The evolution of the TIM network: from the PSTN to the Cloud

Funded by
the European Union
NextGenerationEU

Agenda

Network architecture

The Public Switched Telephone Network

Terminology

The term “Public switched Telephone Network” refers to the traditional public telephone service (aka POTS: Plain Old Telephone Service).
This kind of service was operated along the twentieth century and progressively replaced by Voice over IP (VoIP) technology starting from 2000s.

Properties of Telecom Italia PSTN network

- Full mesh among Class4 telephone switches
- Hierarchical-only interconnections allowed in Gateways Areas (tree-topology)
  - Dual-homed interconnections between Class5 and Class4 switches (direct connections between Class5 switches in special cases)
  - Single-homed interconnection between branch offices and Class5 switches

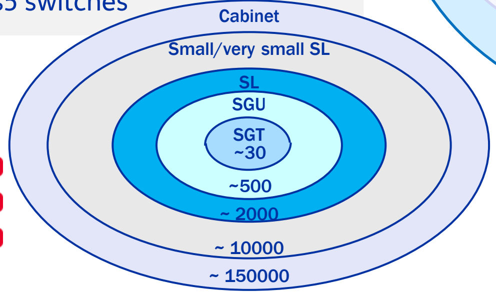

General properties of network architectures

- Fault tolerance guaranteed by duplicated network devices and multihoming
- Hierarchical structure
- The number of network devices and links increases from the centre to the edge
- Devices at the centre of the network have much higher switching capacity than devices at the edge
- Links at the centre of the network have higher bandwidth than those at the edge
- The number of interconnections between devices at the centre is higher than at the edge
- “Service” network devices are interconnected by means of transmission networks

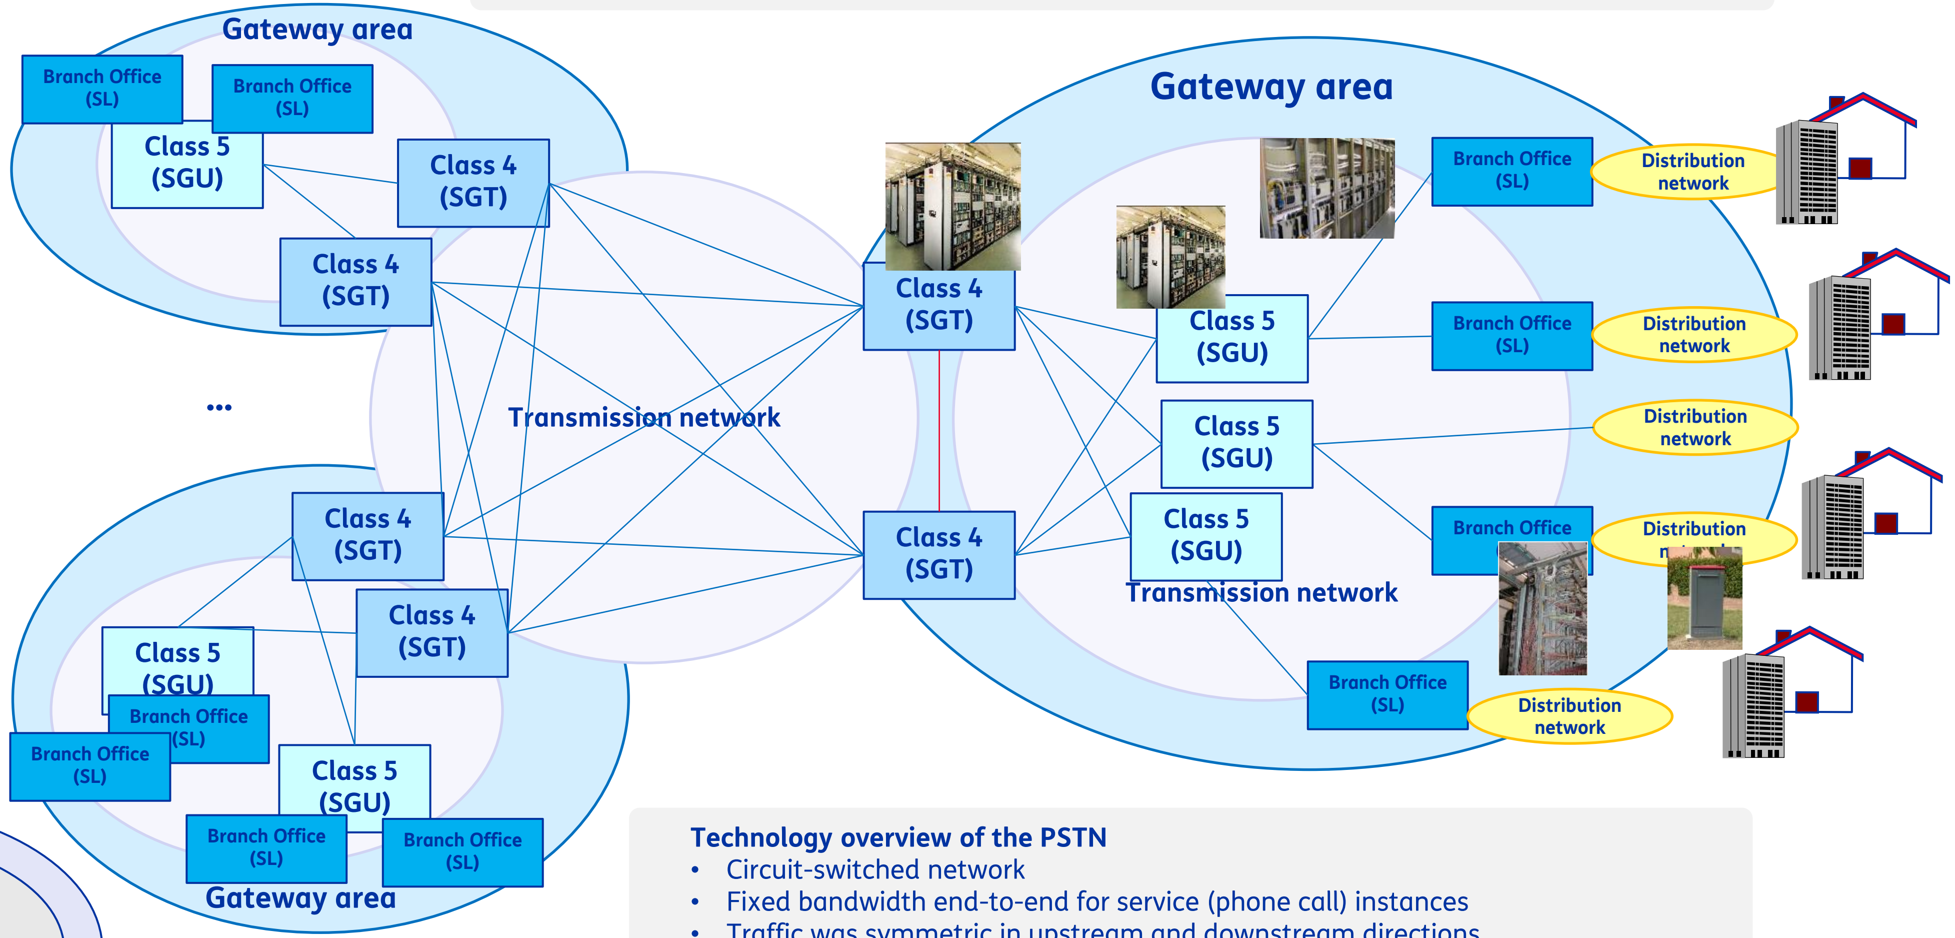

Technology overview of the PSTN

- Circuit-switched network
- Fixed bandwidth end-to-end for service (phone call) instances
- Traffic was symmetric in upstream and downstream directions
- Traffic demand (hourly, daily, etc.) was highly predictable
- Traffic relationships were mostly local
- Transmission networks based on SDH technology (Time Division Multiplexing)

TIM Data Network

Terminology
"Data Network" is a general term that refers to the evolution (or complete replacement) of the PSTN to support modern data services: web browsing, video streaming, Virtual Private Networks (VPN), etc. Data networks encompass both fixed and mobile access connections, for both residential and business customers

Properties of TIM data network
- Hierarchical and symmetric backbone topology
- Hierarchical-only interconnections allowed in Metro-regional network (tree-topology)
  - Dual-homed interconnections between aggregation routers and backbone routers
  - Single-homed interconnection between access network devices and aggregation routers

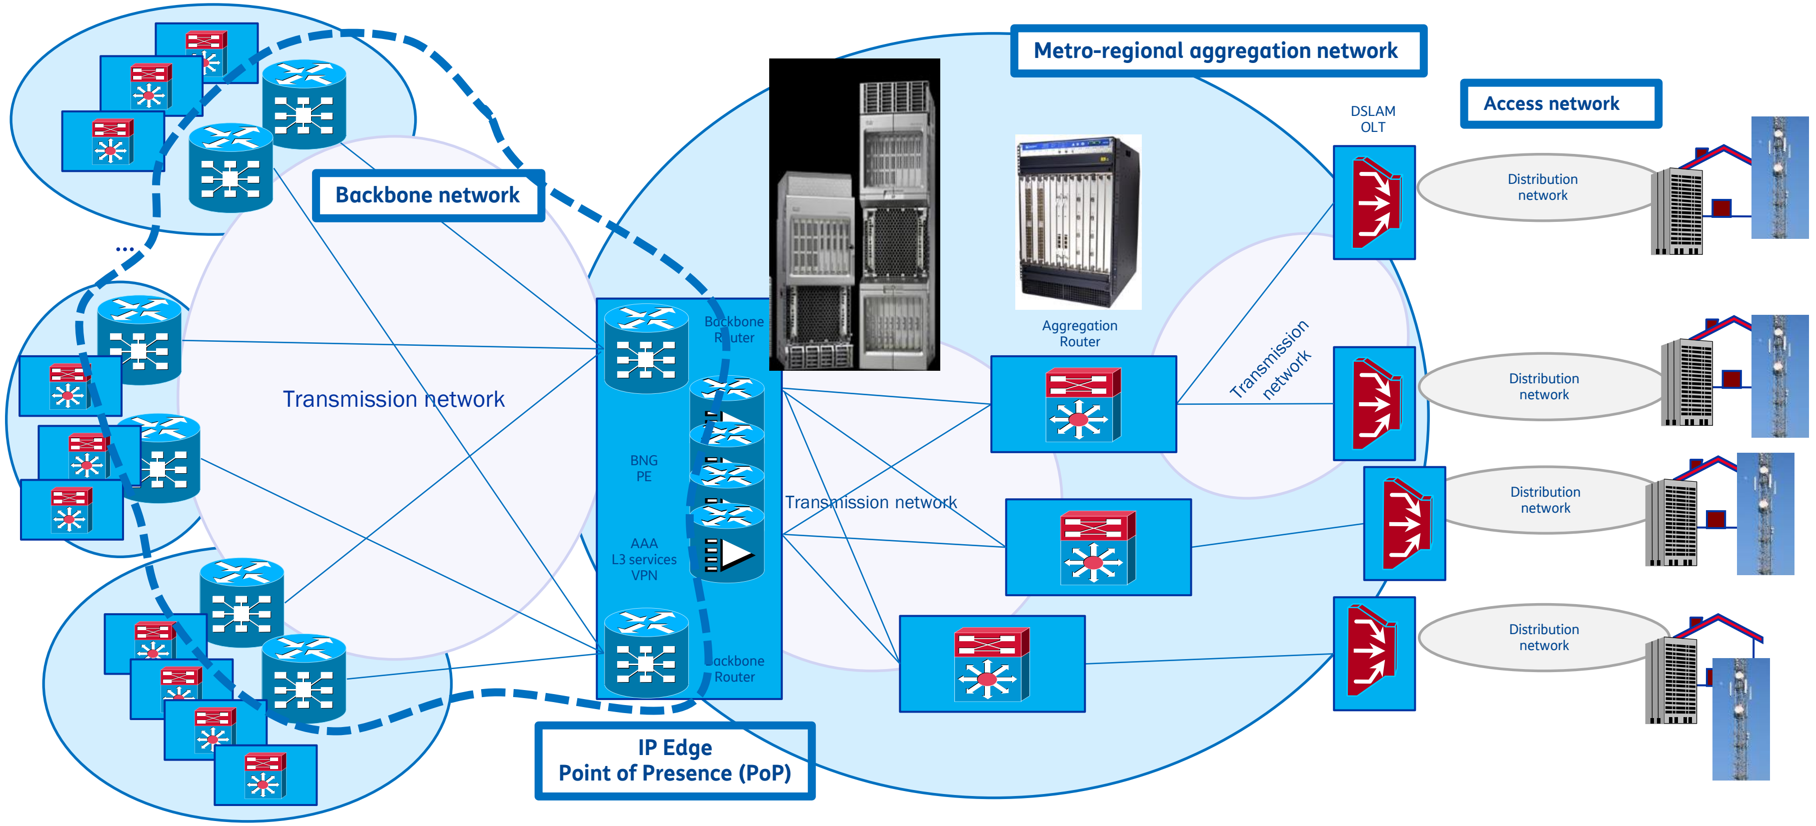

Technology overview of modern data networks
- Packet-switched network (statistical multiplexing)
- End-to-end traffic relationships have variable bandwidth
- Traffic is usually highly asymmetric: downstream traffic is much higher than upstream
- Traffic demand (hourly, daily, etc.) can have spikes not easy to predict
- Traffic relationships are not local, usually towards destinations outside national network
- Transmission networks are based on pure optical technology (Dense Wavelength Division Multiplexing)
- Links between routers are n x 100 Gigabit Ethernet network interfaces
- Backbone routers can forward hundreds of terabits per second (Tbps)
- Routers can have a pure forwarding role or be configured to offer additional services: virtual private networks, encryption, fine-grained Quality of Service management

Interconnection with other Internet Service Providers

Interconnection models
Internet Service Providers exchange traffic by means of either private peerings or public peerings. In both cases, the ISPs must provide a physical interconnection between their network to exchange traffic and also establish a routing relationship in order to guarantee the correct propagation of routing information.
Private (direct) peerings are direct interconnections between routers of two ISPs. The peering can be regulated by a commercial agreement if traffic directions are highly unbalanced or if one of the two ISPs is providing reachability to the Internet (in this case it is called an **Upstream Provider**)
Public peerings take place in ad hoc facilities, usually called Internet Exchange Points (IXP), where many ISPs install their own routers, extending the network up to that site. The IXP is usually run by a private or public company and provides the networking infrastructure to interconnect the routers.
In iXPs, excluding service fees, the peering relationships are usually free of charge

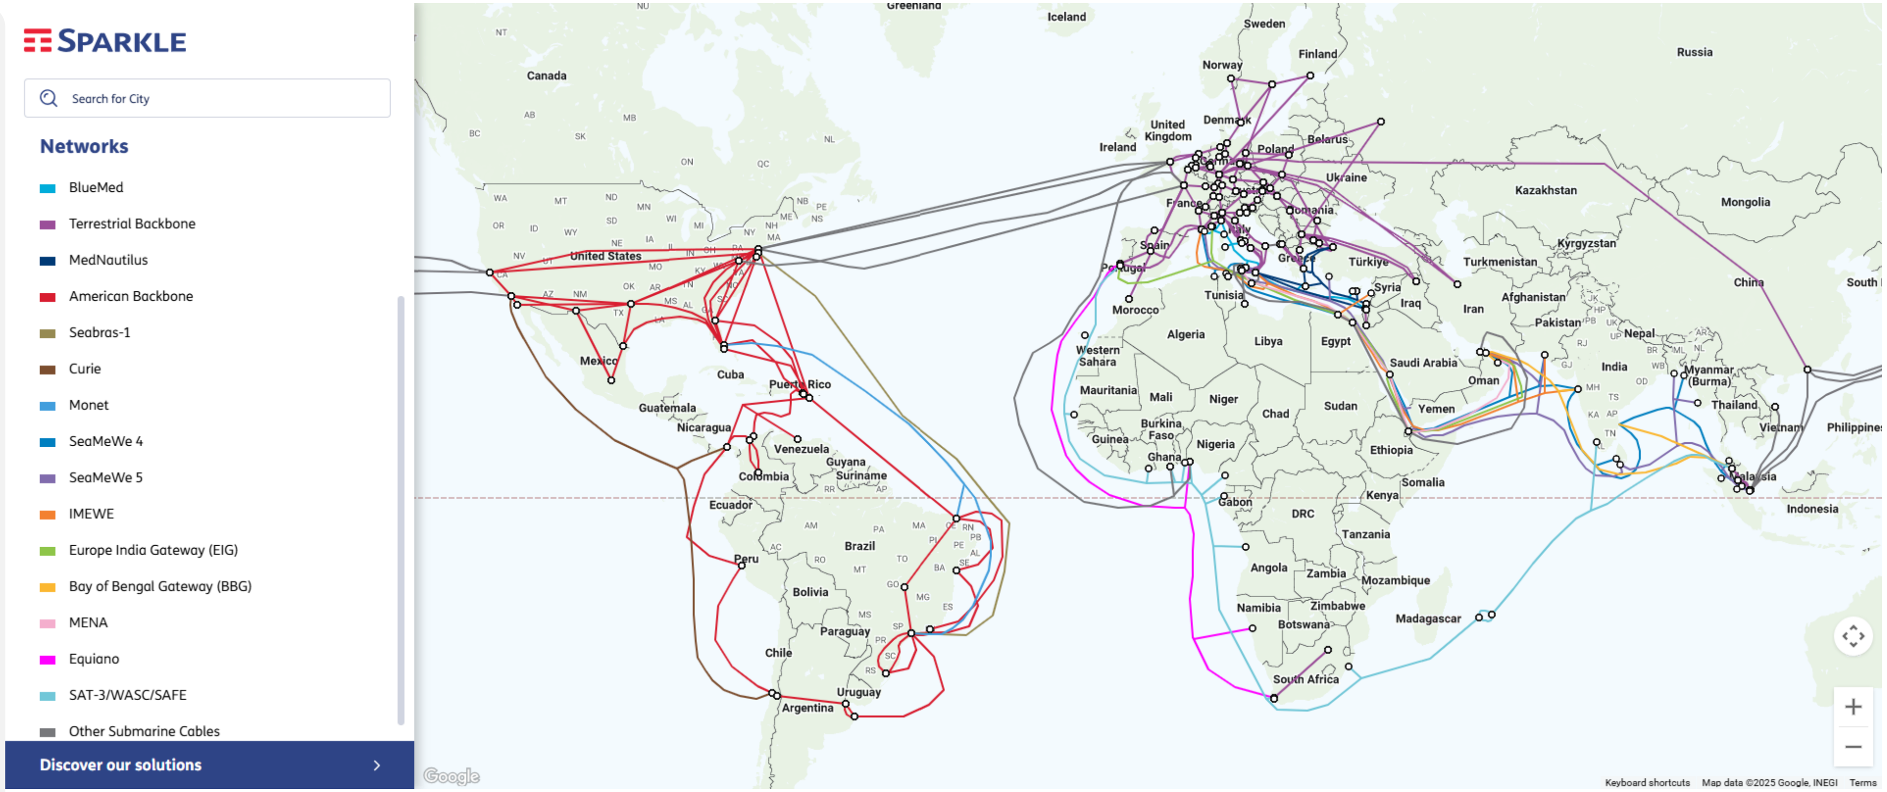

https://www.globalbackbone.tisparkle.com/

International networks
Largest ISPs have international networks to be present in the most important IXPs all over the world.
This makes possible to reach most of the Internet without paying fees to upstream providers

Backbone network

Multi-layer network
The optical transport network provides physical connectivity for IP routers located in different central offices.
Therefore, the optical network is the serving layer for the IP network (client layer).
Optical devices, like ROADM (Reconfigurable Optical Add-Drop Multiplexers) are connected to local IP routers by means of short-reach transceivers and optical fibres. ROADMs are interconnected by means of long-reach transceivers and optical fibres.

Technology considerations
The connectivity between IP routers is usually N x 100 Gbit/s interfaces.
Backbone routers offer an aggregated throughput in order of tens or hundreds of Terabits per second.
The underlying optical network is able to multiplex tens of 400 Gbit/s channels.

IP Layer topology

- The IP topology is purely logical: a network engineer can design it mainly considering traffic patterns and redundancy principles
- In principle, the optical network enables the possibility to establish adjacencies between routers located in any site
- But far sites require long optical paths that need expensive optical regenerations, making them inappropriate from an economic and engineering perspective

Optical Layer topology

- The optical topology is usually highly meshed for redundancy
- Long distance ducts carrying cables with optical fibres are often installed along existing civil infrastructures, like highways or railroads
- Long site-to-site hops may require optical regeneration that is usually expensive for very high-speed interfaces
- Optical circuits can be established between non-adjacent sites: the intermediate ROAMDS provide optical switching and signal regeneration

The backbone IP topology

Hierarchical topology
The backbone central offices (Points of Presence, are interconnected in a hierarchical shape.
This hierarchy is organized in 3 levels:

1. Inner core sites: 2 sites in Milan and two in Rome host routers that represent the central hub of the entire topology
2. Primary PoPs: the second level is connected to a pair of Inner Core routers (dual-homing) of the same geographical area (north or south) but in different central offices
3. Secondary PoPs: the third level is connected to two different Primary PoPs

Each pair of central offices is interconnected by means of two completely disjoint paths. The Inner Core has four disjoint paths. utilization.

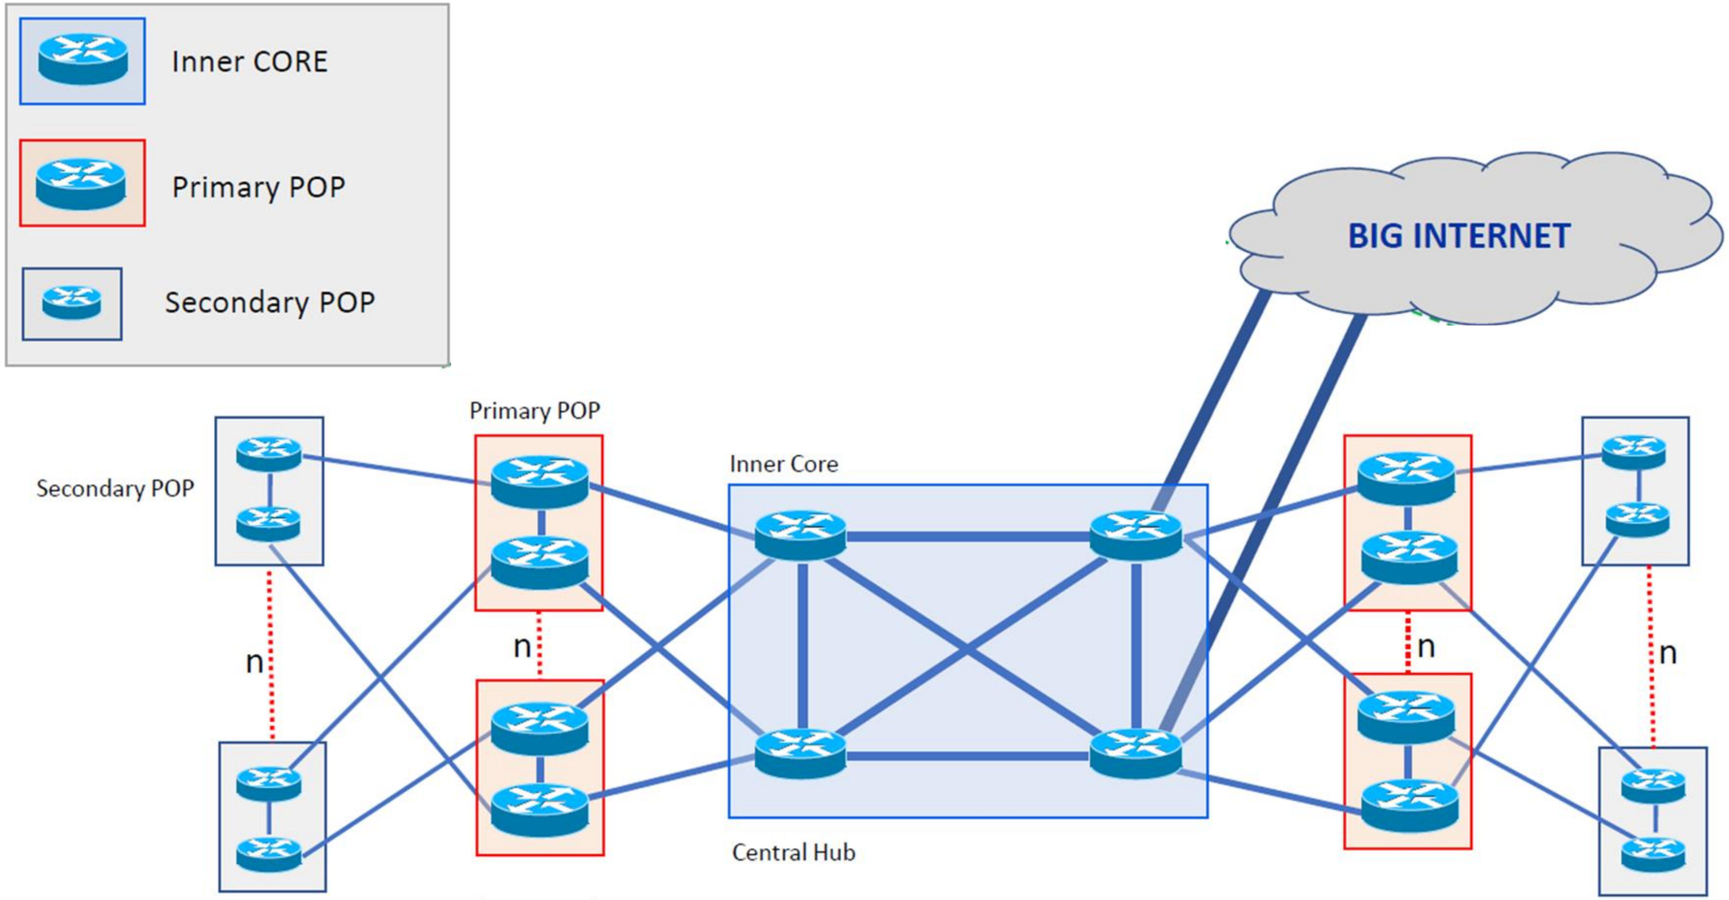

Routing symmetry
The cost of a path, from the routing protocols perspective, is usually the number of traversed sites (hops).
Equal-cost paths can be used by routing protocols to load-balance traffic flows.
The symmetry of the topology originates two equal-cost paths between any pair of PoPs (four in the Inner Core).
This means that all links can be used (almost) evenly, maximizing network utilization.

The IP Edge Point of Presence (PoP)

Role and internal structure
The IP Edge Points of Presence (PoP) are the main central offices of the network. They interconnect the backbone to the local metro-regional aggregation network and provide most of network services, for both fixed and mobile access.
In TIM network, the IP Edge PoPs usually host:

- Backbone routers
- Top level aggregation routers
- Service routers for residential (Broadband Network Gateway) and business (Provider Edge) services
- Content delivery appliances (e.g. caching systems)
- Mobile Core devices (MME, Serving and Packets gateways)
- A small datacentre hosting virtualized network functions.

The internal interconnections among the devices is provided by a couple of high throughput and high fan-out routers, called Service Hubs.

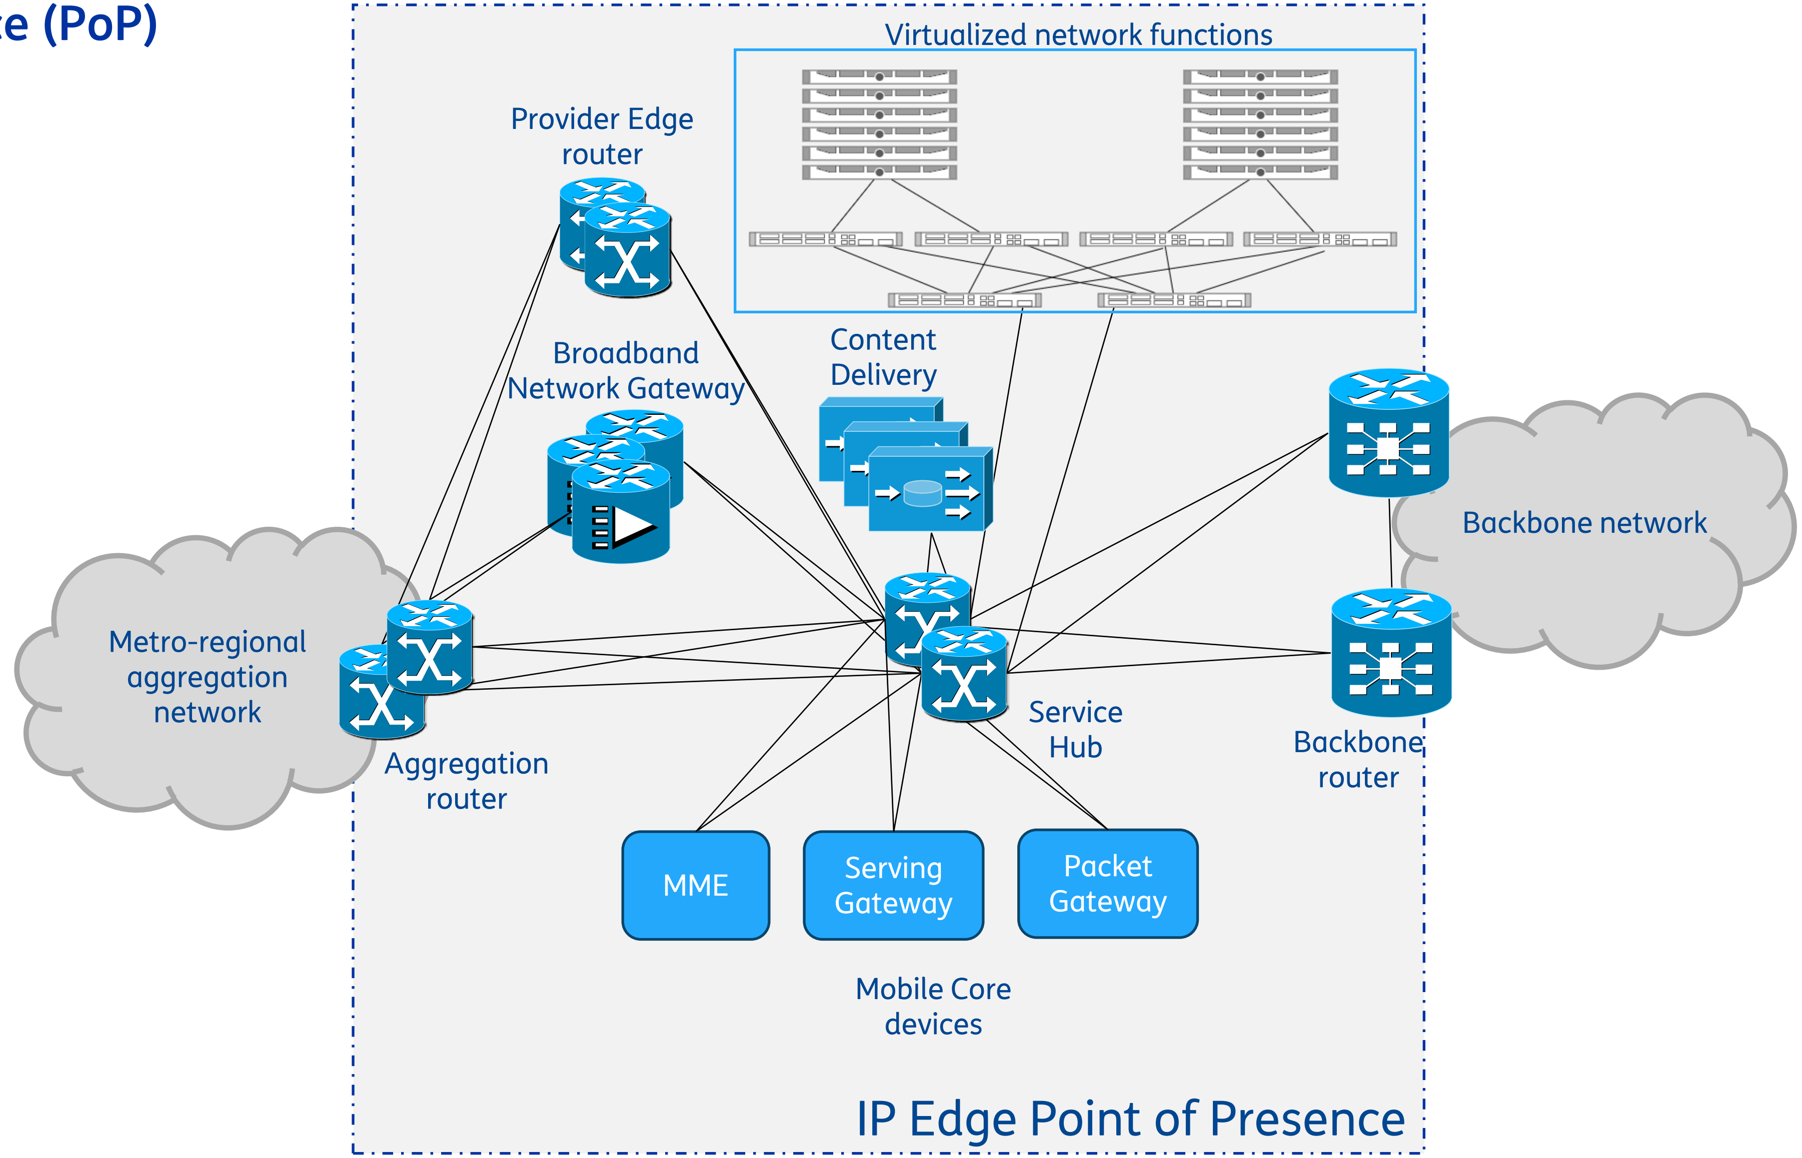

The Metro-Regional Aggregation Network

Role and internal structure
The Metro-regional aggregation network is responsible for aggregating and distributing traffic between the wired/wireless access network and the IP Edge PoPs.
Each IP Edge PoP is backed by an aggregation network.
The aggregation network has a hierarchical tree architecture, to improve traffic grooming, and leverages dual-homing for redundancy as far as it is economically sustainable.
Aggregation devices are IP routers, with 10/100 Gbps interfaces.
The aggregation network is a multi-layer network, like the backbone: the physical interconnections between aggregation routers are provided by means of the underlying Regional Optical Network.
An IP Edge PoPs, besides the edge device described in the previous slide, also hosts aggregation routers and access devices.
Smaller central offices host aggregation routers and access devices.

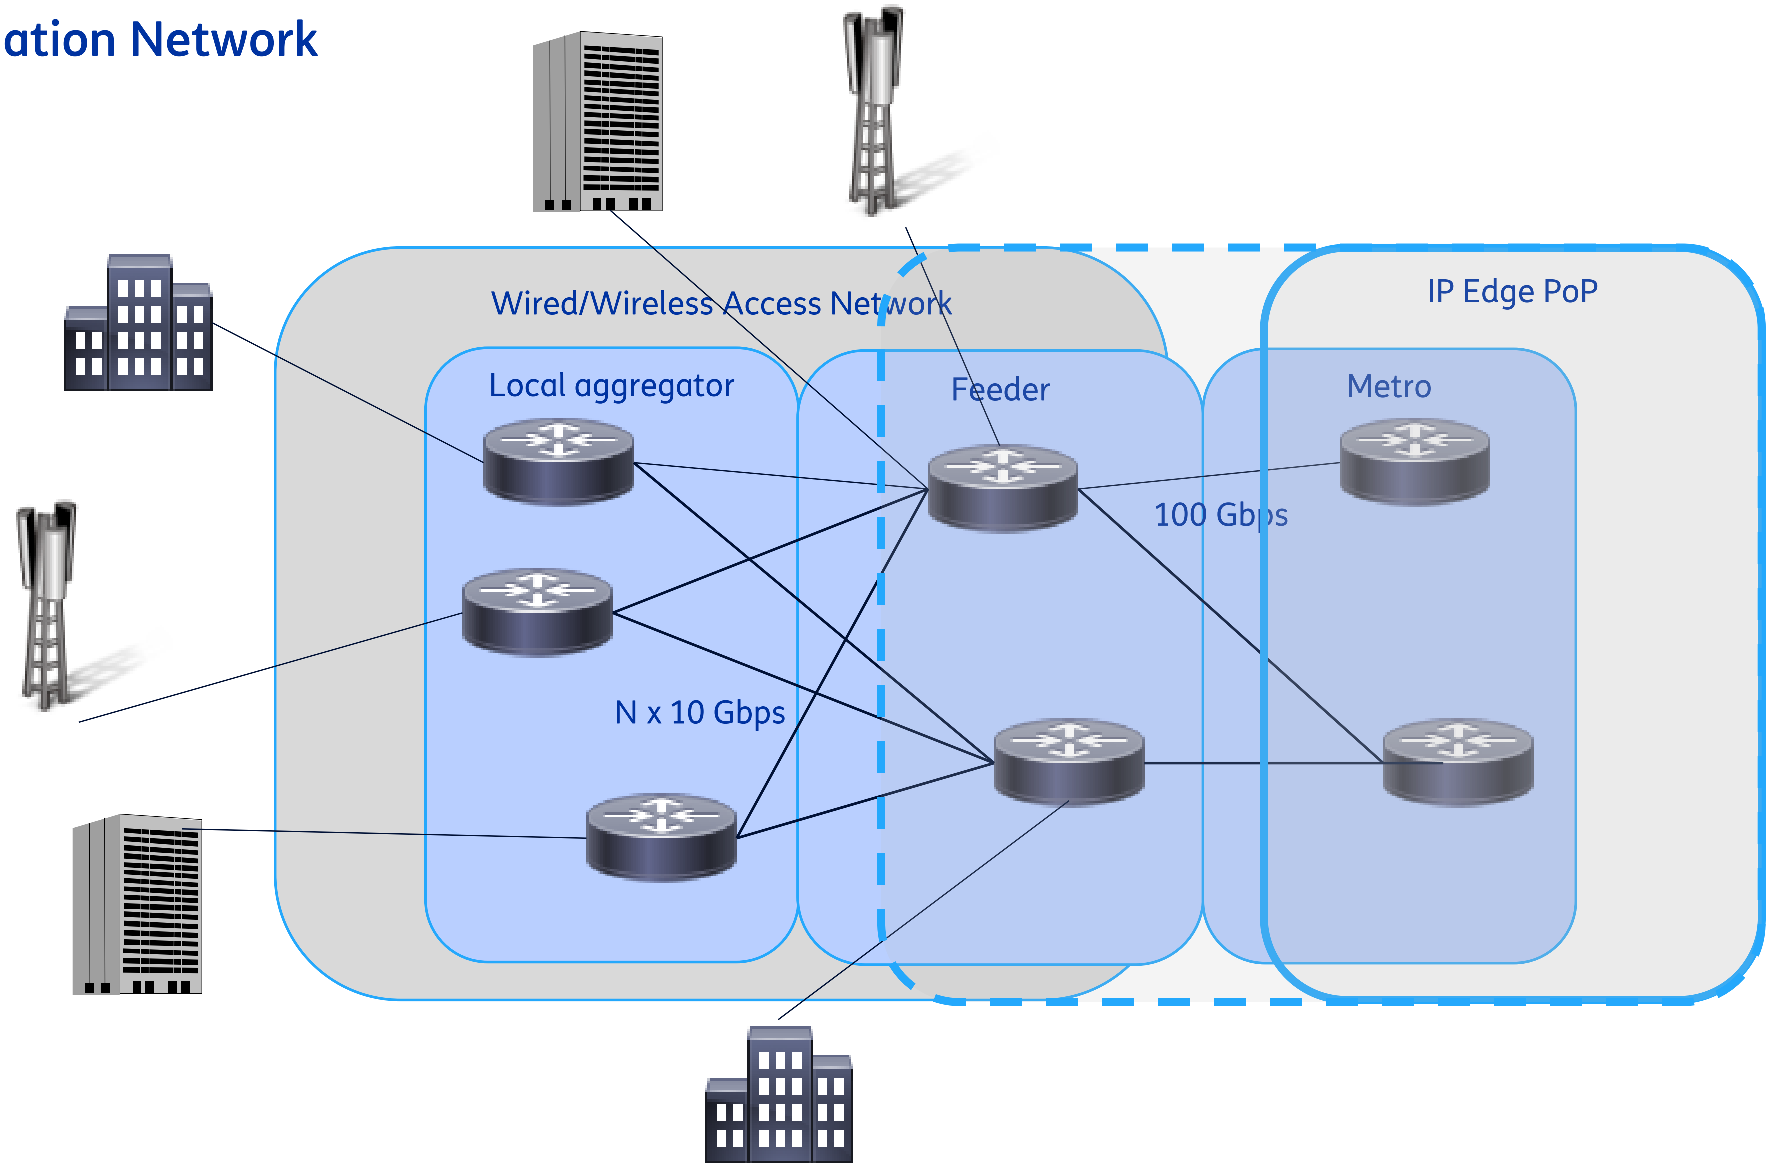

Network operations

Network operations

Network resource management

- It's the set of operational activities aimed at optimizing network performance and reliability.
- Performance is usually expressed in terms of packet loss and latency. Both are impacted by congestion and failure events.
- Congestion prevention is a primary goal of resource management.
- Reducing the impact of network failures requires synergies between resource management and topology design

Fault management

- A fault is a failure of a network component that negatively alter network behaviour and performance
- Failures may involve network device hardware, cabling, or software
- The goal of fault management is to recognize, isolate, correct and log faults that occur in the network
- One of the performance indicators of fault management is the Time To Repair

Traffic growth

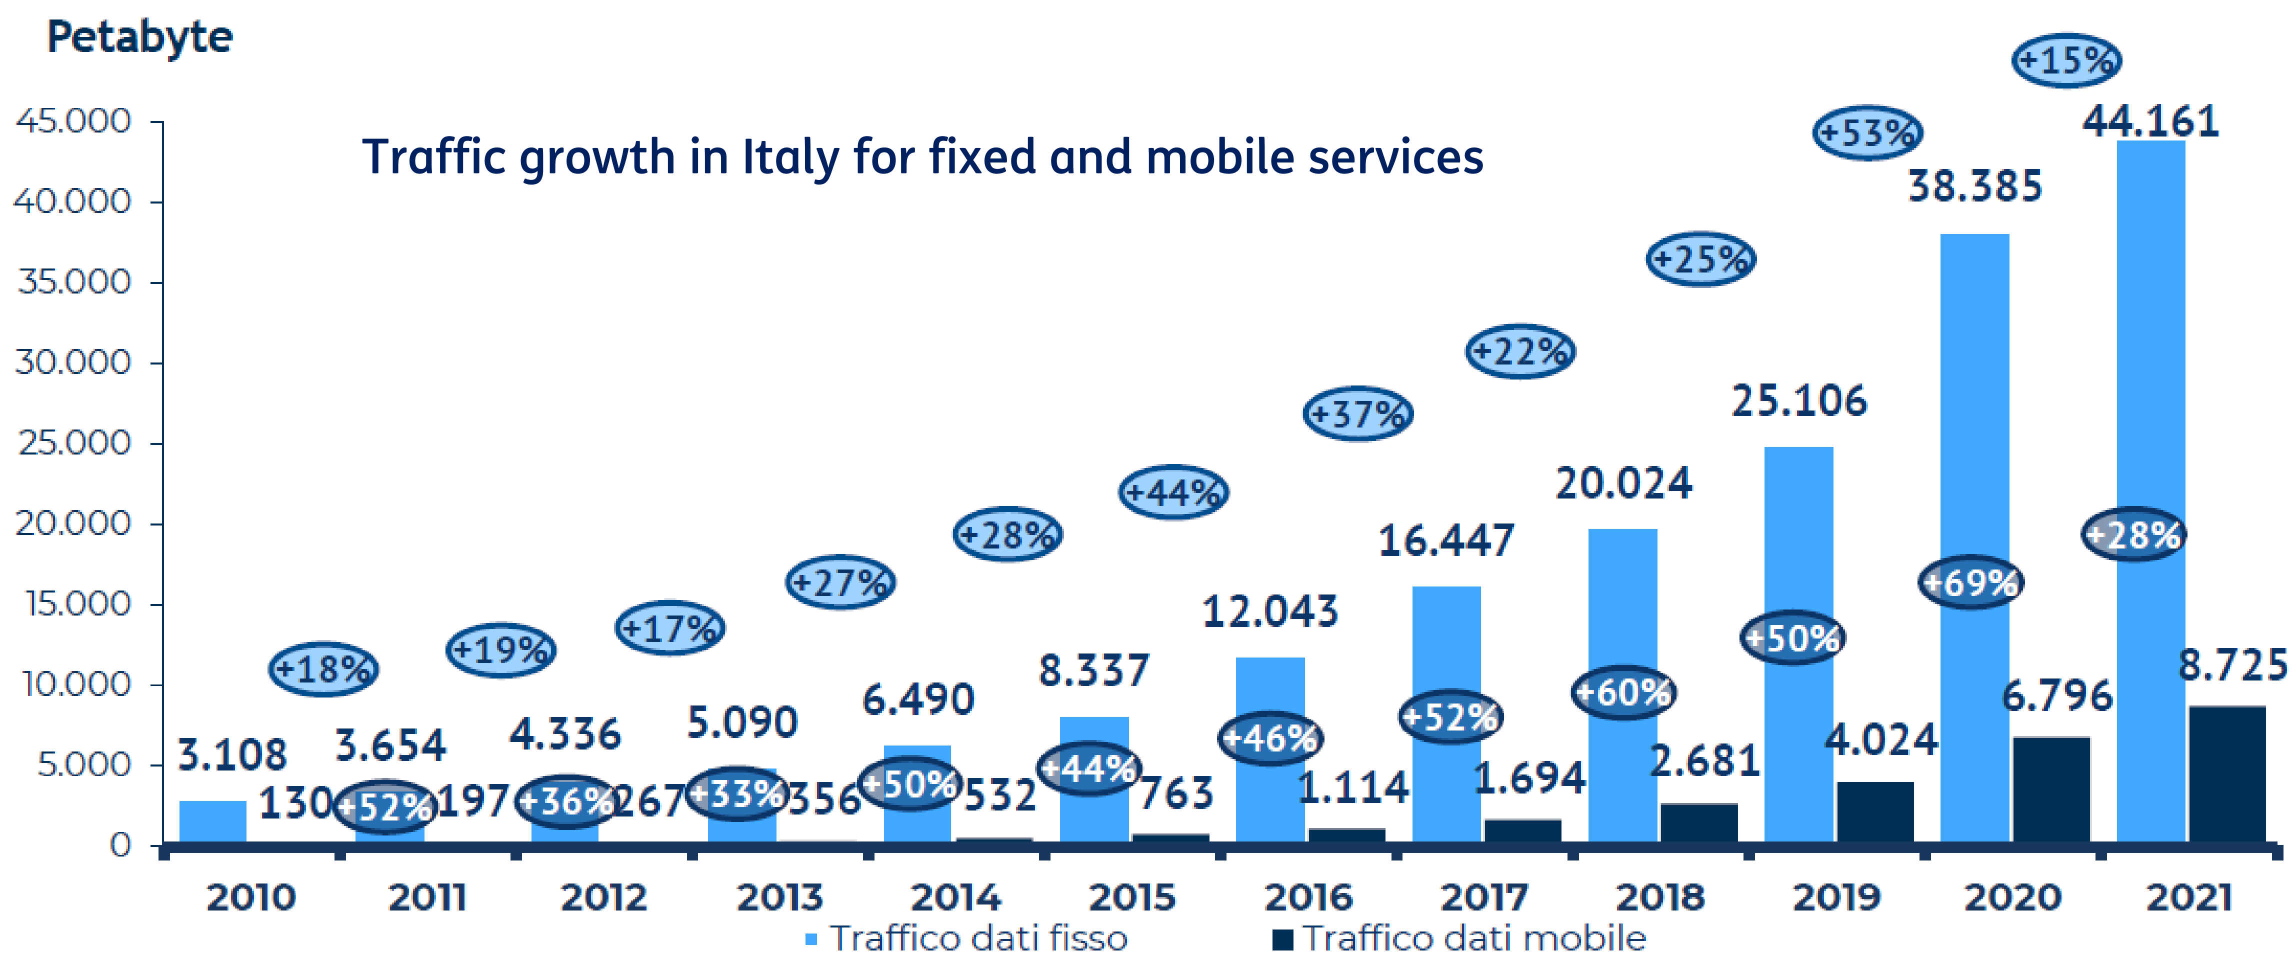

FONTE ELABORAZIONE OSSERVATORI DIGITAL INNOVATION POLITECNICO DI MILANO SU DATI AZIENDALI

Network resource management

Capacity planning

- Capacity planning estimates network upgrades to cope with forecasted traffic demand.
- Upgrades may refer to the IP and/or the optical network.
- Traffic forecasts are mathematical models fed by traffic measures.
- Capacity planning has a temporal horizon of months/years (budget cycle)

Traffic engineering

- Traffic engineering refers to mechanisms to route the traffic along optimal paths, to increase network utilization and satisfy traffic requirements.
- Routing protocols are the default technique
- Advanced tools are needed to granularly steer selected flows over customized paths (MPLS Traffic Engineering, Segment Routing)
- The temporal horizon is hours/days/weeks

Traffic management

- Traffic management reacts to congestion events
- Interface buffers saturate and packets are discarded
- Queue scheduling gives priority to selected flows
- Active Queue Management techniques discard packets in advance when they're experiencing long queueing delays
- Timeframe is short, seconds/minutes

Obsolescence of hardware and software components requires frequent replacements and updates.

TIM strategy for network resource management

Take advantage of topology properties

- The network design avoids single points of failure
- The symmetry of the topology enables load-balancing of traffic over all equal cost paths
- All internal links are dimensioned following the 50% rule: each link can't be loaded more than 50% of its capacity @ peak hour (**overprovisioning**). This protects in case of fault and in case of unexpected traffic surges.

No need for complex traffic engineering solutions

Overprovisioning protects from unexpected events

Traffic engineering powered by Software Defined Network concepts

- Traffic engineering techniques rely on a centralized computing element able to calculate optimal paths for traffic flows
- This is very close to the “controller” concept introduced by SDN but usually it complements (without replacing) the distributed routing algorithms implemented by routers’ control plane
- The controller must receive the complete topology of the network and collect link utilization data
- The controller needs networking tools to route specific traffic flows over specific paths that don’t necessarily are the shortest in terms of hops (MPLS Traffic Engineering, Segment Routing)
- Use cases:
  1. TE as the default technique (at least for a portion of total traffic) if the network utilization is strongly unbalanced
  2. TE used only to reroute (some) traffic over complex paths in case of network failures

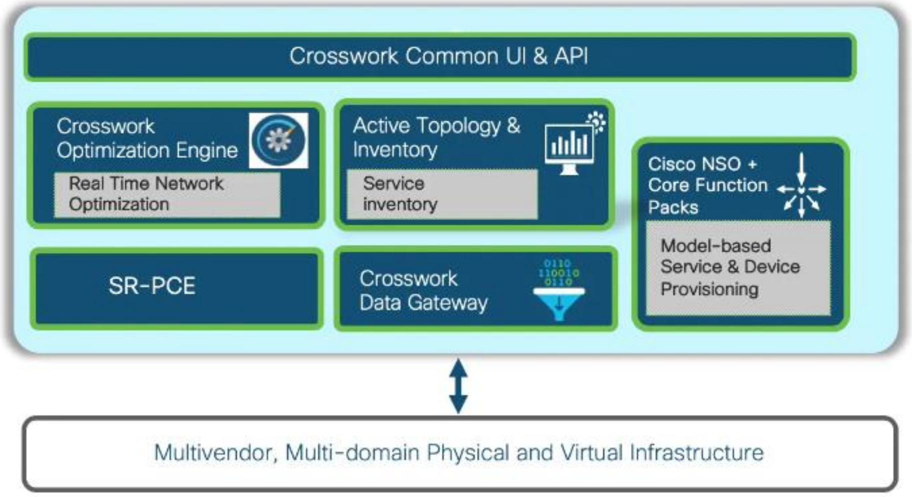

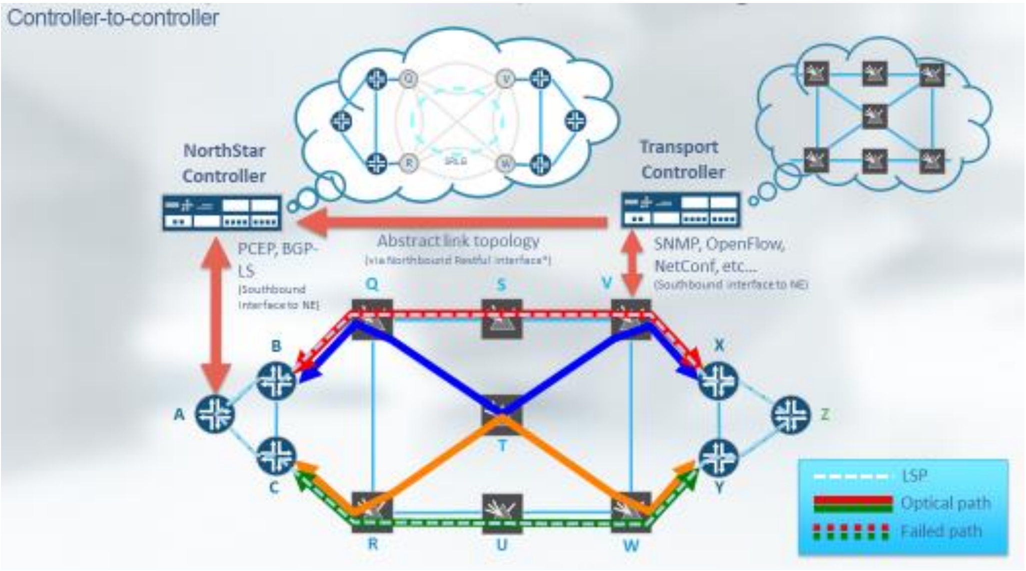

Traffic management: DiffServ QoS management

General idea

Traffic management includes techniques to reacts to congestion events. These techniques usually take the form of scheduling mechanisms at interface queue level to prioritize some traffic flows, according to the DiffServ model.

The goal is to protect critical traffic and penalize traffic that is less impacted from a performance point of view.

When capacity planning strategy uses overprovisioning, congestion events usually occur in case of network failures. If the requirements of the 50% rule are satisfied, at least two simultaneous failures on the same end-to-end path should happen to cause congestion.

Implementation

Traffic is classified using IP precedence, MPLS EXP or Ethernet CoS fields.

Usually 3/4 queues are configured: a strict priority (typically for voice) queue and 2, 3 weighted queues for different data classes. Calibration of correct configuration requires many tests and deep understanding of traffic composition.

Traffic marking is also critical: you can't trust IP precedence of Internet IP packets… most of the traffic in the backbone is in the default class

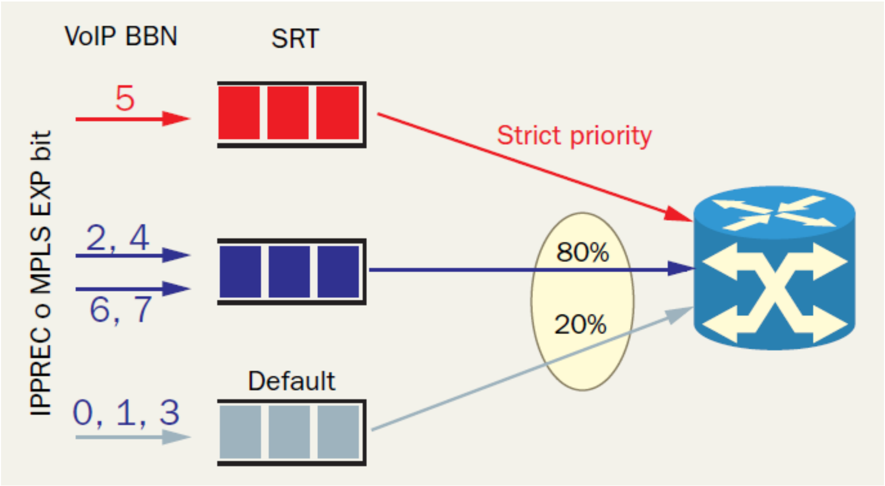

Traffic management: hierarchical QoS management for business and residential customers

Business customers

- Three queues for each customer
  - LLQ + 2 weighted queues (Mission Critical and Default)
  - Example of weights is 30:70
- Hierarchical queuing
  - multiple customers are multiplexed on a single physical interface

Residential customers

For a residential access, only a peak bandwidth is defined

The network applies a limitation of the peak bandwidth at the maximum value defined by the selected offer

For each customer, 1 or 2 queues (LLQ and default) are configured

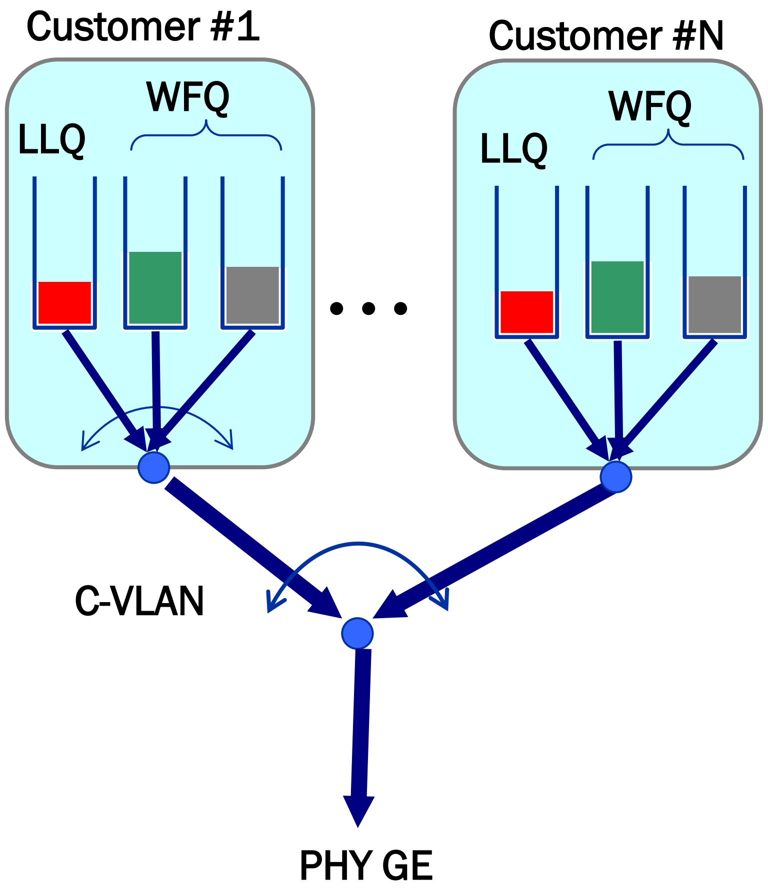

Configuration management

Configuration management includes tools and procedures to manage the configuration of network devices
Two categories of devices:

1. Network devices with a proprietary Network Management System (usually optical and mobile network devices): it's a software that is used to configure the devices and collect metrics and alarms; usually it has to be integrated into the Service Provider's own management systems
2. Network devices without an NMS (usually IP routers): they used to be configured using the Command Line Interface; now their configuration is managed automatically by means of NetConf/RestConf protocols and APIs. Blocks of configuration can be described using modelling languages (e.g. YANG) and translated to the specific syntax of the vendors

Day 0 / Day 1 configuration: refer to the basic configuration to make the device communicate with the SP's management systems (authentication, metrics and alarms collection, etc.) and with the other network devices (protocols configuration); it is usually applied at the beginning and changed only in case of updates or reconfigurations
Day 2 configuration: refer to the configuration provided to add a new customer (e.g. a new VPN, a new QoS profile, etc.); it is applied asynchronously (when needed) during the lifetime of the device

Common (complex) operations on network devices

Hardware upgrades

- Replace a linecard (e.g. a linecard with 10 Gigabit Ethernet Interfaces with a new one with 100 Gigabit Ethernet interfaces)
- Replace a device (e.g. for obsolescence or to improve performance)

Software upgrades

- Upgrade the device’s operating system: if In Service Software Upgrade is supported by the device, the operation can be done during normal operativity, otherwise the device must be put in maintenance state, out of traffic paths

- These activities usually involve a **lot of devices**
- Procedures must be carefully engineered and **validated in lab**
- Traffic disruptions must be avoided: usually the most critical activities must be performed **by night** when traffic is lower
- Planning must be defined accordingly, considering the number of devices to update, the complexity of the operation, the number of people involved, etc.: it may take **months** in case of tens or hundreds of devices

Fault management

Performance data is continuously collected from network devices: interfaces load, discarded packets, CPU load, memory occupancy, etc.
Additional performance measurements are performed using active probes that send artificial traffic and collect results: end-to-end latency, round-trip time, packet loss, …

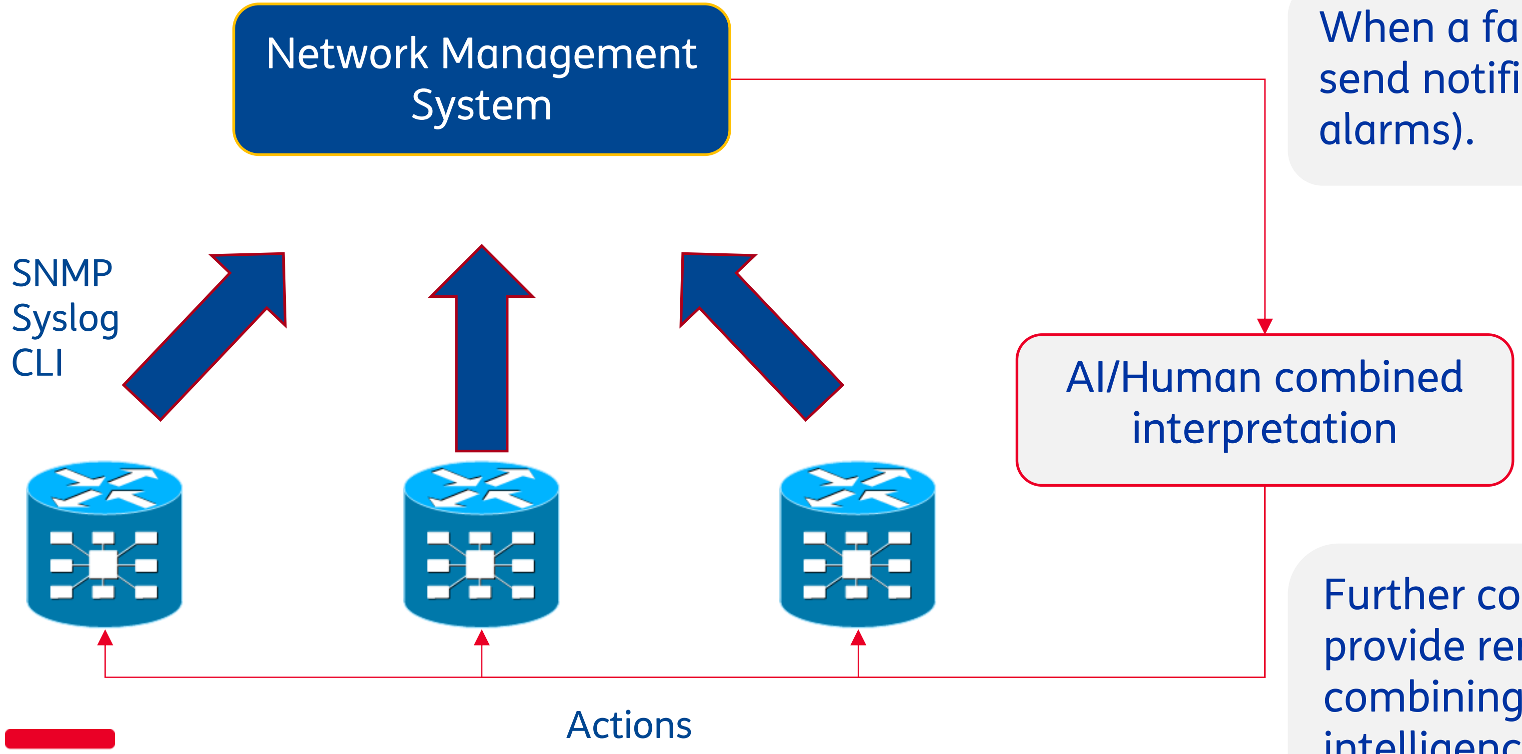

When a fault event occurs, impacted network devices send notifications to the management systems (traps, alarms).

In turn, management systems filter and perform a first events correlation, making the network operating centre (NOC) aware of the problem and allowing appropriate action to be taken.

Further correlations to identify the root cause and provide remediation actions are usually performed by combining software tools, guided by artificial intelligence algorithms, and human expertise.

Closed loop

Analytics

Provide insights based on collected data (e.g. machine learning)

Intelligence

Provide specific decisions and recommendations
AI models, policies and intents

Analysis

Decision

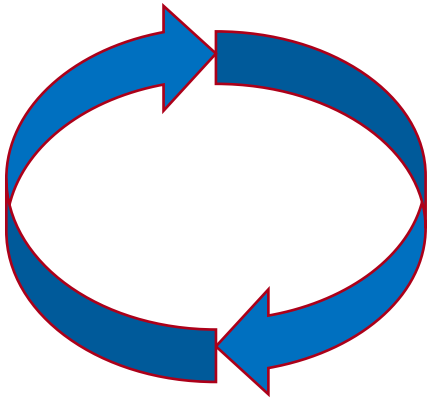

Collection

Execution

Data collection

Monitor the manages entities and provide live fault and performance data

Orchestration & Control

Automate workflows to handle lifecycle management entities
Individually steer the state of managed entities

Network function virtualization

Disaggregation of network devices

A common technology trend is the shift from monolithic systems to disaggregated architectures. Considering the network domain, this means that hardware and software components can be designed, developed and built by different companies, as long as they follow standardized interfaces that ensure interoperability.

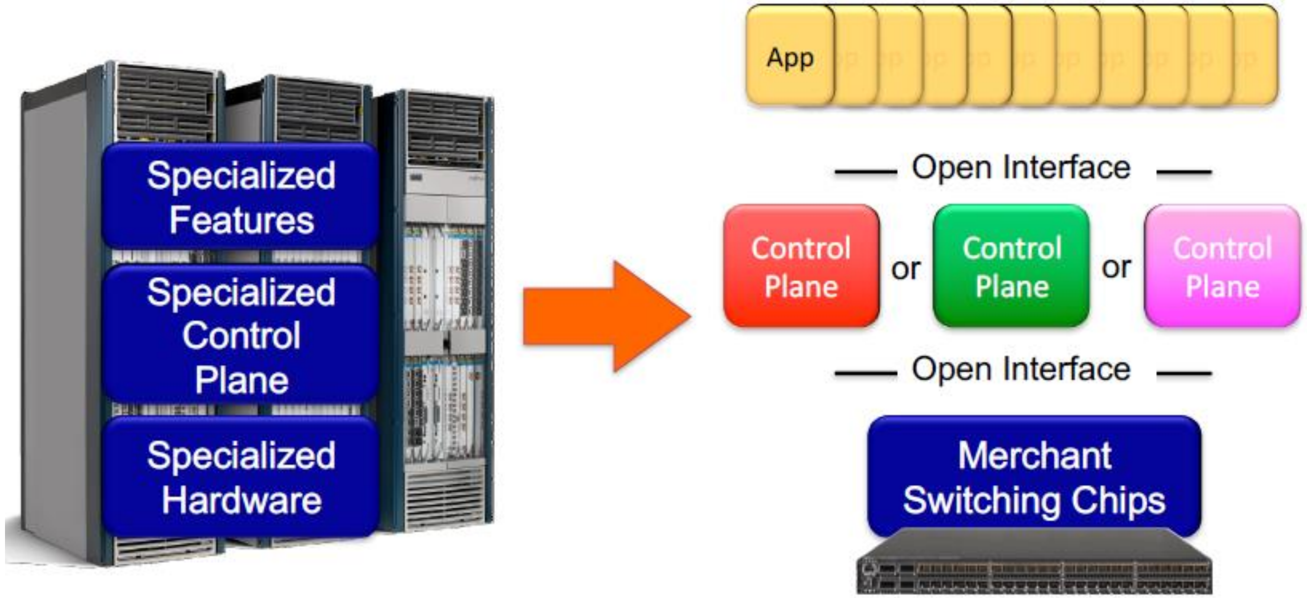

There are vendors that build so called "whitebox" switches, integrating networking chipset provided by a small number of large manufacturers (e.g. Broadcom). The CPUs of these switches can host operating systems developed by different companies.
Capital costs are usually lower, but the availability and quality of support services are often limited.

Whitebox switches can fit use cases where performances are not extreme (e.g. in enterprise networks) and configurations are standard.
In large Service Provider networks, IP routers are usually still vertically integrated systems.

Images from Nick McKeown, Stanford University

Network function virtualization

Network Function Virtualization can be interpreted as a further step in the disaggregation process:

- Specialized chipsets for networking applications are replaced by servers with general-purpose CPUs
- Software components implementing control plane and/or data plane features run in virtualized environments, such as virtual machines or containers.

This is possible because modern multi-core CPUs, even if not specifically designed to forward packets, can provide a reasonable throughput, usually in the order of a few ten of Gigabit per second.
Virtualization technologies help optimizing the use of computing and storage resources, ensuring isolation among different workloads.

The first wave of NFV relied on virtualization infrastructures built on **virtual machines**, like **Openstack** or **VMware** (now Broadcom) vSphere. The network functions developed for such virtualized environments are usually referred to as Virtual Network Functions (**VNF**).
The current wave uses **containers** instead of virtual machines and leverages orchestration platform like **Kubernetes**. In this case network functions are usually called Cloud (or Containerized) Network Functions (**CNF**).

Cloud network functions

While virtualization based on virtual machines typically preserves the traditional architecture—essentially replacing physical appliances with virtualized equivalents—container-based solutions enable a more fundamental rethinking of system design.
With containers and cloud-native practices, network functions are increasingly being disaggregated into smaller software components aligned with microservices principles.

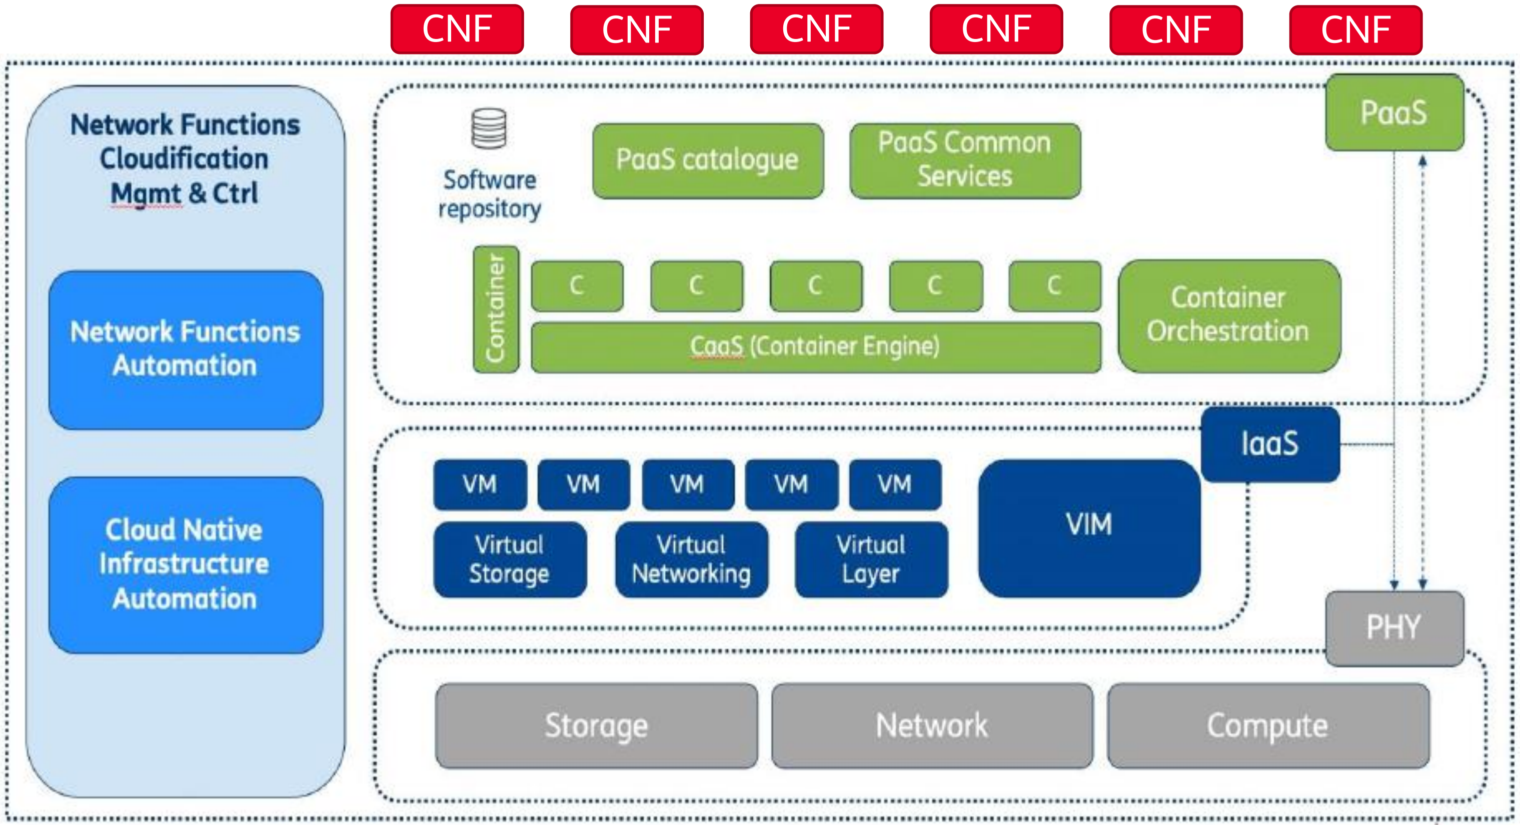

Not all network functions can be broken down this way, but the general trend is toward decomposing monolithic VNFs into autonomous units that can be developed, deployed, and scaled independently. These components commonly communicate through RESTful APIs, although modern implementations increasingly adopt alternatives such as gRPC or messaging systems when higher performance or asynchronous exchanges are required.

Cloud Edge Continuum

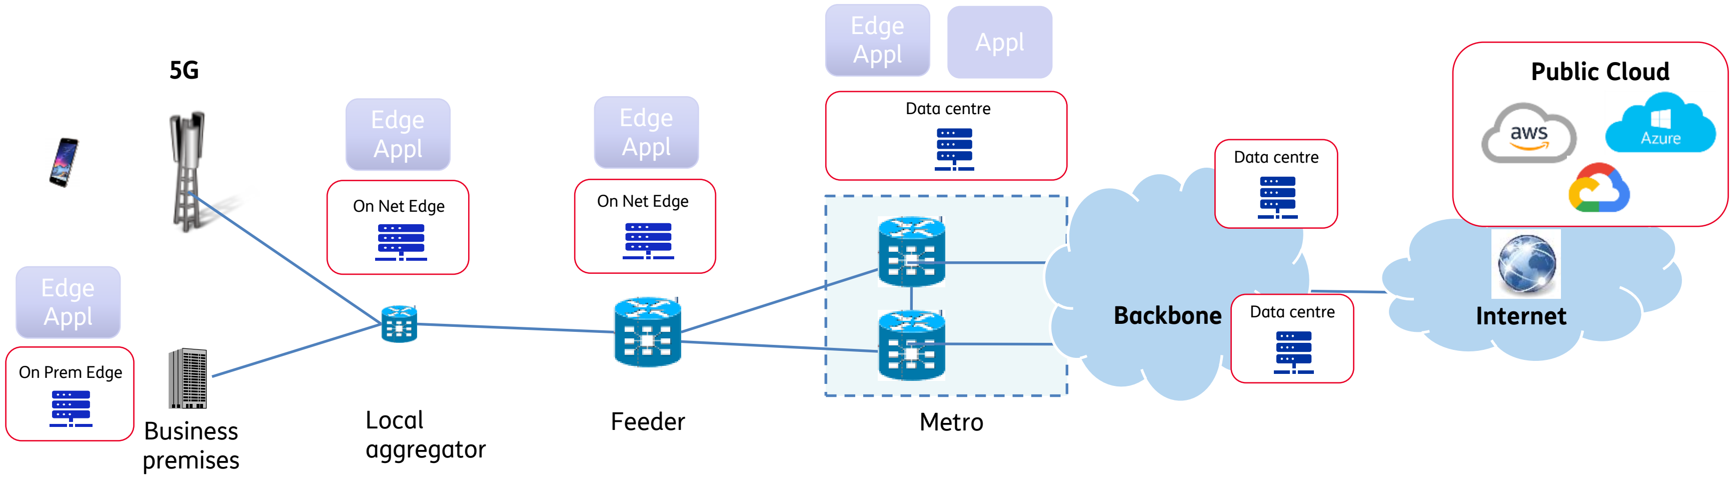

Edge Solution Capability	Example of use case
Low Latency Ultra-low latency (≈ ms) is critical for real-time type of application	**Robotic (Industry 4.0)** **AR/VR (cloud) network centric** **Mobile gaming**
Local Processing Elaboration of data produced locally to extract information and/or take decision avoiding large data transfer toward central D.C.	**Video surveillance** **IoT** **Content Caching** **Sport Event Experience**
Privacy / Security	**Data localization**
Limited Autonomy Ability to continue to run also when disconnected from central Data Center	**Private network services** **Enterprise / Campus Network**

In addition to its technical advantages, virtualization also enables significant architectural innovation. Edge Computing is a cloud-computing paradigm that moves data processing and storage closer to where data is generated or consumed—by devices, applications, and users.
On-Net Edge Cloud sites are small datacentres deployed in selected central offices within the metro-regional network. These sites can host CNFs for the 5G Core and/or IP Edge functions, enabling local breakout of access traffic toward nearby edge applications and services.
Together, all these datacentres, ranging from small edge sites to large regional and national facilities, constitute what is commonly referred to as the **Cloud Edge Continuum**.

Telco-specific requirements for cloud infrastructures

Cloud technologies were not originally developed to cope with the requirements of Telco applications:

- Multiple interfaces
- CPU optimizations
- Direct access to network interfaces (SR-IOV, etc.)
- Real-time kernels
- ...

While CPU performances are increasing fast, the packet throughput they offer is still much lower than what networking chipset or specialized ASICs are able to do.
The operating systems need to be fine tuned to optimize performance, often bypassing the OS kernel.
In many cases, optimizations are not enough and general purpose CPUs need to be assisted by smartNICs, hardware accelerators, etc.

At the same time, Telco vendors are progressively improving their CNFs in order to make them really cloud-native and able to take the full benefits from cloud technologies

Sylva project

The main carriers in Europe, together with network function providers, in **November 2022 initiated the Sylva project** to address Telco and Edge use cases

The project objectives are:

- To release a **cloud software framework tailored for Telco functionalities and edge requirements** that address the technical challenges of the industry layer of this ecosystem
- To develop a reference implementation of the cloud software framework and **create a validation program for such implementations**

https://sylvaproject.org/

IPCEI-CIS

Important Project of Common European Interest - Cloud Infrastructure and Services

IPCEI-CIS creates a "**Multi-Provider Cloud-Edge Continuum**" without being tied to a single provider. In terms of interoperability, portability and scalability, an advanced digital infrastructure is being developed that will significantly revolutionize current offerings on cloud and edge markets, which are dominated by third-party providers. This will enable the **seamless integration** of a variety of different new data processing solutions in the European Union in just a few years to support **novel applications** such as autonomous driving.

3.5
billion euros
will be made available for projects across Europe

113
projects by companies & research institutions
are involved in IPCEI-CIS throughout Europe

12
EU countries
participate in the IPCEI-CIS

Funded by the European Union
NextGenerationEU

Thanks

For any questions, feel free to contact me at

alessandro.capello@telecomitalia.it