Il Tier-1 e la sua evoluzione: risorse opportunistiche, remote, commerciali

Luca dell'Agnello, Gaetano Maron
INFN - CNAF

Workshop CCR - June 14 2018

The Tier-1 at INFN-CNAF

- Started in 2003 as computing center for BaBar, CDF, Virgo and LHC experiments (ATLAS, CMS, LHCb, ALICE)
  - Nowadays provides services and resources to more than 30 scientific collaborations

  1.000 WNs , ~32.000 computing slots, ~340 kHS06
  - Also small (~33 TFlops) HPC cluster available
- 29 PB of storage on disk (~33 PB soon), ~48 PB on tape (~85 PB foreseen at the end of this year)
- Dedicated network channel (60 Gb/s) for LHC OPN + LHC ONE
  - 20 Gb/s reserved for LHC ONE
  - Ready for upgrade to (2x)100 Gb/s (before the end of Summer 2018)

The flood

The Tier-1 location

Transformers

Electrical room

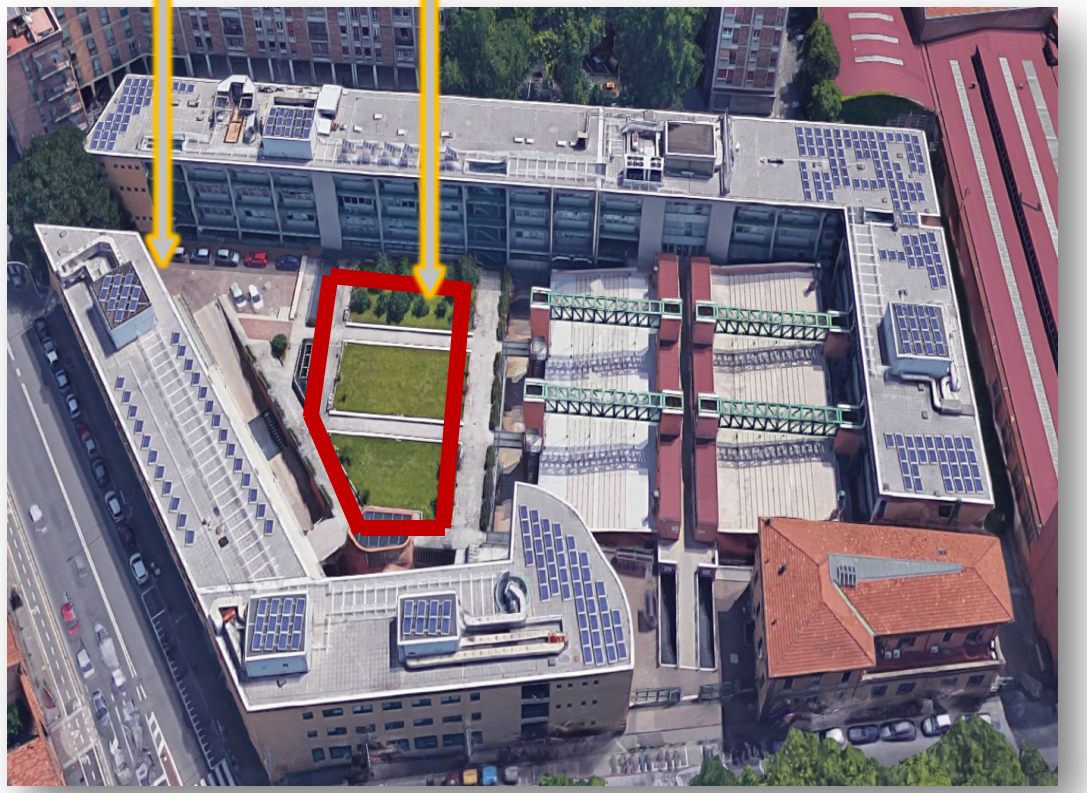

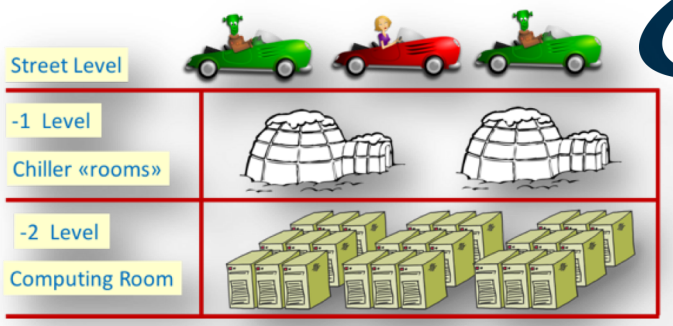

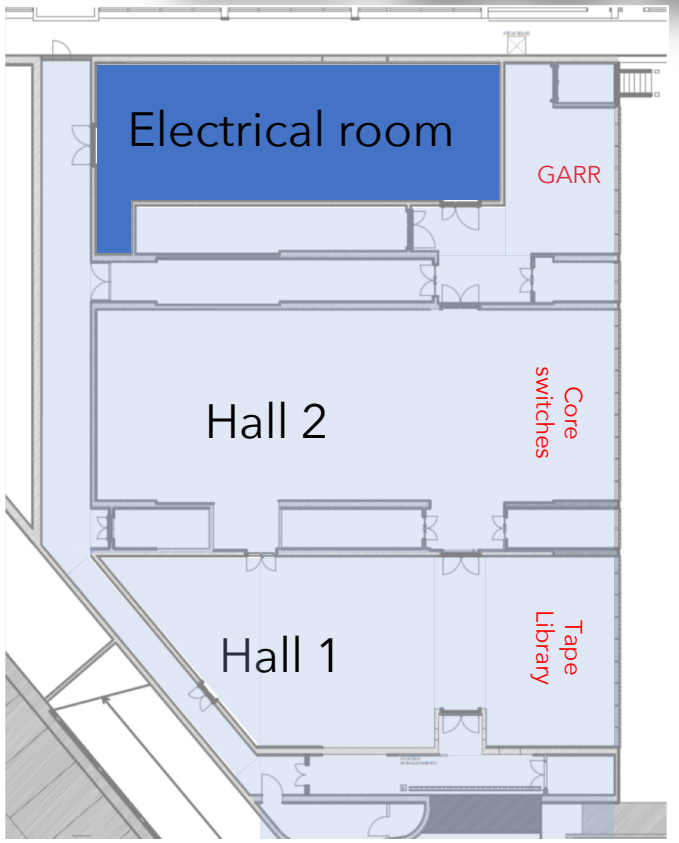

11/9: the flood

- The flood happened on November 9 early in the morning
  - Breaking of one of the main water pipelines in Bologna
  - Also the road near CNAF seriously damaged

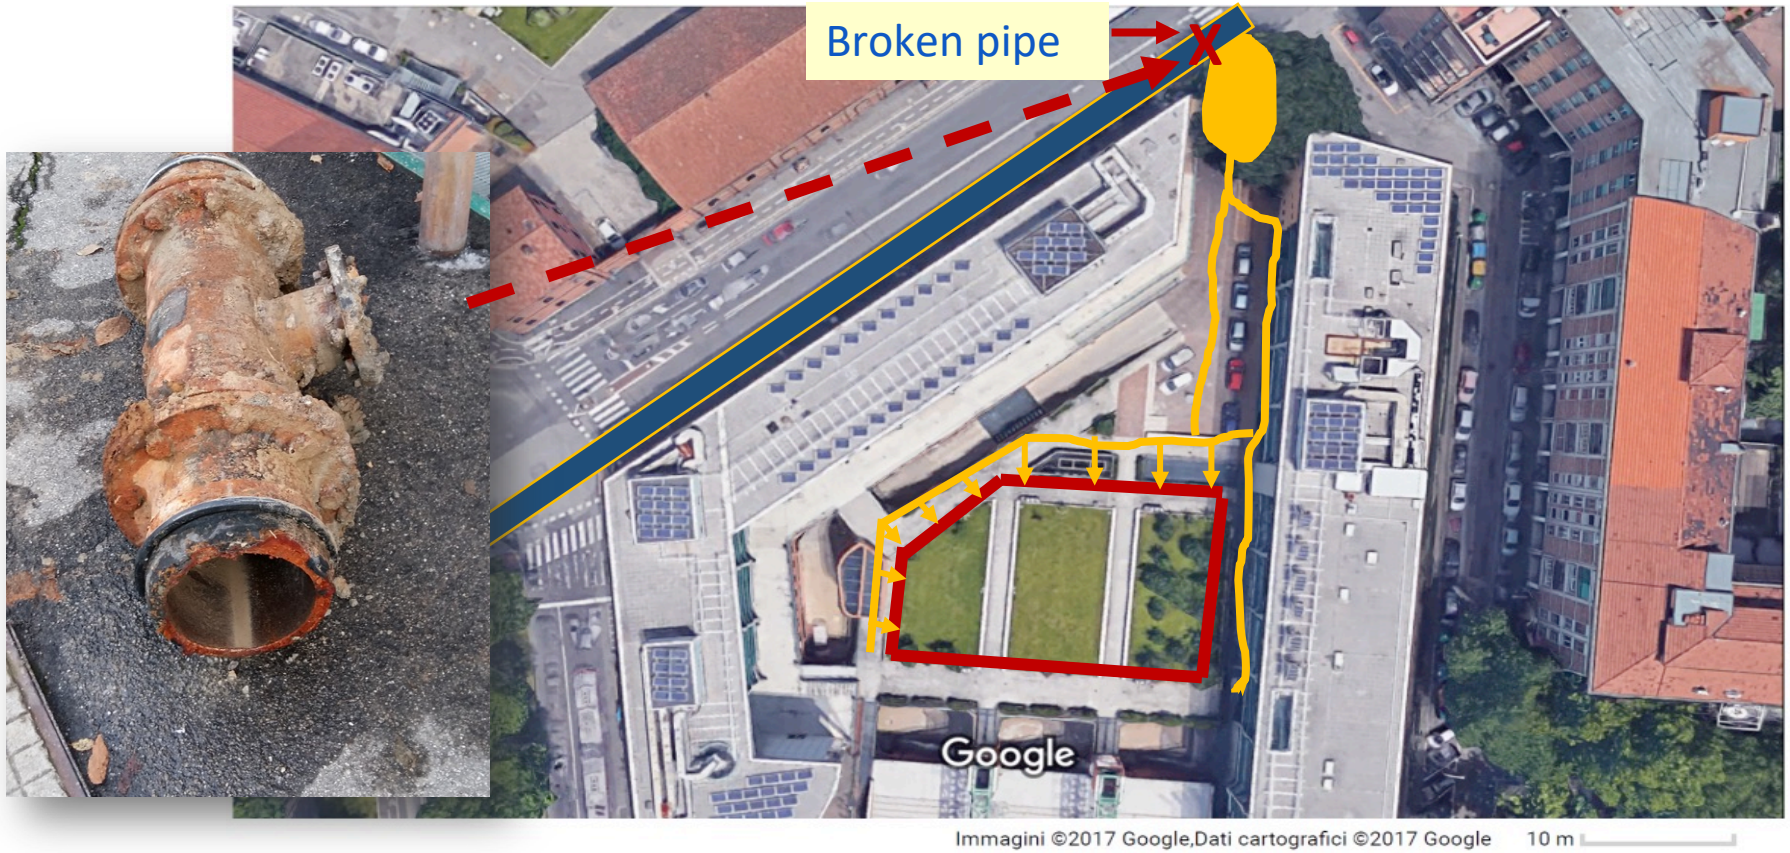

11/9: the flood

ovember 9 early in the morning

n water pipelines in Bologna

eriously damaged

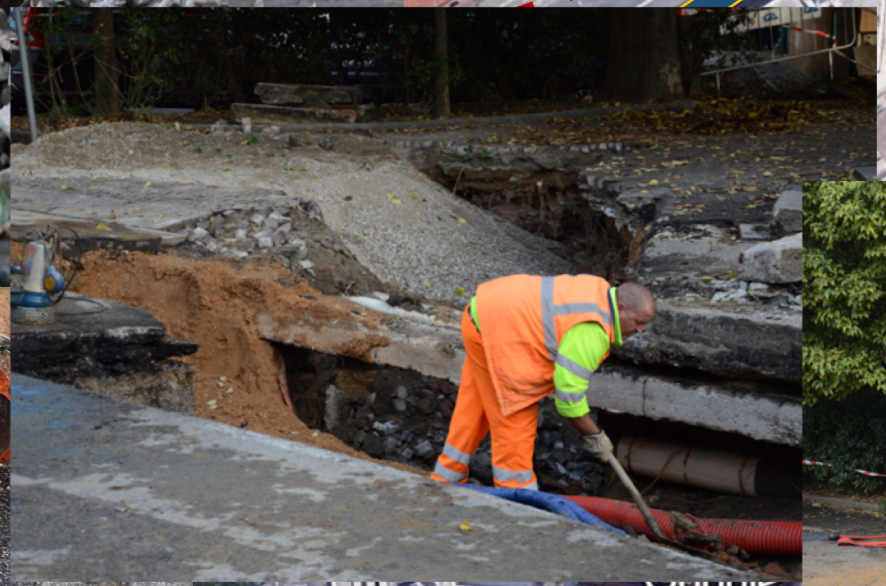

The Tier-1 entrance that morning

All Tier-1 doors are watertight
Height of water outside: 50 cm
Height of water inside: 10 cm (on floating floor)
for a total volume of ~500 m^3

After the flood

- The main issue was to recover the power center
  - Both 1.4 MW power lines compromised
- Temporary power line from the beginning
- First operations: data center dried over the first week-end and cleaned from dust and mud completed during the first week of December
- First line recovered before Xmas
  - (Small) UPS only from mid January
  - Full UPS + Diesel Engine only from 20/2
- Chillers + air conditioning in the IT halls from mid January
  - Only half chillers can be powered on with one line

Power Center configuration before the flood

Electrical Transformers (Energy Supplier)

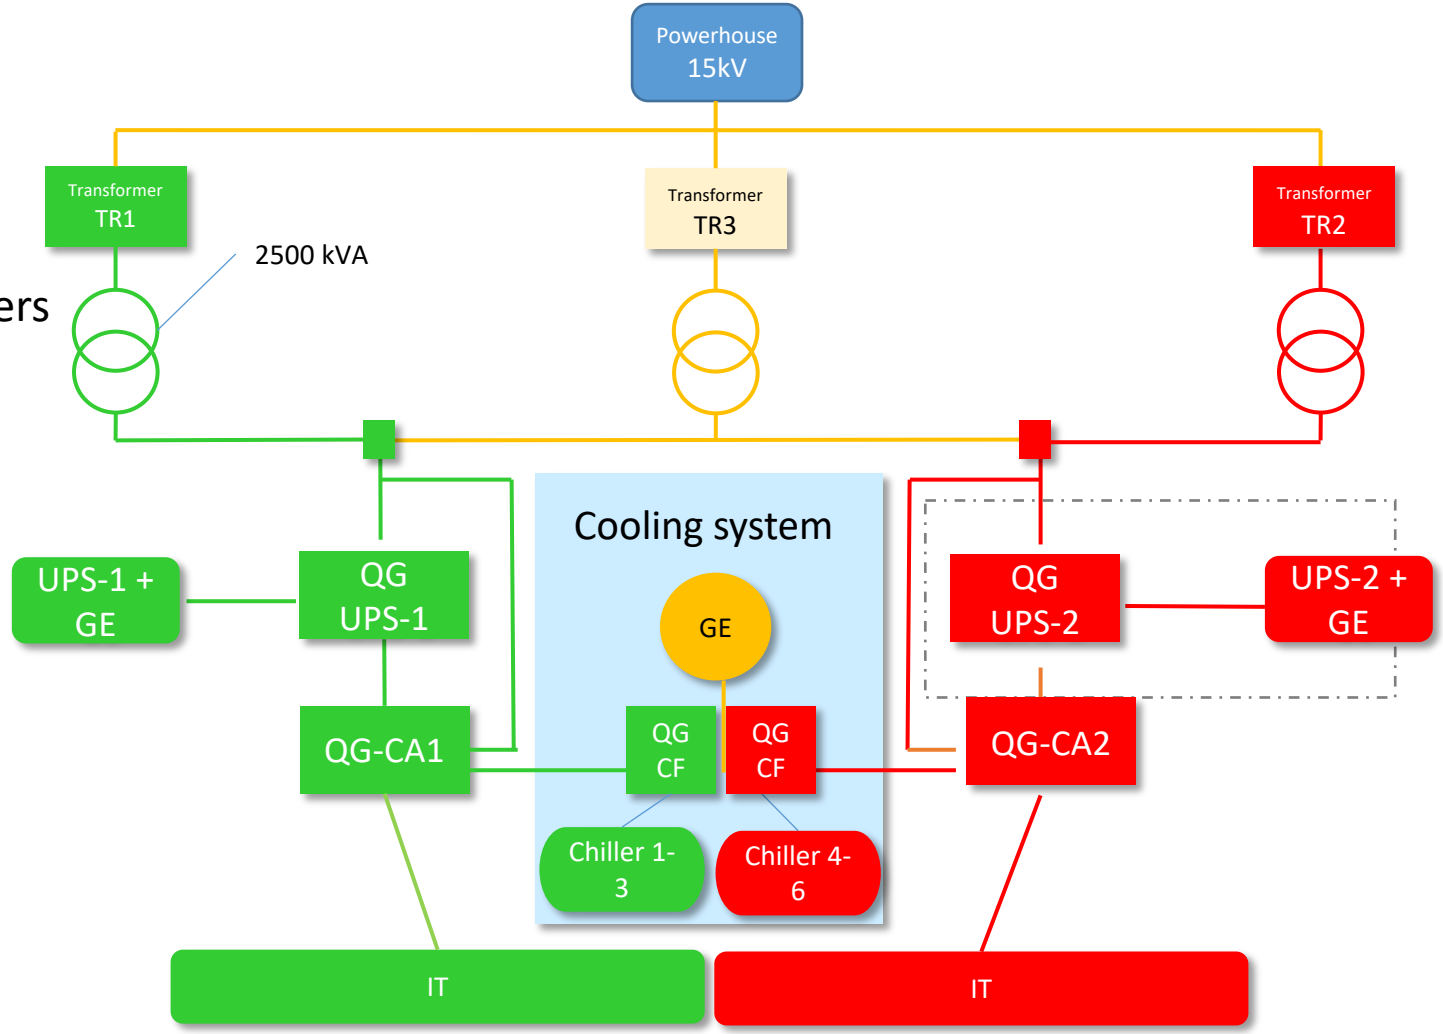

QG=Electric Panel
GE=Generator

Power Center after the flood

Electrical Transformers (Energy Supplier)

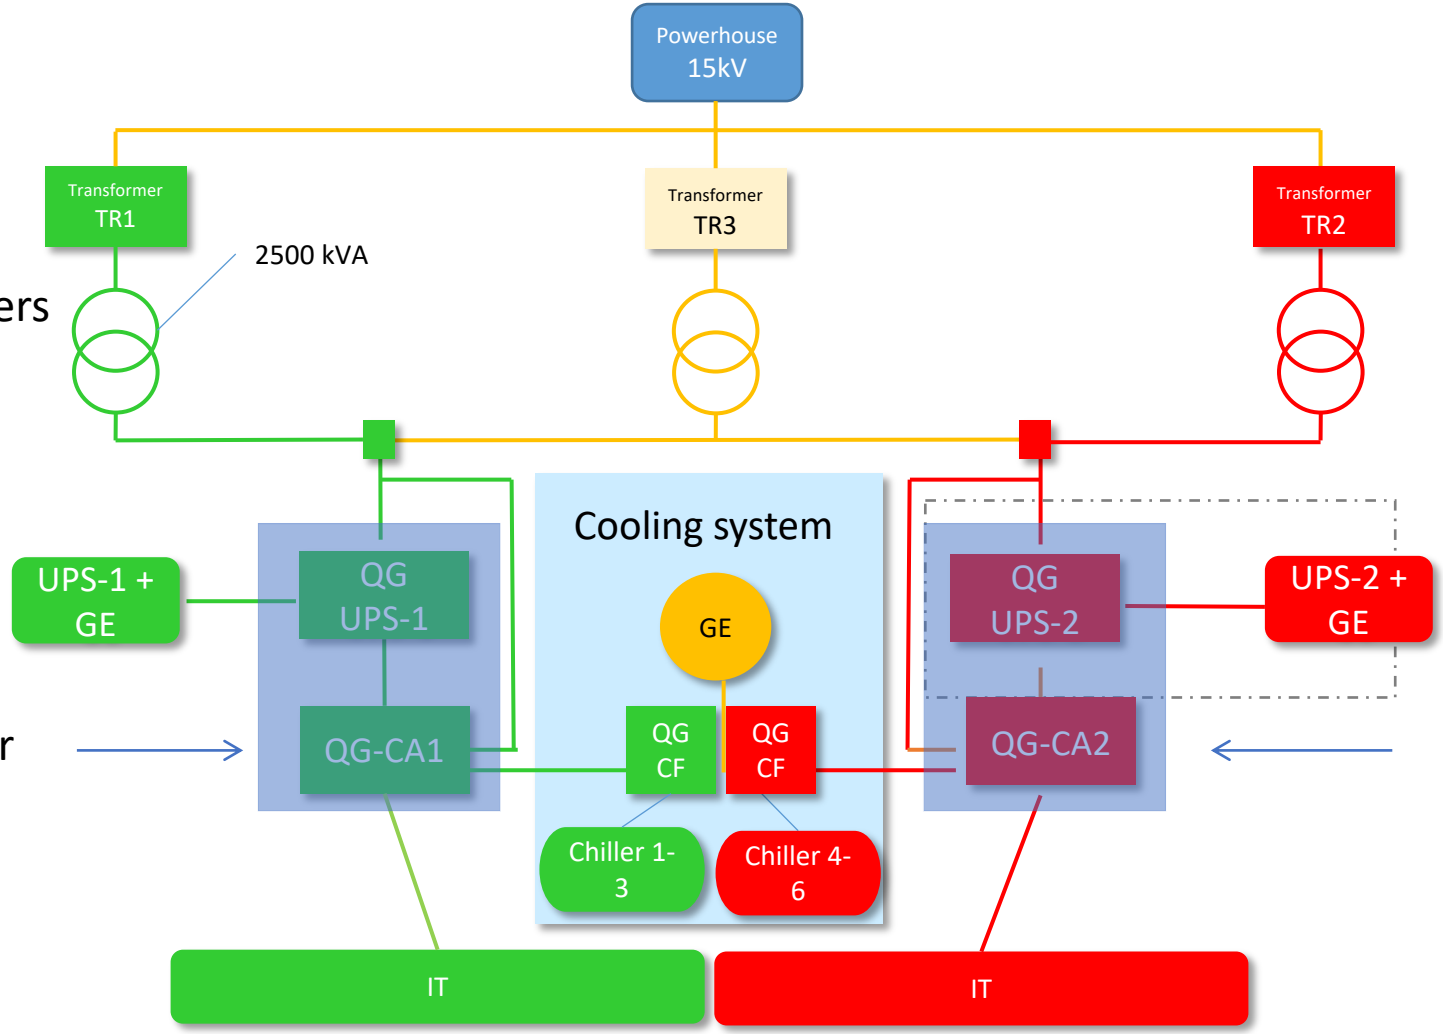

Damaged by water

Damaged by water

QG=Electric Panel
GE=Generator

Power Center after the flood

Electrical Transformers (Energy Supplier)

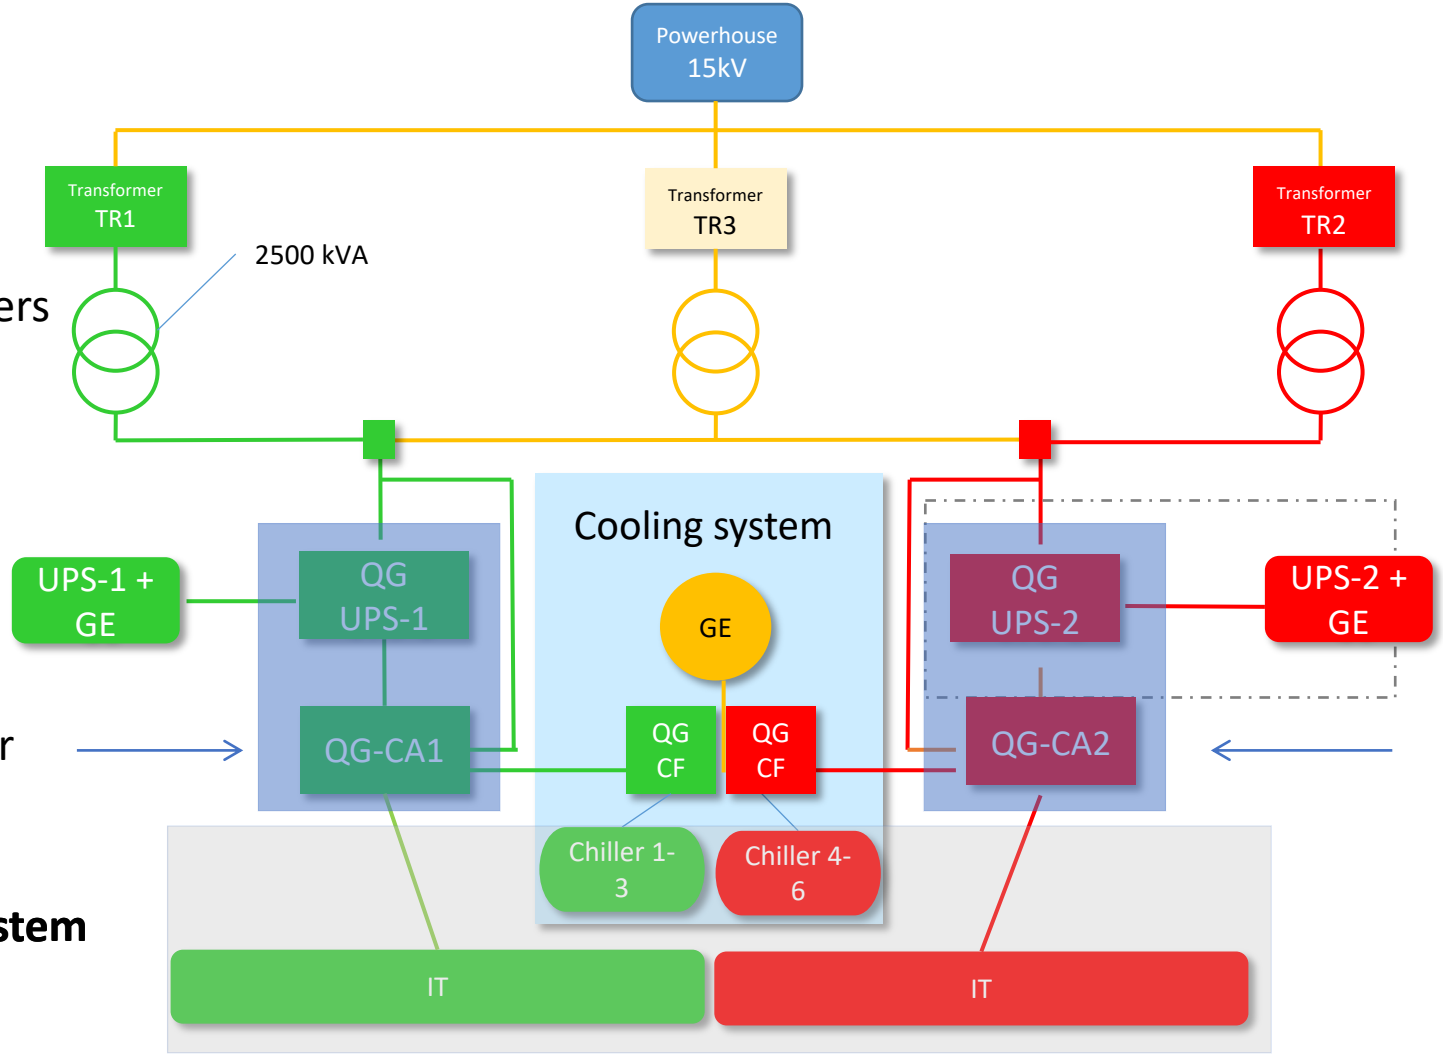

Damaged by water

Damaged by water

All the electrical system Affected !

QG=Electric Panel
GE=Generator

Present Power Center status

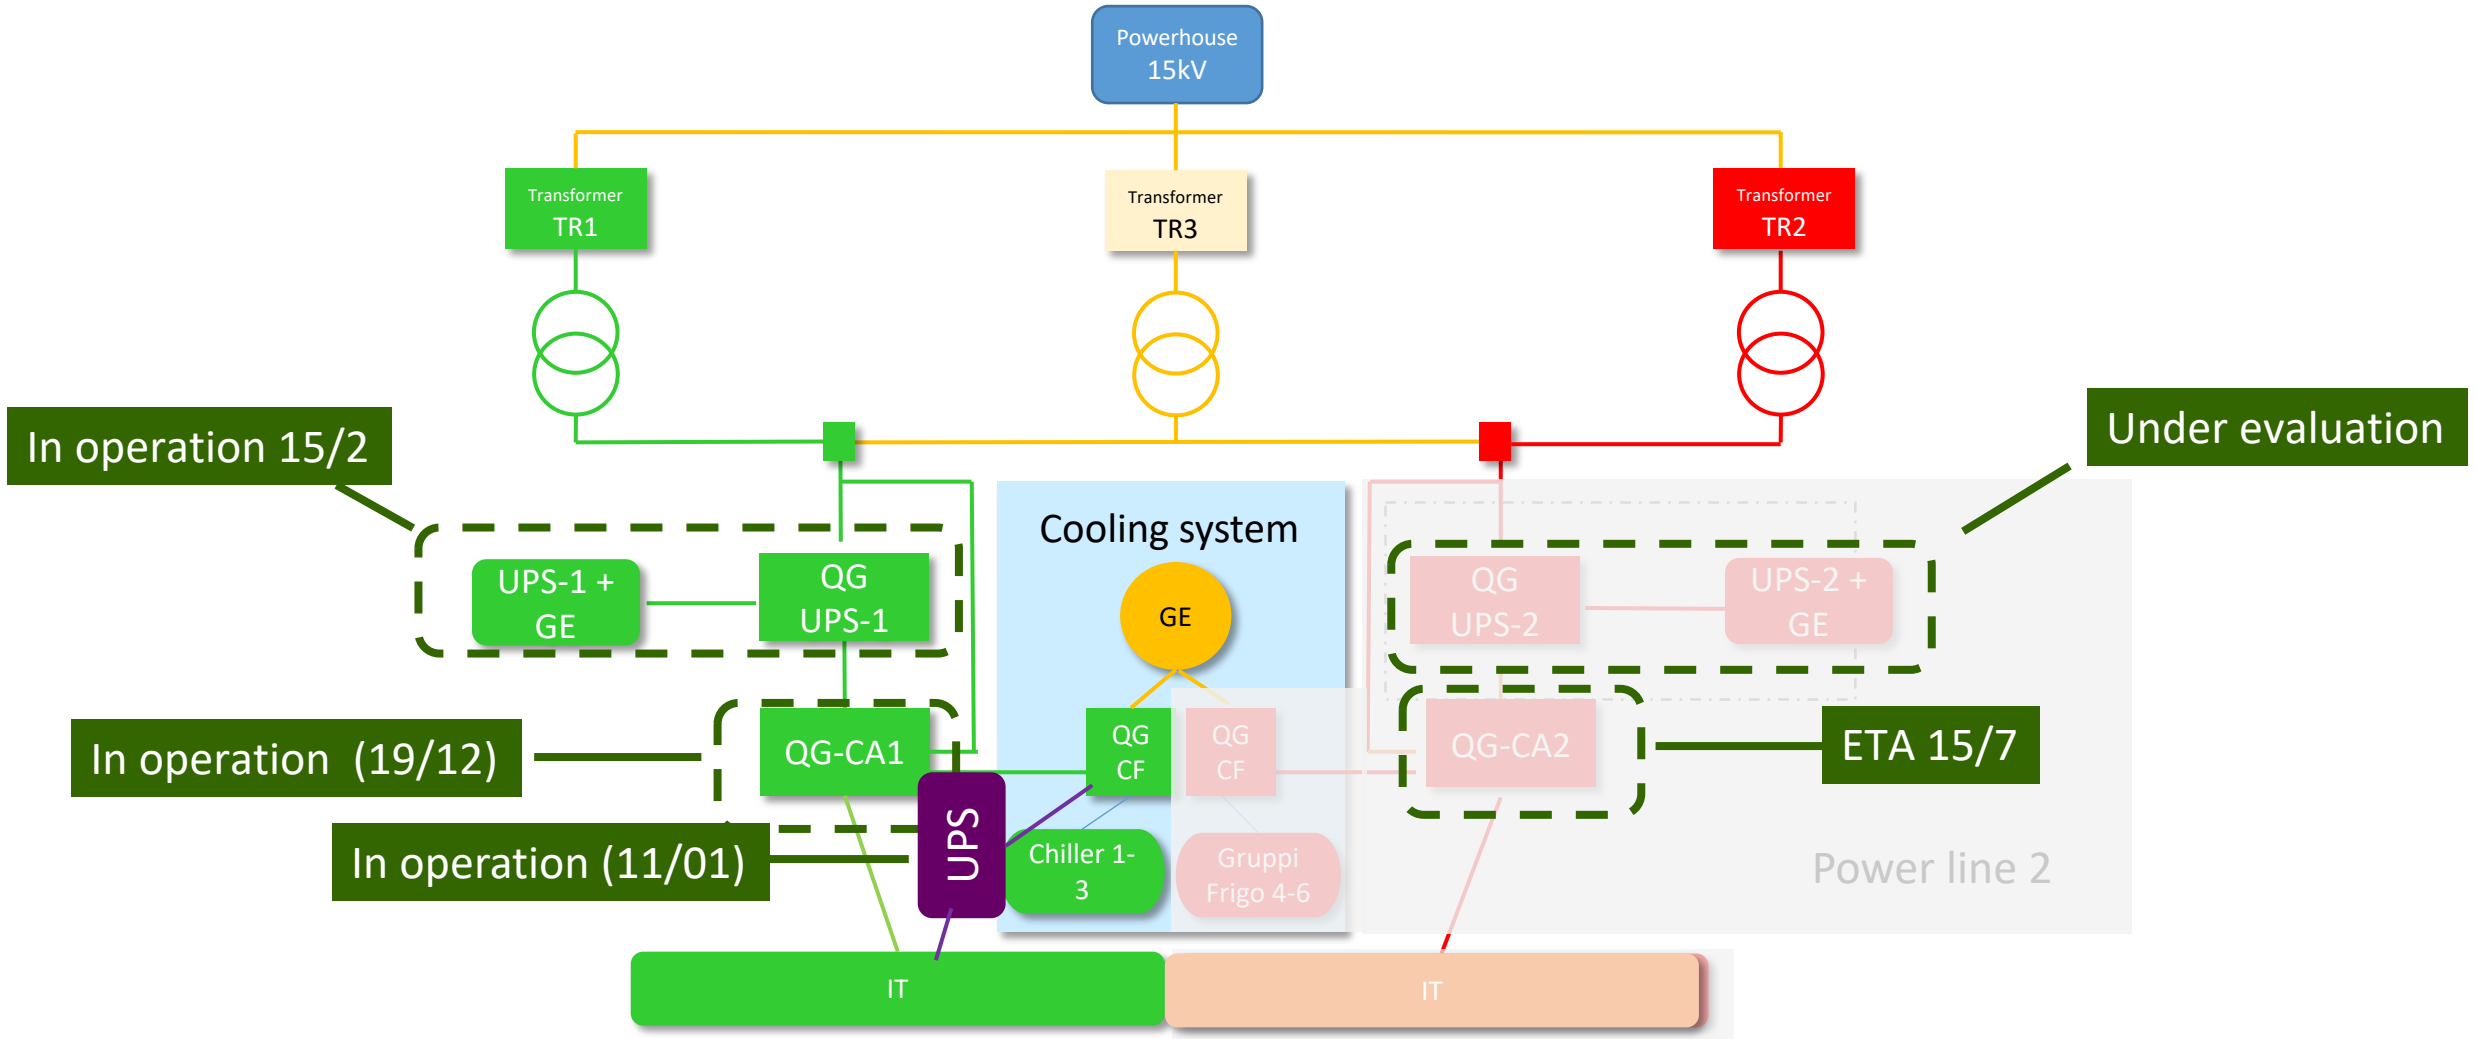

Damage to IT equipment

- The two lower units of all racks in the IT halls were submerged
  - Including the two lowest rows of tapes in the library
  - All storage systems involved
- The 3 Core Switch/Routers and the General IP Router were installed above the 3rd rack unit (safe for few centimeters)
- In parallel with the recovery of the power system, various activities done to recover wet IT equipment
  - Cleaning and drying disks, servers, switches (using oven when appropriate)
  - Damaged IT components ordered and replaced (with a variable timeline depending on the brand….)

Damage to IT equipment: the list

- Computing farm
  - ~34 kHS06 are now lost (~14% of the total capacity)
  - No special action taken (replaced)
- Library and HSM system
  - 1 drive damaged
  - Several non critical components
  - Library recertified in January
  - 4 TSM-HSM servers
- Tapes
  - 136 tapes damaged
  - Tapes being recovered in lab
    - 40 +22 tapes recovered
    - 1 tape partially recovered (LHCb)
    - 3 (CMS) + 3 (LHCb) tapes undergoing second round
    - 6 tapes (CMS) to be recovered
- Nearly all storage disk systems involved
  - 11 DDN JBODs (2 for CMS)
    - *RAID parity affected*
  - 2 Huawei JBODs (non-LHC experiments)
  - 2 Dell JBODs including controllers
  - 4 disk-servers

System	PB	JBODs	Involved experiments
Huawei	3.4	2	All CSN2 and 3 experiments excepting AMS, Darkside e Virgo
Dell	2.2	2	Darkside and Virgo
DDN 1,2	1.8	*4*	ATLAS, Alice and LHCb
DDN 8	2.7	*2*	LHCb
DDN 9	3.8	2	CMS
DDN 10, 11	10	3+2	ATLAS, Alice and AMS
Total	**23.9**	**9**	

Storage recovery

- Replacement components ordered only for systems under support in 2018
  - DDN8 (LHCb) to be phased out in Q1 2018
- Moreover, some components not available for bulk replacement for old systems
  - i.e. disks for DDN8 (LHCb) and DDN9 (CMS)
  - Other older systems (DDN1, DDN2) could be repaired with spare parts we had in house
- Need to move LHCb data to a new storage systems ASAP (lost parity)
- Key element was the installation of 2017 tender storage
  - Installed in January but still not accepted ☹
  - We used it anyway to move LHCb data from the damaged storage system
  - Later on, “good” disks from DDN8 used to replace wet disks of DDN9 (CMS)
- Unfortunately we could recover only 1/3 of data on Huawei system (astro-particle experiments)
  - ~2.2 PB of data lost (retransferred or regenerated)
  - We suspect an erroneous strategy from the support

Network (1)

- For the installation of the 2017 storage tender we had to upgrade our network infrastructure to support 100G Ethernet connections
- The core switches (2x Cisco Nexus 9516) upgraded in December
  - New Fabric modules and each core expanded with 32 x 100G Ethernet port modules

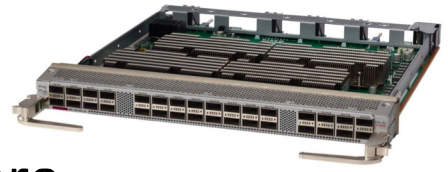

- Disk Servers of the 2017 tender are connected at 2x100Gb/s to the core switches
  - Most probably also disk servers of the new tender (next week!)
- 100 Gb/s connections needed also for DCI to CINECA (remote farm extension)
  - 4x100Gb/s Ethernet extension (upgradable to 12x100Gb/s)
- Scheduled upgrade of OPN/ONE access from 6x10Gb to 2x100Gb– (Q1-Q2 2018)

Network (2)

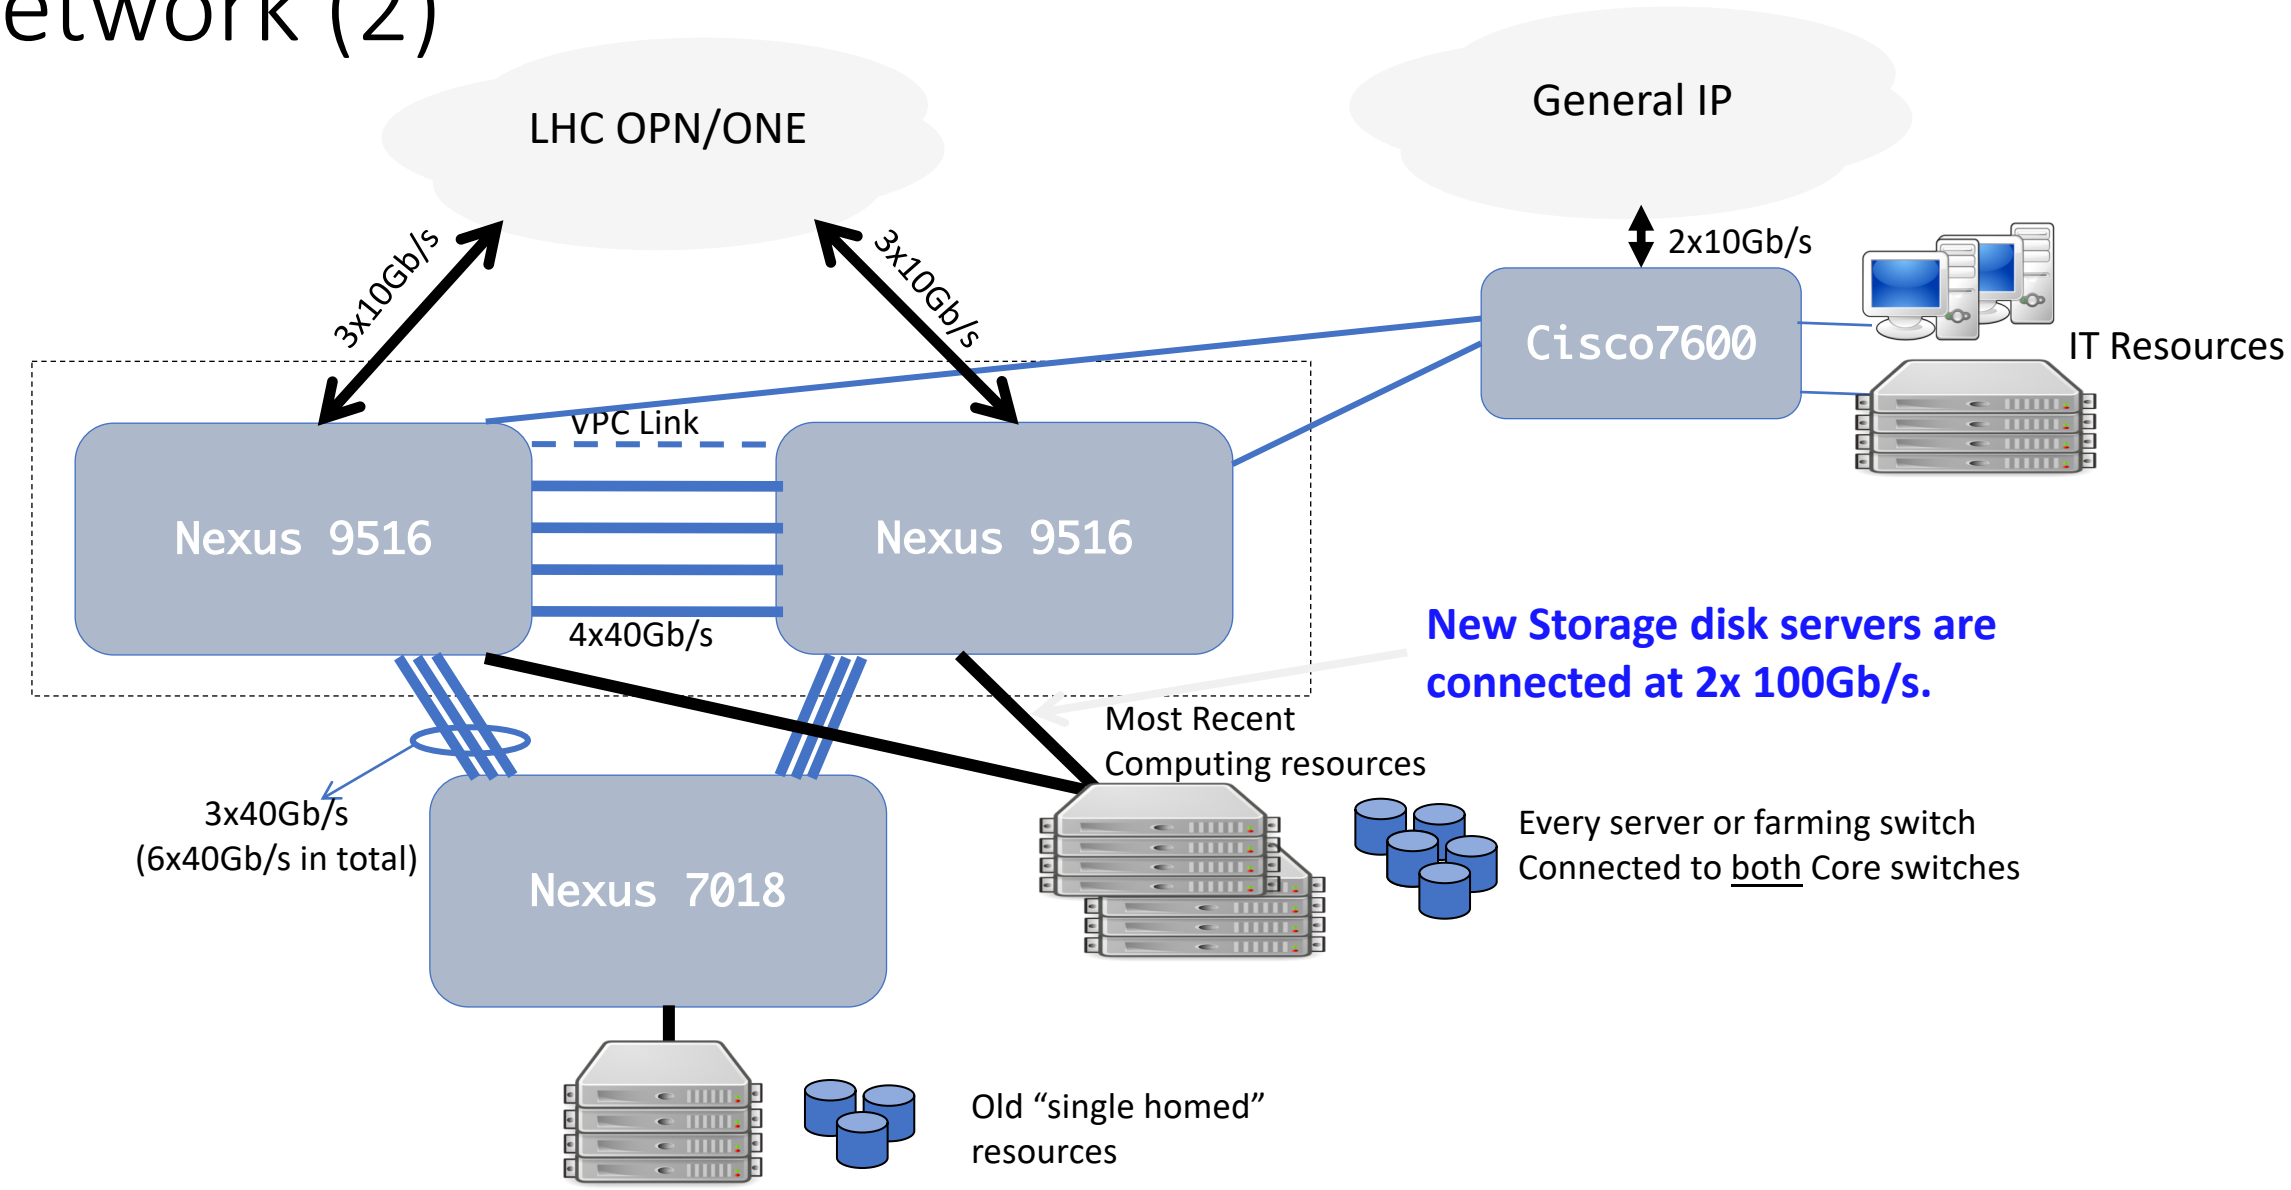
LHC OPN/ONE
General IP
3x10Gb/s
3x10Gb/s
2x10Gb/s
Cisco7600
IT Resources
VPC Link
Nexus 9516
Nexus 9516
4x40Gb/s
New Storage disk servers are connected at 2x 100Gb/s.
Most Recent Computing resources
3x40Gb/s
(6x40Gb/s in total)
Nexus 7018
Every server or farming switch
Connected to both Core switches
Old “single homed” resources

Farm recovery

- During February we started reopening the services
  - LSF masters, CEs, squids etc…
- Not all experiments at the same time (depending on storage availability)
- Performed upgrade of WNs
  - Middleware, security patches (i.e. meltdown etc..)
- Only part of the local farm powered on (only 3 chillers in production)
  - ~150 kHS06 (out of ~200kHS06 available)
- But exploiting the CNAF farm elastic extension to provide more computing power
  - Remote farm partition in Bari-RECAS (~22 kHS06)
  - Remote extension farm (~ 180 kHS06) at CINECA – In production since March

Tier-1 remote extensions

Various types of extensions…

- Static allocation of remote resources
  - Bari-RECAS (2017-)
  - CINECA (2018 -)
- Resources on commercial clouds
  - Aruba (2015)
  - Azure (2017)
  - HNSciCloud (2018)
- Opportunistic on HPC
  - Scheduled test on CINECA HPC resources

Farm remote extensions (1)

- Some functional tests on cloud providers (Aruba, Azure)
  - No cache, xrootd access
- In 2017 ~13% of CPU pledged resources to WLCG experiments located in Bari-RECAS data center
  - Transparent access for WLCG experiments
    - CNAF CEs and LSF as entry-point
    - Auxiliary services (i.e. squids) in Bari
  - Similar to CERN/Wigner extension
  - 20 Gbps VPN provided by GARR
    - All traffic with farm in Bari routed via CNAF
  - Disk cache provided via GPFS-AFM
    - “Transparent” extension of CNAF GPFS

Opportunistic computing on Aruba (1)

- One of the main Italian commercial resource providers
  - Web, host, mail, cloud, …
  - Main datacenters in Arezzo and near Bergamo
- Small scale test
  - 10x8 cores VM (160 GHz) managed by VMWare
- Use of idle CPU cycles
  - When a customer requires a resource we are using, the CPU clock speed of “our” VMs is decreased to a few MHz (not destroyed!)
- Only CMS mcore jobs tested
  - No storage for data on site: remote data access via Xrootd
  - Use of GPN (no dedicated NREN infrastructure)

Opportunistic computing on Aruba (2)

- Developed in house solution, dynfarm
  - Authenticates connection requests coming from remote hosts and delivers the information needed to creates a VPN tunnel
    - Communication enabled only through remote WNs and local CEs, LSF and argus
    - All other traffic goes through its default route
- Shared file system access through a R/O GPFS cache (AFM, see later for details)
- Job efficiency (CPT/WCT) depends on type of job
  - Very good for MC
  - Low on average (0.49 vs. 0.80)

Data access in Bari-RECAS

- GPFS AFM
  - A cache providing geographic replica of a file system
  - Manages RW access to cache
- Two sides
  - Home - where the information lives
  - Cache
  - Data written to the cache is copied back to home as quickly as possible
  - Data is copied to the cache when requested
- AFM configured as RO for Bari-ReCaS
  - ~400 TB of cache vs. ~11 PB of data
- Several tunings and reconfigurations required!
- In any case decided to avoid submission of high throughput jobs in Bari (possible for Atlas)
- Alice jobs access data directly through XrootD

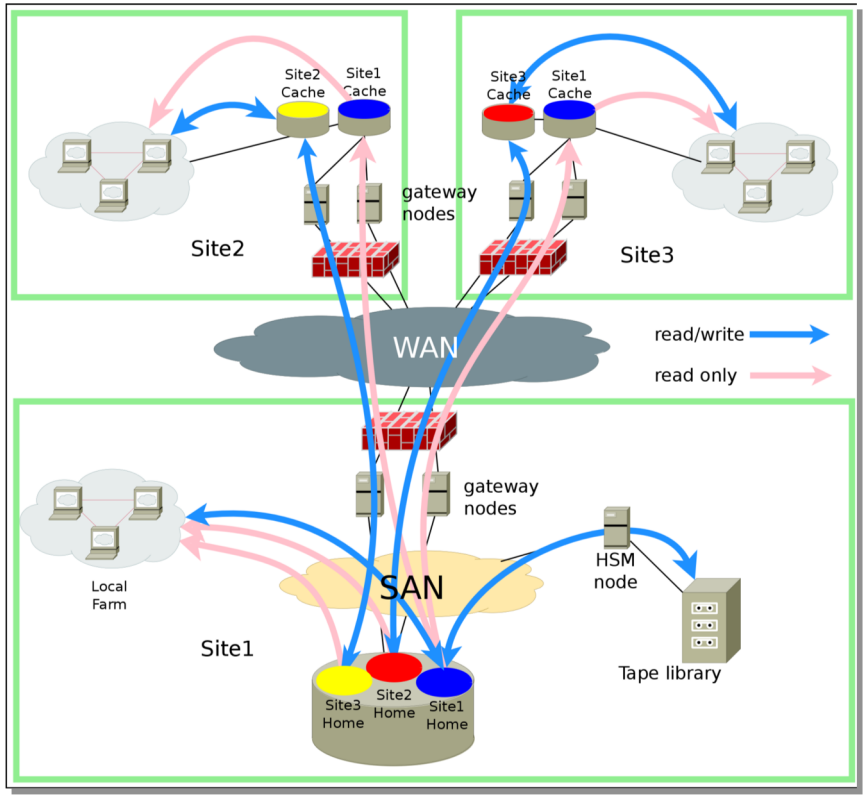

Farm remote extensions (2)

- In 2018 ~180 kHS06 provided by CINECA
  - CINECA, located in Bologna too, is the Italian supercomputing center (~15 Km far from CNAF)
  - 216 WNs (10 Gbit connection to rack switch and then 4x40 to router aggregator) managed by LSF@T1
- Dedicated fiber directly connecting Tier-1 core switches to our aggregation router at CINECA
  - 500 Gbps (upgradable to 1.2 Tbps) on a single fiber couple via Infinera DCI
- No disk cache, direct access to CNAF storage
  - Quasi-LAN situation (RTT: 0.48 ms vs. 0.28 ms on LAN)
- In production since March
  - Need to disentangle effects from migration to CentOs7, singularity etc… to have a definitive assessment on efficiency

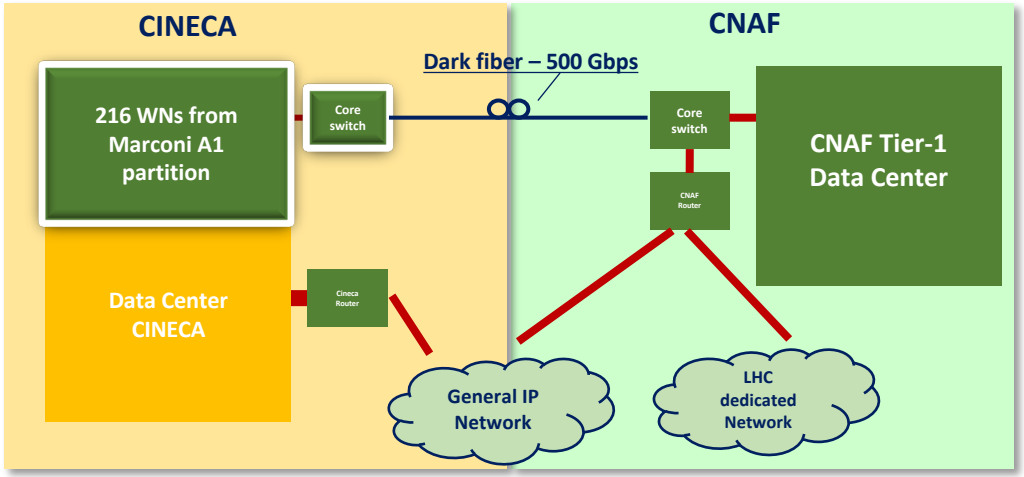

DCI (Software Defined – WAN)

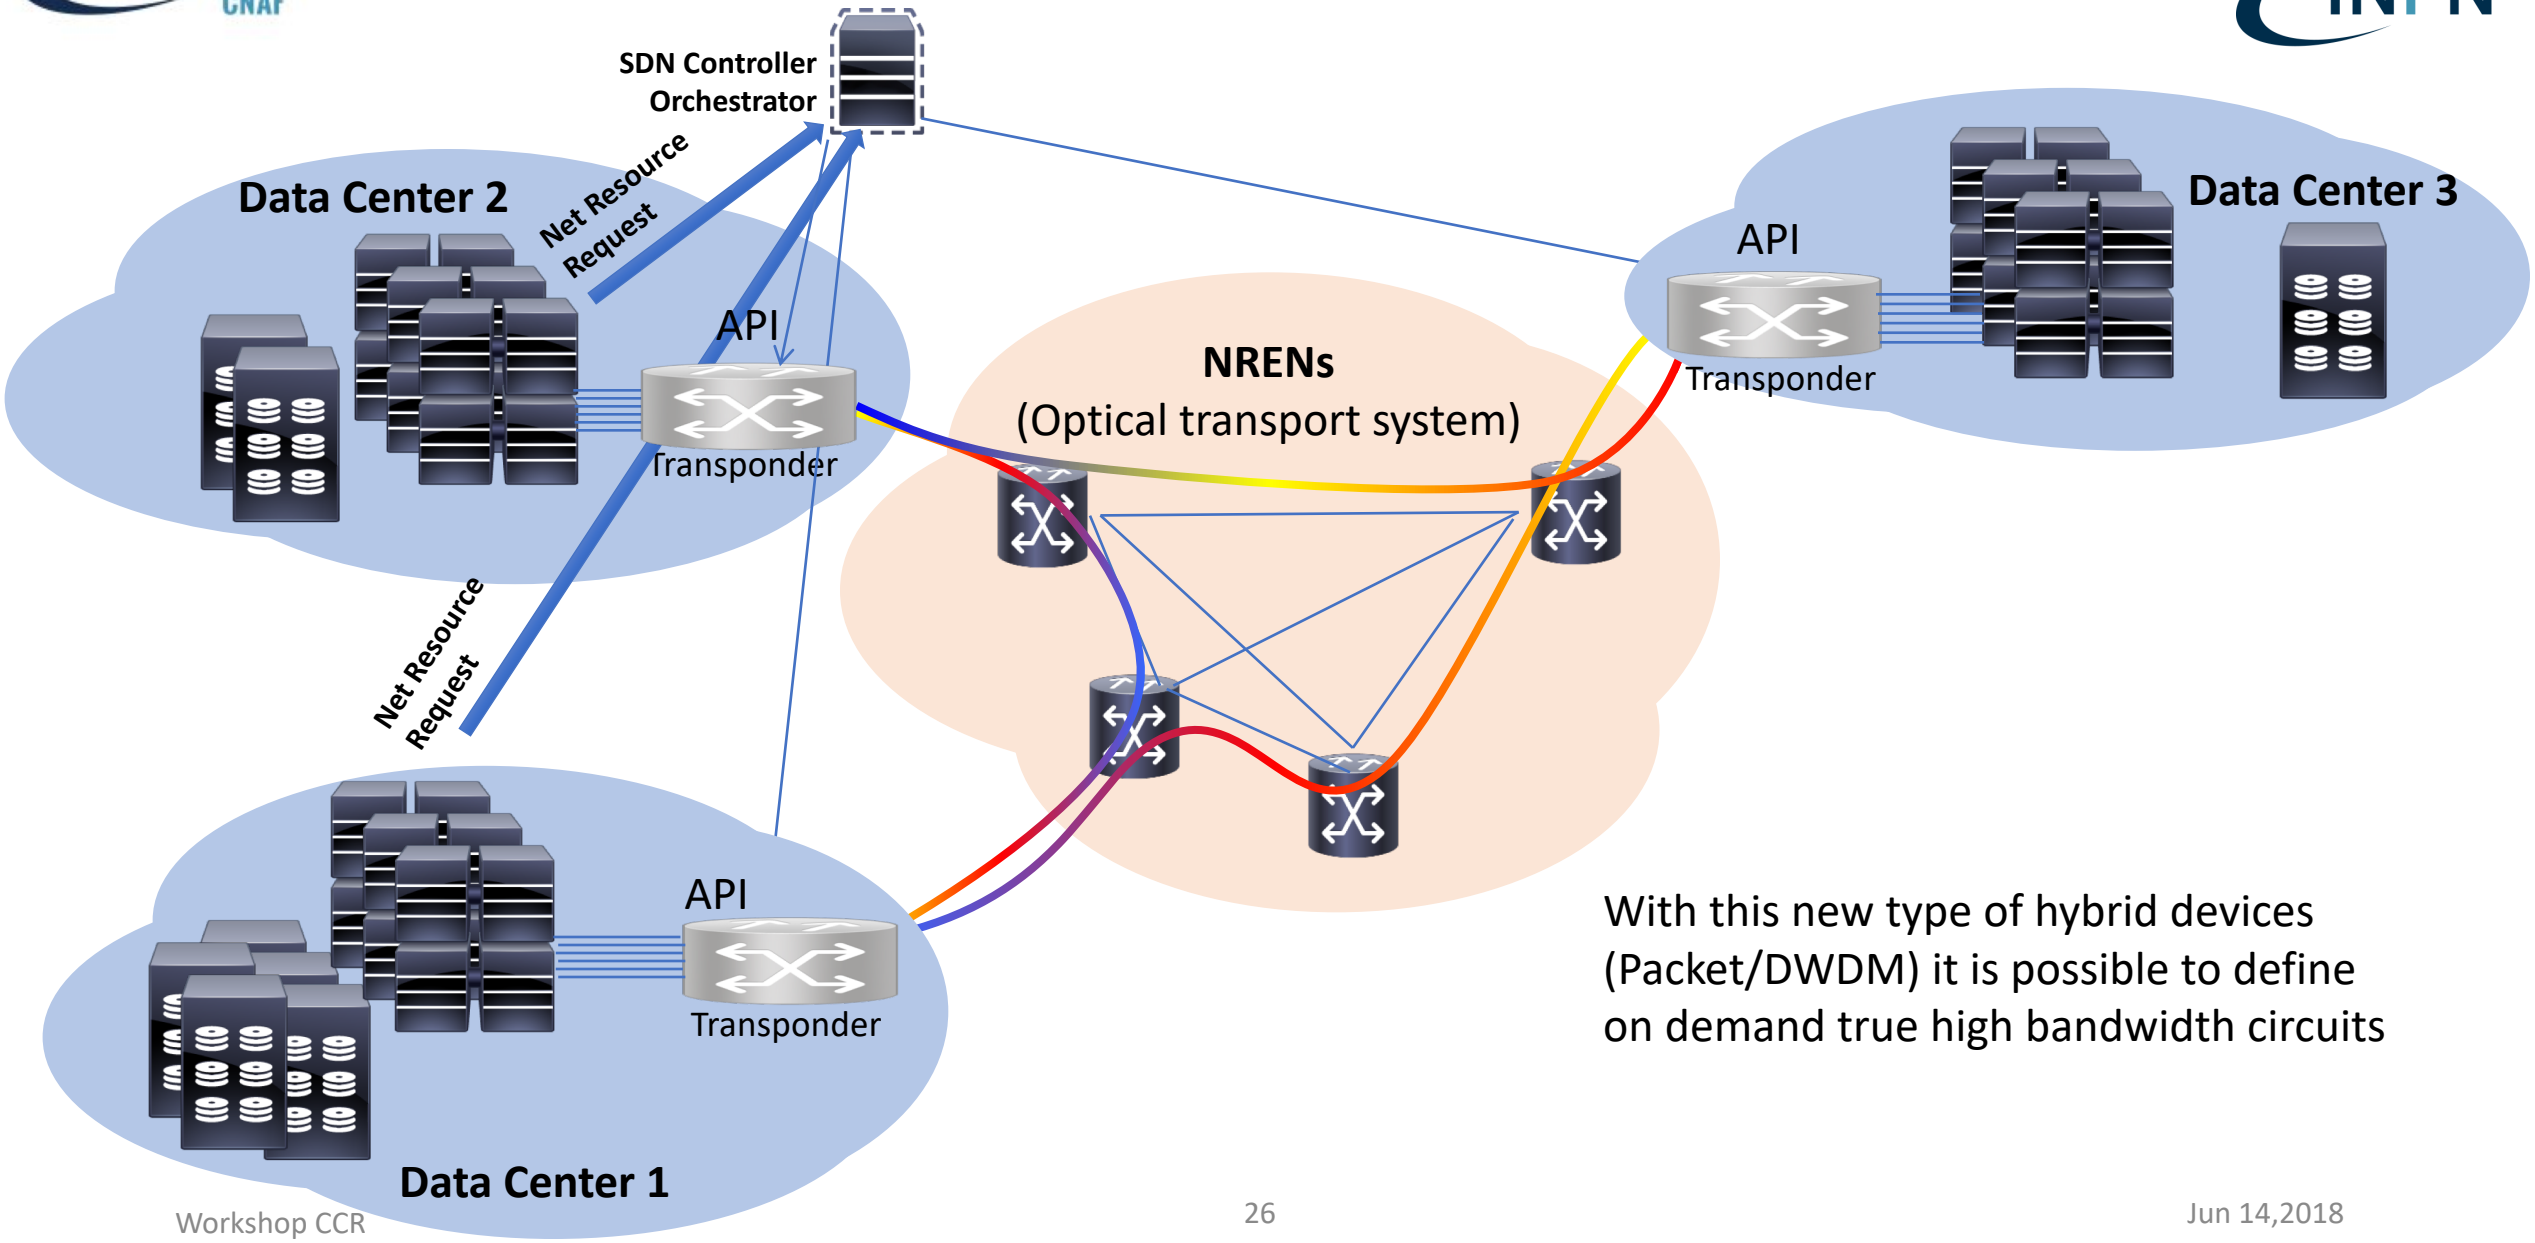

With this new type of hybrid devices (Packet/DWDM) it is possible to define on demand true high bandwidth circuits

Conclusions (?)

- INFN Tier-1 fully operational since March
  - Some hiccups at the restart
- Some systems not completely recovered yet
  - Still working on 2nd power line (needed for redundancy)
  - Strategy for continuity on the 2nd line not decided
- Some damaged components on the library ordered but still to be installed (not critical)
- In the meanwhile activity ongoing to improve the isolation of the data center perimeter (water broke into through various “sources” on the wall)

A possible future for the CNAF Tier1 towards the HL-LHC Data Lake

Looking for a new location for the Tier-1

The goal: take into account the needs for HL-LHC (i.e.data lake) and expansions due to astro-particle experiments

ECMWF center will be hosted in Bologna from 2019 in the Technopole area.

Possibility to host in the area also:
- INFN Tier-1
- CINECA computing center

Allocated 40 M€ from the Italian government to refurbish the area (works should be completed in 2019)
Extra budget 19 M€ promised for INFN & CINECA

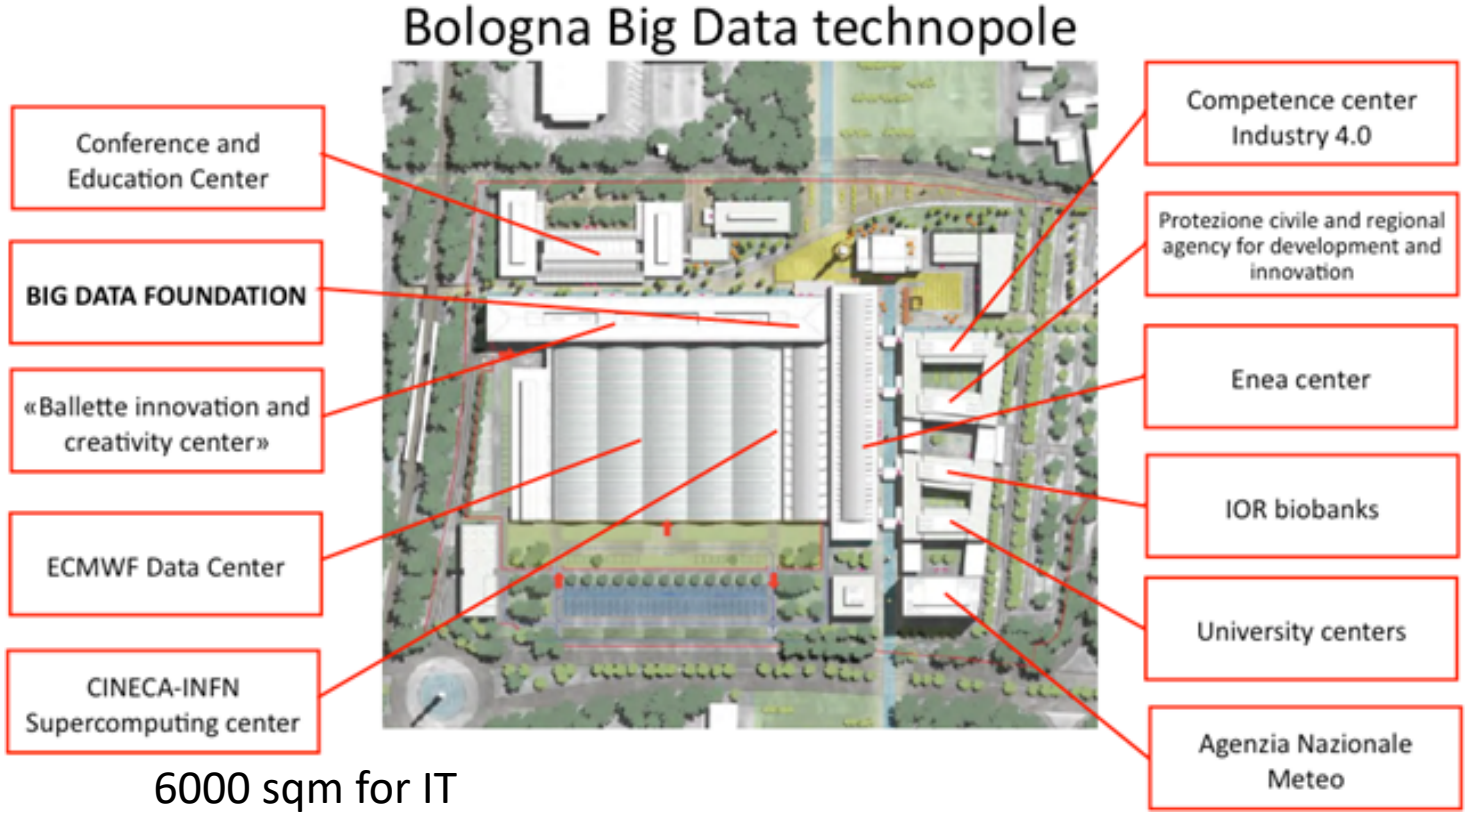


Supermarket
Via Ferrarese
Tecnopolo di Bologna
(ECMWF Candidate)
SS64
Sprint Gas S.P.A
McDonald's
Bologna Stalingrado
Via Stalingrado
Auto Parts Store

Technopole

The data centers at the Science Park

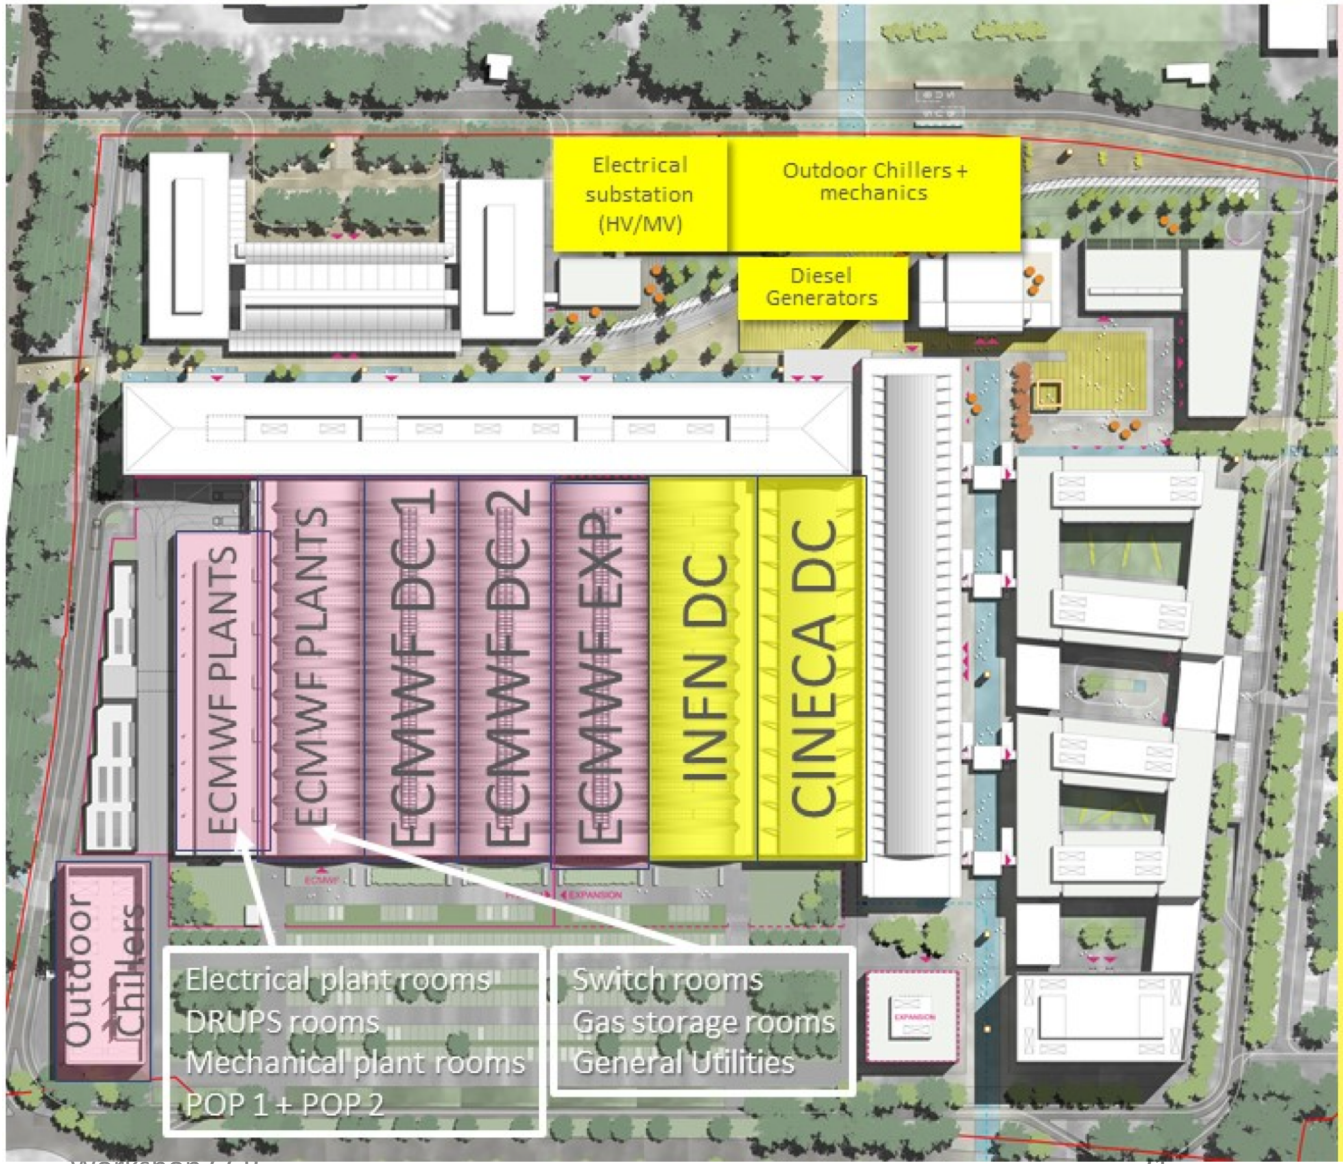

ECMWF DC main characteristics

- 2 power line up to 10 MW (one bck up of the other)
- Expansion to 20 MW
- Photovoltaic cells on the roofs (500 MWh/year)
- Redundancy N+1 (mechanics and electrical)
- 5 x 2 MW DRUPS
- Cooling
  - 4 dry coolers (1850 kW each)
  - 4 groundwater welles
  - 5 refrigerator units (1400 kW each)
- Peak PUE 1.35 / Maximum annualized PUE 1.18

INFN – CINECA DC main characteristics

- up to 20 MW (one bck up of the other)
- Possible use of Combined Heat and Power Fuel Cells Technology
- Redundancy strategy under study
- Cooling, still under study
  - dry coolers
  - groundwater welles
  - refrigerator units
- PUE < 1.2 – 1.3

Current status & activity

- Writing requirements for data center infrastructure
  - Joint effort with CINECA
- Available space is divided in 2 buildings (“botti”)
  - 3000 mq. Each (1500 mq for IT)
- Power ramping up to 20 MW for IT in 2028
  - $PUE_{mean} \leq 1.18$; $PUE_{peak} \leq 1,25$
- Tight schedule
  - First “botte” must be ready before the end of 2020!
  - Second “botte” ready for 2023
- This data center would fit perfectly the requirements for WLCG data lake!

IT power [MW]	2020	2021	2024	2028
INFN	2	2	3	10
CINECA	2.5	8	10	10
Totale	4.5	10	13	20