Network Architecture and IPv6 Deployment at CERN

CHEP Oct 2013
David.Gutierrez@cern.ch

Co-authors: Edoardo.Martelli@cern.ch, Carles.Kishimoto@cern.ch
IT/Communication Systems

Agenda

1. Network Infrastructures
2. LCG Architecture
3. Management
4. IPv6 Deployment Status

astruct

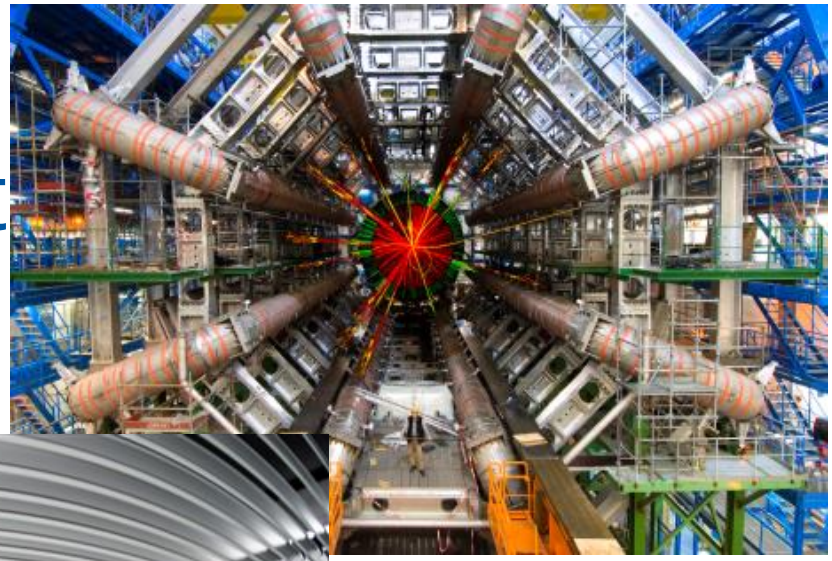

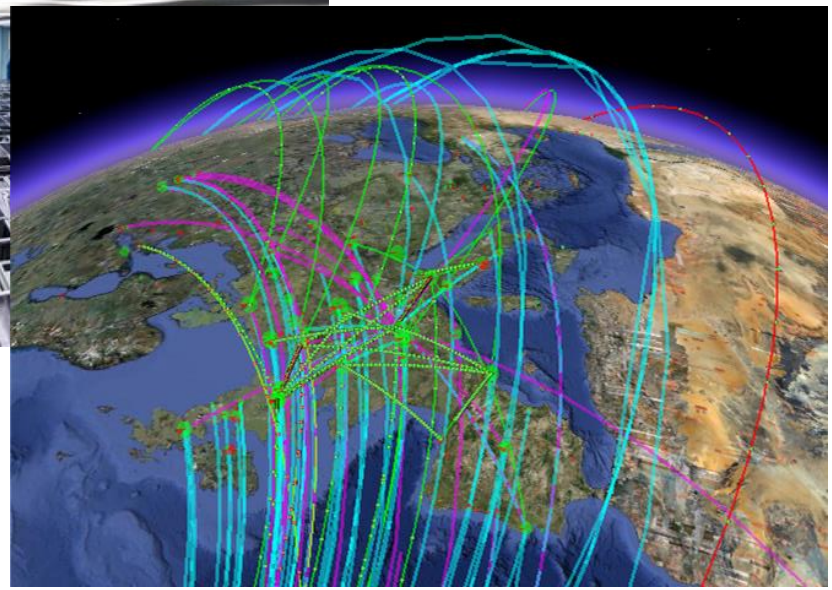

CORE Network

- Interconnects all infrastructures
- Extends between Geneva-Wigner
- Non-blocking 1Tbps
- **IPv6 Ready**
- OSPF backbone
- Security policies
- Problem isolation

Technical Network

- Technical services
  - Safety, electricity, cooling,…
- LHC Operation
  - Cryogenics, vacuum,…
- Industrial systems

Devices	**8,073**
Switches	**485**
HP Routers	**24**
IPv4/IPv6 Dual Stack	**NO**

Experiments

- IT/CS provides
  - Detector Control Network
    - ALICE
    - ATLAS
    - CMS
  - Data Recording to T0
    - 20Gbps, up to 50Gbps
    - 25 PB/year

Devices	**6,111**
Switches	**390**
HP Routers	**15**
IPv4/IPv6 Dual Stack	**NO**

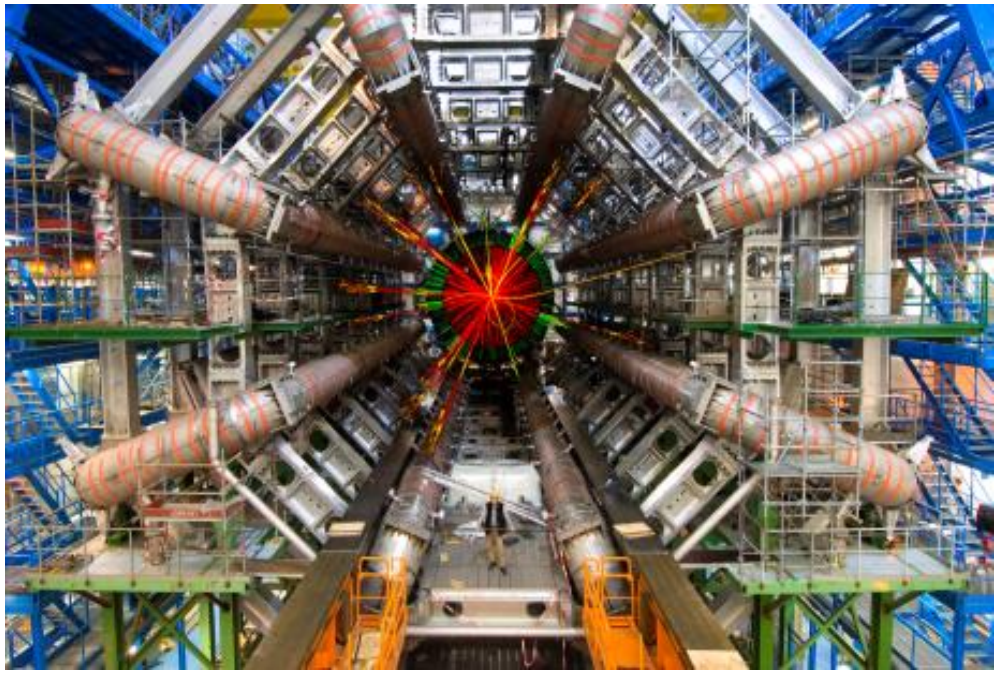

LHC Computing Grid

- High Performance Network
- 5.28 Tbps Non-blocking Switching Fabric
- WLCG Tier0
  - Long-term storage
  - Distribution WLCG
  - 1.5 PB/day

Devices	**8,902**
Switches	**588**
Brocade Routers	**13**
IPv4/IPv6 Dual Stack	**YES**

External Network

- Public general purpose connections
  - Full BGP Internet routing table
  - Geant, CIXP, ISPs
- Private WLCG
  - LHCOPN
    - 70Gbps peaks to T1
  - LHCONE

Brocade Routers	**8**
BGP Peerings	**86**
Aggregated BW	**232 Gbps**
IPv4/IPv6 Dual Stack	**YES**

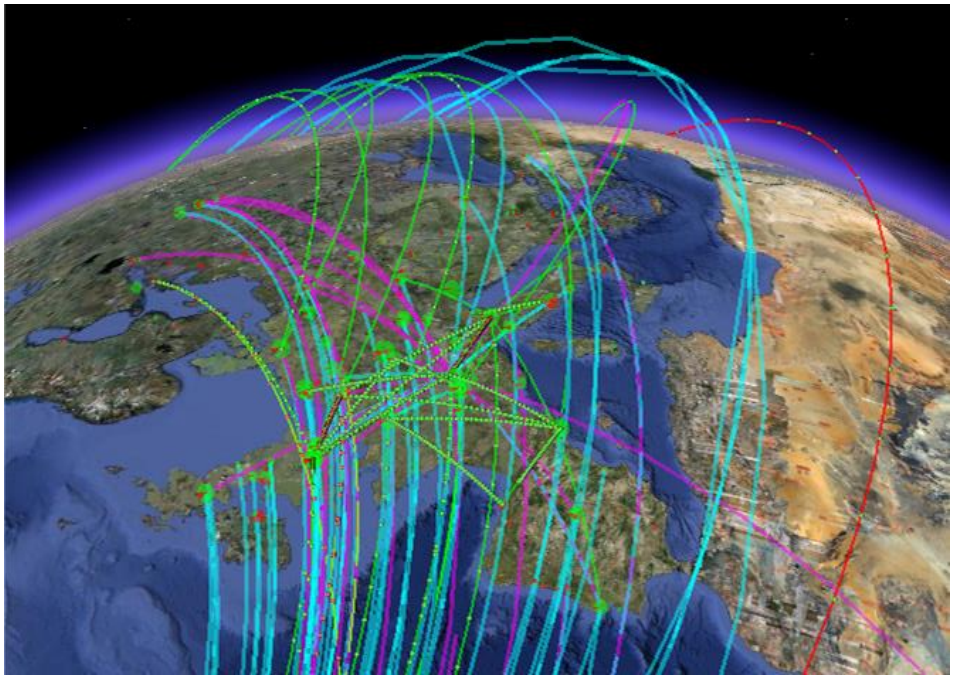

General Purpose Network

- Desktop computing
- Wired and wireless
- Central Services
  - AFS, www, mail, databases

Users	**14,592**
Buildings	**650**
Devices	**114,061**
Switches	**1,550**
WiFi Access Points	**1,514**
HP Routers	**100**
Brocade Routers	**7**
IPv4/IPv6 Dual Stack	**YES***

2. LCG Network Architecture

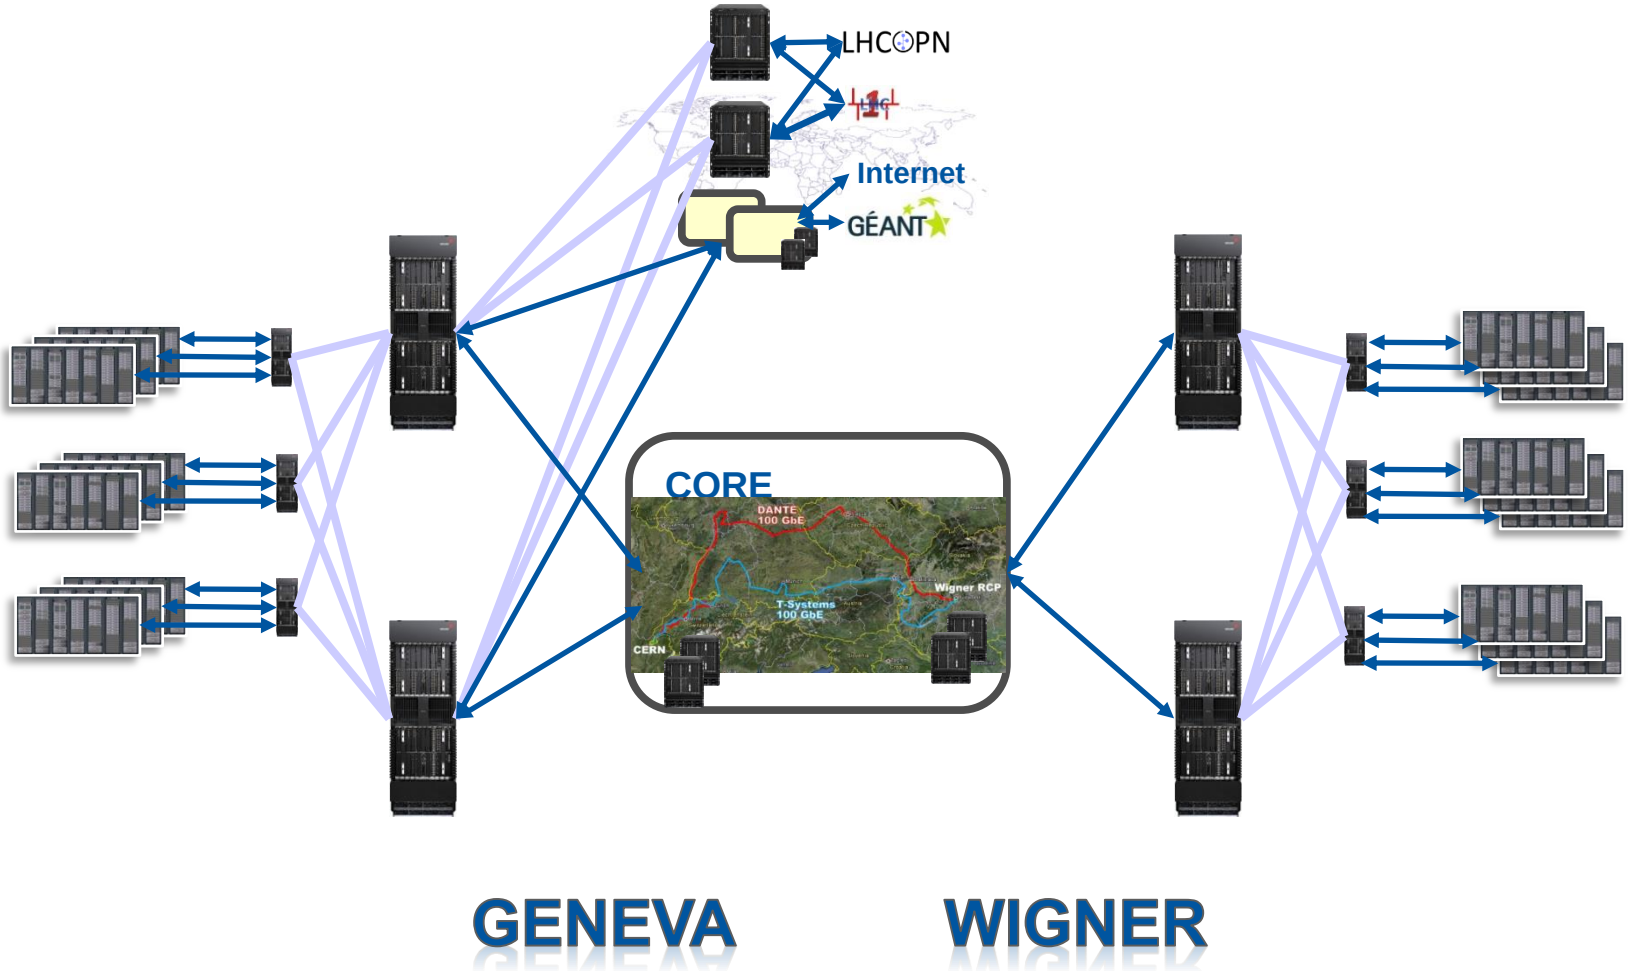

Building block: Service

- Network Access modeled as **Service**
- A **Service** is a broadcast domain
  - Providing access to end systems
  - **IPv6/IPv4 Dual Stack**
- Features are defined by a Service Type
  - Jumbo, LACP, tagged, loop protect, access-list protect, OSPF, VRRP, …
- A ToR Switch provides a **Service**
- A Router interconnects multiple **Services**

Service Bandwidth

- Service capacity depends on Service purpose
- Blocking Factor: 2 for CPUs, 5 for Storage

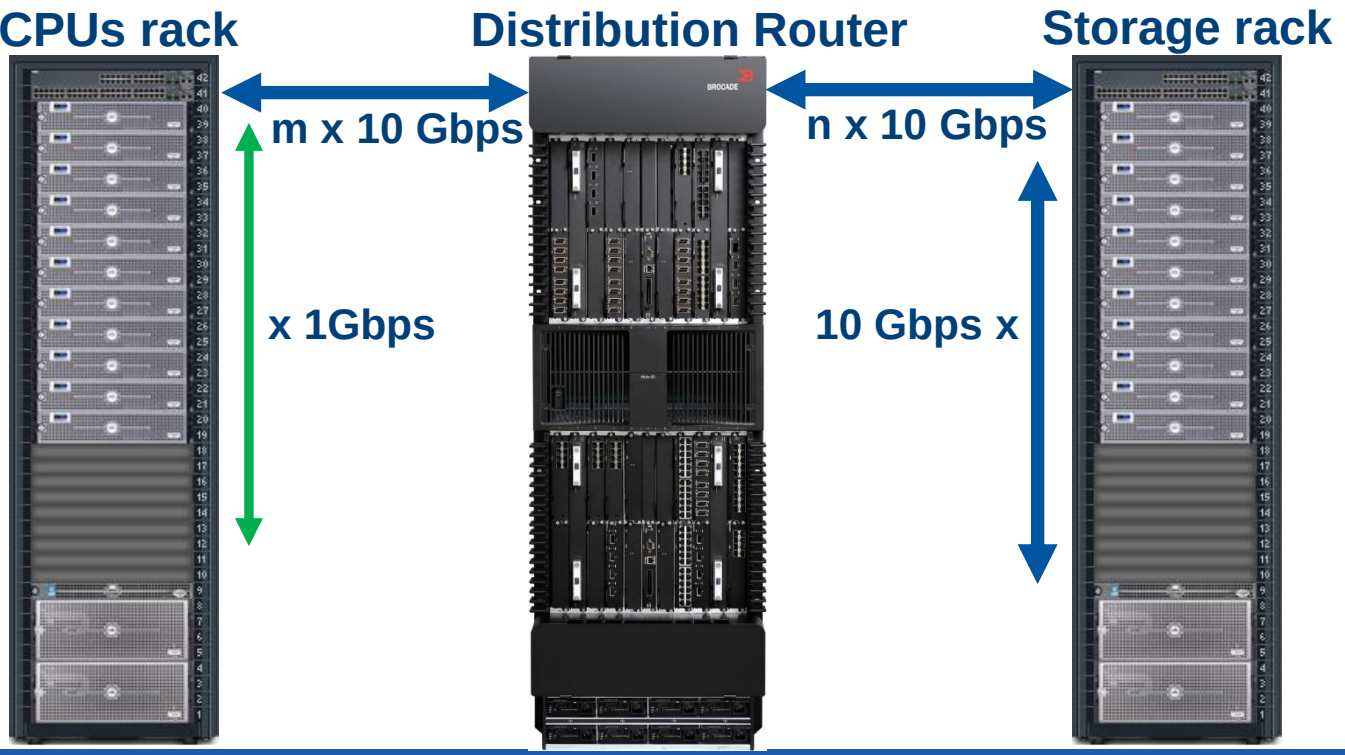

Scaling the Data Center

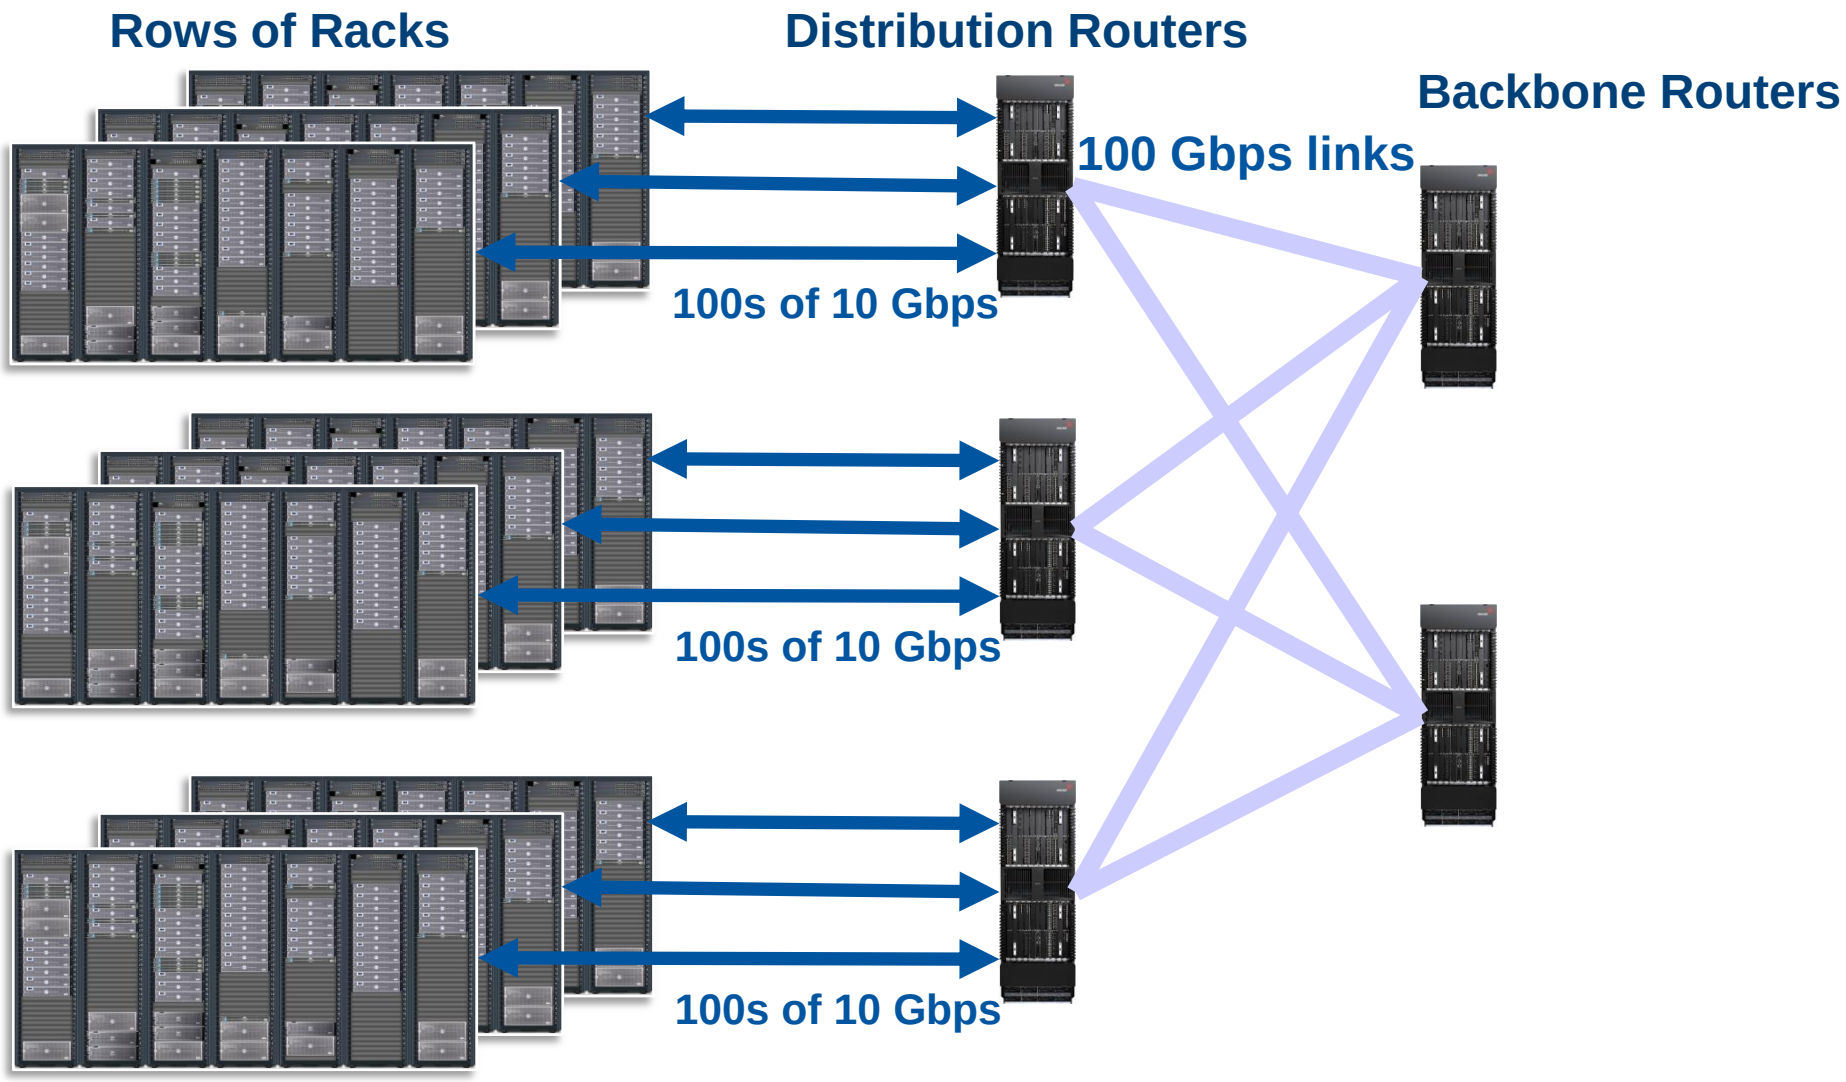

Worldwide LCG

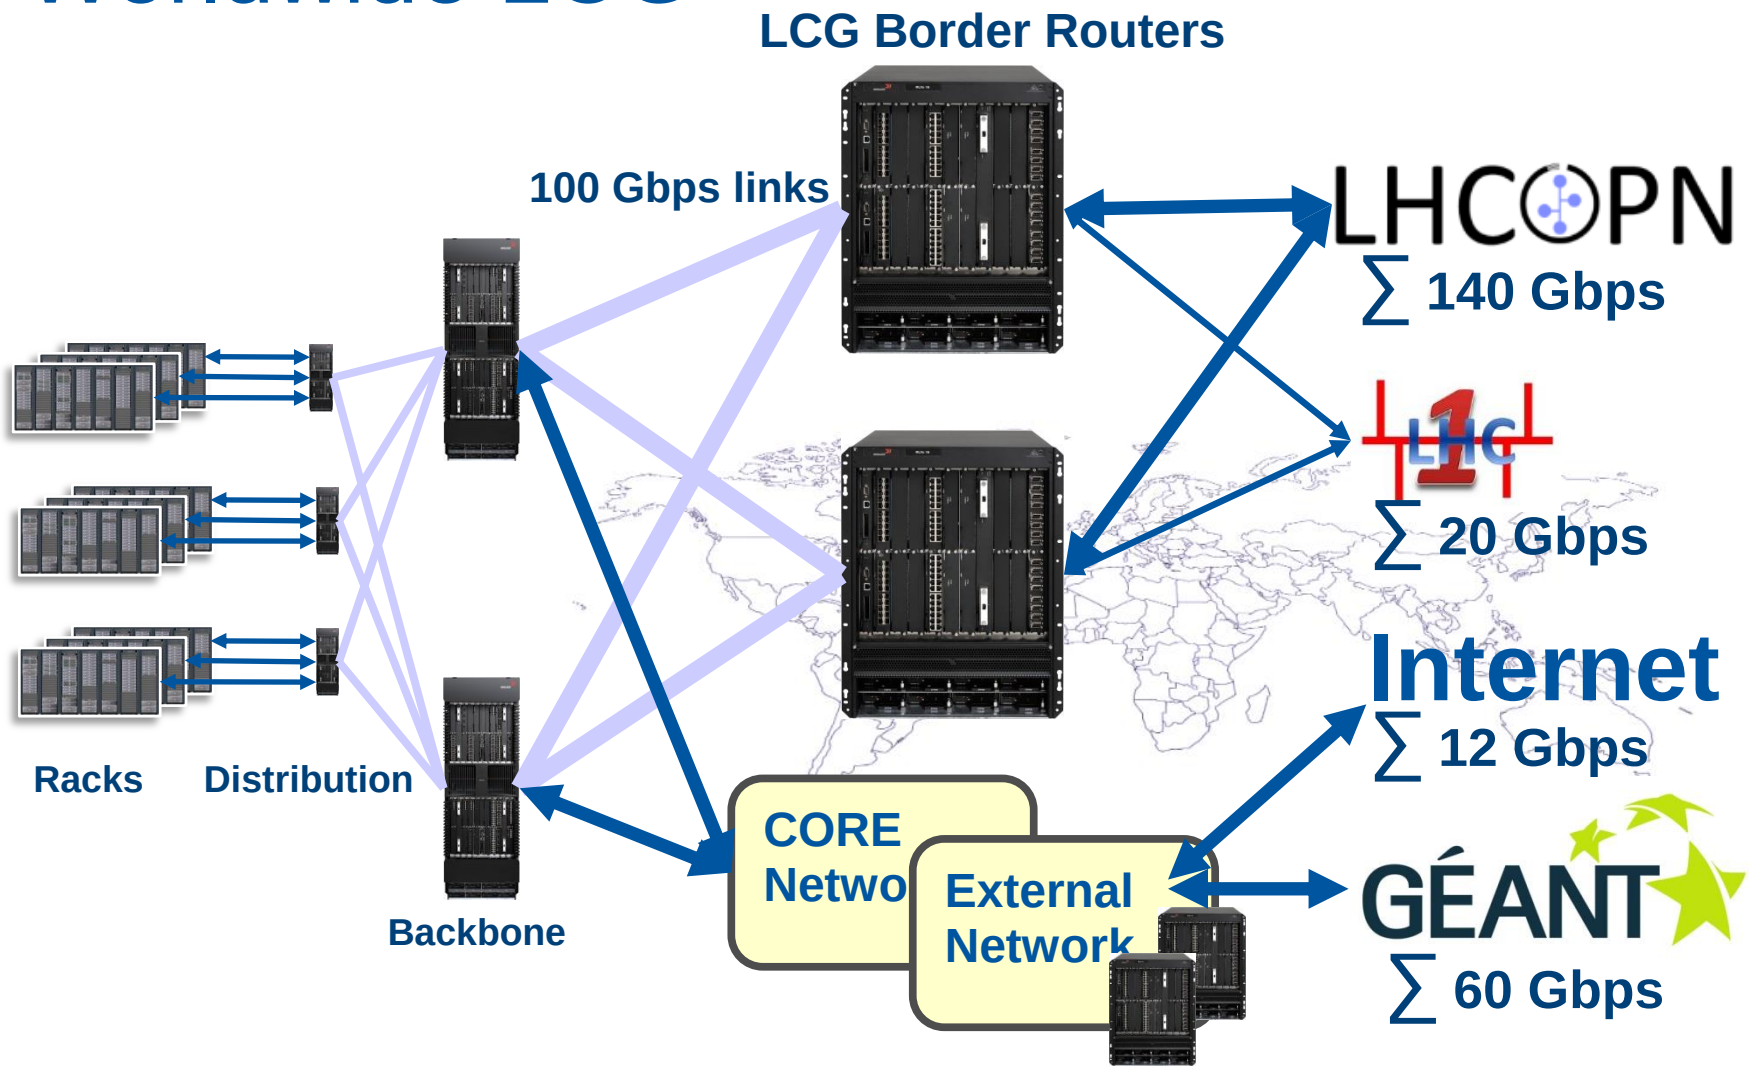

Extending the Tier0 to Wigner

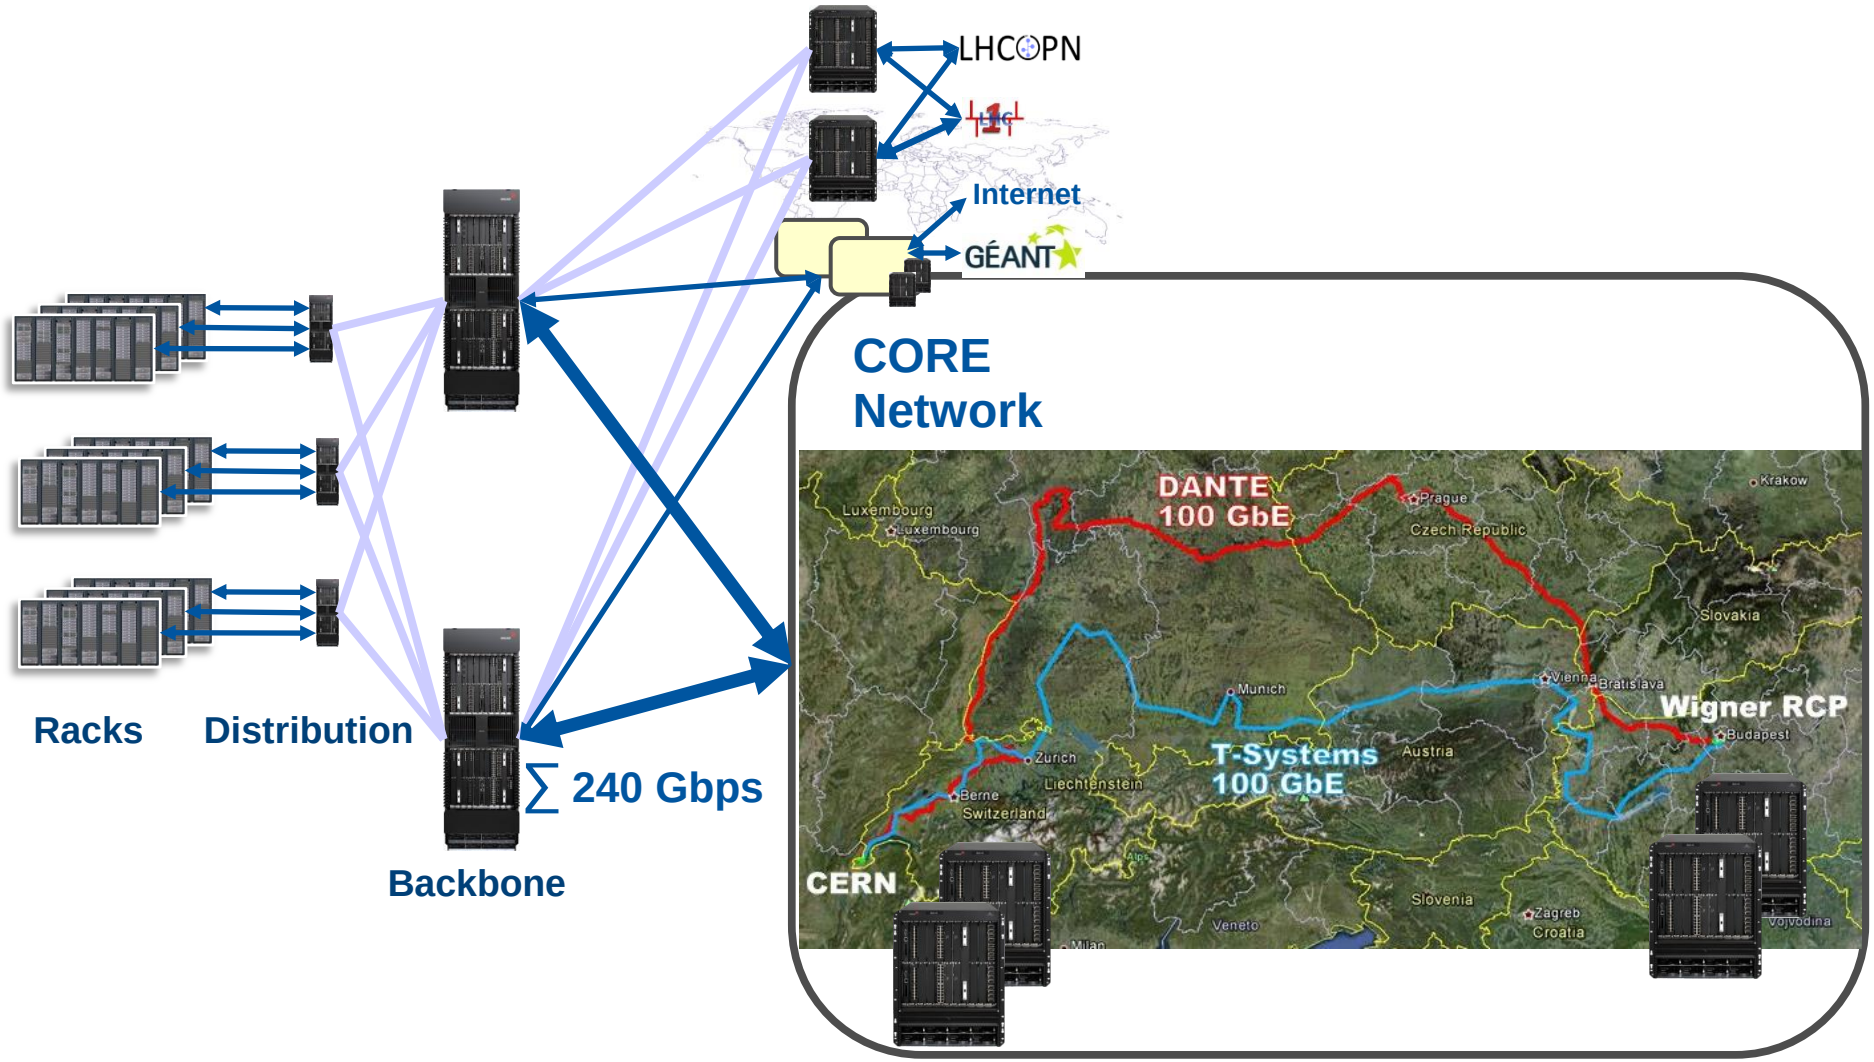

WLCG Tier0

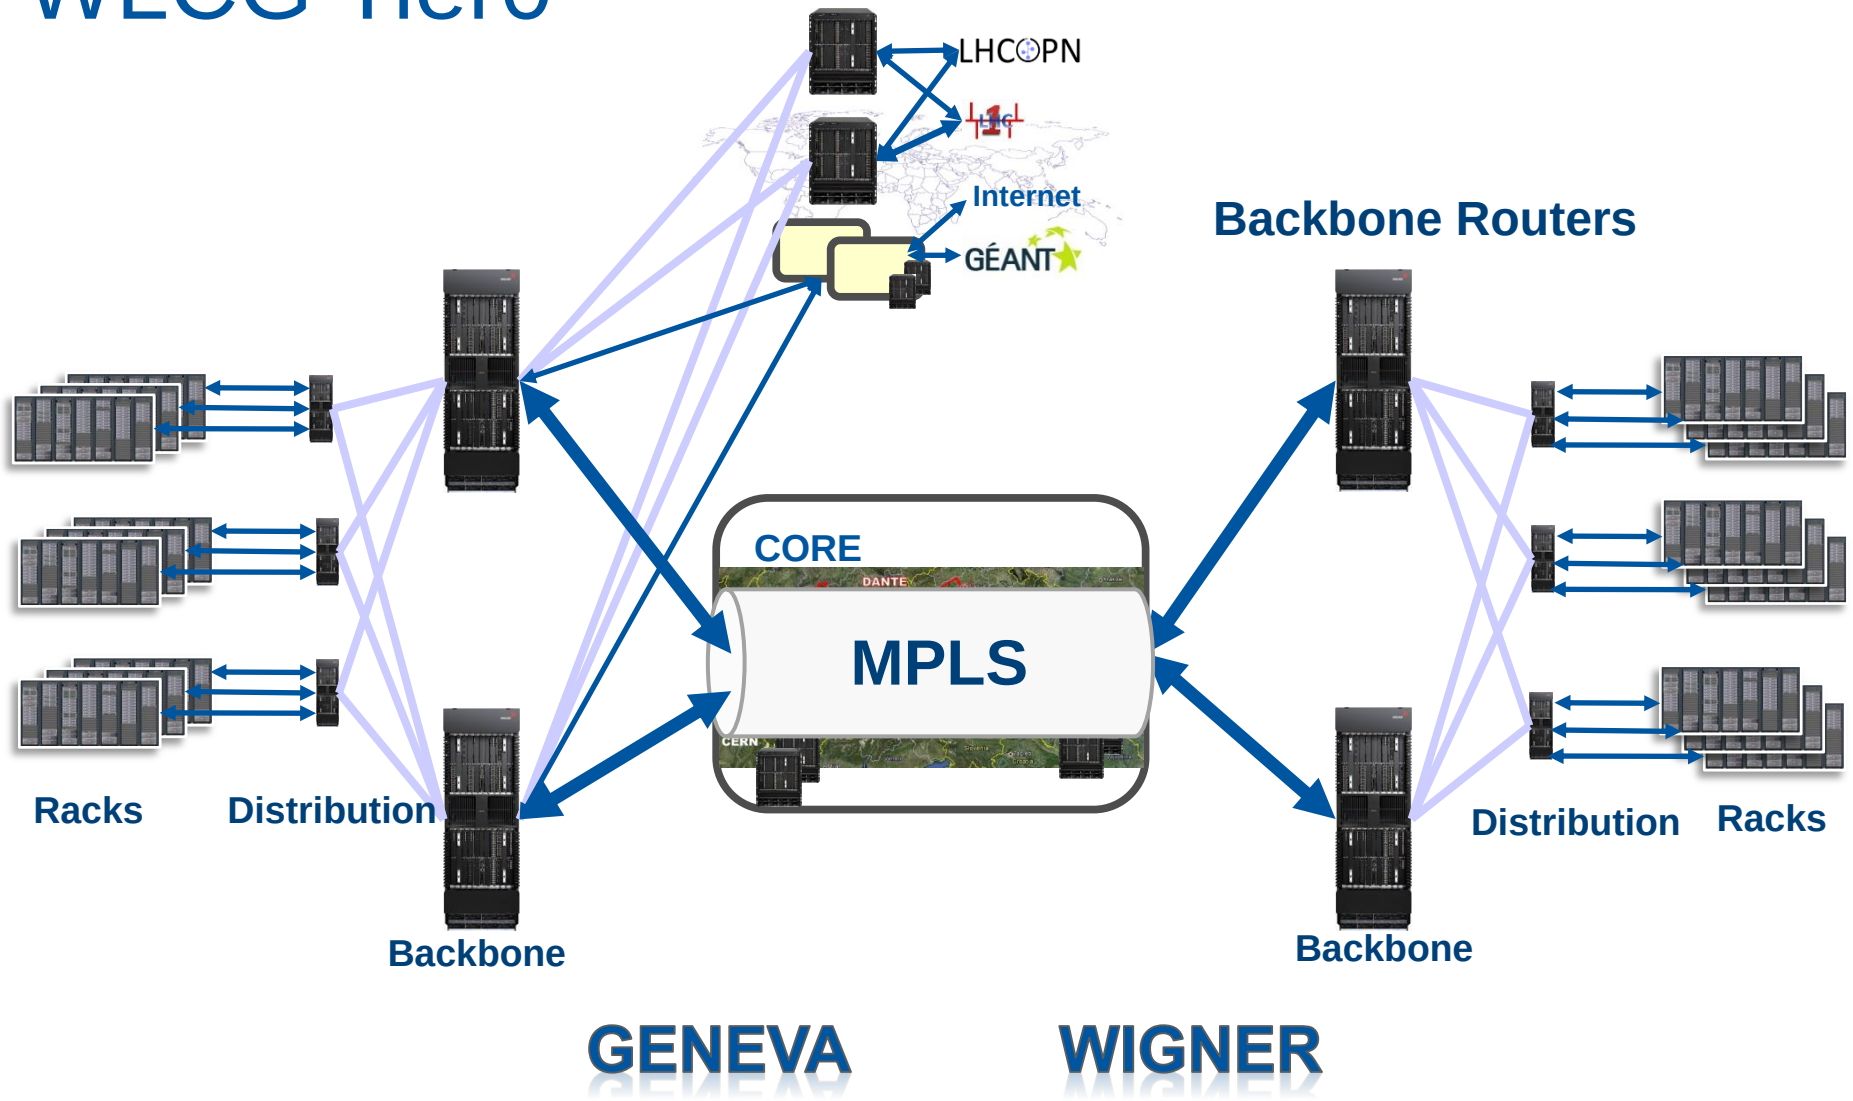
LHCOPN
Internet
GÉANT
Backbone Routers
CORE
MPLS
Racks
Distribution
Backbone
Backbone
Distribution
Racks
GENEVA
WIGNER

3. Network Management

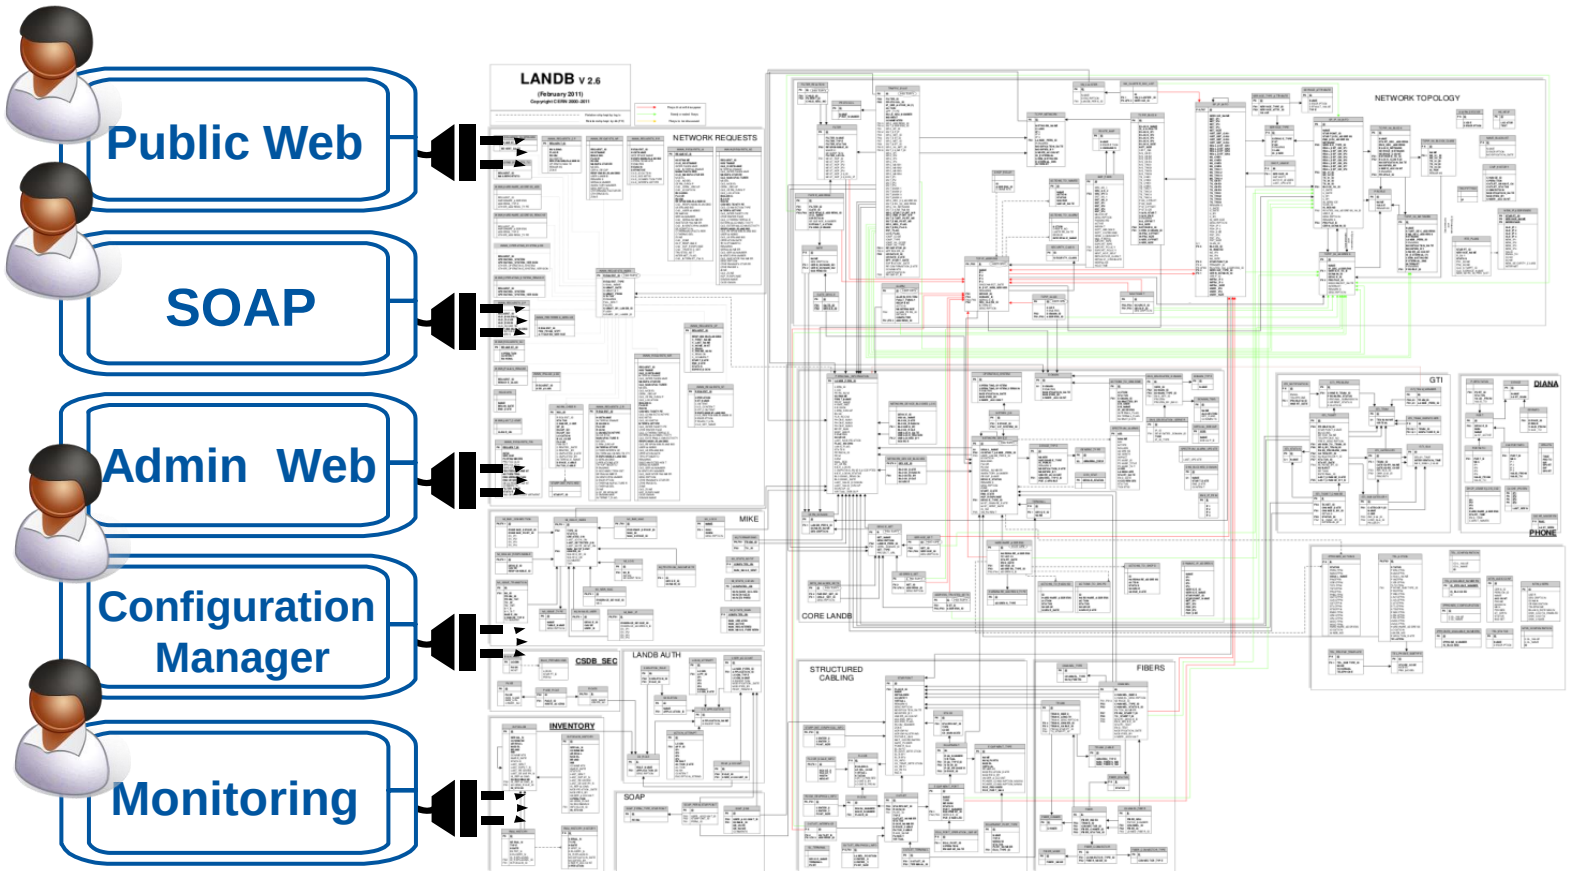

Dyn-A4 three columns example of a distribution router configuration

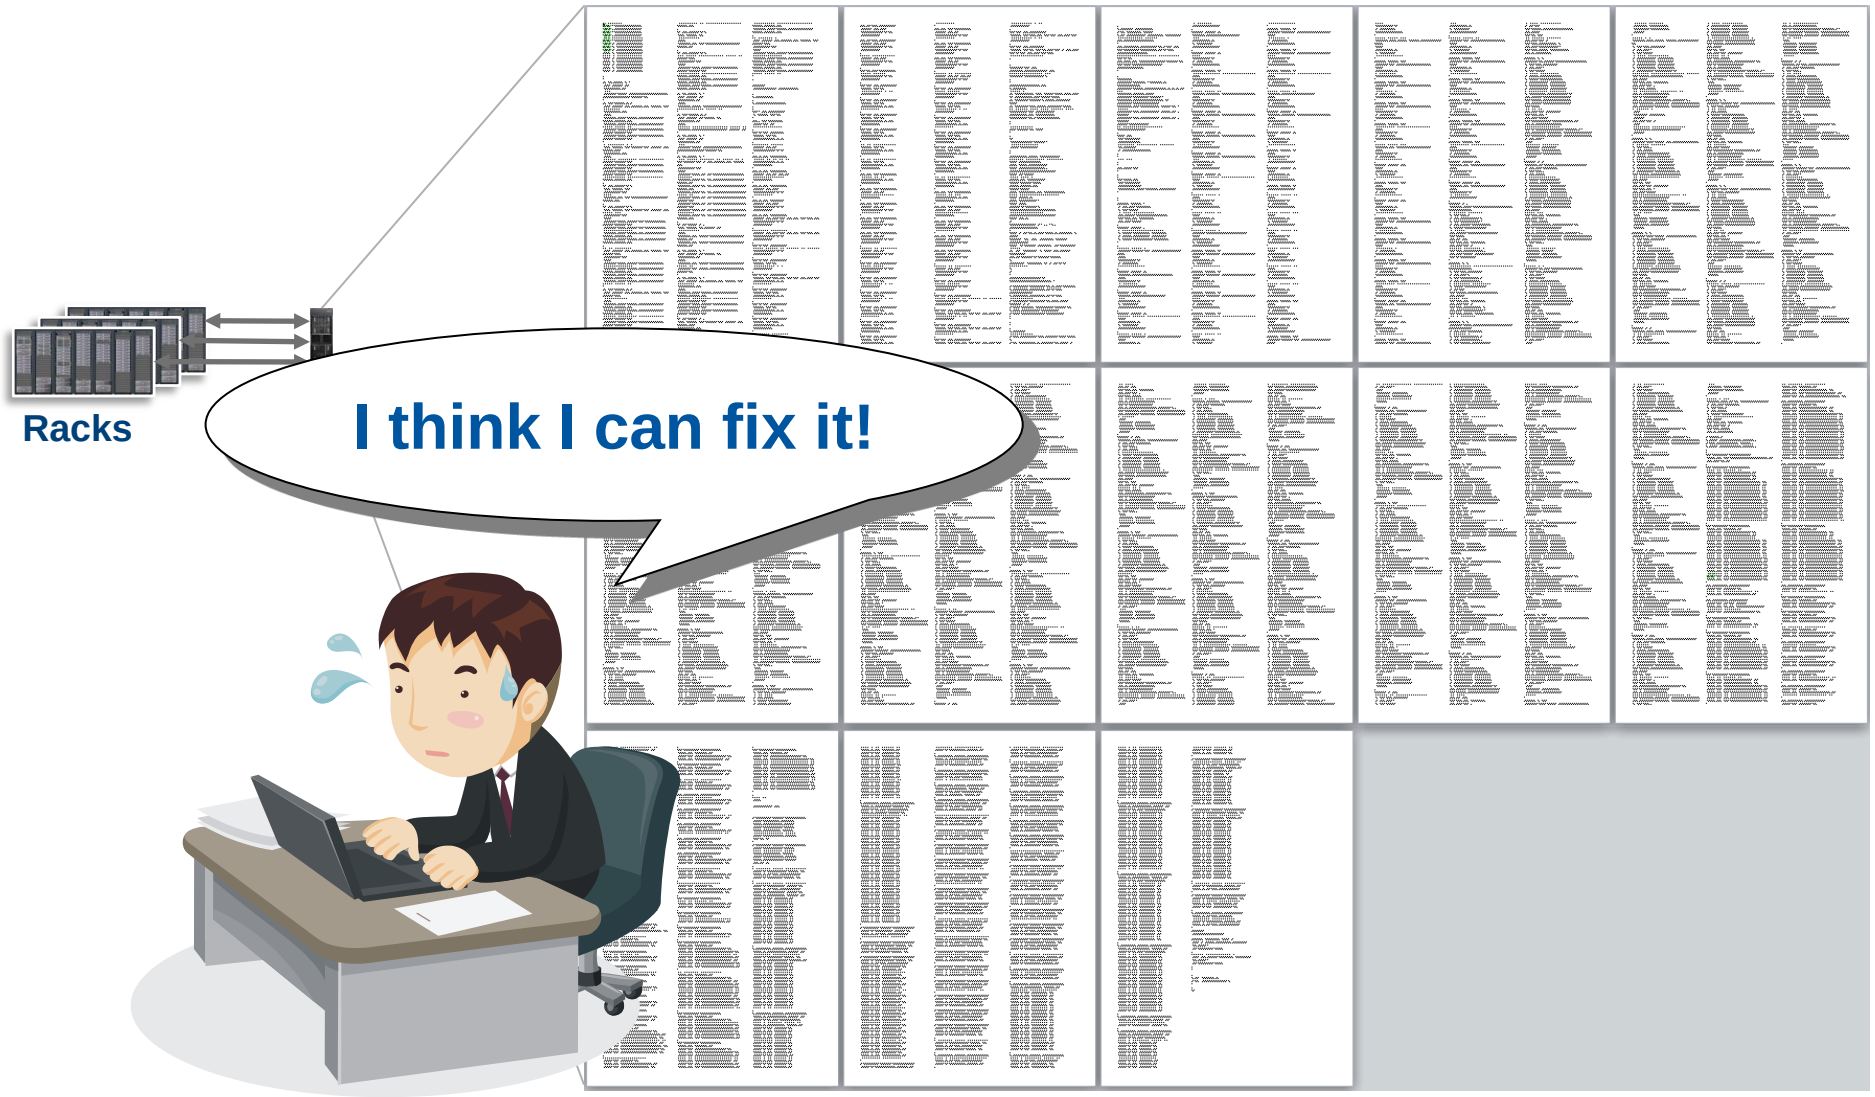

Network Database

- Buildings and cabling infrastructure
- IPv4/IPv6 addressing schemas
- Services and service features
- Devices (Servers, VMs, desktops, smartphones, …)
- Switches
- Routers/Firewalls
- Security for the Infrastructure (firewall rules)

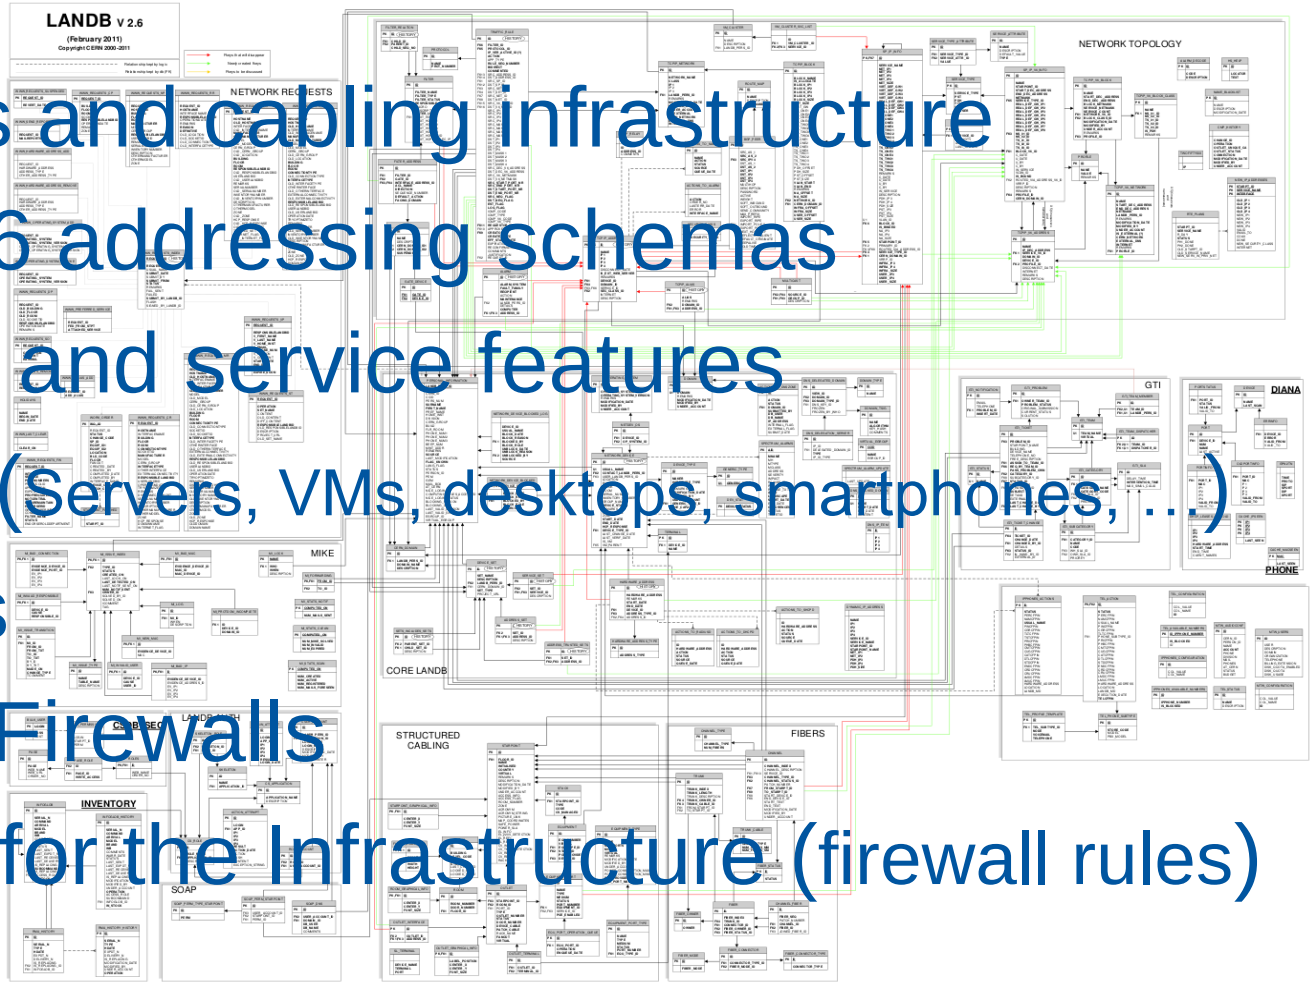

Interfaces to the Network Database

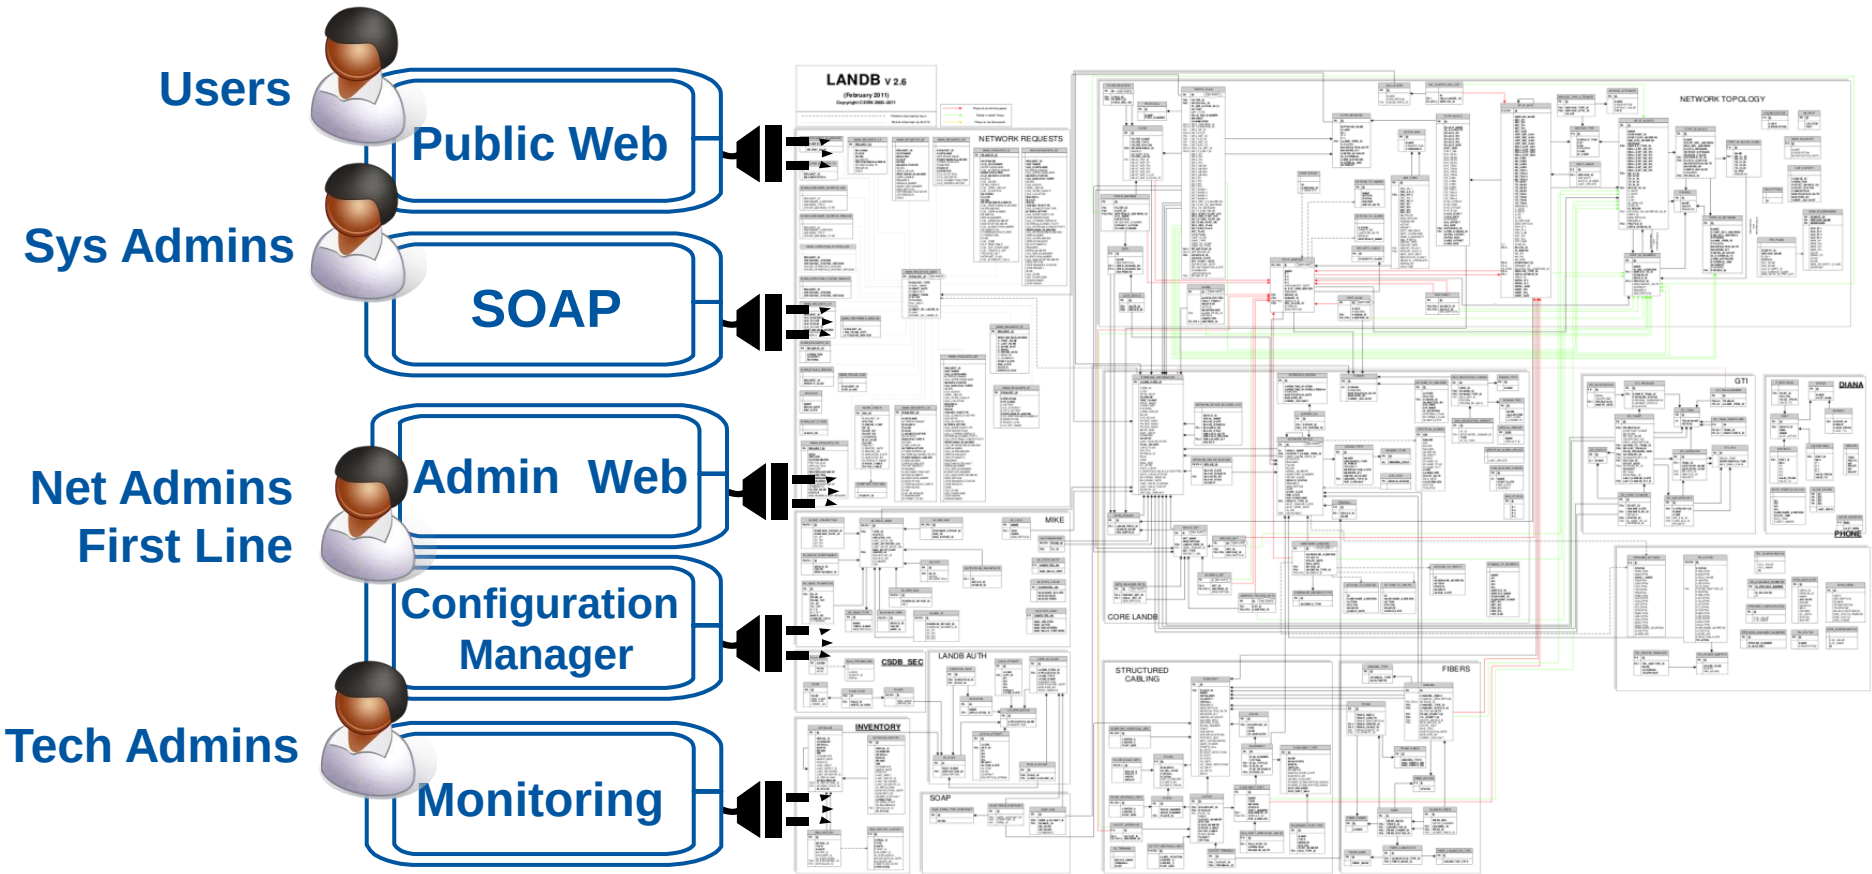

Software-based Network configuration

Net Admins
First Line

Automatic Updates

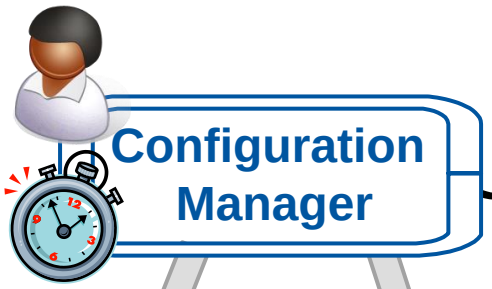

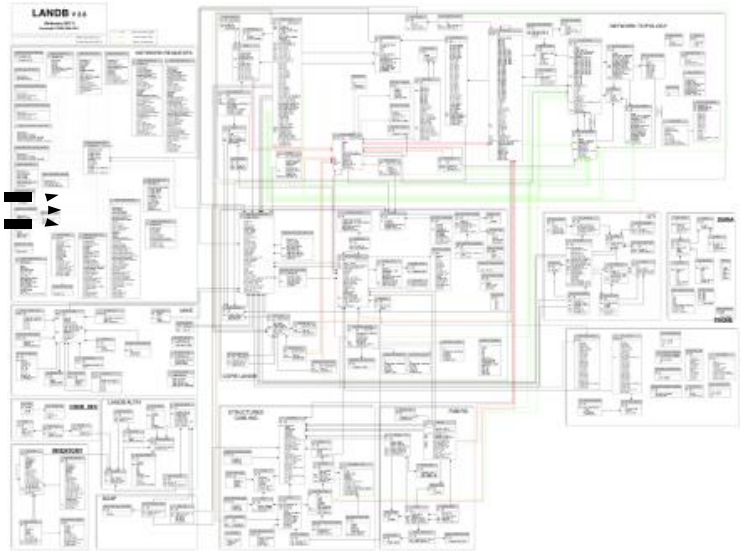

Network Services

- DNS
- DHCP
- RADIUS

4. IPv6 Deployment Status

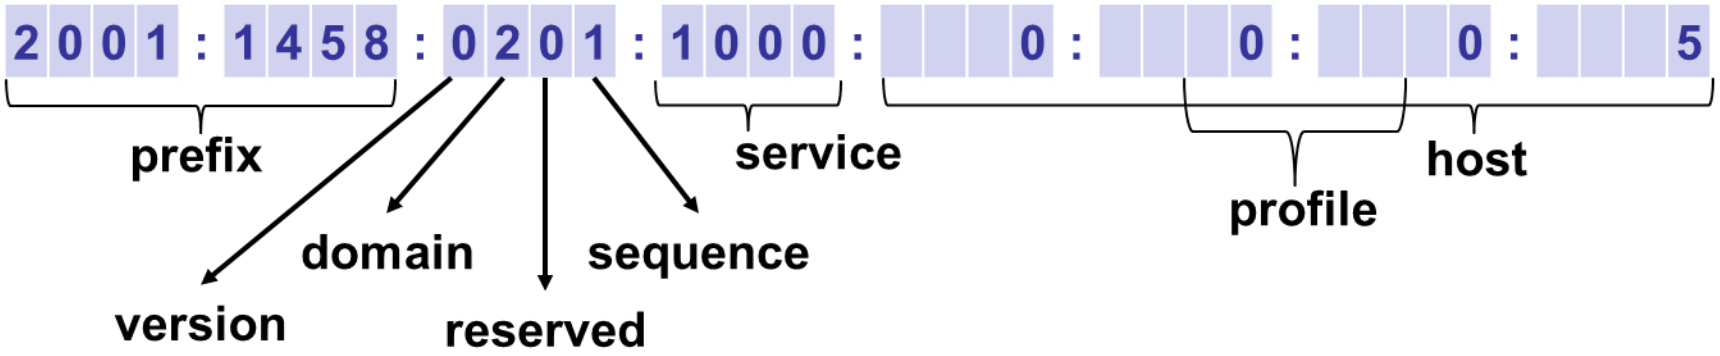

GVA prefixes
2001:1458::/32
fd01:1458::/32

WIGNER prefixes
2001:1459::/32
fd01:1459::/32

Network Domains	
0	EXTNET and Firewall
1	CORE
2	General Purpose Net
3	LHC Computing Grid
5	ALICE

Network Profiles	
fffe	EUI64
0000	Net Equipment
0001	User device

Well known hosts	
x::1	Gateway
x::2	VRRP backup

2012

- Network Database: Schema and Data IPv6 Ready
- Admin Web: IPv6 integrated
- Configuration Manager supports IPv6 routing

2013

- Gradual deployment on the routing infrastructure starts
- The Data Center is Dual-Stack
- NTPv6 and DNSv6
- DHCPv6

Today

- Infrastructure is Dual-Stack
- Firewallv6 automated configuration
- User Web and SOAP integrate IPv6
- Automatic DNS AAAA configuration

IPv4 / IPv6 same portfolio

- Identical performance, common tools and services
- Dual Stack, dual routing
  - OSPFv2/OSPFv3
  - BGP ipv4 and ipv6 peers
- Service managers decide when ready for IPv6
- Devices must be registered
  - SLAAC disabled
  - RAs: Default Gateway + IPv6 prefixes no-autoconfig
  - DHCPv6
    - MAC addresses as DUIDs: painful without RFC6939
    - **ISC has helped a lot** (βcode implementing classes for ipv6)
    - DHCPv6 clients might not work ‘out of the box’

Conclusions

- The Network is ready to accommodate the new demands after Long Shutdown 1
- Before Eo2013, IPv6 will be fully deployed and available to the CERN community

Thank you!

Extra Slides

Data Centers	Geneva	Wigner 2013
Power	3,500KW	~900KW
Racks	828	90
Servers	10,173	~1,200
Routers	22	6
100Gbps ports	60	18
ToR Switches	662	140
ToR Switching		
1Gbps ports	22,776	3,072
10Gbps ports	4,284	528

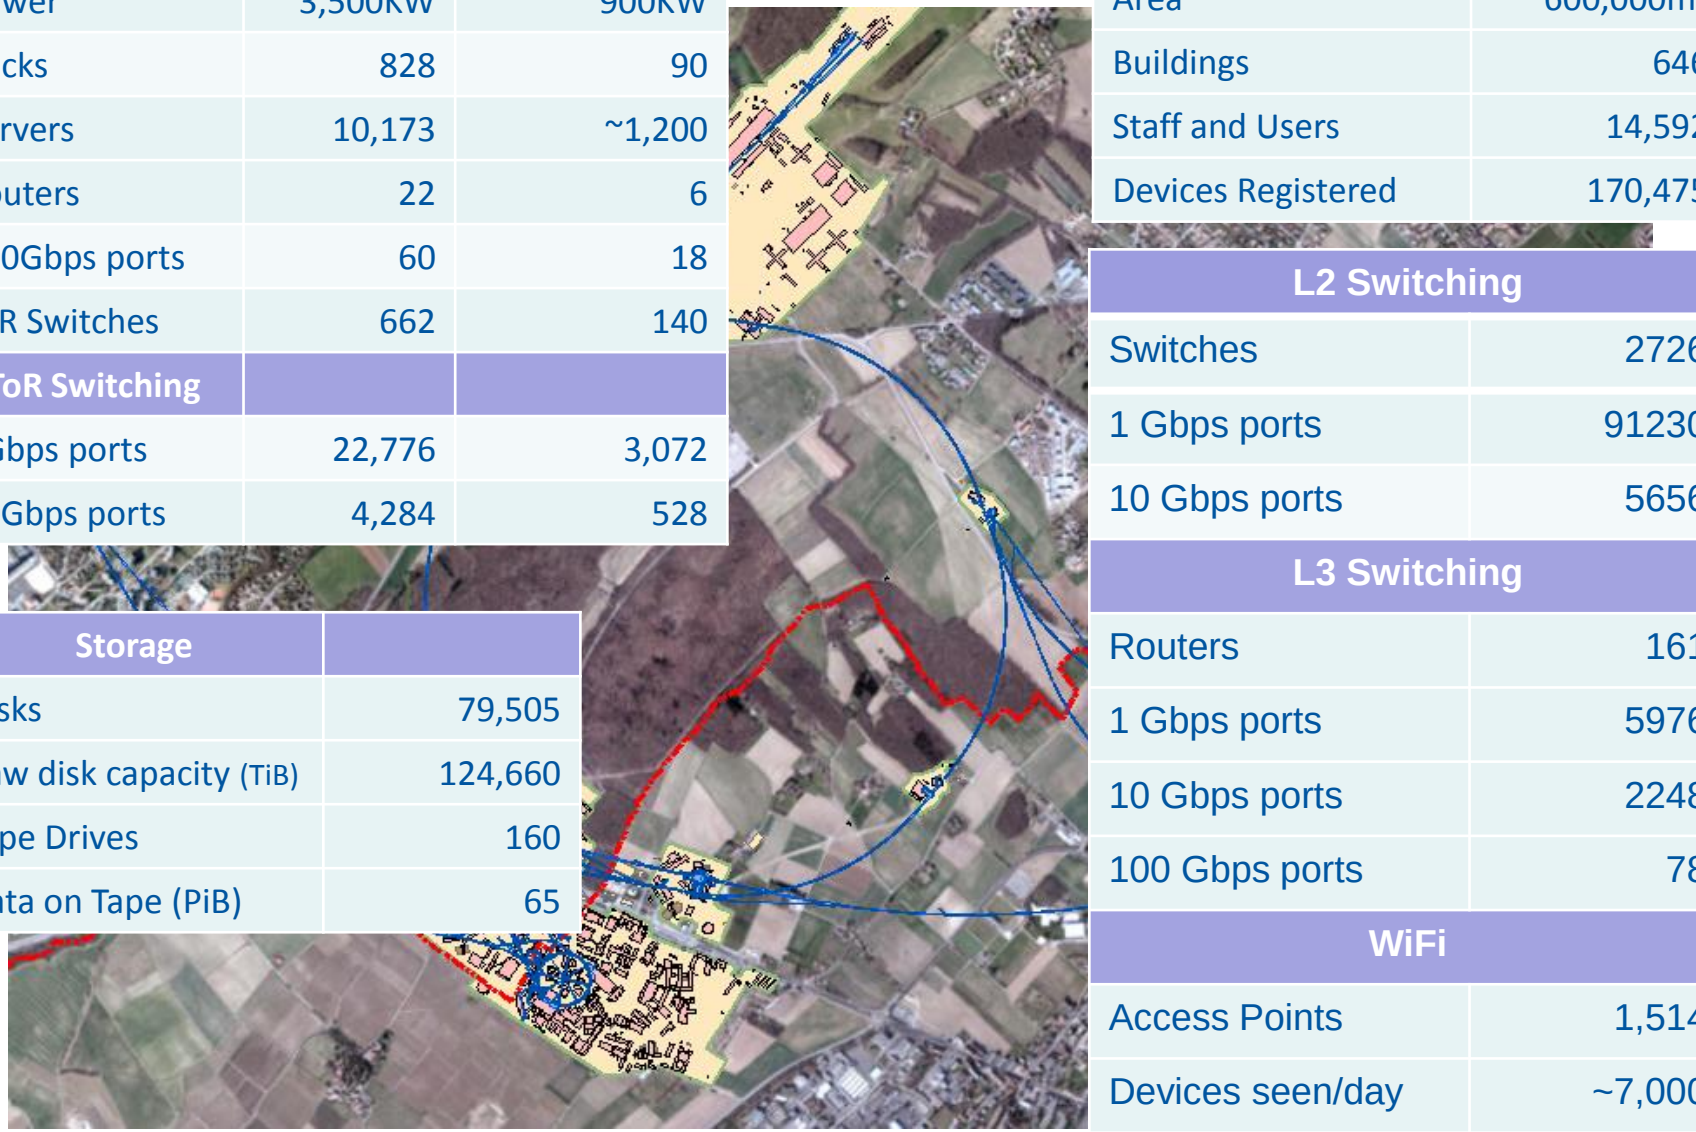

Storage	
Disks	79,505
Raw disk capacity (TiB)	124,660
Tape Drives	160
Data on Tape (PiB)	65

CERN	
Area	~600,000m^2
Buildings	646
Staff and Users	14,592
Devices Registered	170,475

L2 Switching	
Switches	2726
1 Gbps ports	91230
10 Gbps ports	5656

L3 Switching	
Routers	161
1 Gbps ports	5976
10 Gbps ports	2248
100 Gbps ports	78

WiFi	
Access Points	1,514
Devices seen/day	~7,000