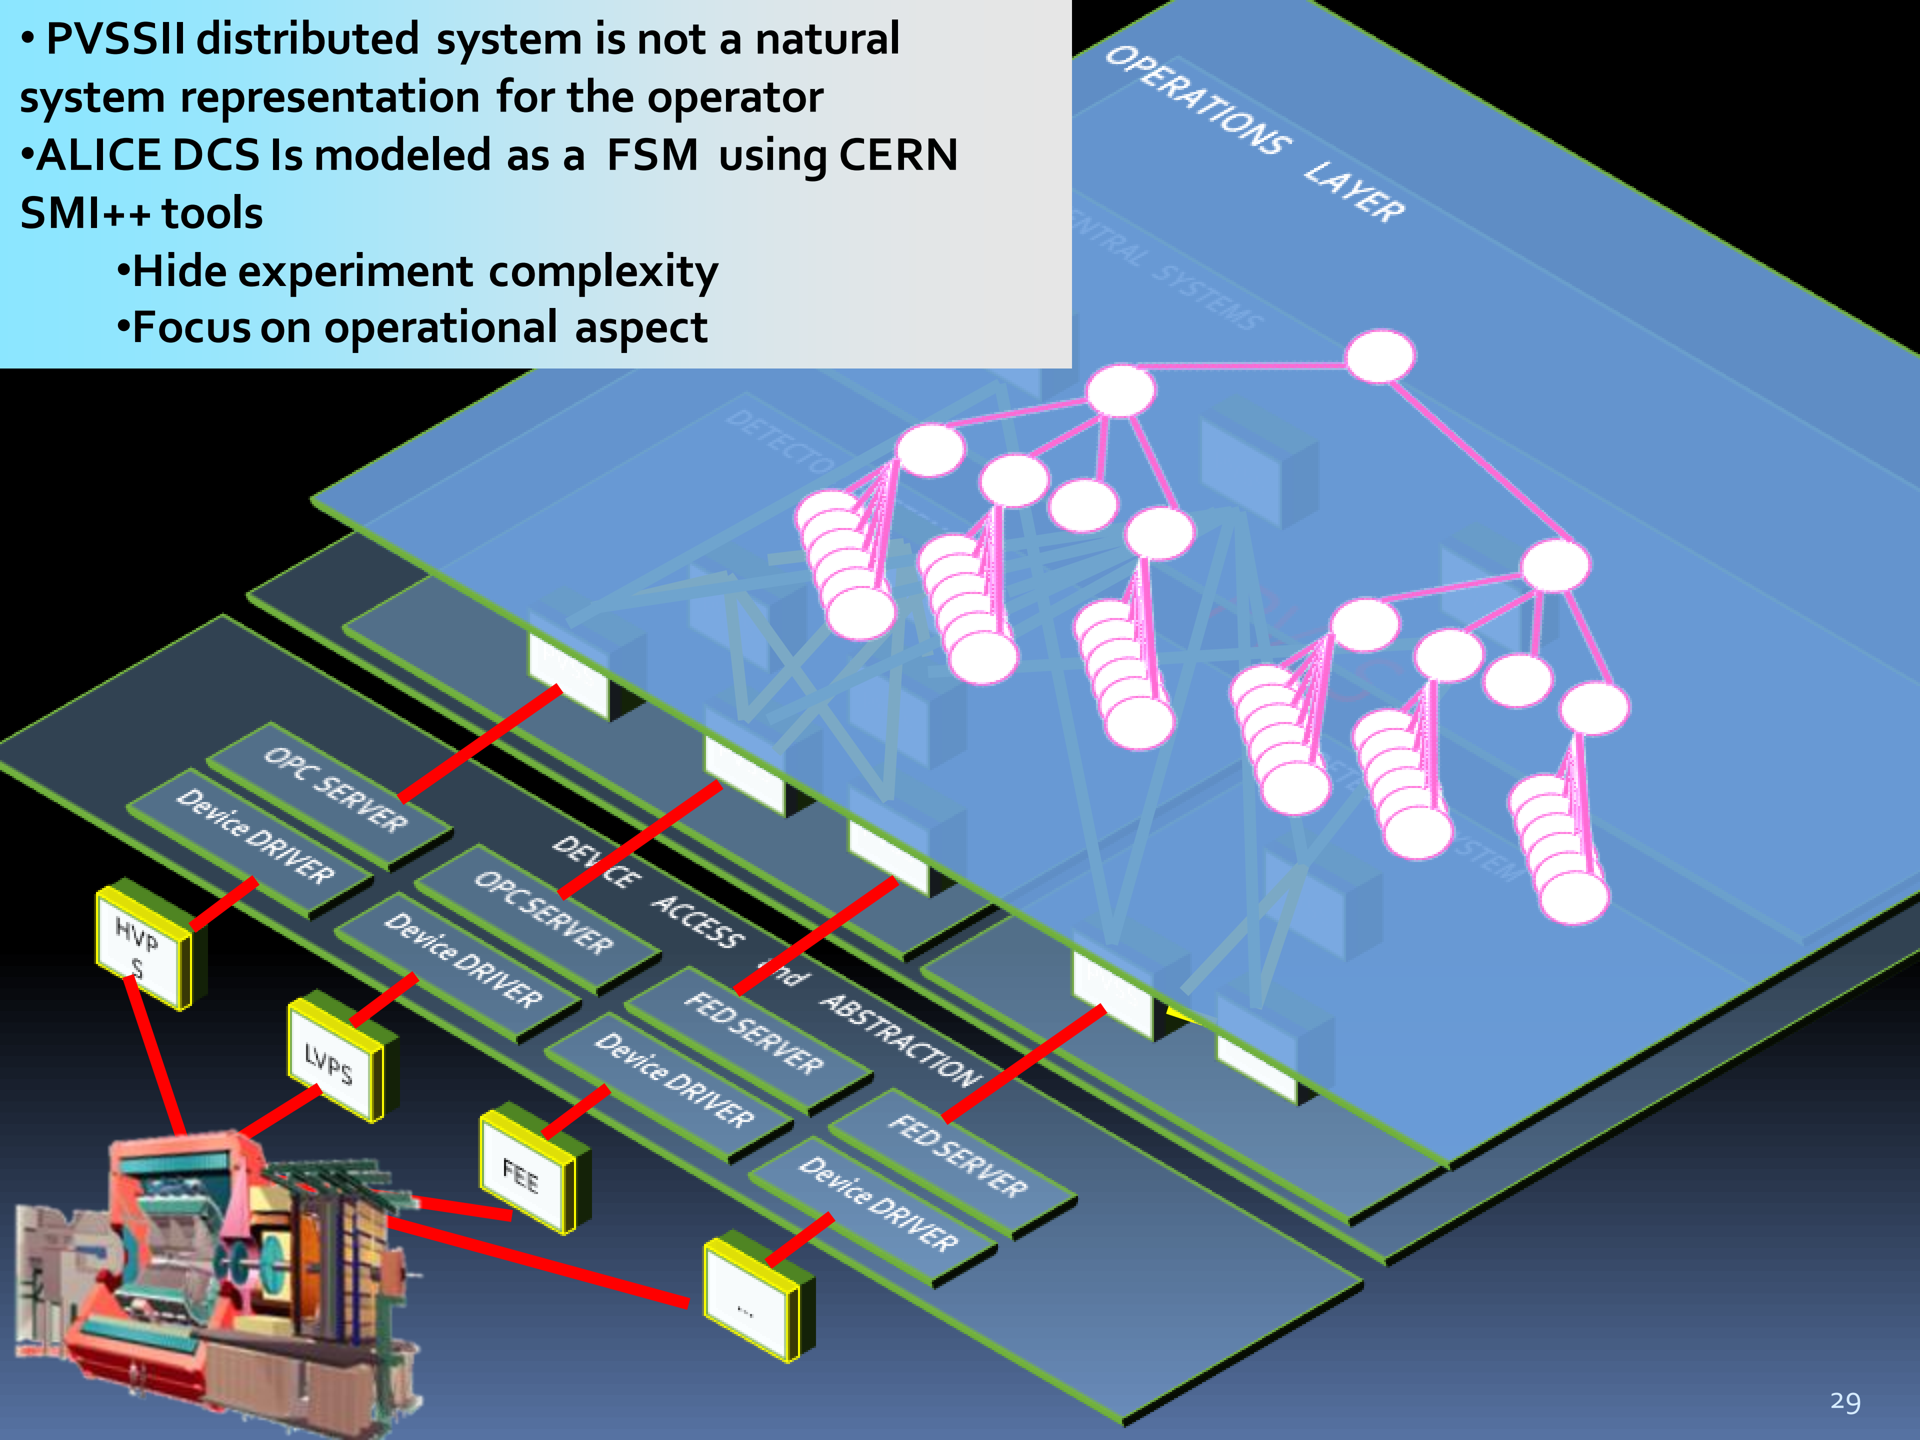

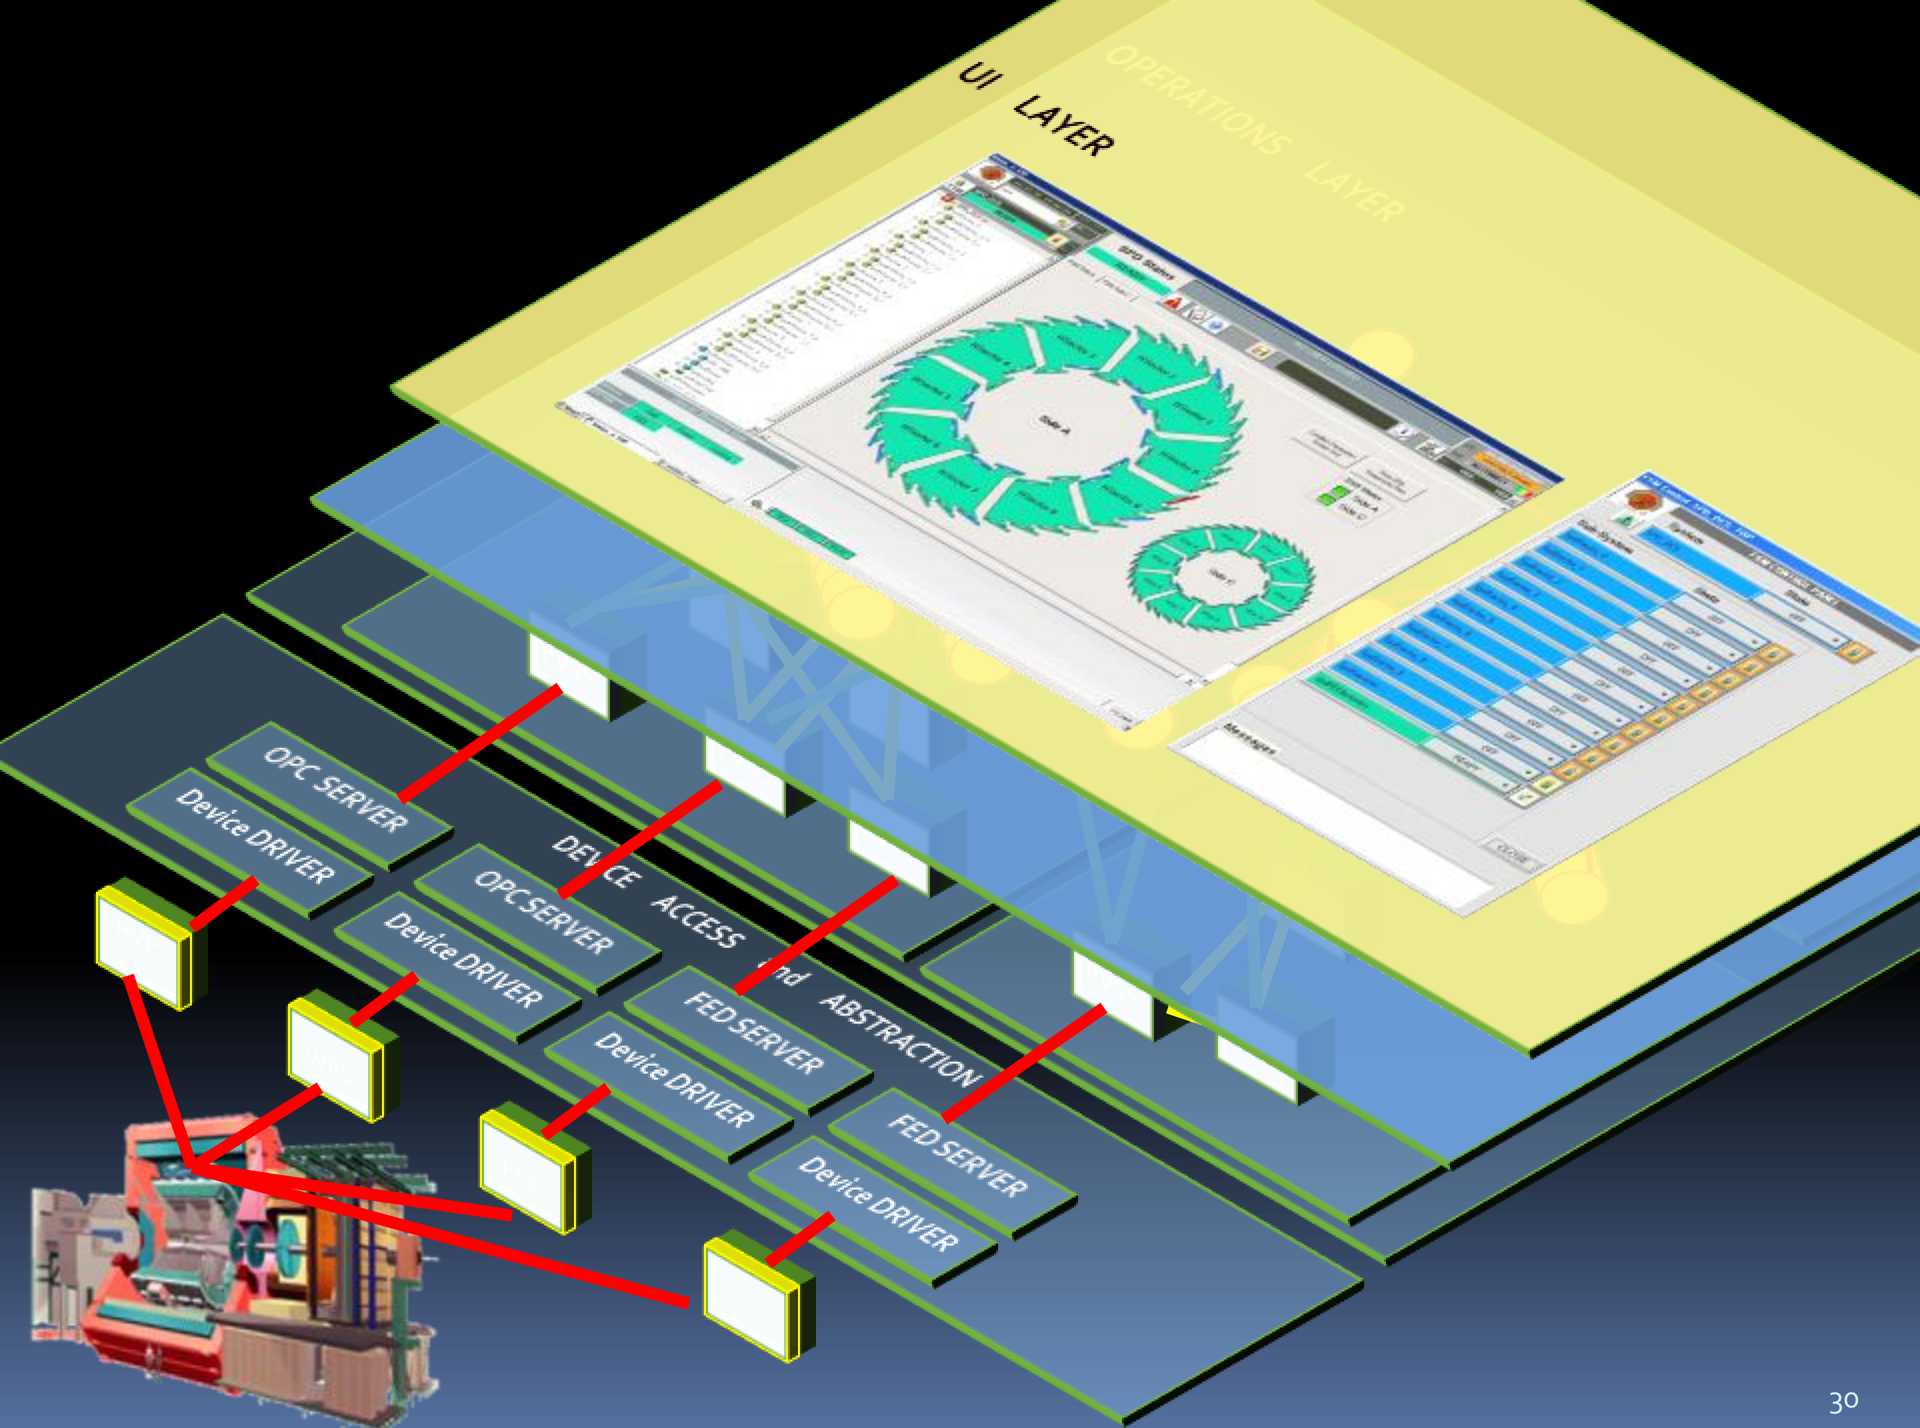
UI LAYER
OPERATIONS LAYER
DEVICE ACCESS and ABSTRACTION
OPC SERVER
Device DRIVER
OPC SERVER
Device DRIVER
FED SERVER
Device DRIVER
FED SERVER
Device DRIVER

DCS Computing model

- Two categories of DCS computers:
  - Worker nodes – executing the controls tasks and running detector specific software
  - Operator node – used by operators to interact with the system

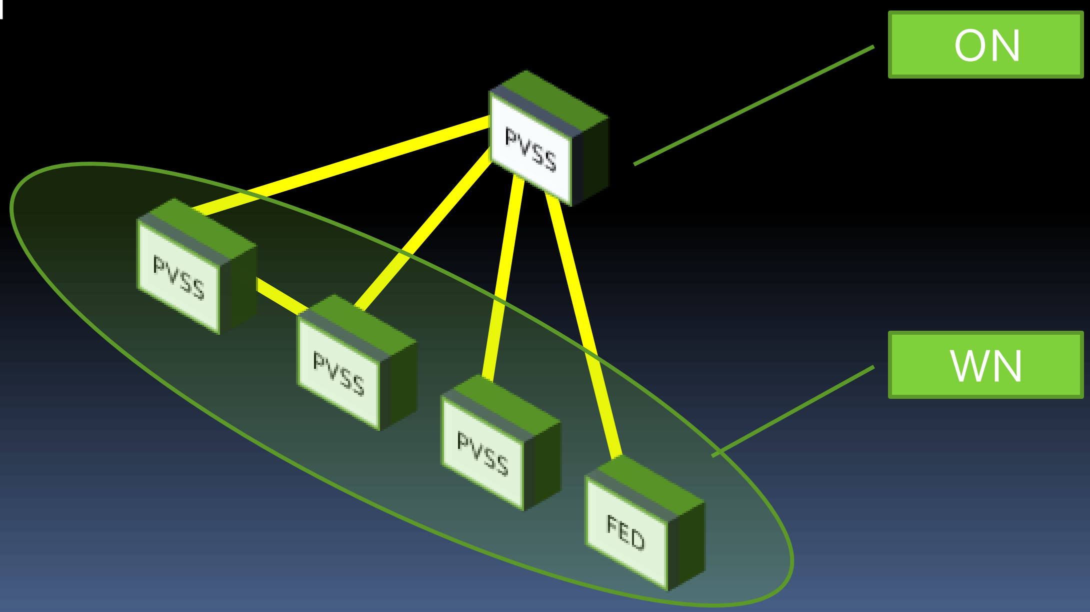

- ALICE network architecture

Remote access to the DCS network

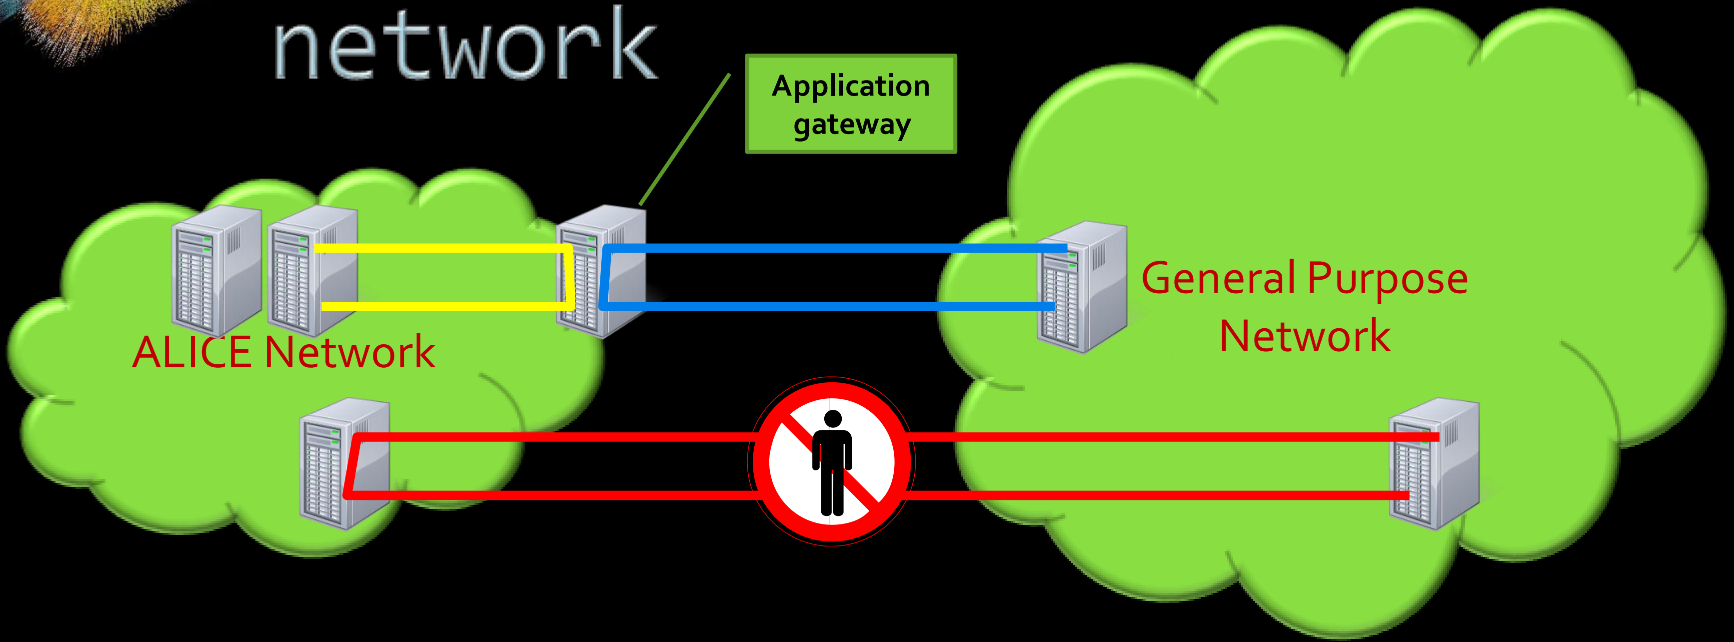

- No direct user access to the ALICE network
- Remote access to ALICE network Is possible via the application gateways
  - User makes RDP connection to the gateway
  - From the gateway further connection is granted to the network

- ALICE host exposed to NetA:
- Can see all NetA and ALICE hosts
- Can be seen by all NetA hosts

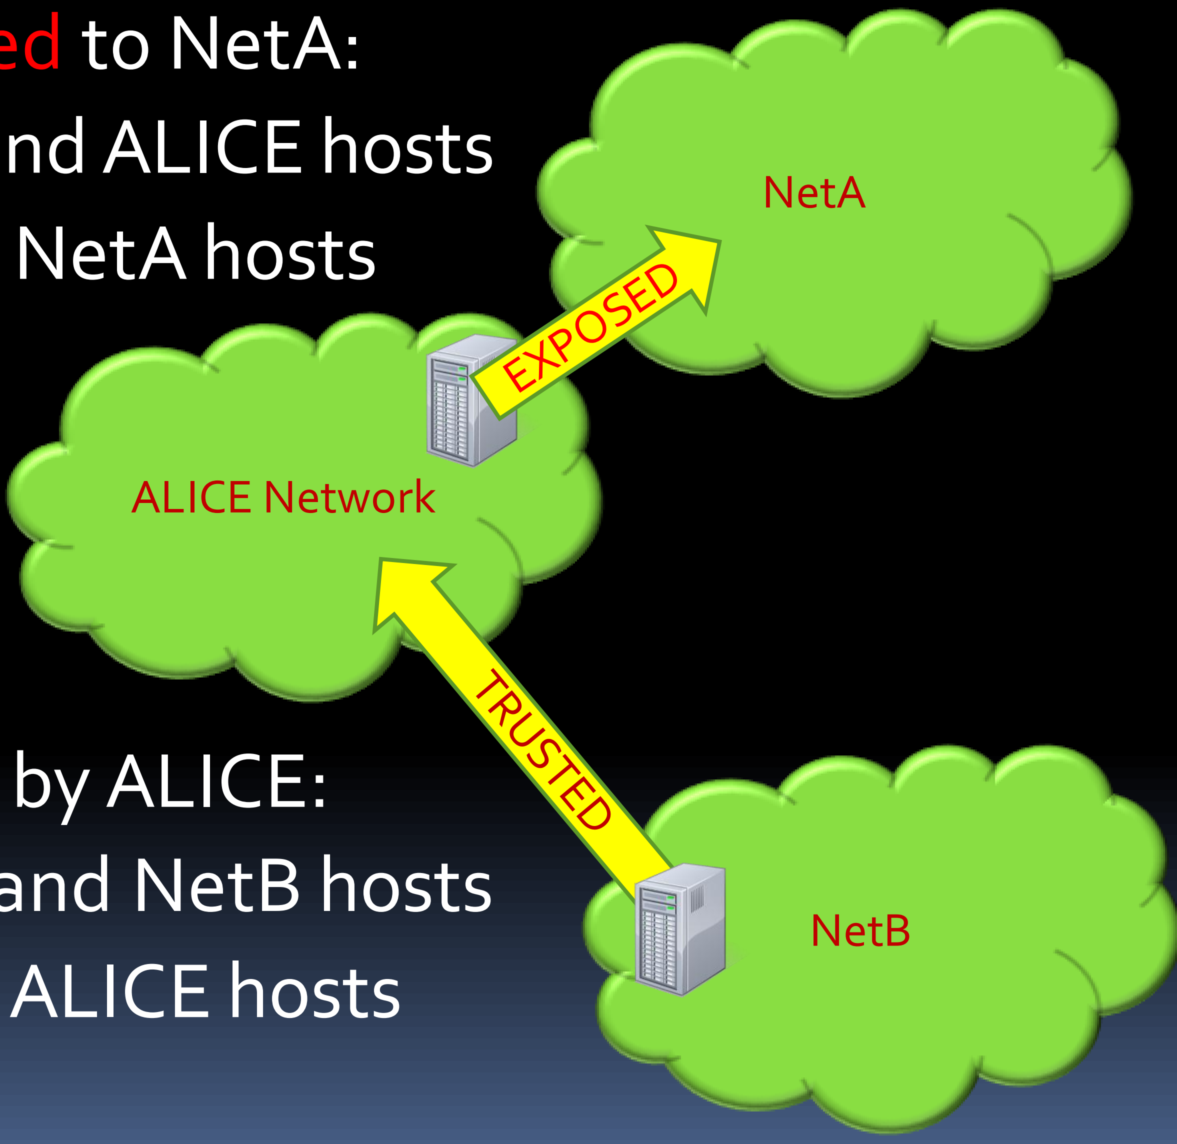

- NetB host trusted by ALICE:
- Can see all ALICE and NetB hosts
- Can be seen by all ALICE hosts

Are we there?

- The simple security cookbook recipe seems to be:
  - Use the described network isolation
  - Implement secure remote access
  - Add firewalls and antivirus
  - Restrict the number of remote users to absolute minimum
  - Control the installed software and keep the systems up to date
- Are we there?
  - No, this is the point, where the story starts to be interesting

Remote access

- Why would we need to access systems remotely?

- ALICE is still under construction, but experts are based in the collaborating institutes
  - Detector groups need DCS to develop the detectors directly in situ
    - There are no test benches with realistic systems in the institutes, the scale matters
- ALICE takes physics and calibration data
  - On-call service and maintenance for detector systems are provided remotely

The user challenge

- Natural expectation would be that there are few users requiring access to the controls system
- The today's reality is more than 400 authorized accounts…
  - Rotation of experts in the institutes is very frequent
    - Many tasks are carried out by students (graduate or PhD)
  - Commitments to collaboration expect shift coverage
    - Shifters come to CERN to cover 1-2 weeks and then are replaced by colleagues

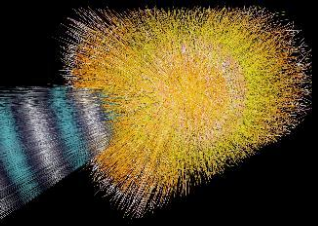

- How do we manage the users?

Authorization and authentication

- User authentication is based on CERN domain credentials
  - No local DCS accounts
  - All users must have CERN account (no external accounts allowed)
- Authorization is managed via groups
  - Operators have rights to logon to operator nodes and use PVSS
  - Experts have access to all computers belonging to their detectors
  - Super experts have access everywhere
- Fine granularity of user privileges can be managed by detectors at the PVSS level
  - Only certain people are for example allowed to manipulate very high voltage system etc.

Operator access to computers

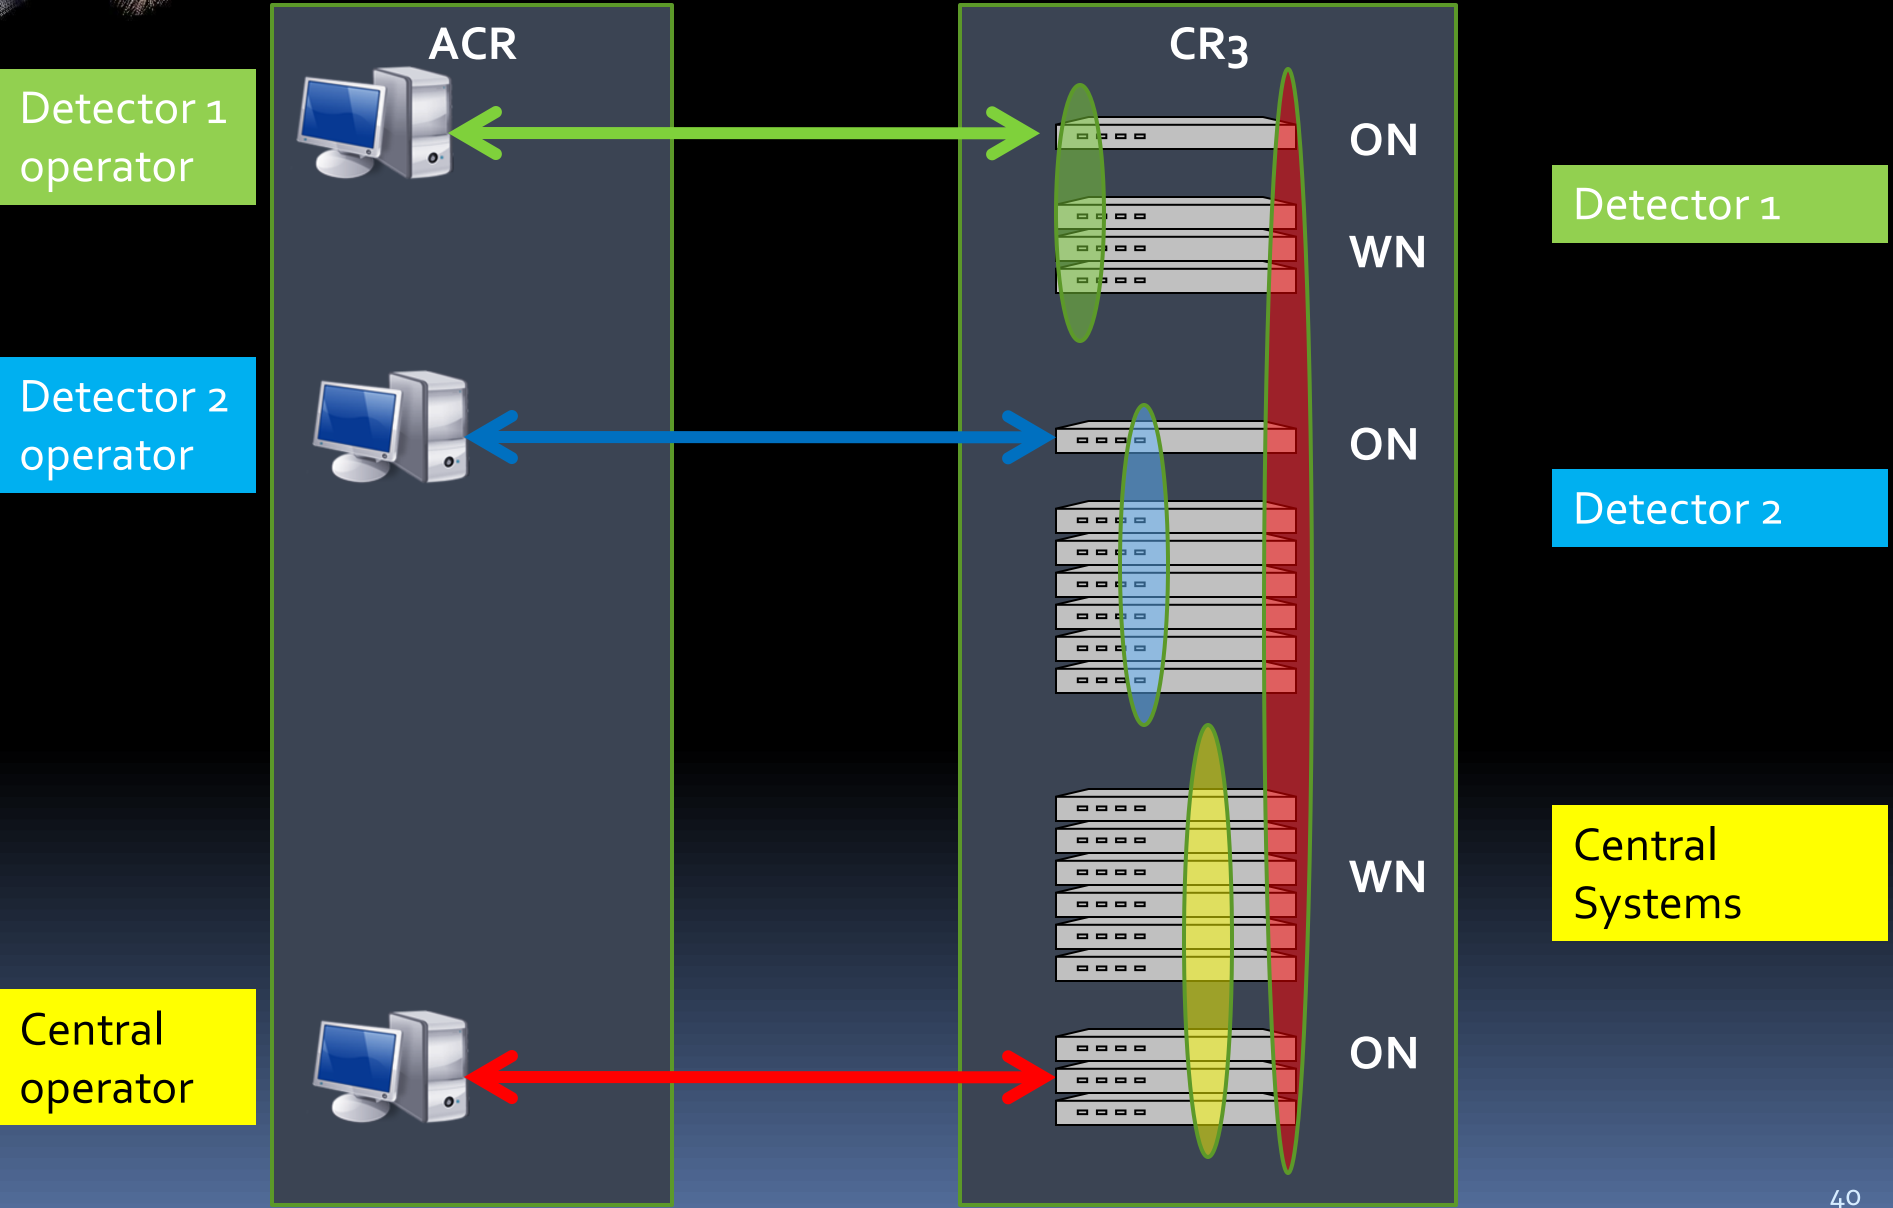

- Could there be an issue?

Authentication trap

- During the operation, the detector operator uses many windows, displaying several parts of the controlled system
  - Sometimes many ssh sessions to electronic cards are opened and devices are operated interactively
- At shift switchover old operator is supposed to logoff and new operator to logon
  - In certain cases the re-opening of all screens and navigating to components to be controlled can take 10-20 minutes, during this time the systems would run unattended
    - During beam injections, detector tests, etc. the running procedures may not be interrupted
- Shall we use shared accounts instead?
  - Can we keep the credentials protected?

Information leaks

- Sensitive information, including credentials, can leak
  - Due to lack of protection
  - Due to negligence/ignorance

….in scripts

```
echo " ----- make the network connections -----"
rem --- net use z: \\alidcsfs002\DCS_Common XXXXXX  /USER:CERN\dcsoper
rem --- net use y: \\alidcscom031\PVSS_Projects XXXXXX  /USER:CERN\dcsoper
echo " ------ done ---------"
rem ---ping 1.1.1.1 -n 1 -w 2000 >NULL
```

```
START C:\Programs\PVSS\bin\PVSS00ui.exe -proj lhc_ui -user operator:XXXXXX
-p lhcACRMonitor/lhcACRDeskTopDisplay.pnl,$panels:BackGround:lhcBackground/
lhcBackgroundMain.pnl;Luminosity_Leveling:lhcLuminosity/
lhcLuminosityLumiLevelling.pnl;Collisions_Schedule:BPTX/
lhc_bptxMonitor.pnl;V0_Control:lhcV00Control/lhcV00ControlMain.
```

These examples are real, original passwords in clear text are replaced by XXXXXX in this presentation

```
# Startup Batch Program for the LHC Interface Desktop
#
# Auth : deleted v1.0 4/8/2011
# - rdesktop -z -f -a 16 -k en-us -d CERN -u dcsoper -p XXXXXX -s “D:
\PVSS_Profiles\ACRLHCDesk.bat” alidcscom054
rdesktop -z -g2560x1020 -a 16 -k en-us -d CERN -u
```

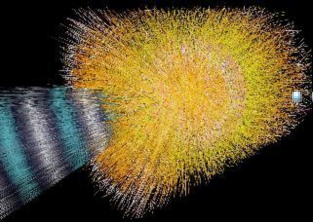

… In documentation

Entries like this :

The relevant parameters are
• Window dimension : 1920x1050;
• RDT credential : host = alidcscom054, user = dcsoper, password = XXXXXX;
• shell command to start : D:\PVSS_Profiles\ACRLHCBigScreen.bat
• panel to reference : lhcACRMonitor/lhcACRMain.pnl

Can be found in

Thesis

Reports

Web pages

…..

Twikies

Printed manuals

We protect our reports and guides, but institutes republish them very often on their unprotected servers

... or even worse!

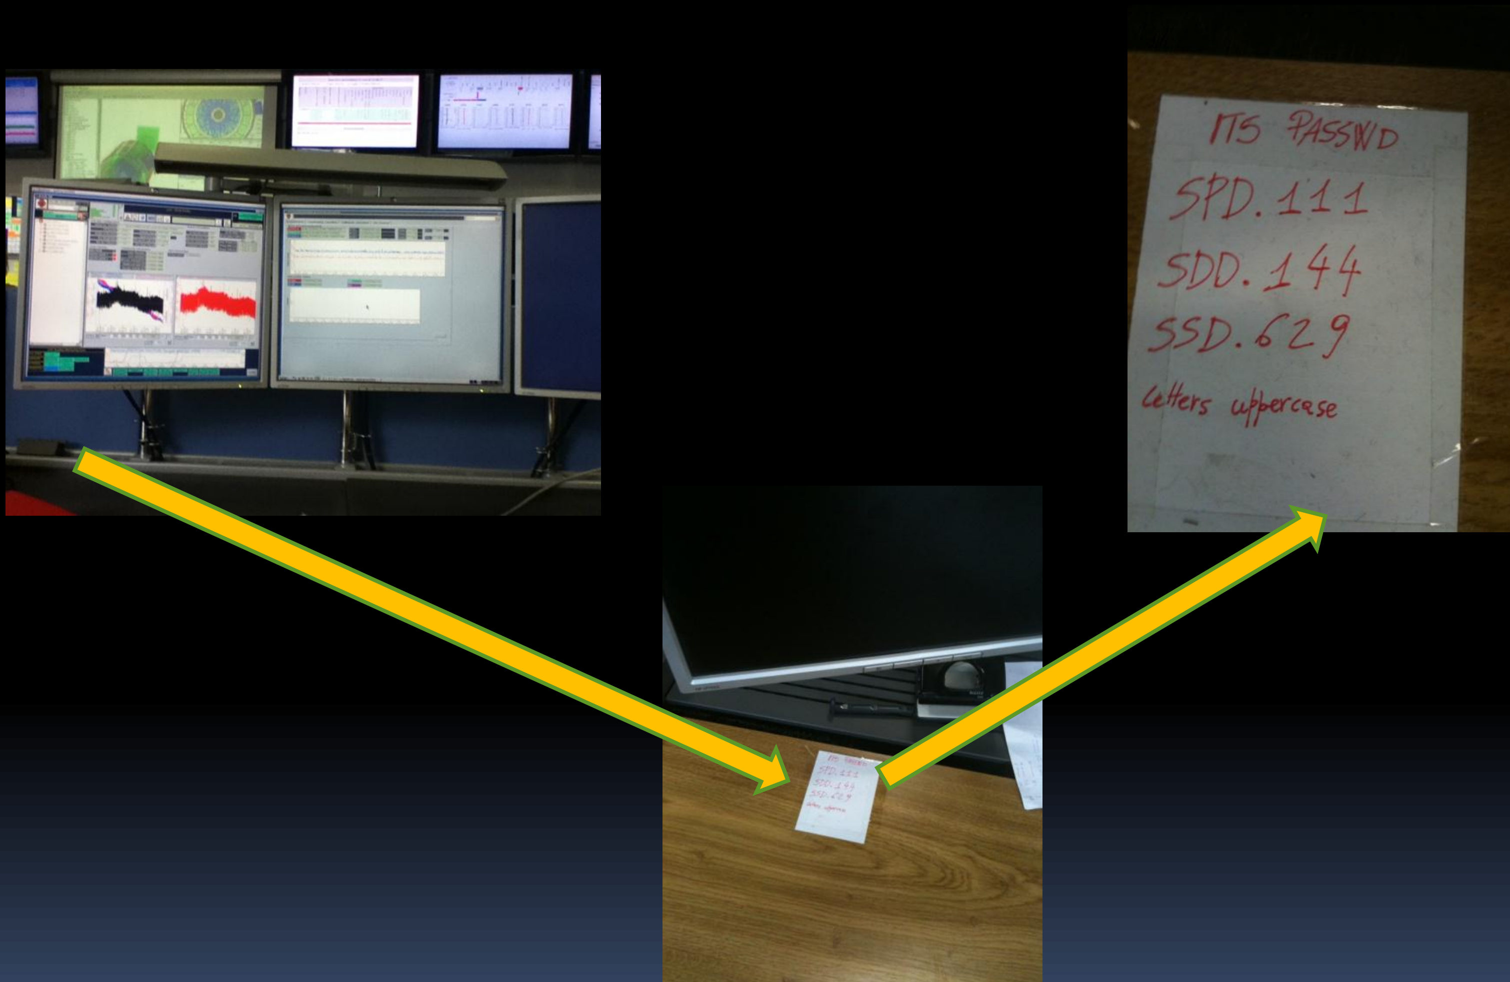

Using shared accounts

- In general, the use of shared accounts is undesired
- However, if we do not allow for it, users start to share their personal credentials

- Solution – use of shared accounts (detector operator, etc.) only in the control room
  - Restricted access to the computers
  - Autologon without the need to enter credentials
  - Logon to remote hosts via scripts using encrypted credentials (like RDP file)
  - Password known only to admins and communicated to experts only in emergency (sealed envelope)
- Remote access to DCS network allows only for physical user credentials

- OK, so we let people to work from the control room and remotely.
- Is this all?

Data exchange

- The DCS data is required for physics reconstructions, so it must be made available to external consumers
- The systems are developed in institutes, and the elaborated software must be uploaded to the network
- Some calibration data is produced in external institutes, using semi-manual procedures
  - Resulting configurations must find a way to the front end electronics
- Daily monitoring tasks require access to the DCS data from any place at any time

- How do we cope with that requests?

Getting data OUT of DCS

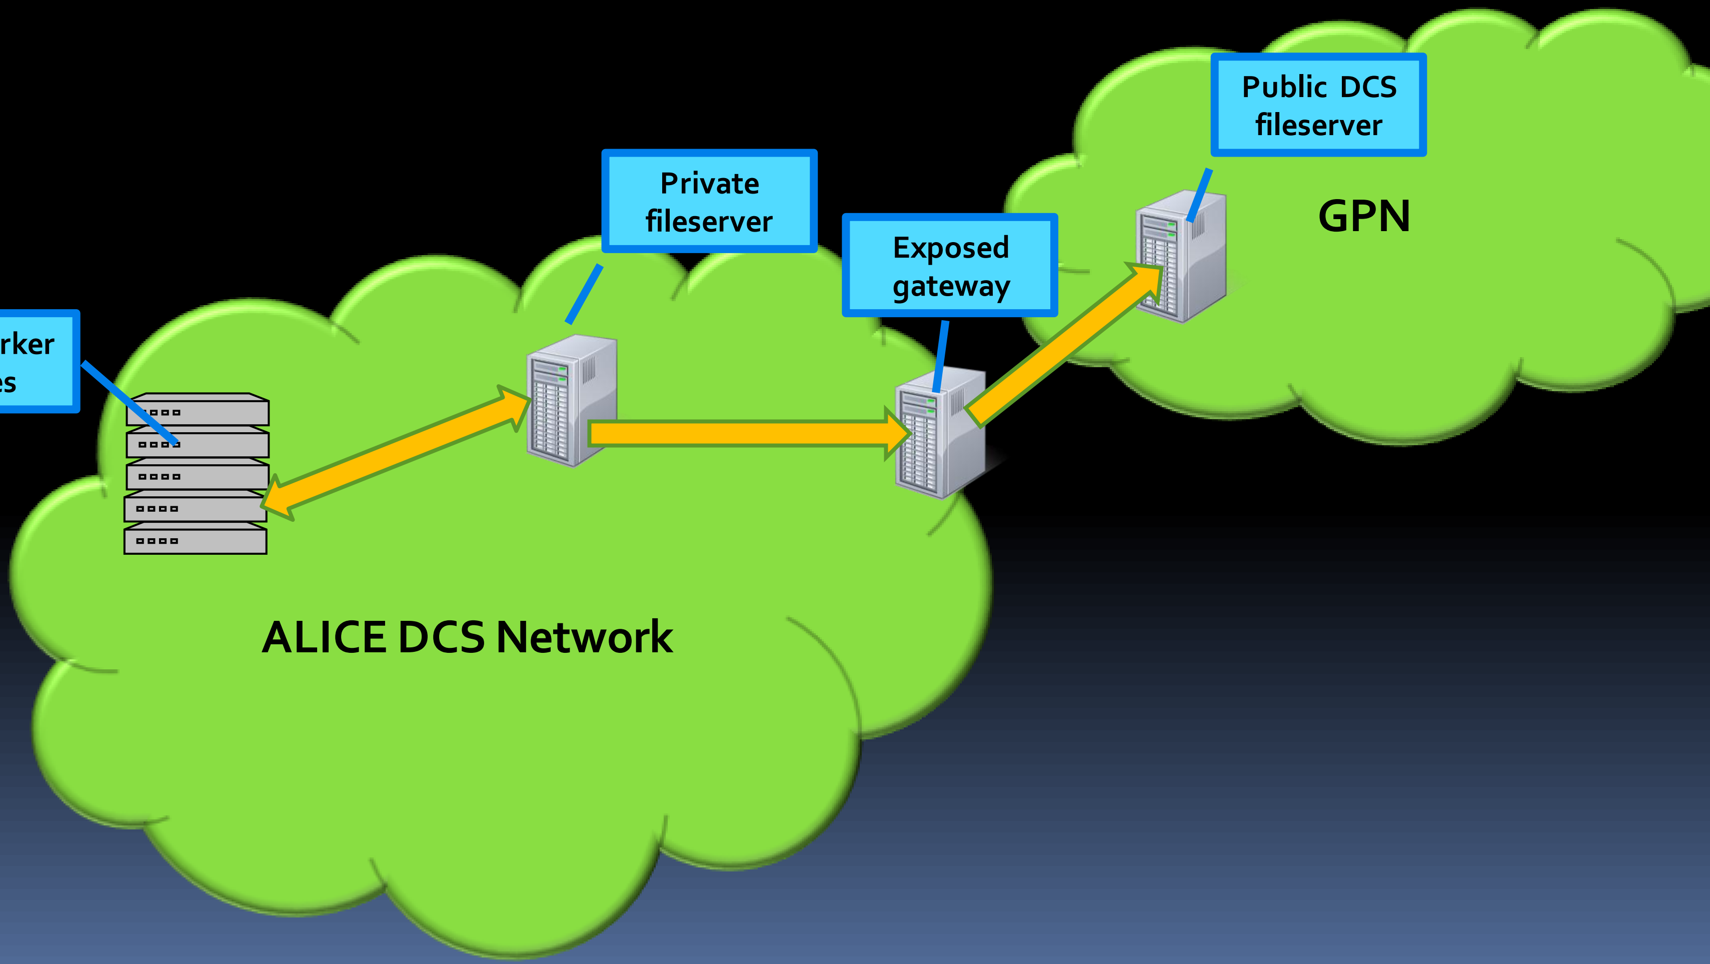

DCS WWW monitoring

- WWW is probably the most attractive target for intruders
- WWW is the most requested service by institutes

- ALICE model:
  - Users are allowed to prepare a limited number of PVSS panels, displaying any information requested by them
  - Dedicated servers opens these panels periodically and creates snaphosts
  - The images are automatically transferred to central Web servers
- Advantage:
  - There is no direct link via the WWW and ALICE DCS, but the web still contains updated information
- Disadvantage/challenges:
  - Many

WWW monitoring

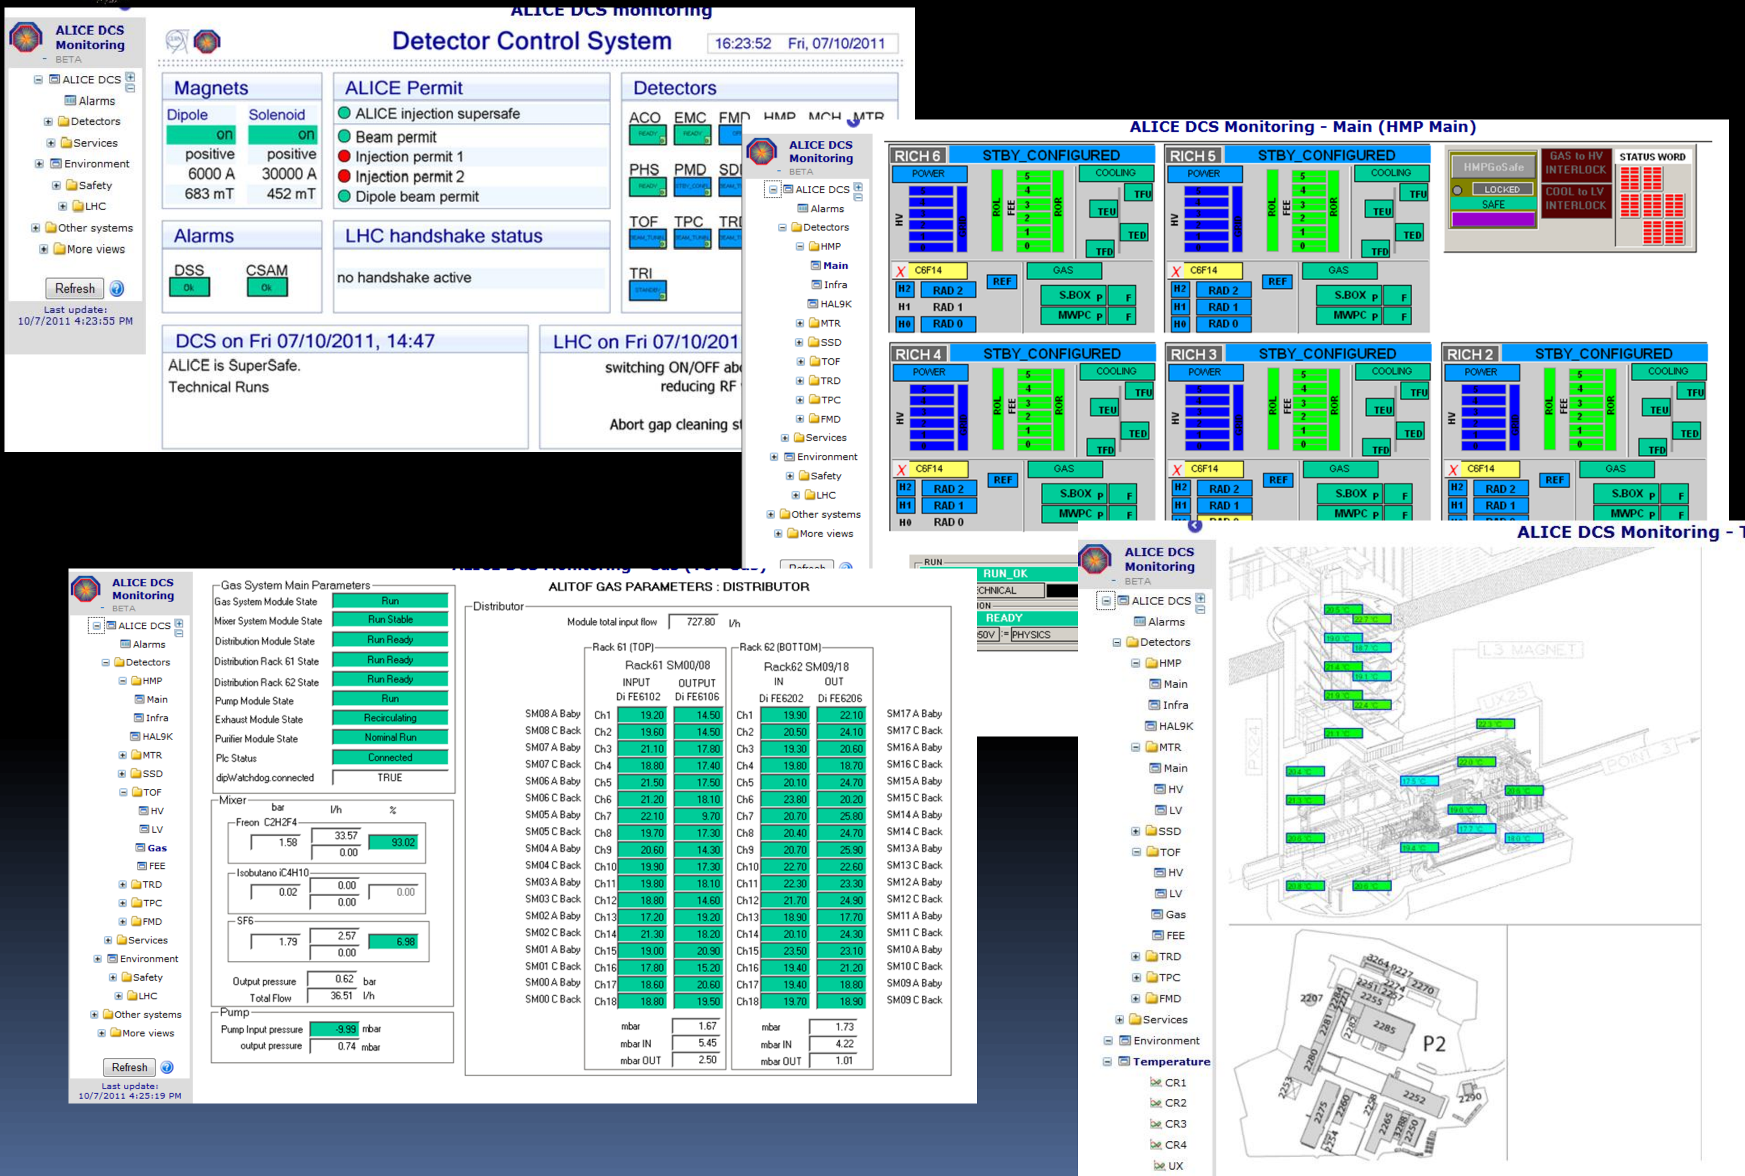

Getting data OUT of DCS

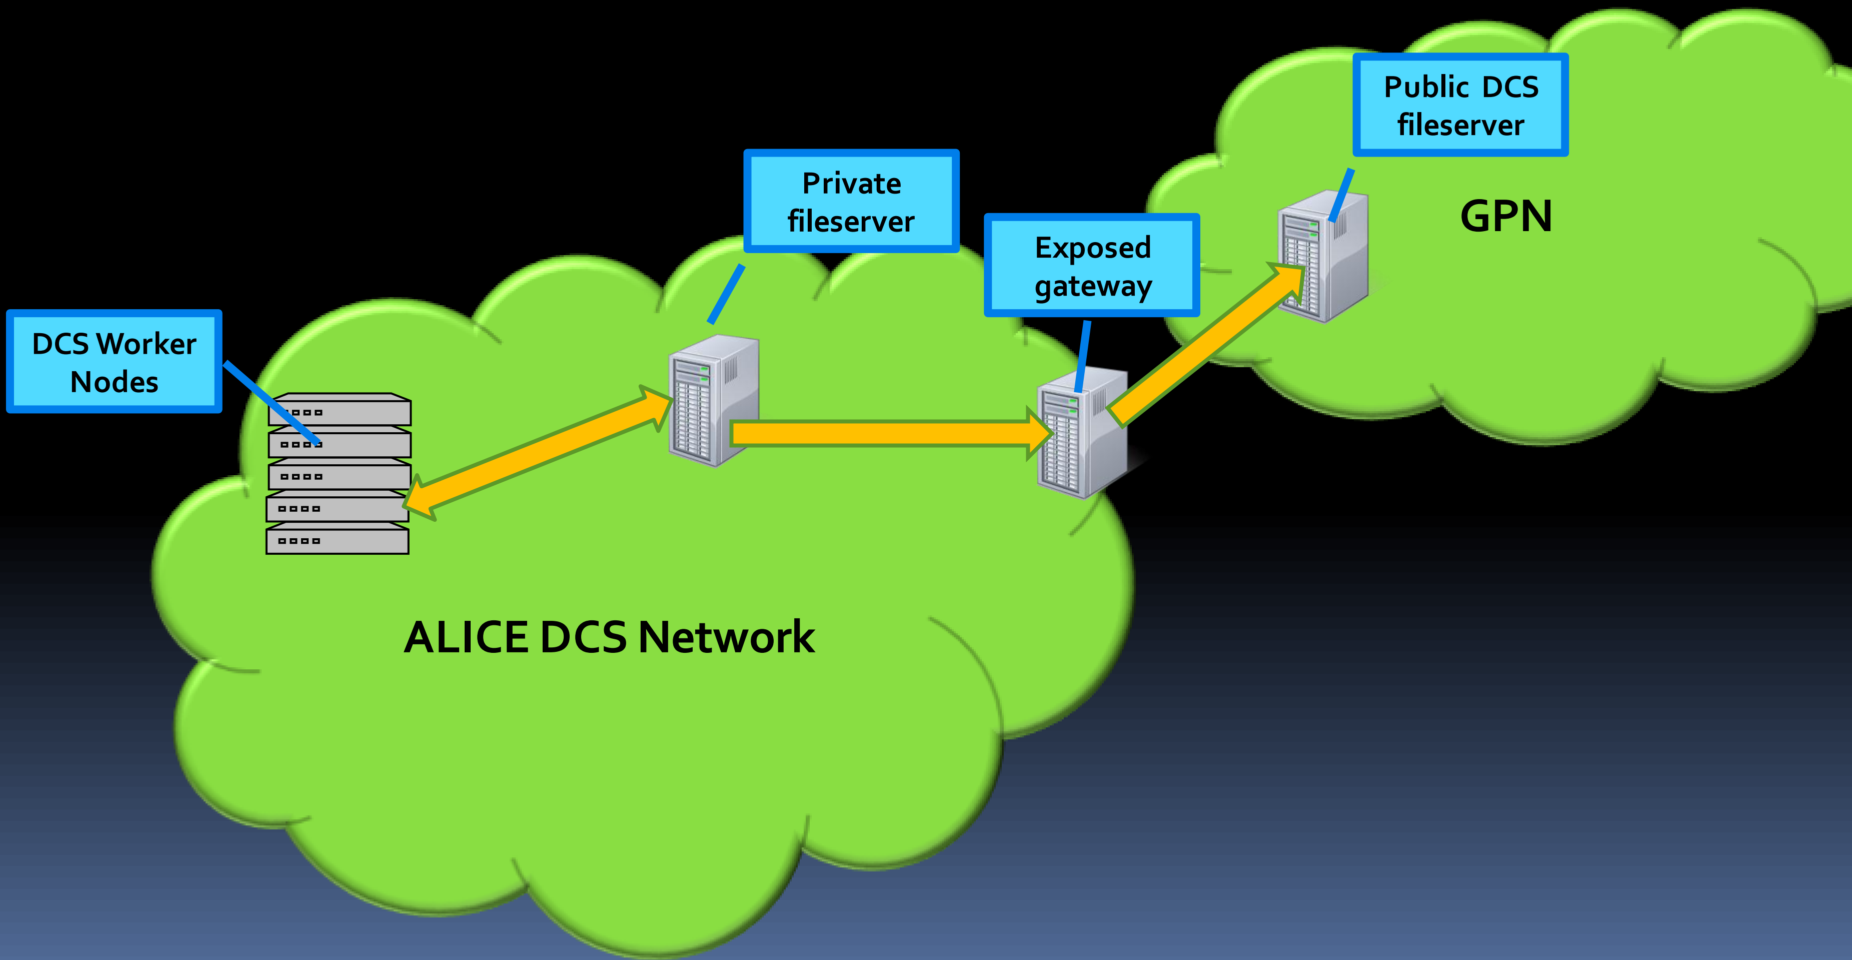

Getting data OUT of DCS

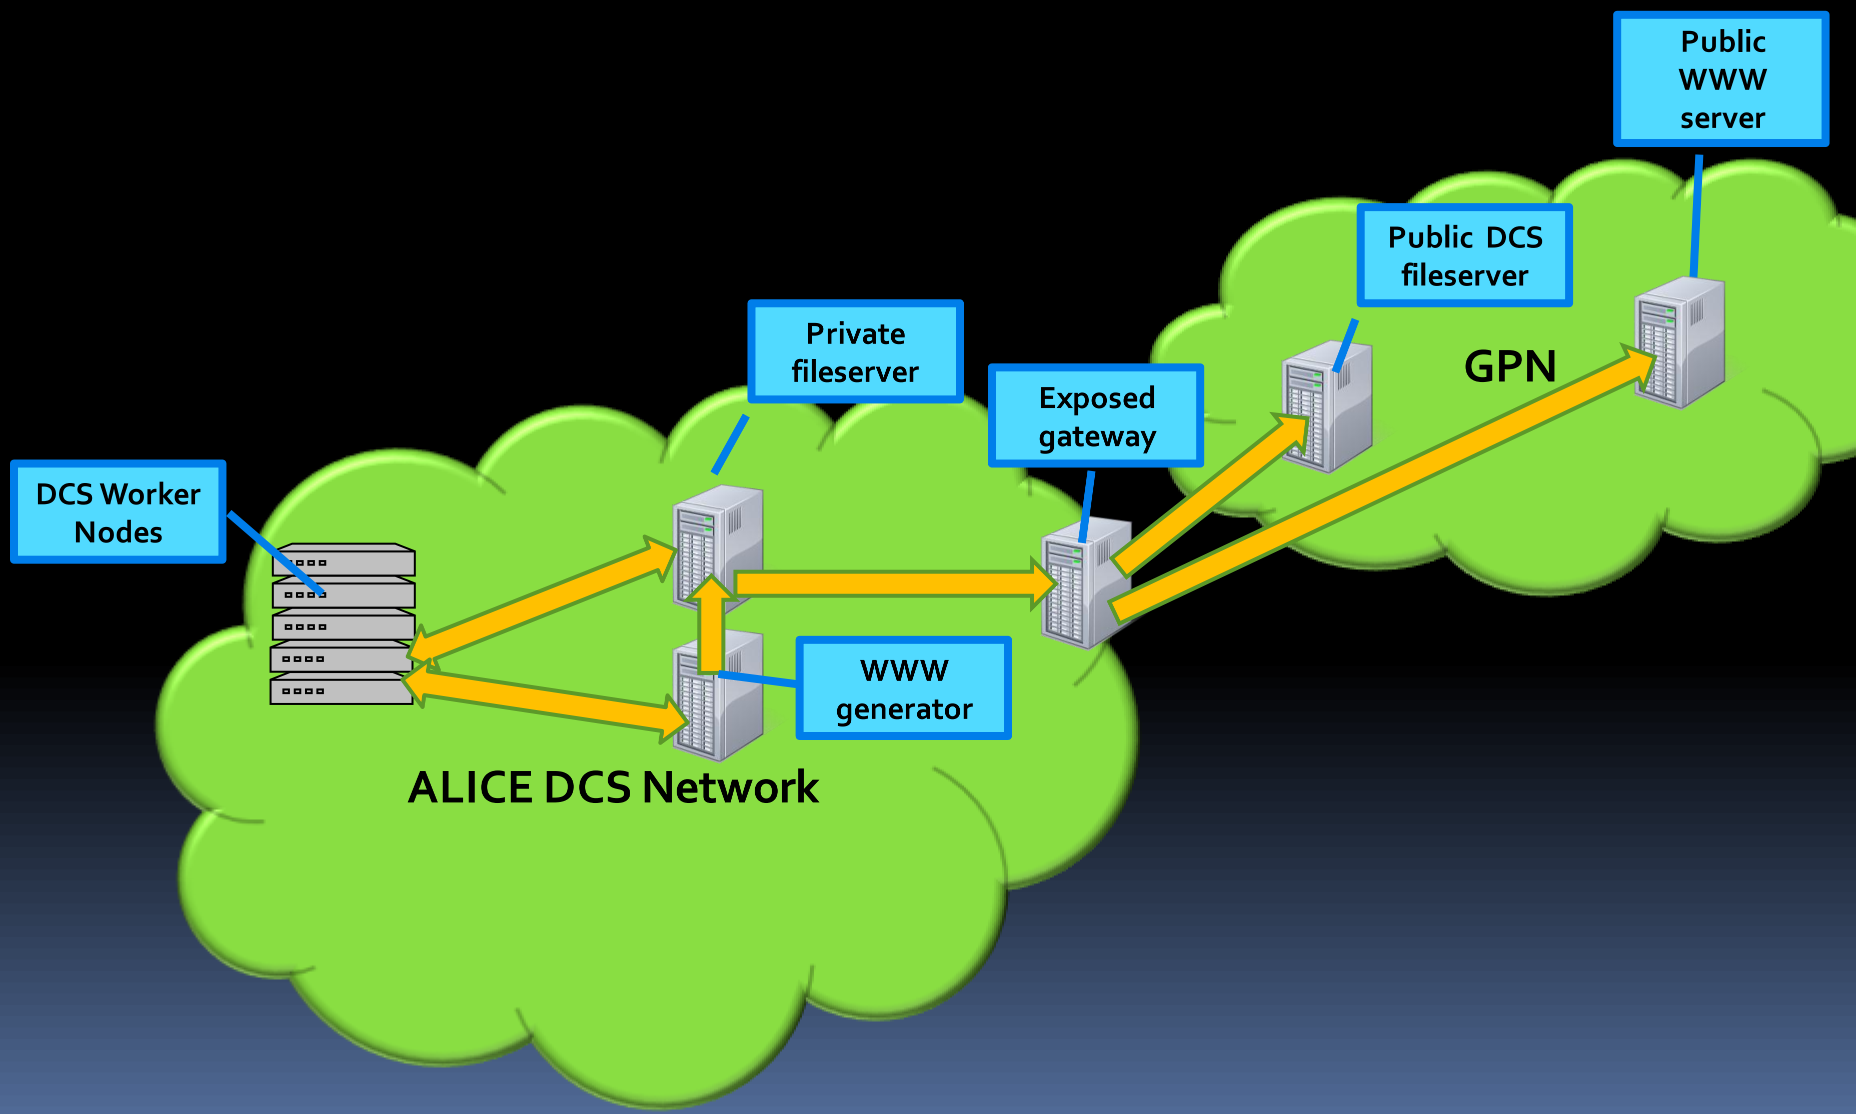

Data for OFFLINE

- Certain DCS data is required for offline reconstruction
  - Conditions data
  - Configuration settings
  - Calibration parameters
- Conditions data is stored in ORACLE and sent to OFFLINE via dedicated client-server machinery
- Calibration, configuration, memory dumps, etc. are stored on private fileserver and provided to offline
- OFFLINE shuttle collects the data at the end of each run

Getting data OUT of DCS

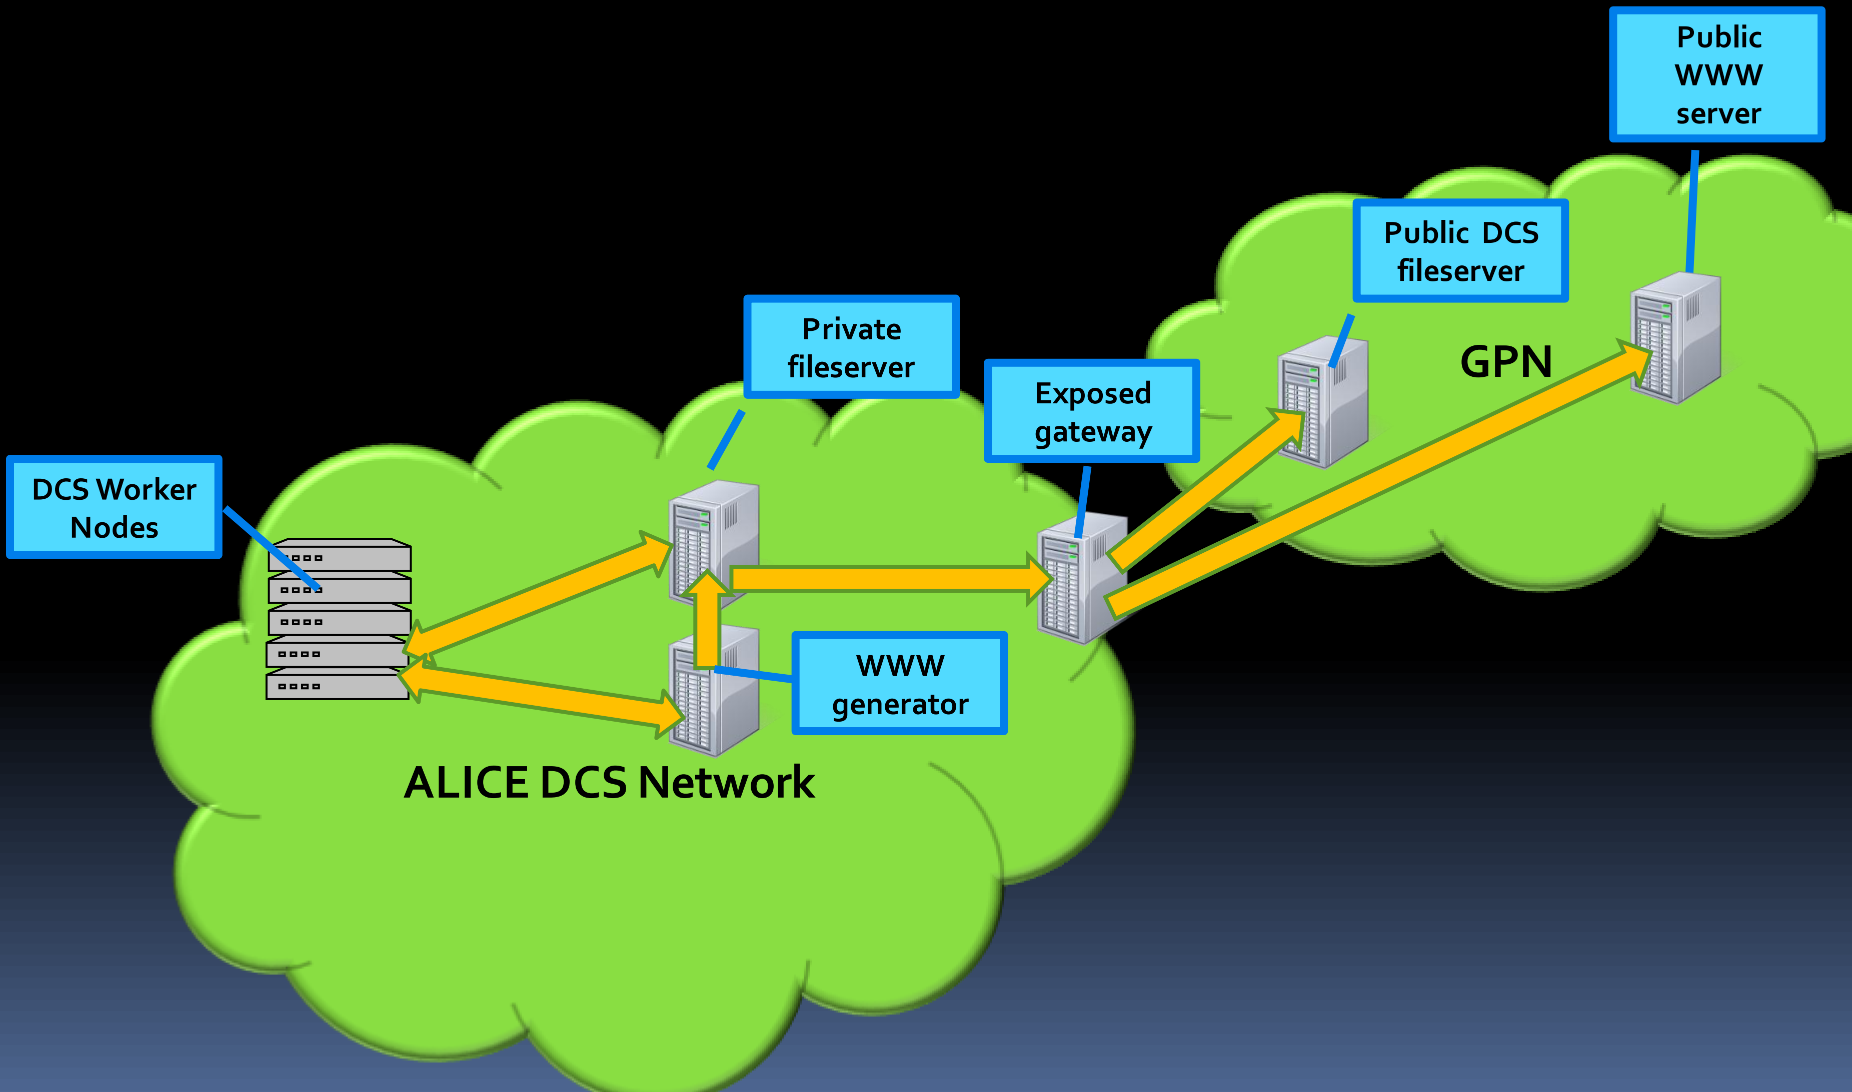

Getting data OUT of DCS

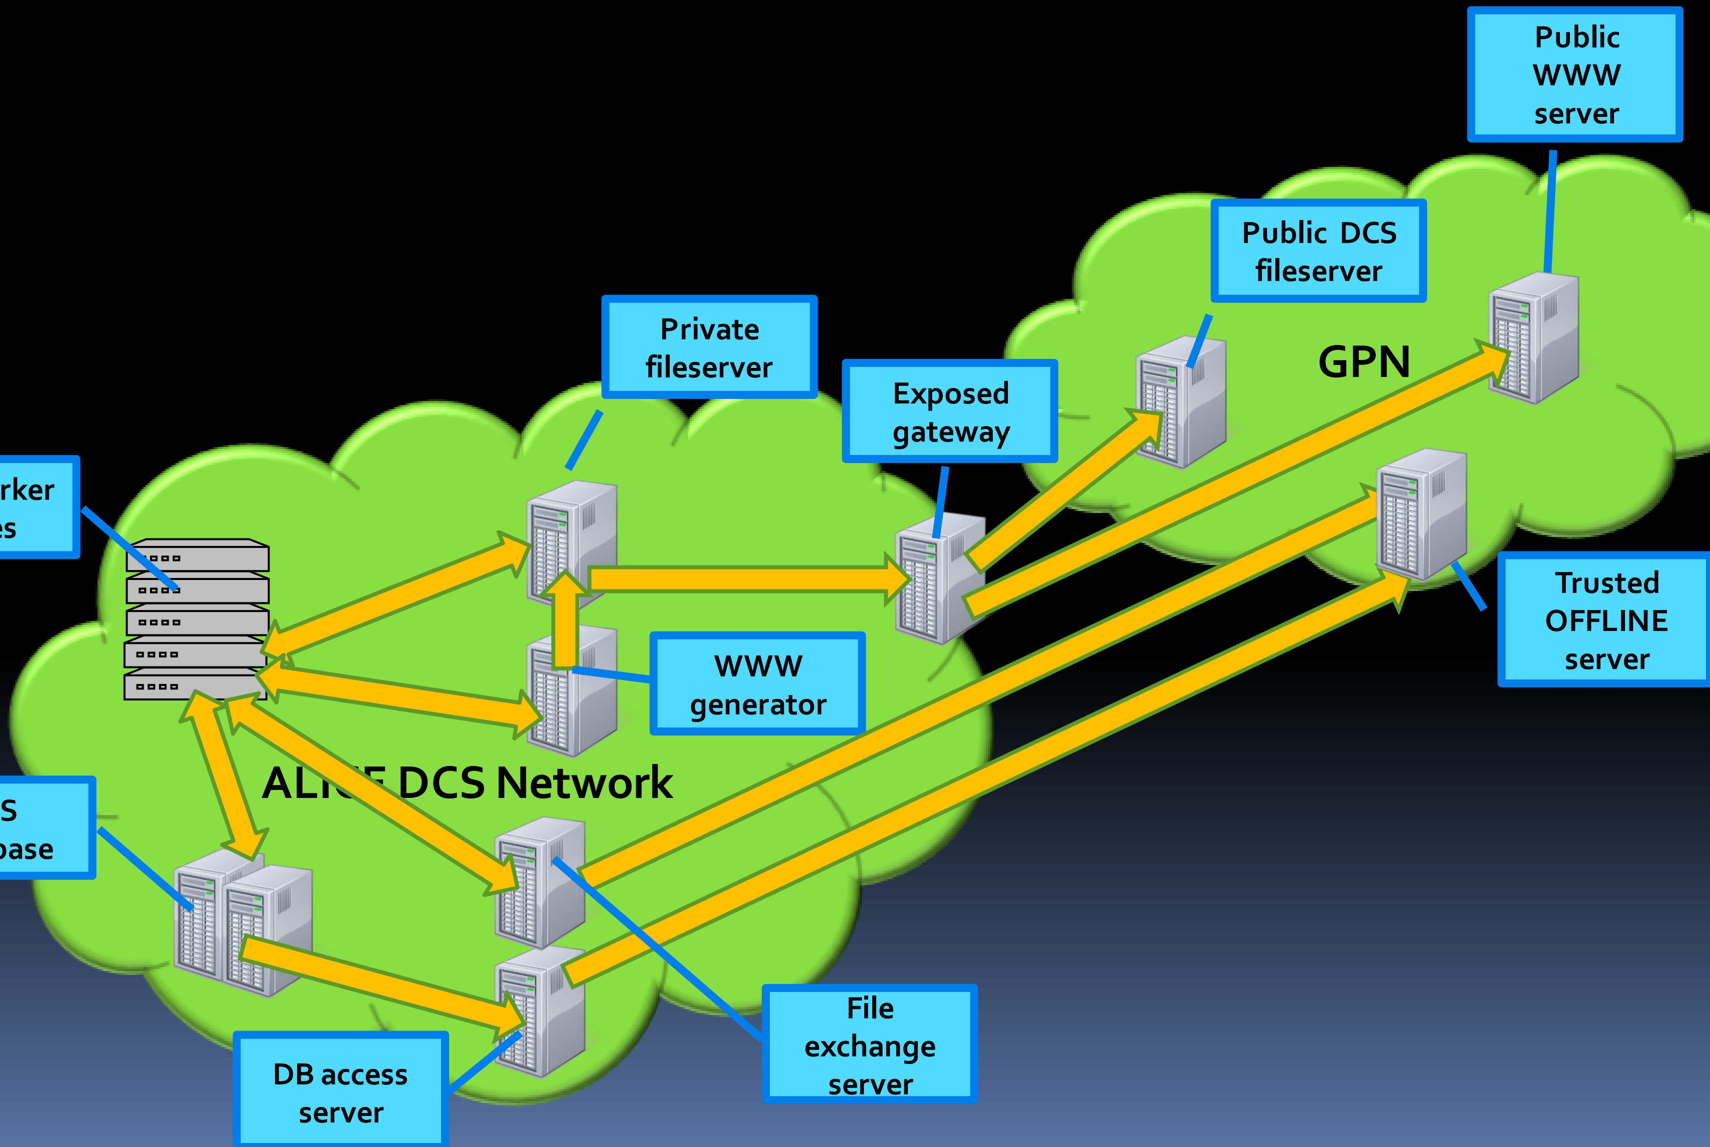

Data to ALICE online systems

- During the runs, DCS status is published to other online systems for synchronization purposes
  - Run can start only if DCS is ready
  - Run must be stopped if DCS needs to perform safety related operations
  - Etc.
- Conditions data is sent to online and quasi-online systems for further processing
  - Data quality monitoring
  - Calibration parameters for HLT
  - Etc.

Getting data OUT of DCS

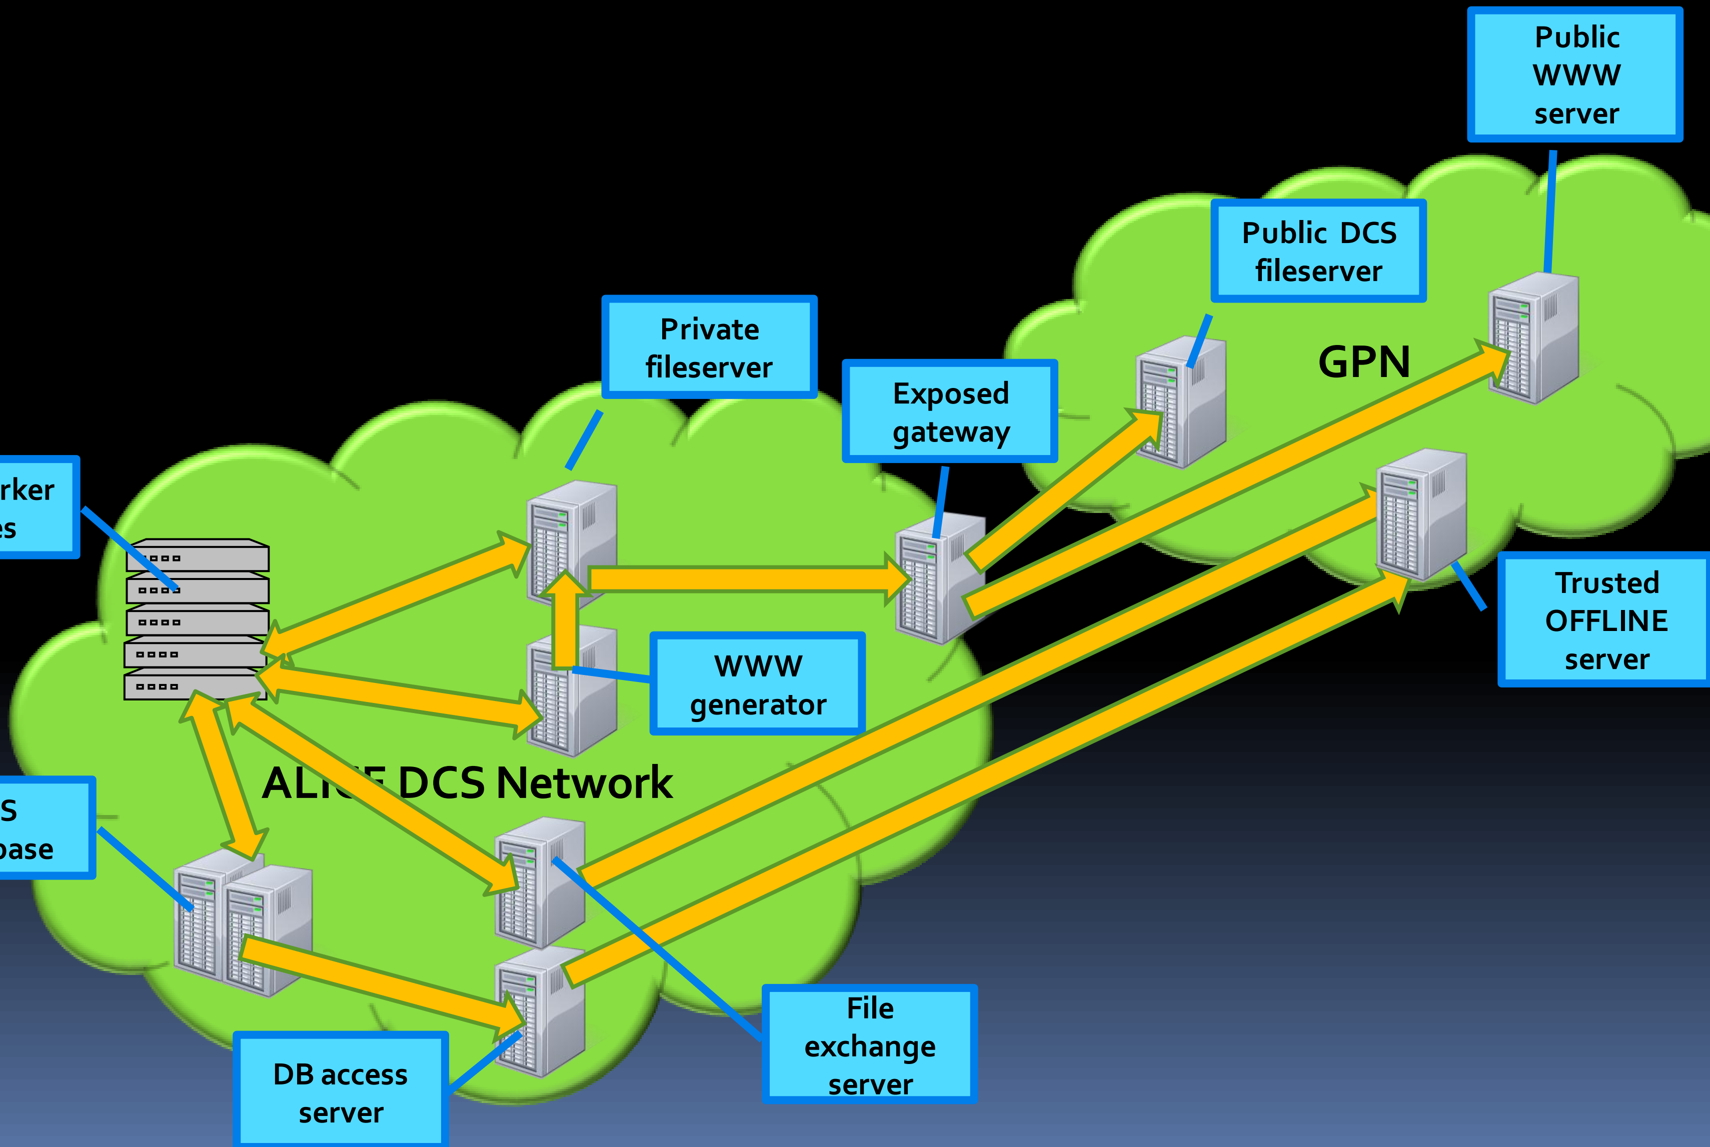

Getting data OUT of DCS

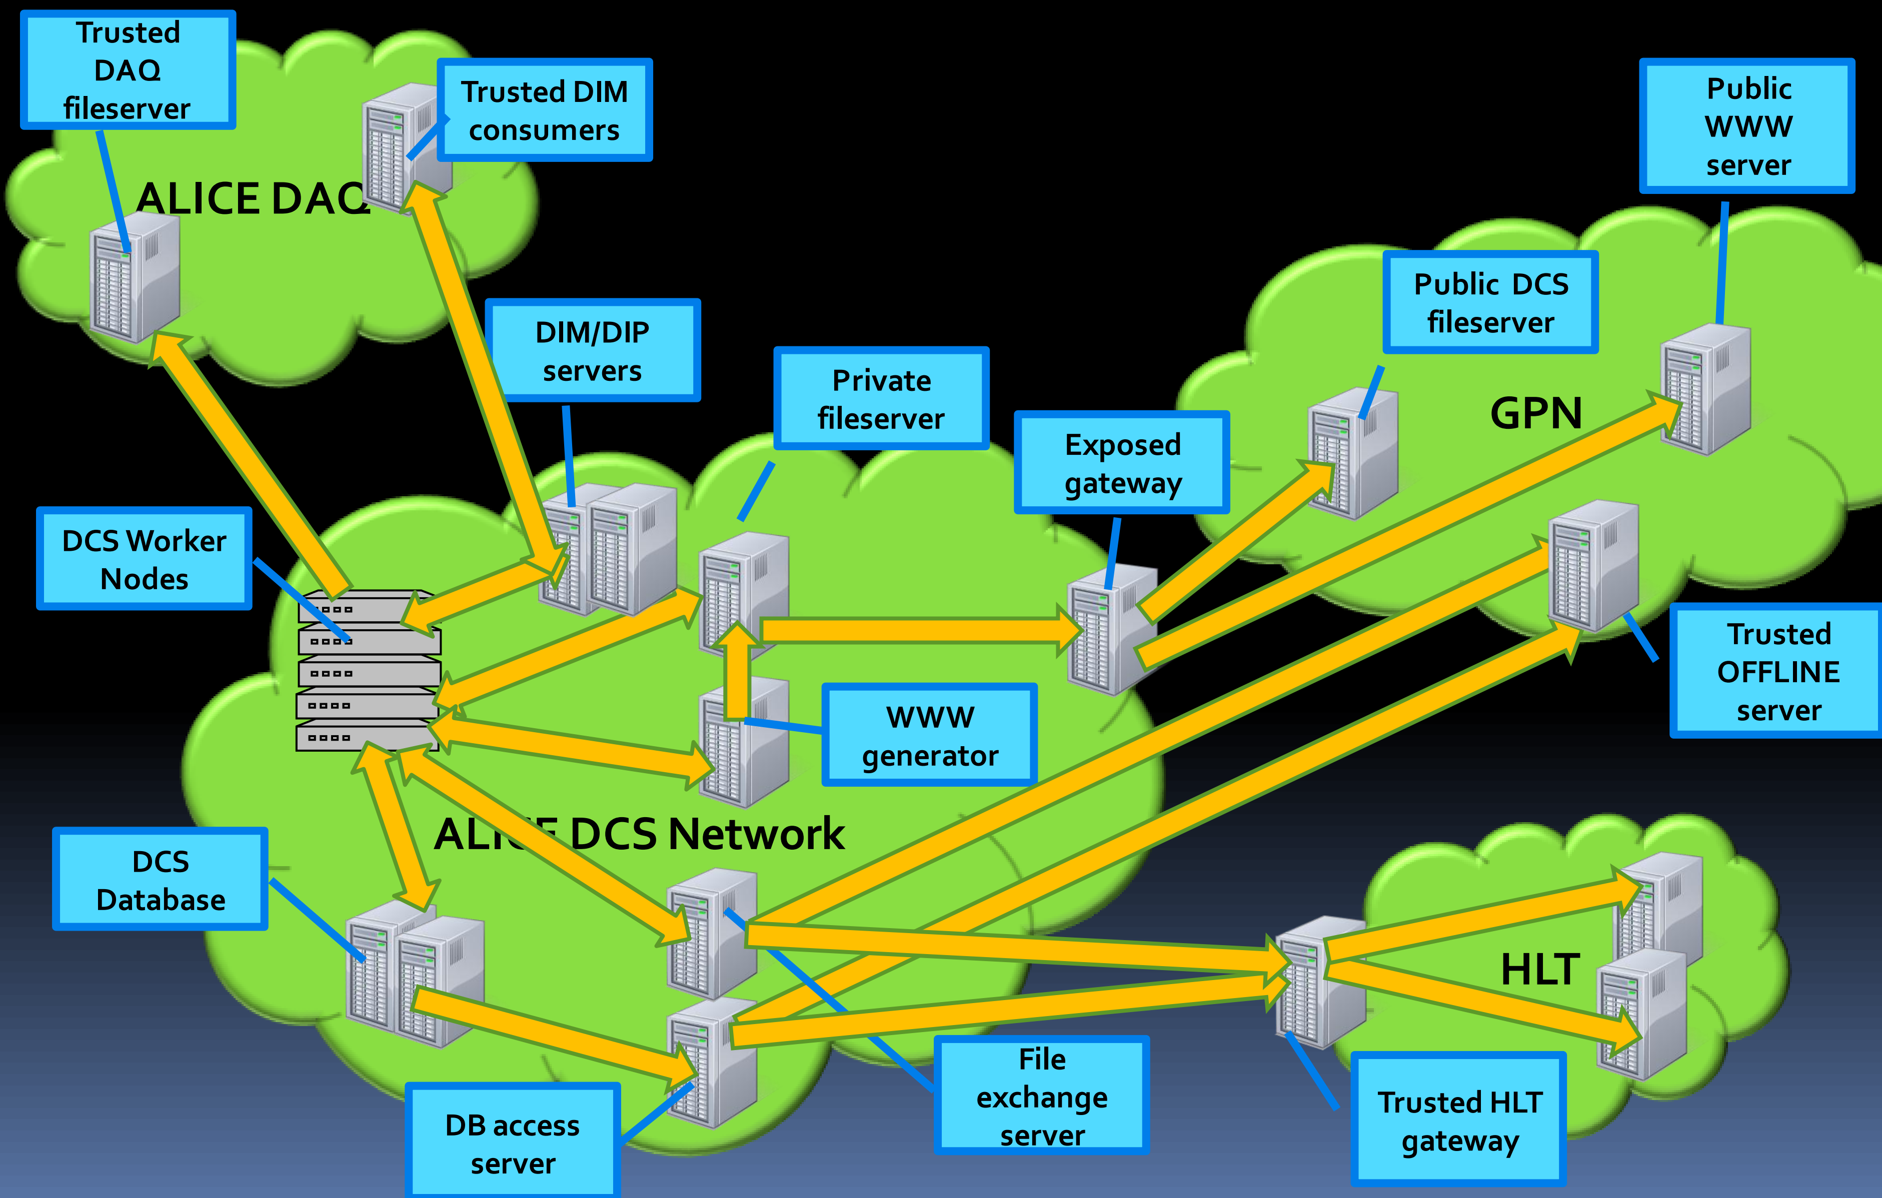

External data published to other sources

- DCS provides feedback to other systems
  - LHC
  - Safety
  - …

Getting data OUT of DCS

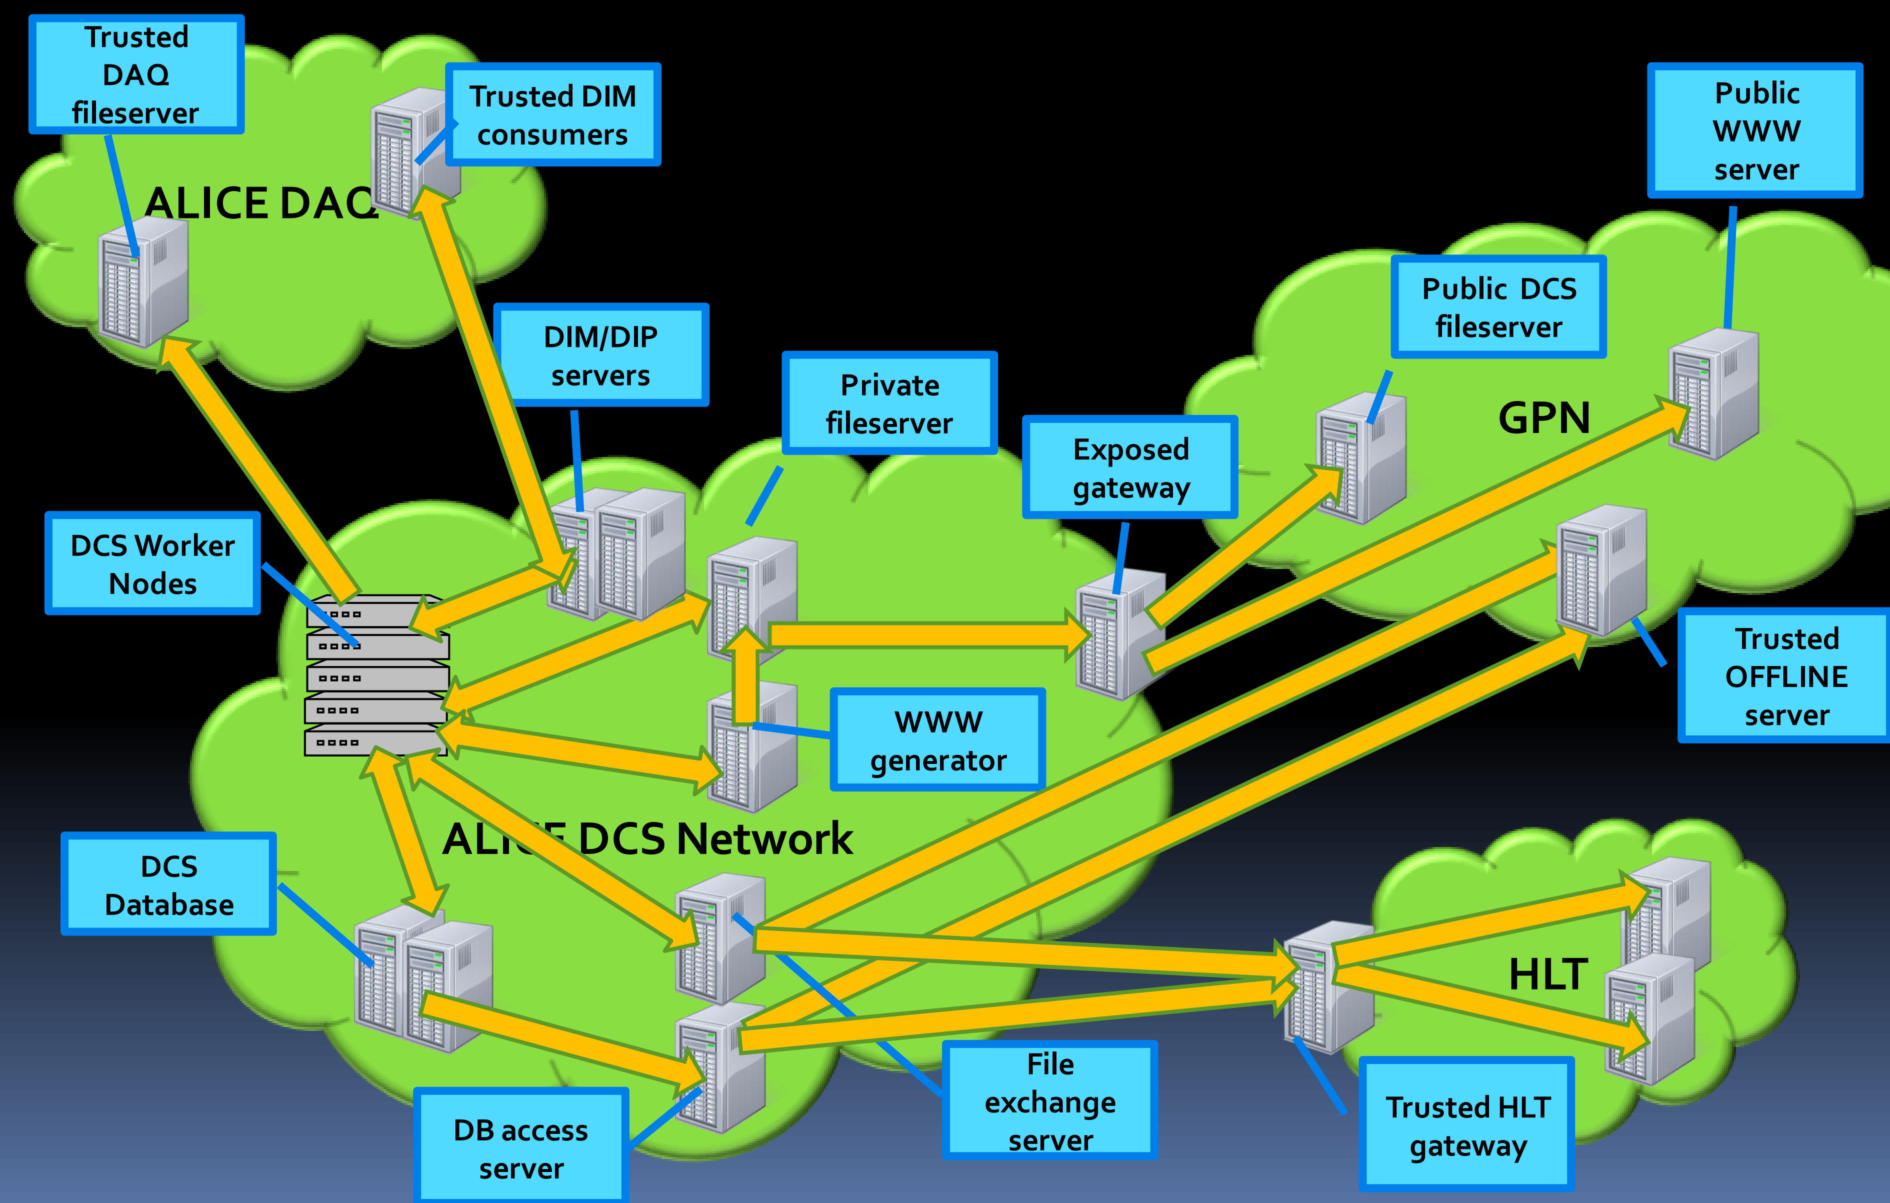

Getting data OUT of DCS

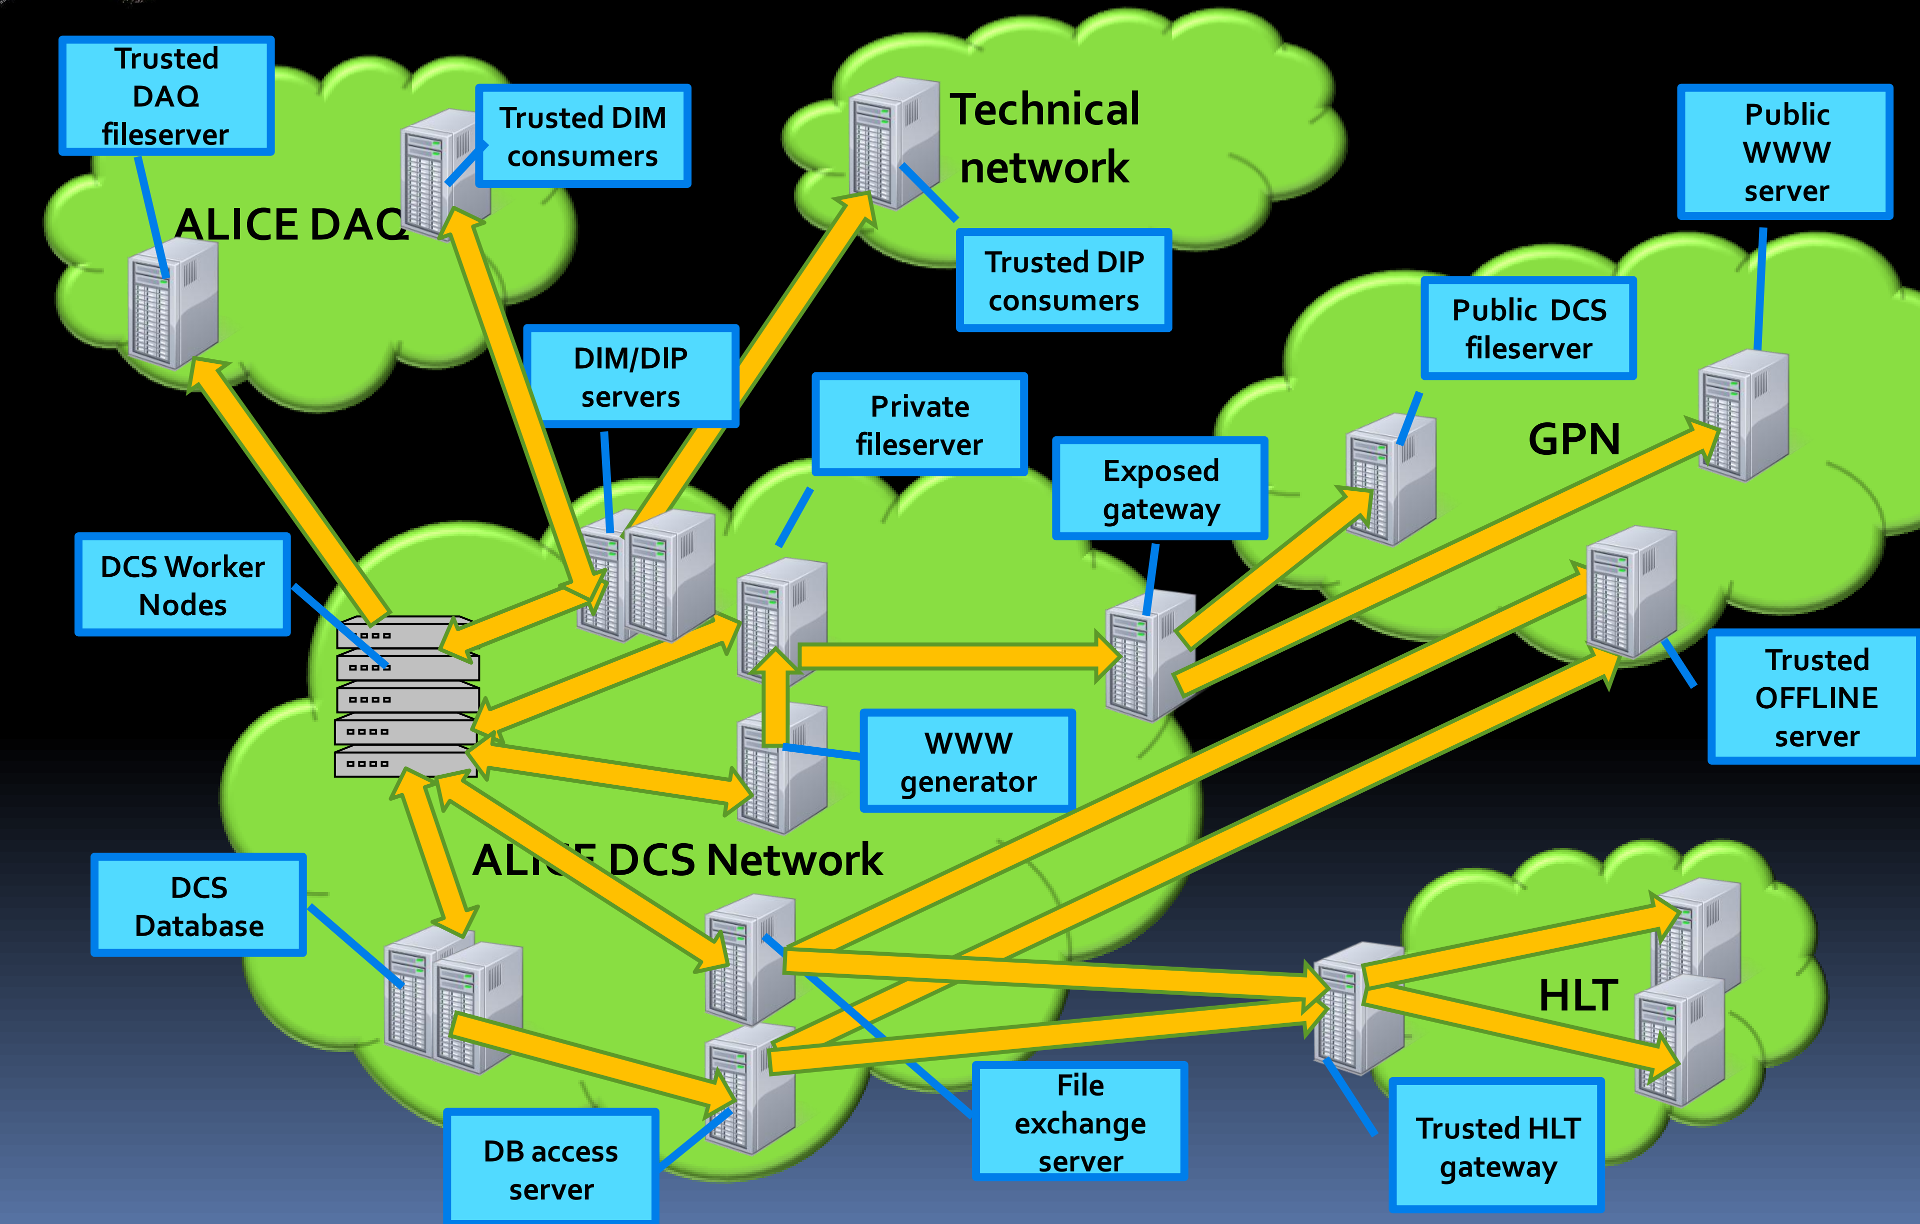

Getting data OUT of DCS

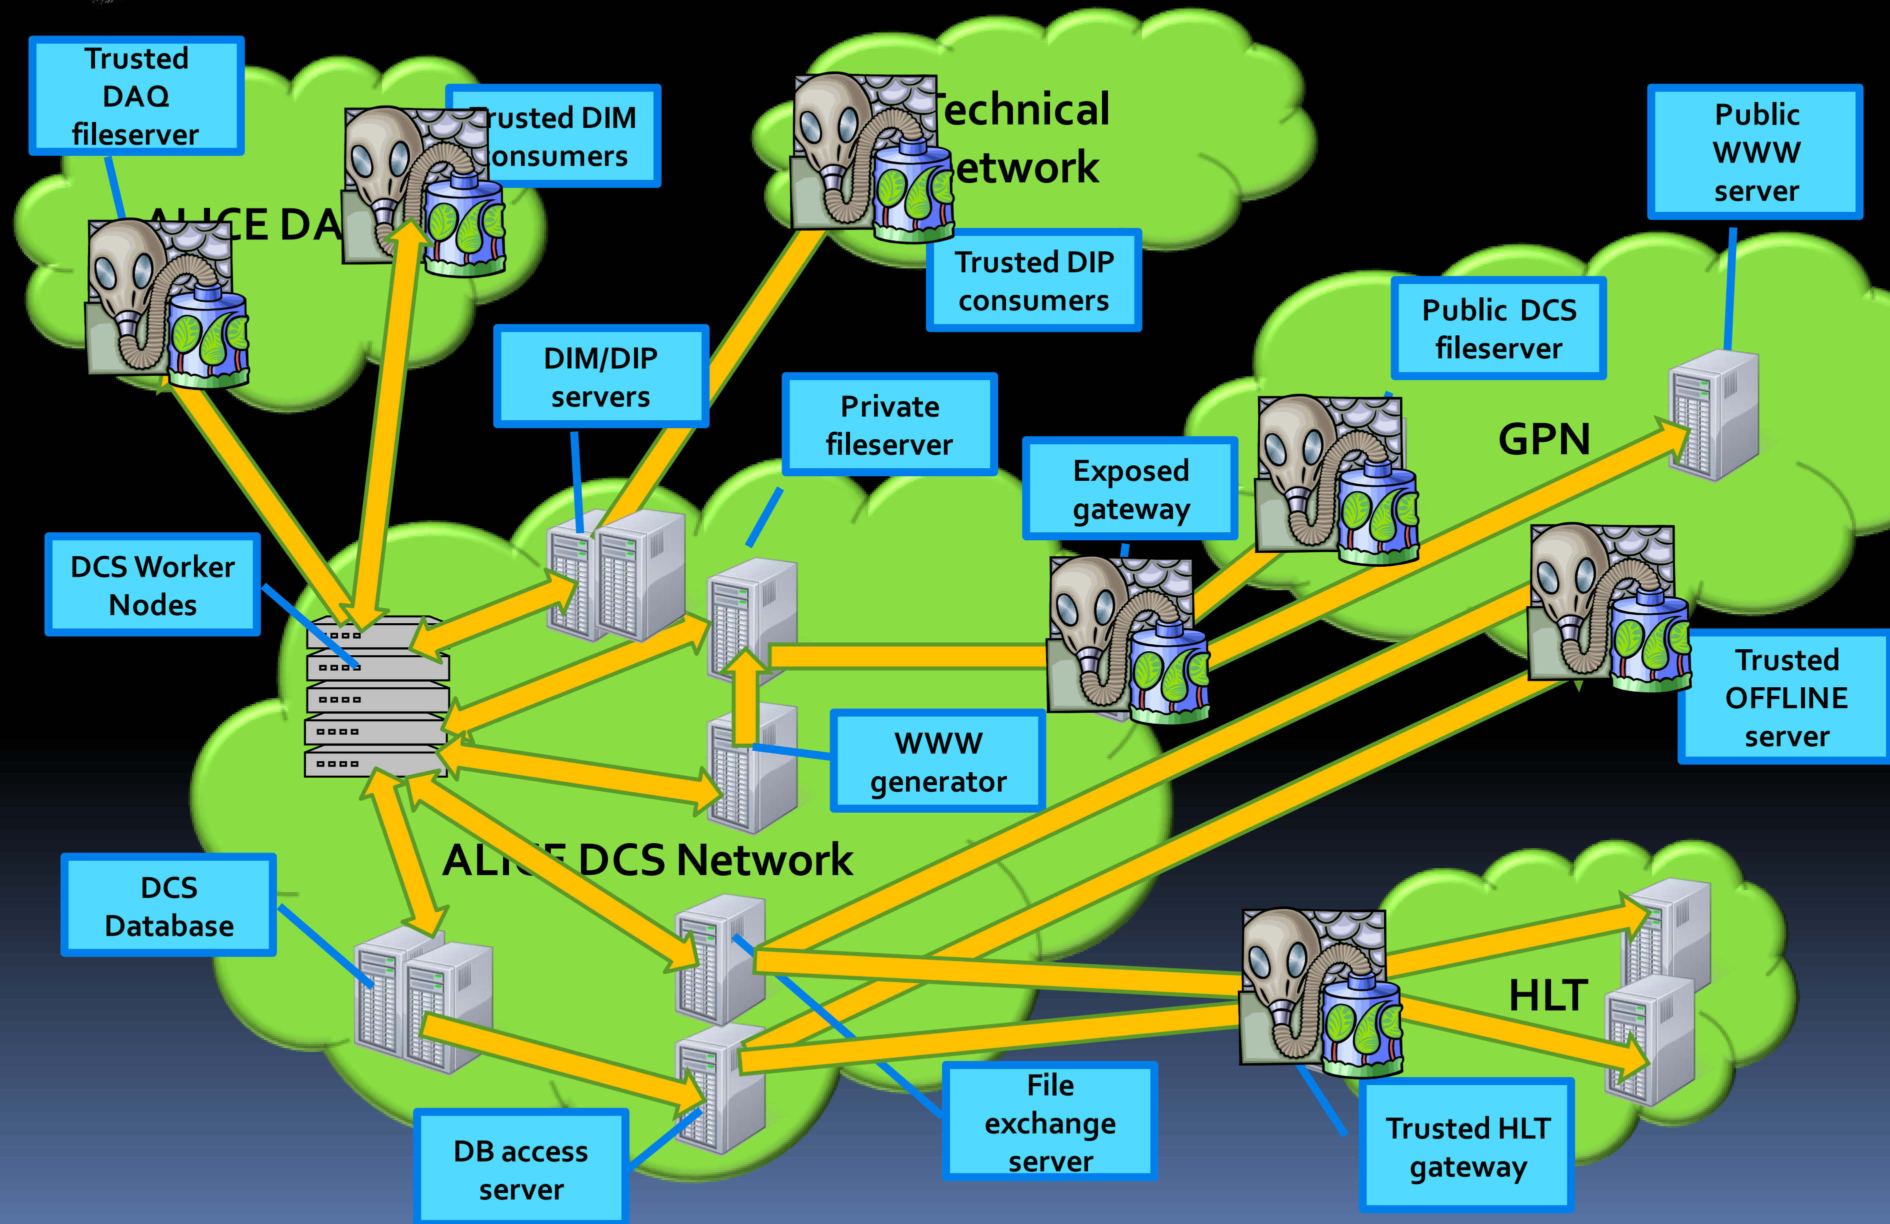

Getting data OUT of ALICE

- A number of servers in different domains need to be trusted
  - The picture still does not contain all the infrastructure needed to get the exchange working (nameserves, etc.)
- Filetransfer OUT of the DCS network is not limited
  - Autotriggered filetransfers
  - Data exchange on client request

Getting data to ALICE DCS

- All file transfers to ALICE DCS are controlled
  - Users upload the data to public fileservers (CERN security apply) and send transfer request
  - After checking the files (antivirus scans), data is uploaded to private DCS fileservers and made visible to DCS computers
- Automatic data flow to ALICE DCS is possible only via publisher/subscriber model
  - DCS clients subscribe to LHC services, environment monitors, safety systems …. and data is injected into the PVSS

Getting data IN to DCS

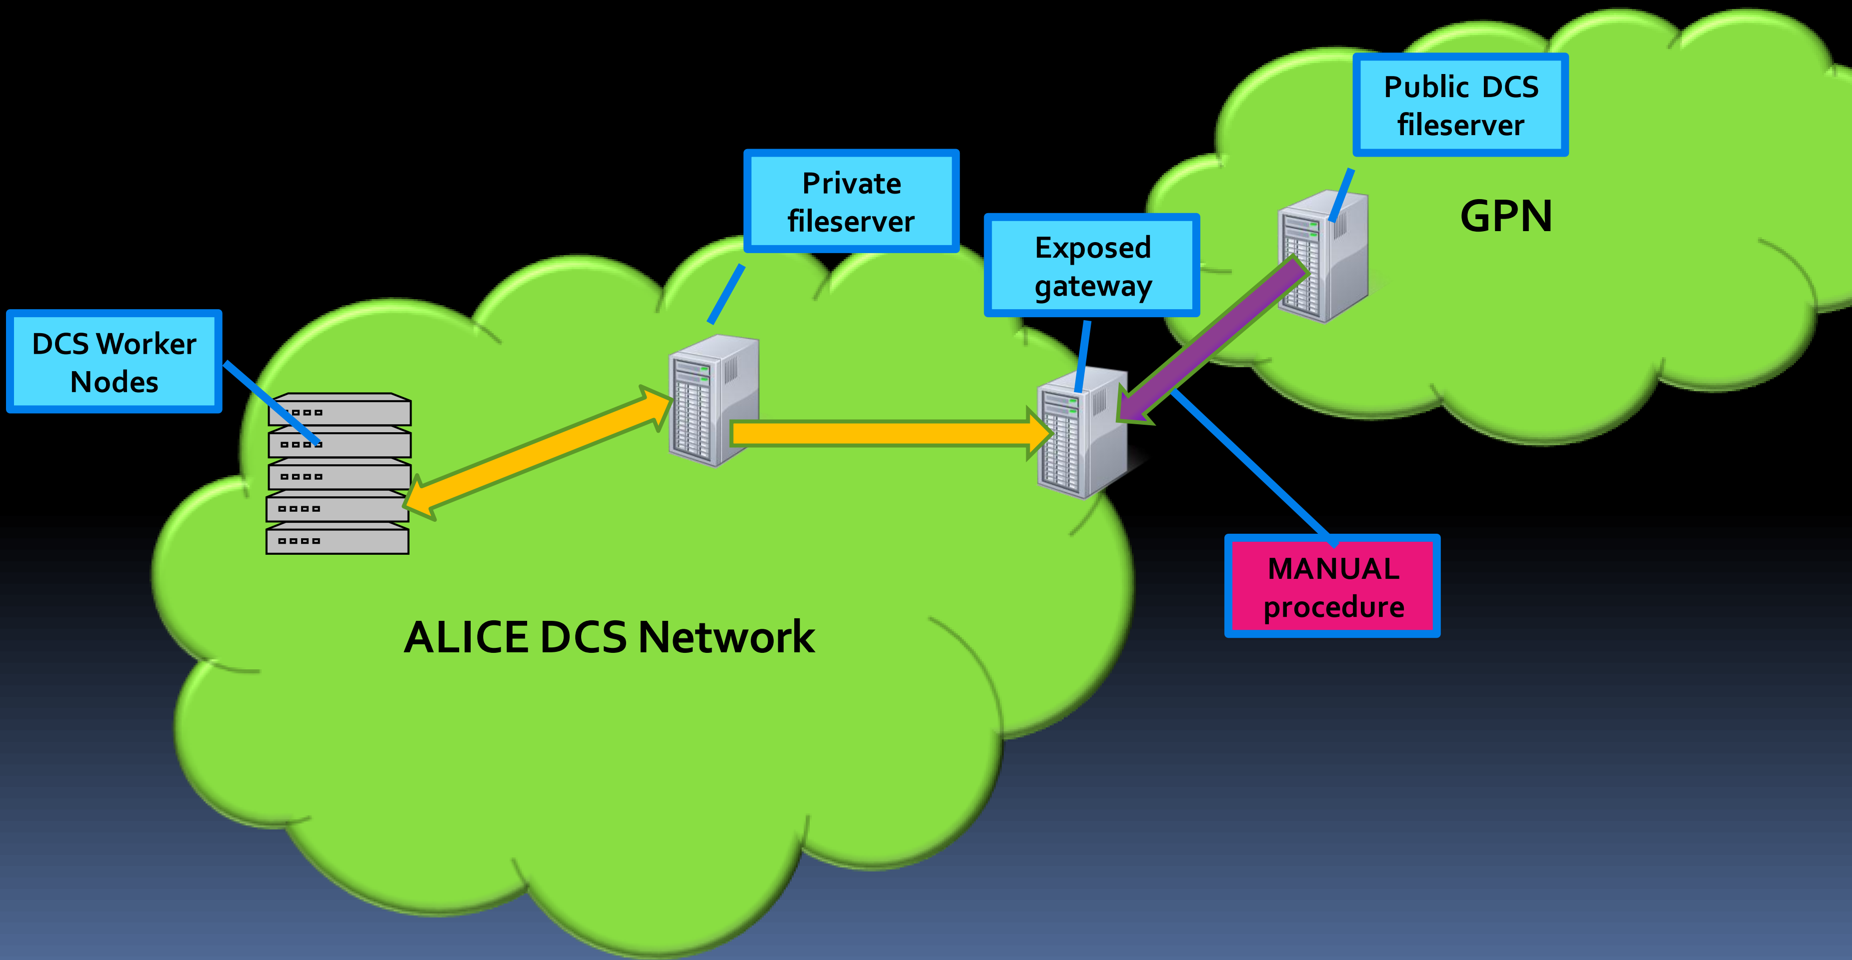

Getting data IN to DCS

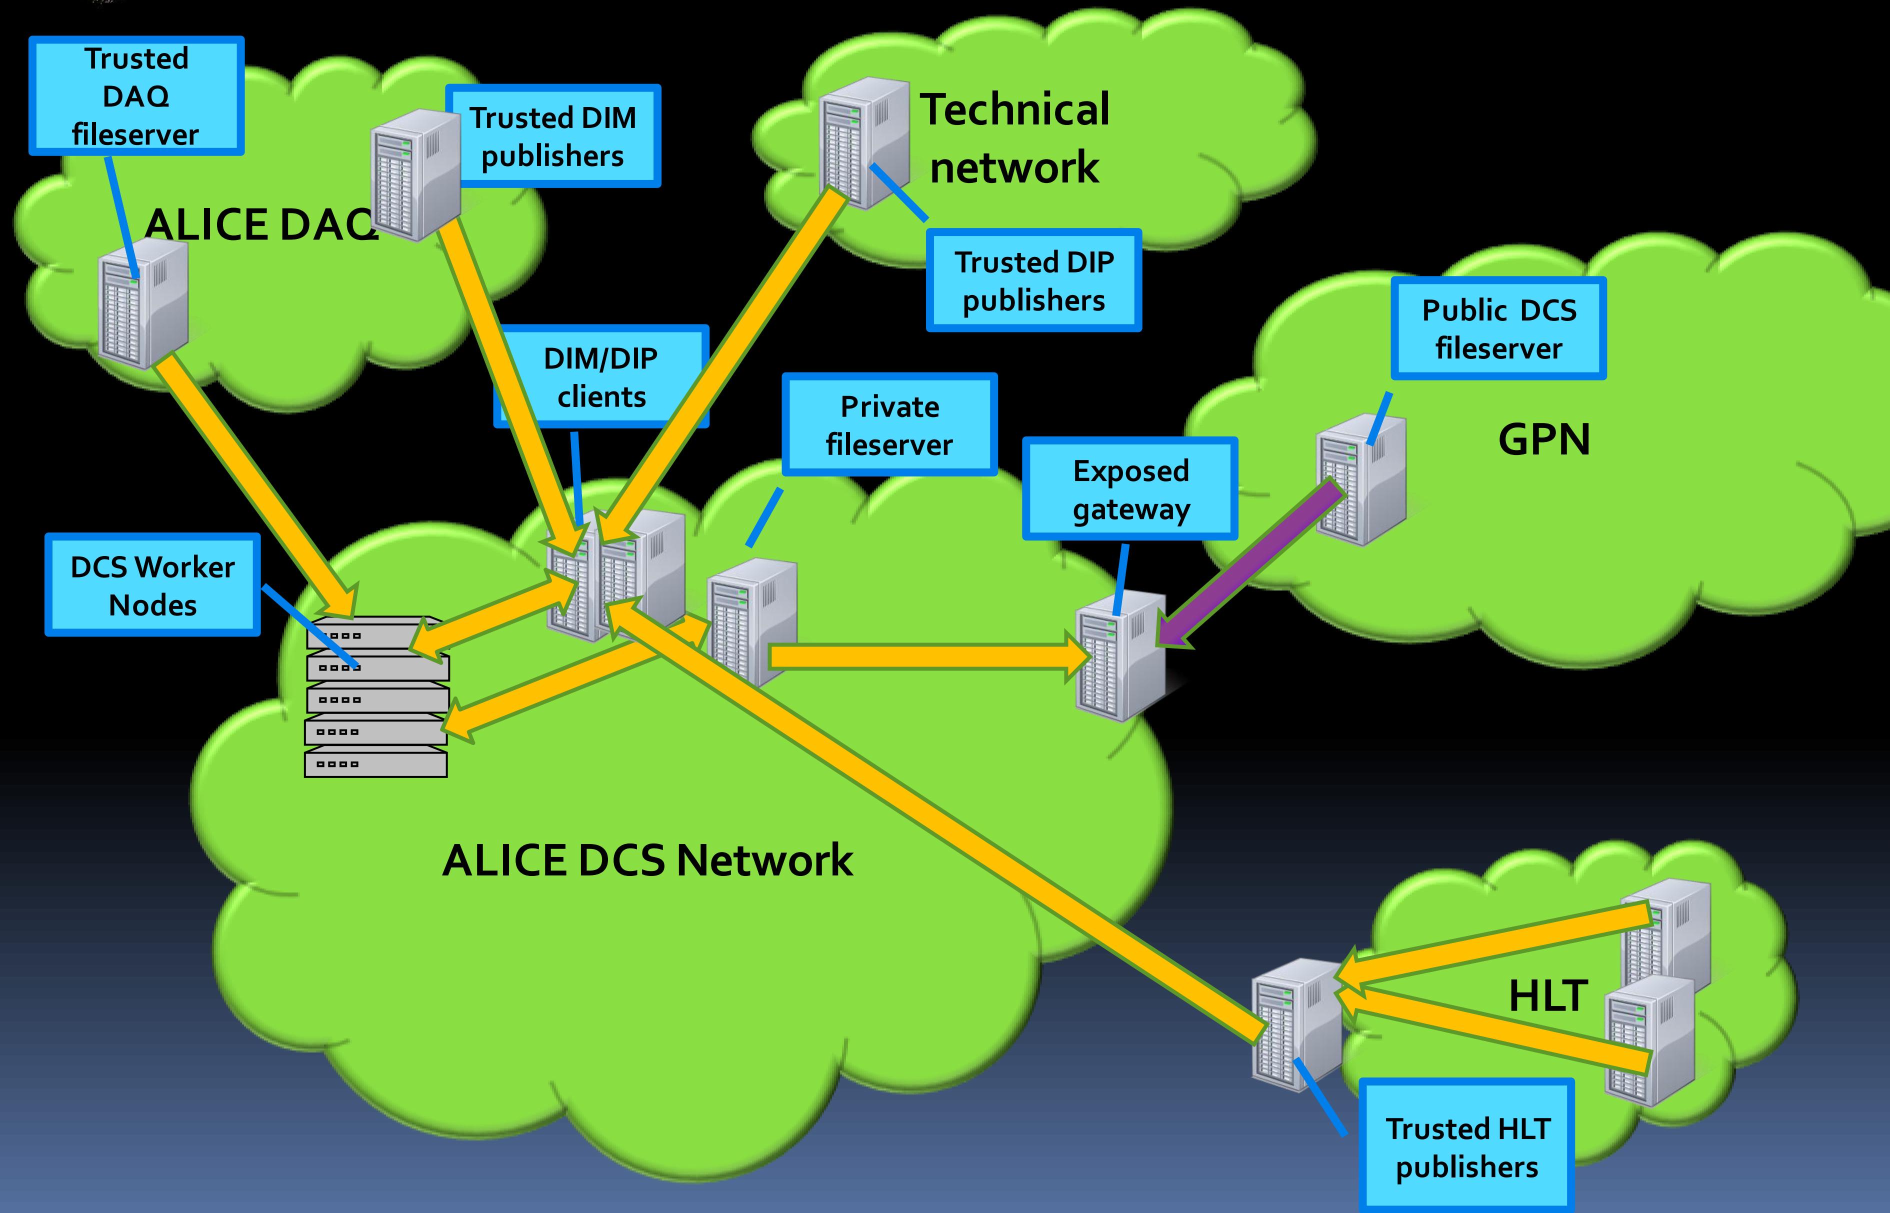

- Are people happy with this system?

- One example for all

From: XXXXXXXX [mailto:XXXXXXX@YYYYYYYY.ZZ]

Sent: Tuesday, February 1, 2011 11:03 PM

To: Peter Chochula

Subject: Putty

Hi Peter

Could you please install Putty on com001? I'd like to bypass this annoying upload procedure

Grazie

UUUUUUUUUUU

- Few more
  - Attempt to upload software via cooling station with embedded OS
  - Software embedded in the frontend calibration data
  - .....

  - We are facing a challenge here
  - ... and of course we follow all cases....
  - The most dangerous issues are critical last minute updates

DCS

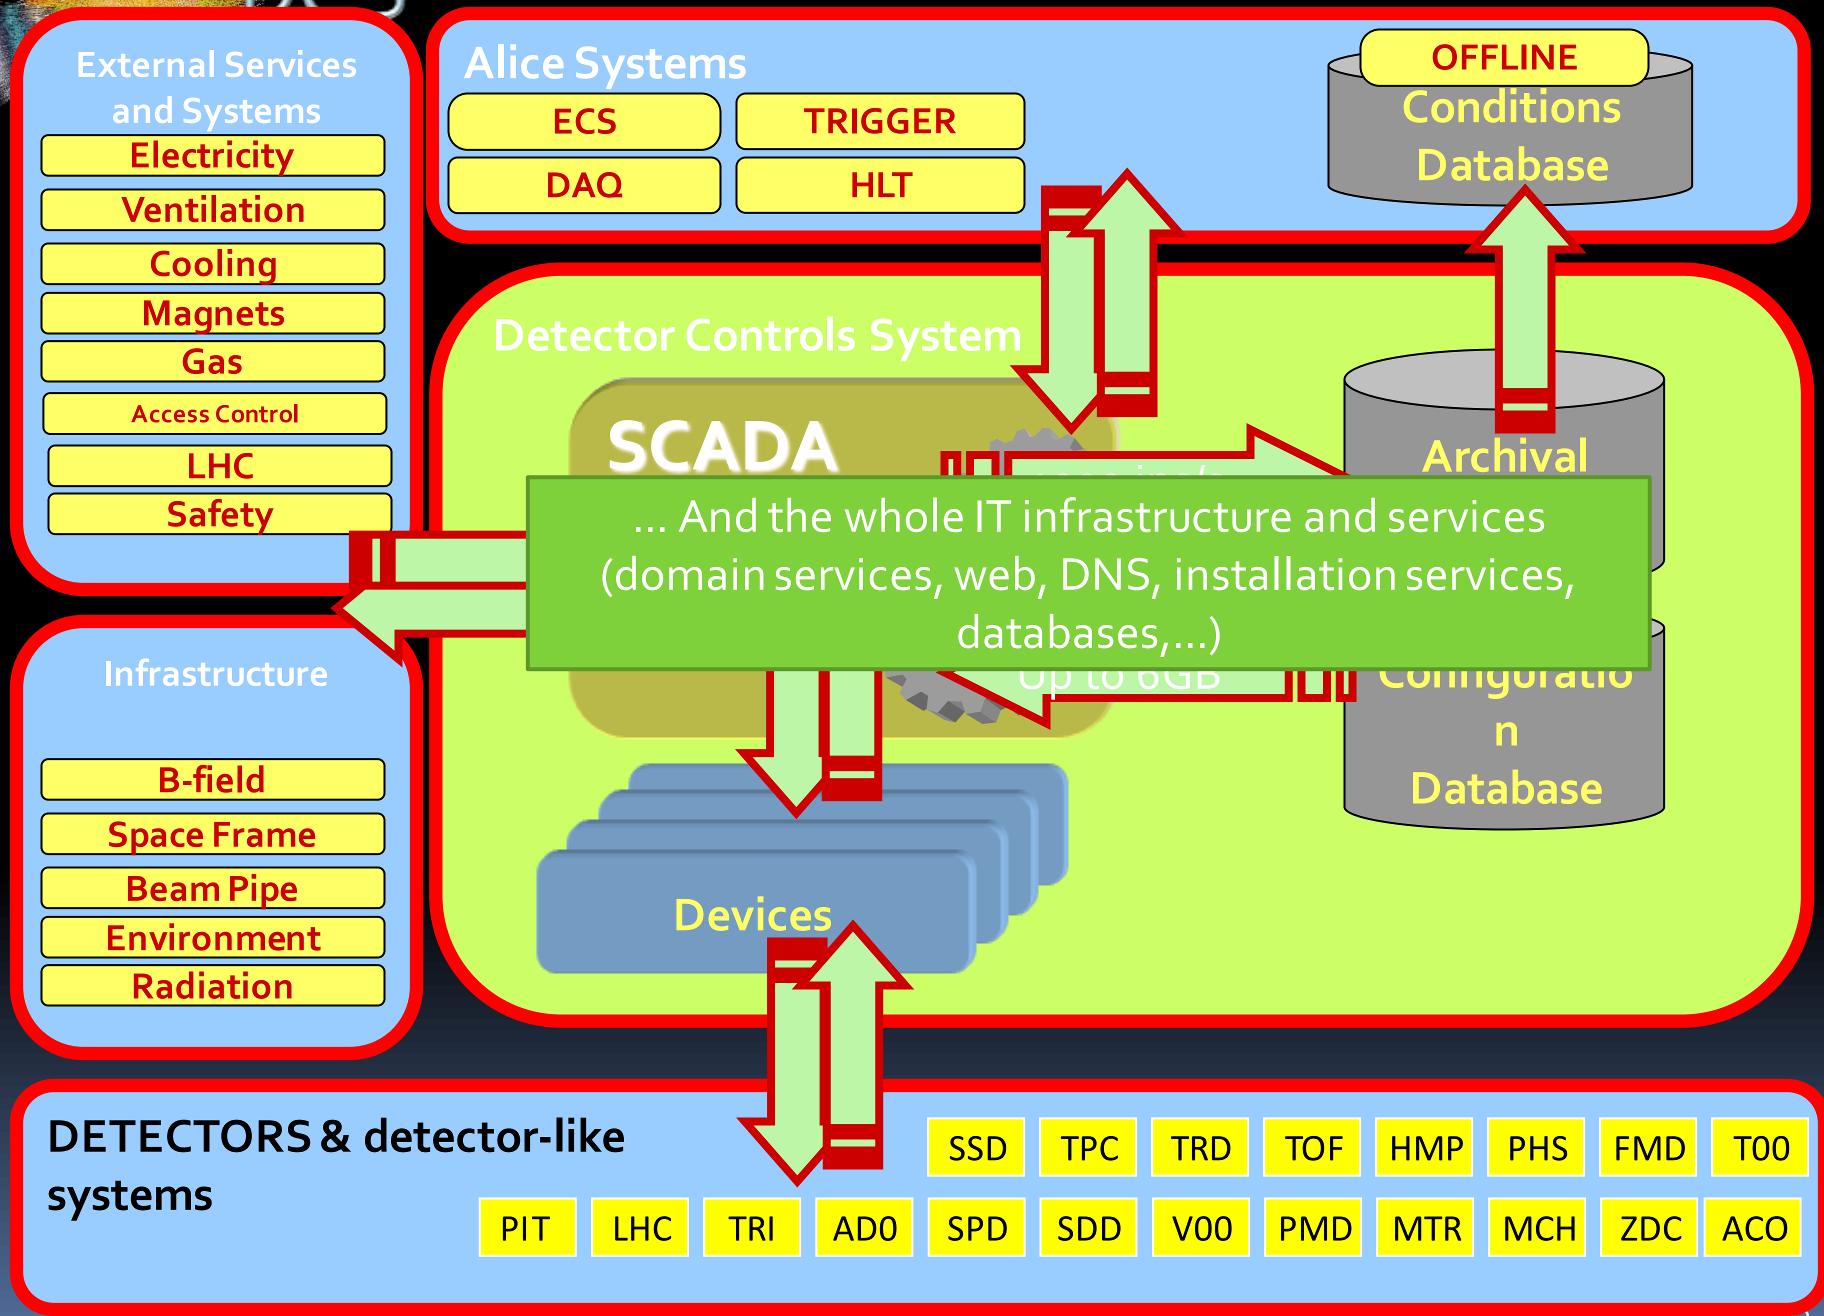
External Services and Systems
Electricity
Ventilation
Cooling
Magnets
Gas
Access Control
LHC
Safety
Alice Systems
ECS
TRIGGER
DAQ
HLT
OFFLINE
Conditions Database
Detector Controls System
SCADA
Archival
... And the whole IT infrastructure and services (domain services, web, DNS, installation services, databases,...)
Up to 6GB
Configuratio
n
Database
Infrastructure
B-field
Space Frame
Beam Pipe
Environment
Radiation
Devices
DETECTORS & detector-like systems
SSD
TPC
TRD
TOF
HMP
PHS
FMD
T00
PIT
LHC
TRI
AD0
SPD
SDD
V00
PMD
MTR
MCH
ZDC
ACO

Firewalls

- In the described complex environment firewalls are a must
  - Can be the firewalls easily deployed on controls computers?

Firewalls

- The firewalls cannot be installed on all devices
  - Majority of controls devices run embedded operating systems
    - PLC, front-end boards, oscilloscopes,…
  - The firewalls are MISSING or IMPOSSIBLE to install on them

Firewalls

- Are (simple) firewalls (simply) manageable on controls computers?

Firewalls

- There is no common firewall rule to be used
- The DCS communication involves many services, components and protocols
  - DNS, DHCP, WWW, NFS, DFS,
  - DIM, DIP, OPC, MODBUS, SSH,
  - ORACLE clients, MySQL clients
  - PVSS internal communication
- Efficient firewalls must be tuned per system

Firewalls

- The DCS configuration is not static
  - Evolution
  - Tuning (involves moving boards and devices across detectors)
  - Replacement of faulty components
- Each modification requires a setup of firewall rules by expert
  - Interventions can happen only during LHC access slots, with limited time for the actions
    - Can the few central admins be available 24/7?

System Complexity

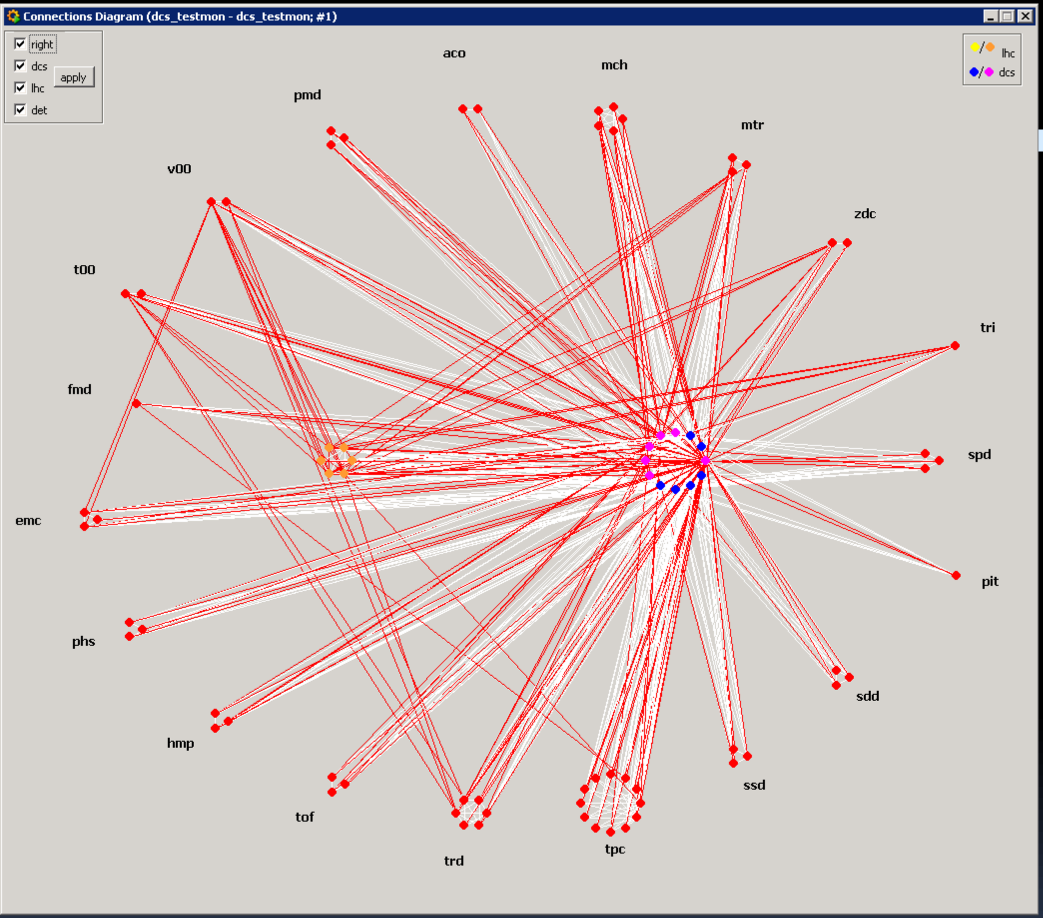

- Example of the cross-system connectivity as seen by monitoring tools
  - Red dots represent PVSS systems

Firewalls

- Firewalls must protect the system but should not prevent its functionality
  - Correct configuration of firewalls on all computers (which can run firewalls) is an administrative challenge
  - Simple firewalls are not manageable and sometimes dangerous
    - for example Windows firewall turns on full protection in case of domain connectivity loss
      - Nice feature for laptops
        - Killing factor for controls system which is running in emergency mode due to restricted connectivity
- And yes, most violent viruses attack the ports, which are vital for the DCS and cannot be closed…

Antivirus

- Antivirus is a must in such complex system
- But can they harm? Do we have resources for them?

Antivirus

- Controls systems were designed 10-15 years ago
  - Large portion of the electronics is obsolete (PCI cards, etc.) and requires obsolete (=slow) computers
- Commercial software is sometimes written inefficiently and takes a lot of resources without taking advantage of modern processors
  - Lack of multithreading forces the system to run on fast cores (i.e. Limited number of cores per CPU)

Antivirus

- Operational experience shows that fully operational antivirus will start interacting with the system preferably in critical periods like the End of Run
  - When systems produce conditions data (create large files)
  - When detectors change the conditions (communicate a lot)
    - adopt voltages as a reaction to beam mode change
    - Recovery from trips causing the EOR…

Antivirus and firewall finetuning

- Even a tuned antivirus typically shows on top 5 resource hungry processes

- CPU core affinity settings require huge effort
  - There are more than 2300 PVSS managers in ALICE DCS, 800 DIM servers, etc.
- The solutions are:
  - Run firewall and antivirus with very limited functionality
  - Run good firewalls and antivirus on the gates to the system

Software versions and updates

- It is a must to run the latest software with current updates and fixes
  - Is this possible?

Software versions and updates

- ALICE operates in 24/7 mode without interruption
- Short technical stops (4 days each 6 weeks) are not enough for large updates
  - DCS supervises the detector also without beams
  - DCS is needed for tests
- Large interventions are possible only during the long technical stops - around Christmas
- Deployment of updates requires testing, which can be done only on the real system
- Most commercial software excludes the use of modern systems (lack of 64 bit support)
- Front-end boards run older OS versions and cannot be easily updated
- ALICE deploys critical patches when operational conditions allow for it
  - Whole system is carefully patched during the long stops

Conclusions

- The cybersecurity importance is well understood in ALICE and is given high priorities
- The nature of a high energy physics experiment excludes a straightforward implementation of all desired features
  - Surprisingly , the commercial software is a significantly limiting factor here

- Implemented procedures and methods are gradually developing in ALICE
- The goal is to keep ALICE safe until 2013 (LHC long technical stop) and even safer afterwards