NET-CENTRIC 2021

# IETF Network Slicing

Adrian Farrel (adrian@olddog.co.uk)

# Slicing the Presentation

- What and Why “Network Slicing”?
- 3GPP Perspective
- Is there anything new here?
- IETF Network Slicing
  - Architecture
  - It’s all about the Terminology
  - Slicing as a Service : the Northbound Interface
- Solution Perspectives
  - ACTN
  - VPN+
  - SR Policy
  - Slicing Anything
  - The TeraFlow Project

# 30,000 Thousand Foot Perspective

- Many new advanced services and applications
  - Holographic telepresence
  - Augmented reality
  - Collaborative gaming
  - Massive IoT
  - Public safety
- Some opinion that these applications and services are “enabled by 5G”
- New applications and services place very different requirements on the network
  - Low latency
  - High throughput
  - Resilience against other users
  - Predictable jitter
- Shared networks struggle to make adequate guarantees
  - Need to reserve/partition the network resources for specific user
- This “partition” is a “network slice”

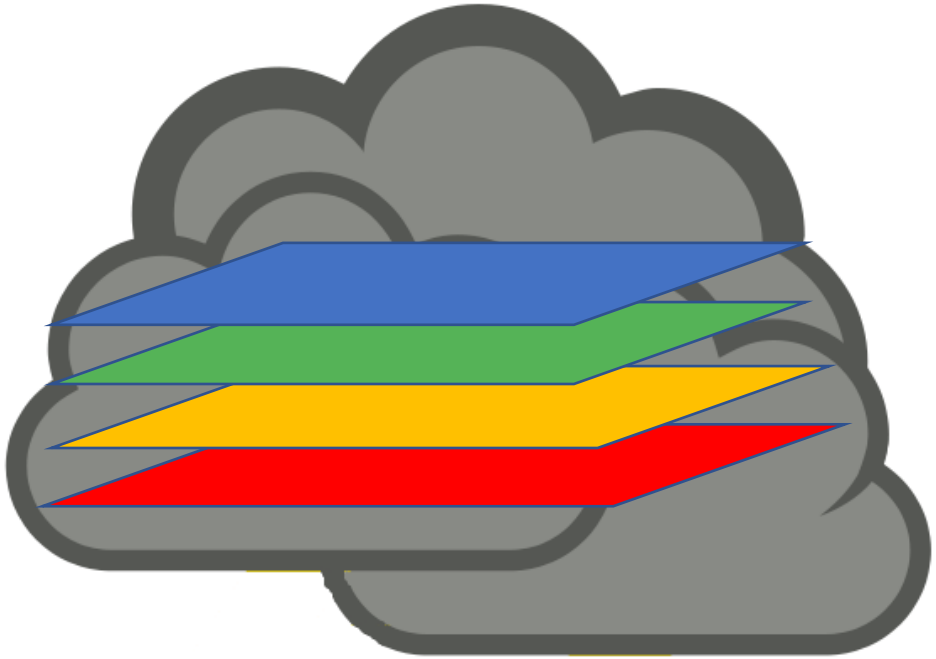

# 3GPP Proposals

- 3GPP focus
  - Slice the radio access network (RAN) and core network (CN)
    - Make sure resources are available for 5G applications
  - Place dependencies on “transport networks” (TN)
    - Deliver quality connectivity for end-to-end (E2E) network slices

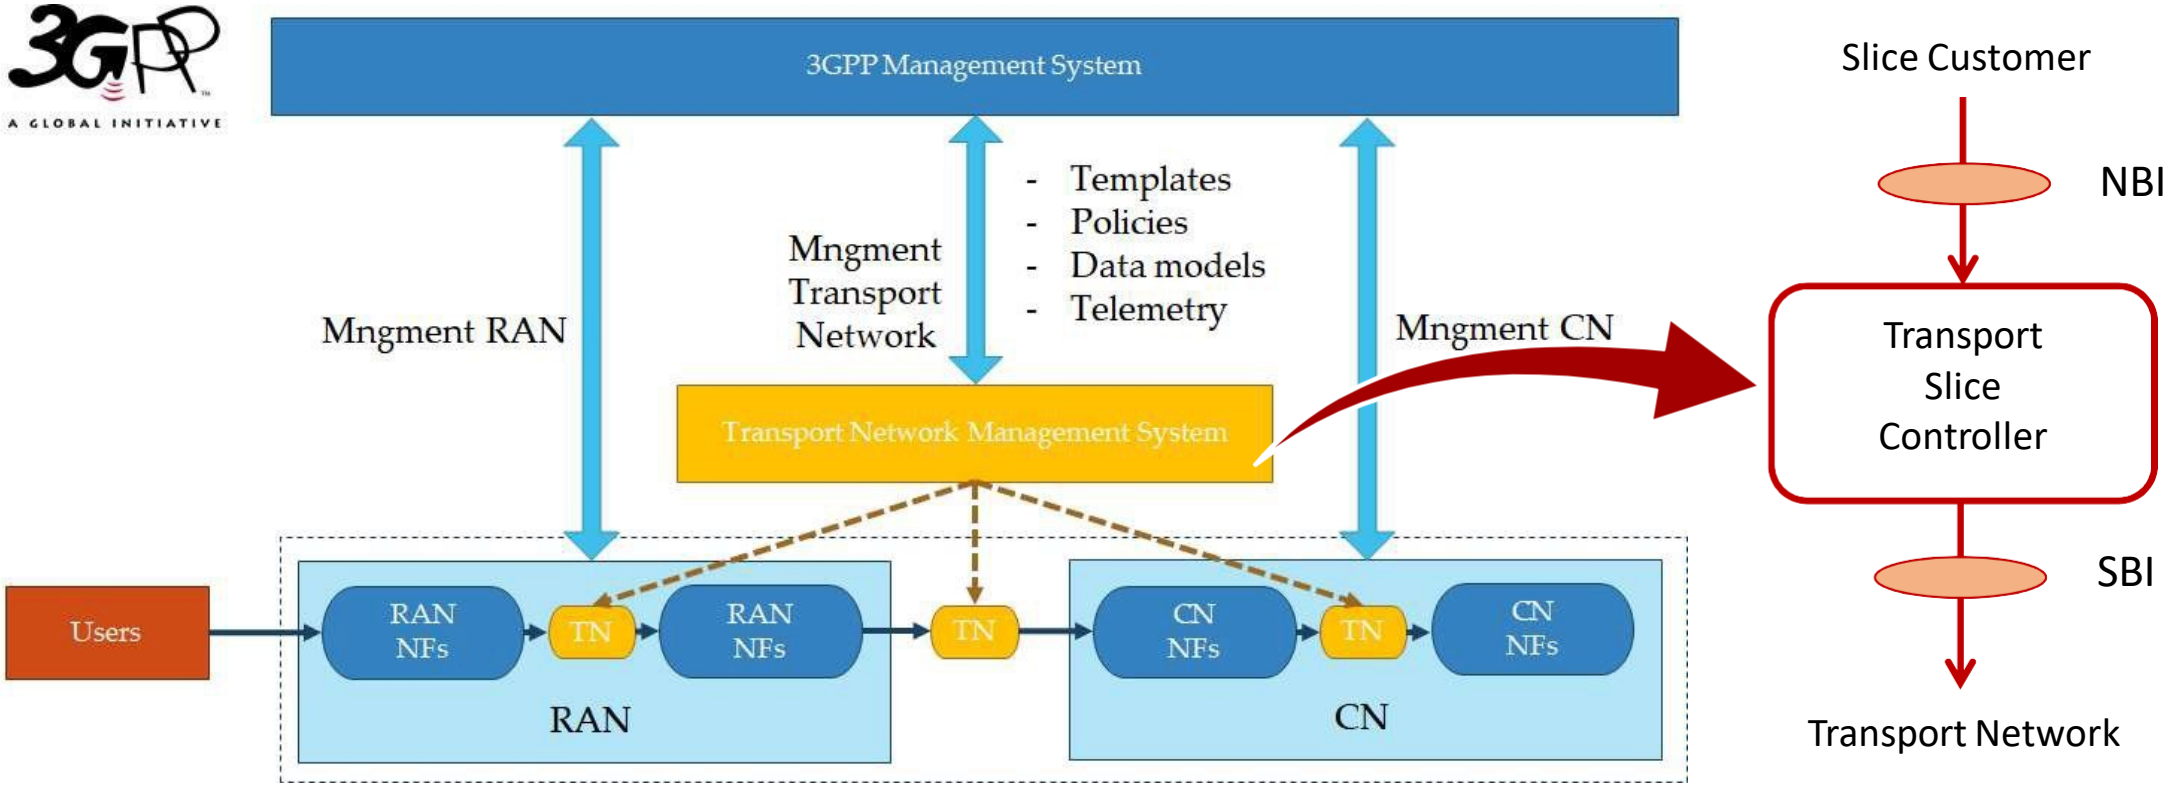

# What Goes Around Comes Around

- We have been building “slices” in the IETF for many years
  - Virtual Networks
  - Abstract Topologies
  - Resource Reservations
  - Traffic Engineering
  - (see draft-ietf-teas-rfc3272bis)
- What’s new this time?
  - Specific request by the “user” for a slice
  - Resource partitioning tied to specific “service level objectives”
  - Collection of multiple “connectivity constructs” into one managed entity

# Some IETF History on Network Slicing

- 3GPP TR 28.801 V1.0.0 (2017-03) gives a generalised definition

  *A network slice instance is: a set of network functions and the resources for these network functions which are arranged and configured, forming a complete logical network to meet certain network characteristics.*

- Side meetings at IETF-97 and 98 (2016/17) for discussion
- BoF at IETF-99 (2017) and at IETF-101 (2018)
- TEAS Working Group Design Team formed 2019
  - draft-ietf-teas-network slice-definition
  - draft-ietf-teas-network-slice-framework
- draft-ietf-teas-ietf-network-slices
  - Merger of Design Team documents
  - Work in progress to polish and develop wider consensus

# IETF Network Slice Service Definition

- Specified in terms of:
  - A set of endpoints (e.g., CEs)
  - A set of connectivity matrices between subsets of these endpoints
    - Think of these as virtual connections or tunnels
    - P2P, P2MP, MP2MP
  - Service level behaviours requested for each sender on each connectivity matrix
- Sounds very simple, but:
  - Result of a huge amount of discussion
  - A lot depends on the understanding of these three key elements

# Endpoints

- The points at which a network slice service is delivered
  - Those things that are connected in connectivity matrices
- There are four possibilities: all are supported
  1. CE is operated by slice provider. The edge is within the CE
  2. CE is operated by customer or provider. Edge is the CE port. Slice includes AC
  3. Edge is the PE port. Similar to a port-based VPN
  4. Edge is within the PE. PE performs classification of traffic onto the slice

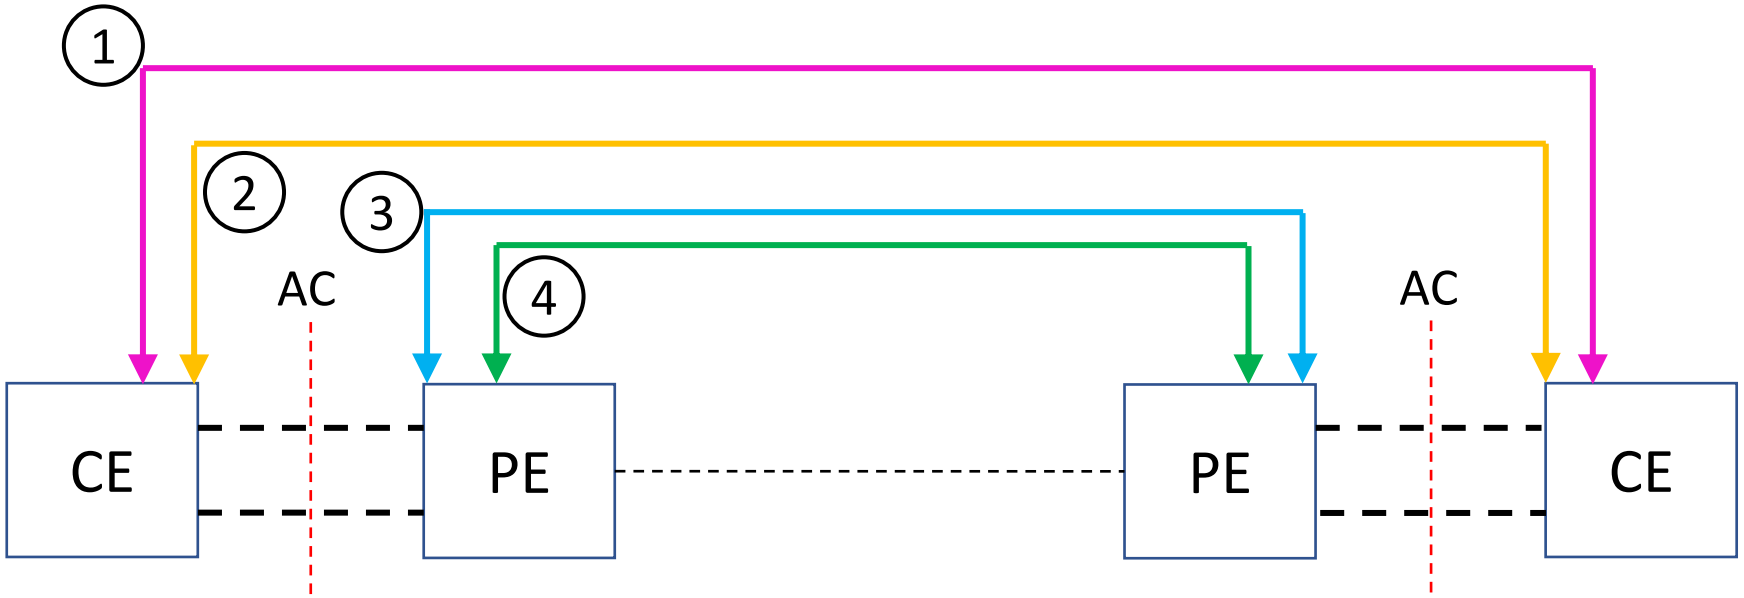

# The Service Level Agreement


- Service Level Indicator (SLI)
  - A quantifiable measure of an  aspect of the performance of a network
  - E.g., a measure of throughput in bits per second, or a measure of latency in milliseconds
- Service Level Objective (SLO)
  - A target value or range for the measurements returned by observation of an SLI
  - E.g., expressed as "SLI <= target", or "lower bound <= SLI <= upper bound“
  - A customer can determine whether the provider is meeting the SLOs by measurements on the traffic
- Service Level Expectation (SLE)
  - An expression of an unmeasurable service-related request that a customer makes of the provider
  - The customer has little or no way of determining whether the SLE is being met
  - But there is still a contract for a service that meets the expectation.
- Service Level Agreement (SLA)
  - An explicit or implicit contract between the customer and the provider of the slice
  - Expressed in terms of a set of SLOs and SLEs to be applied to the connections between endpoints
  - May include commercial terms and consequences of missing/violating the SLOs

# IETF Network Slicing Management Architecture

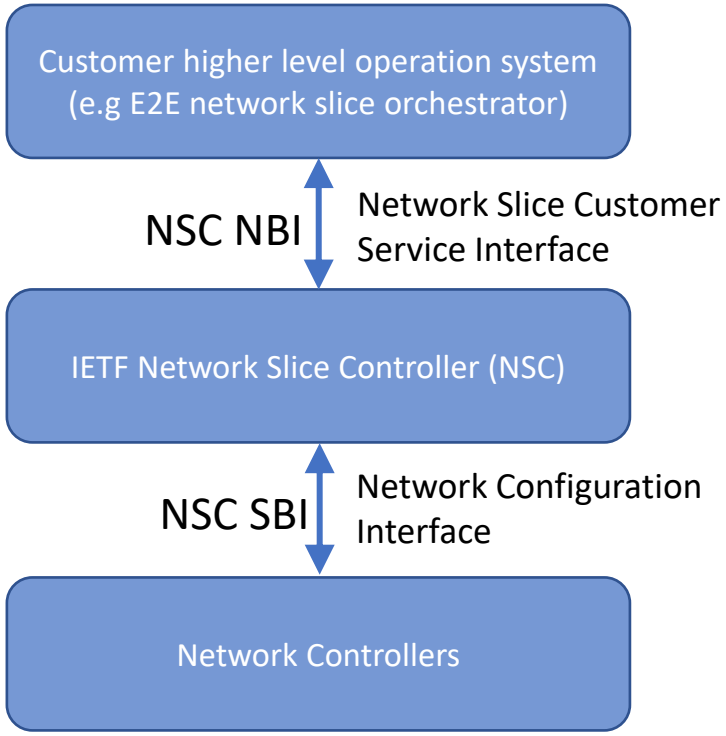


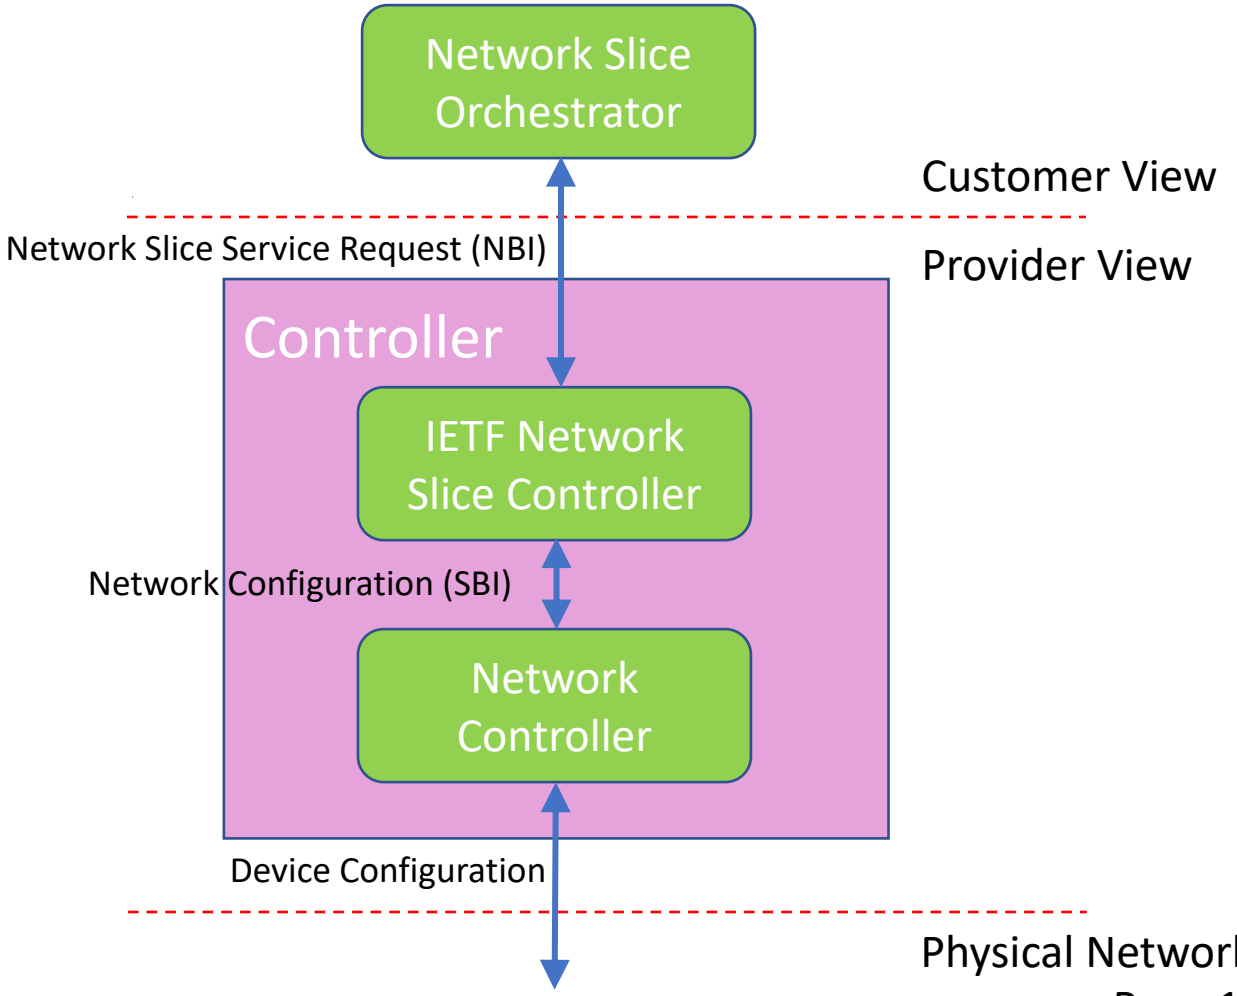

# IETF Network Slice Service Interface

- Customer Service Interface
  - This is how a customer requests a service form a provider
  - Compare with Layer 2 VPN and Layer 3 VPN Service Models
    - See RFC 8466 and RFC 8049
- This is the IETF Network Slicing Northbound Interface (NBI)
- Needs to express
  - Customer details
  - Endpoints
  - Connectivity matrices
  - SLOs and SLEs
    - Complex parameterisation is needed
- draft-wd-teas-ietf-network-slice-nbi-yang
  - Early work on a YANG specification

# IETF Network Slice Generic Architecture

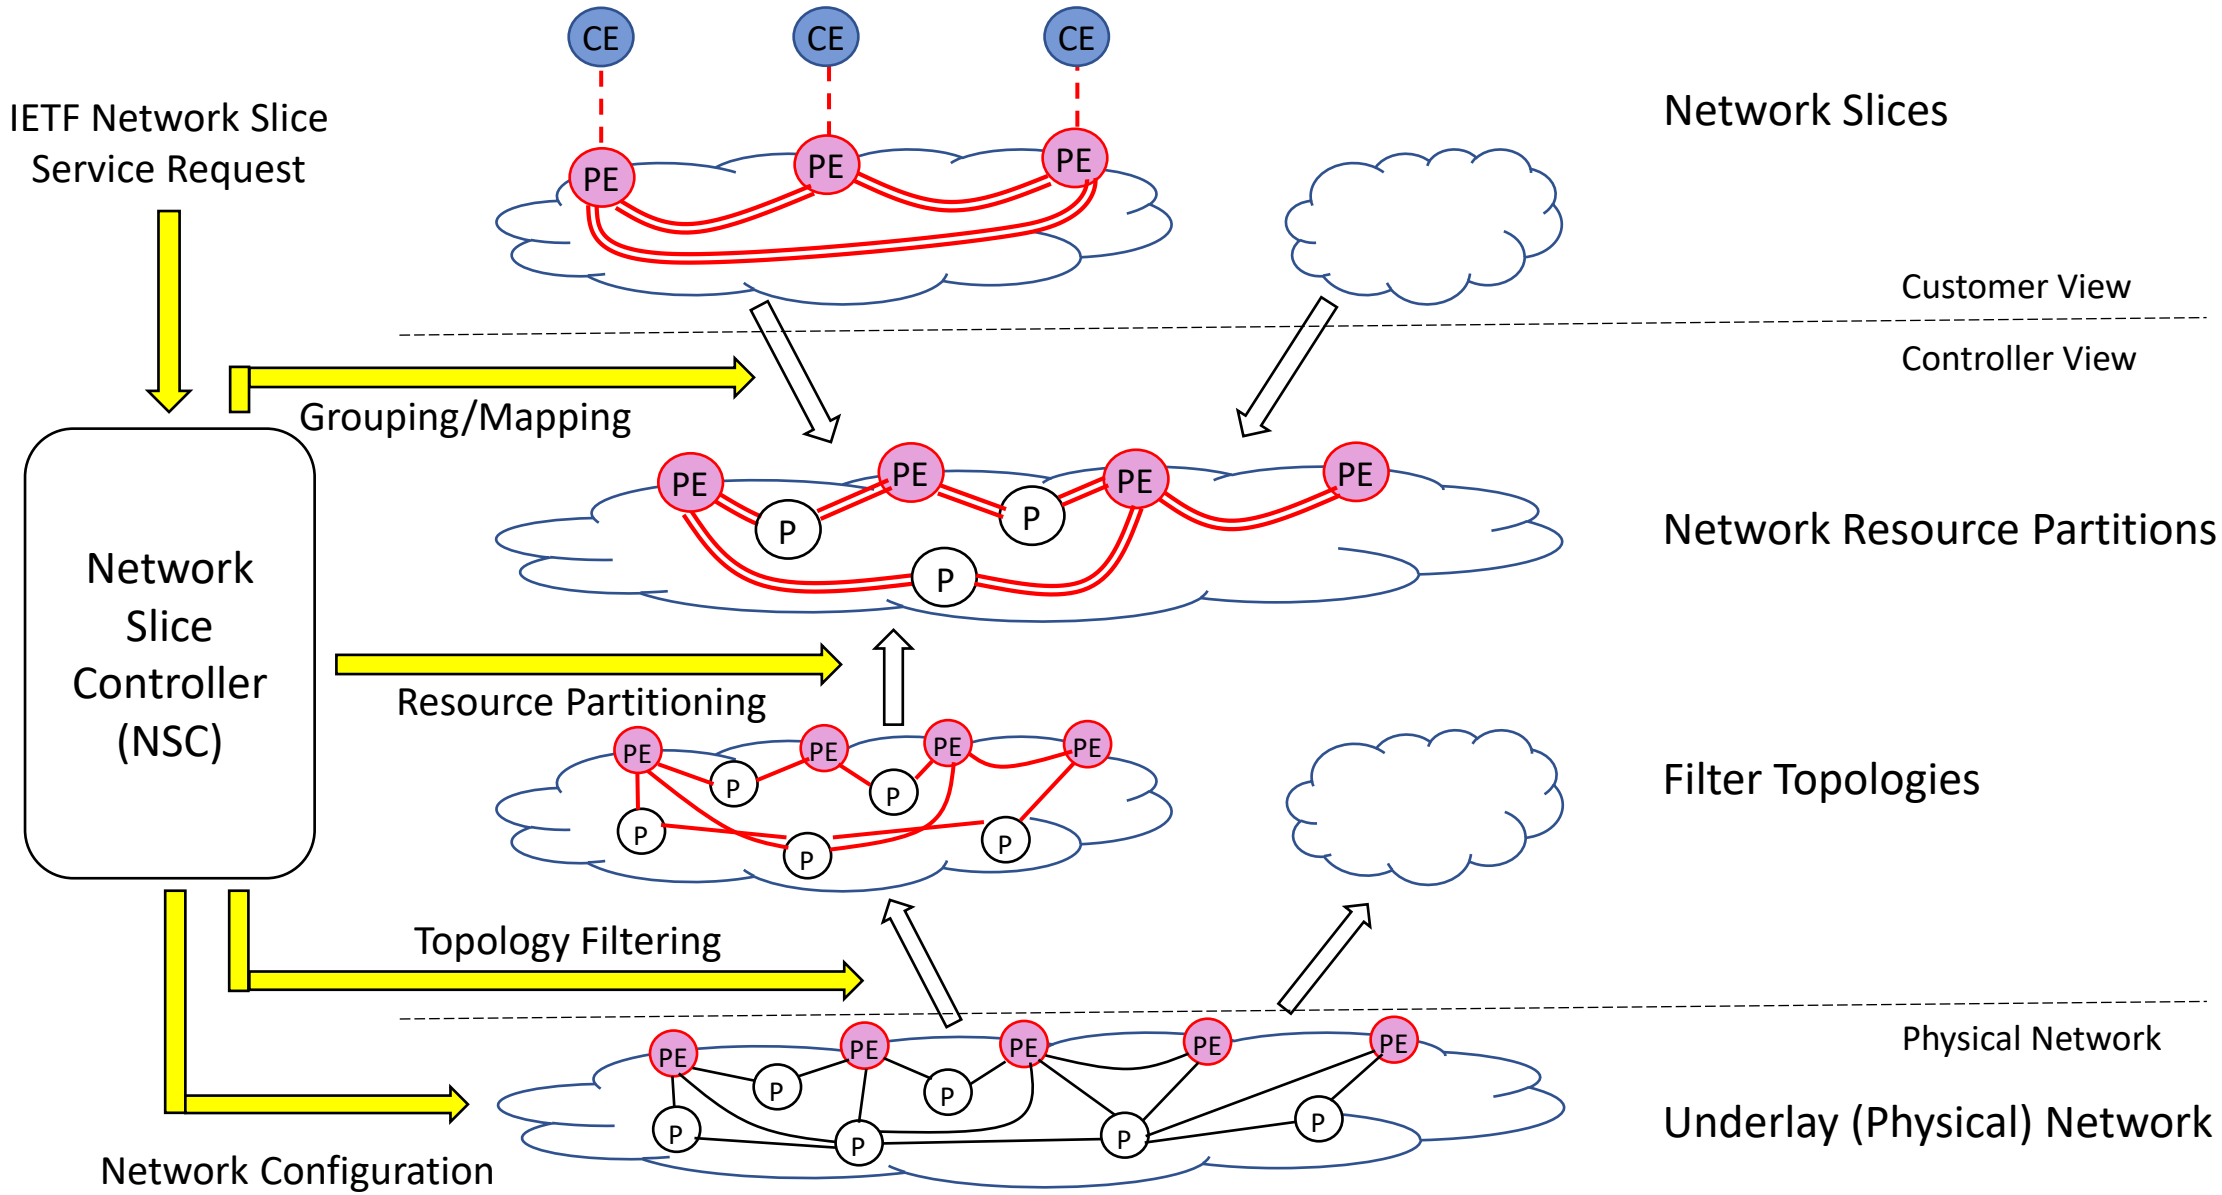

# Applicability of ACTN to IETF Network Slicing

- Abstraction and Control of TE Networks (ACTN) defined in RFC 8453
  - A management architecture and YANG models for building Virtual Network services
- ACTN can be used to deliver IETF Network Slices
  - draft-ietf-teas-applicability-actn-slicing

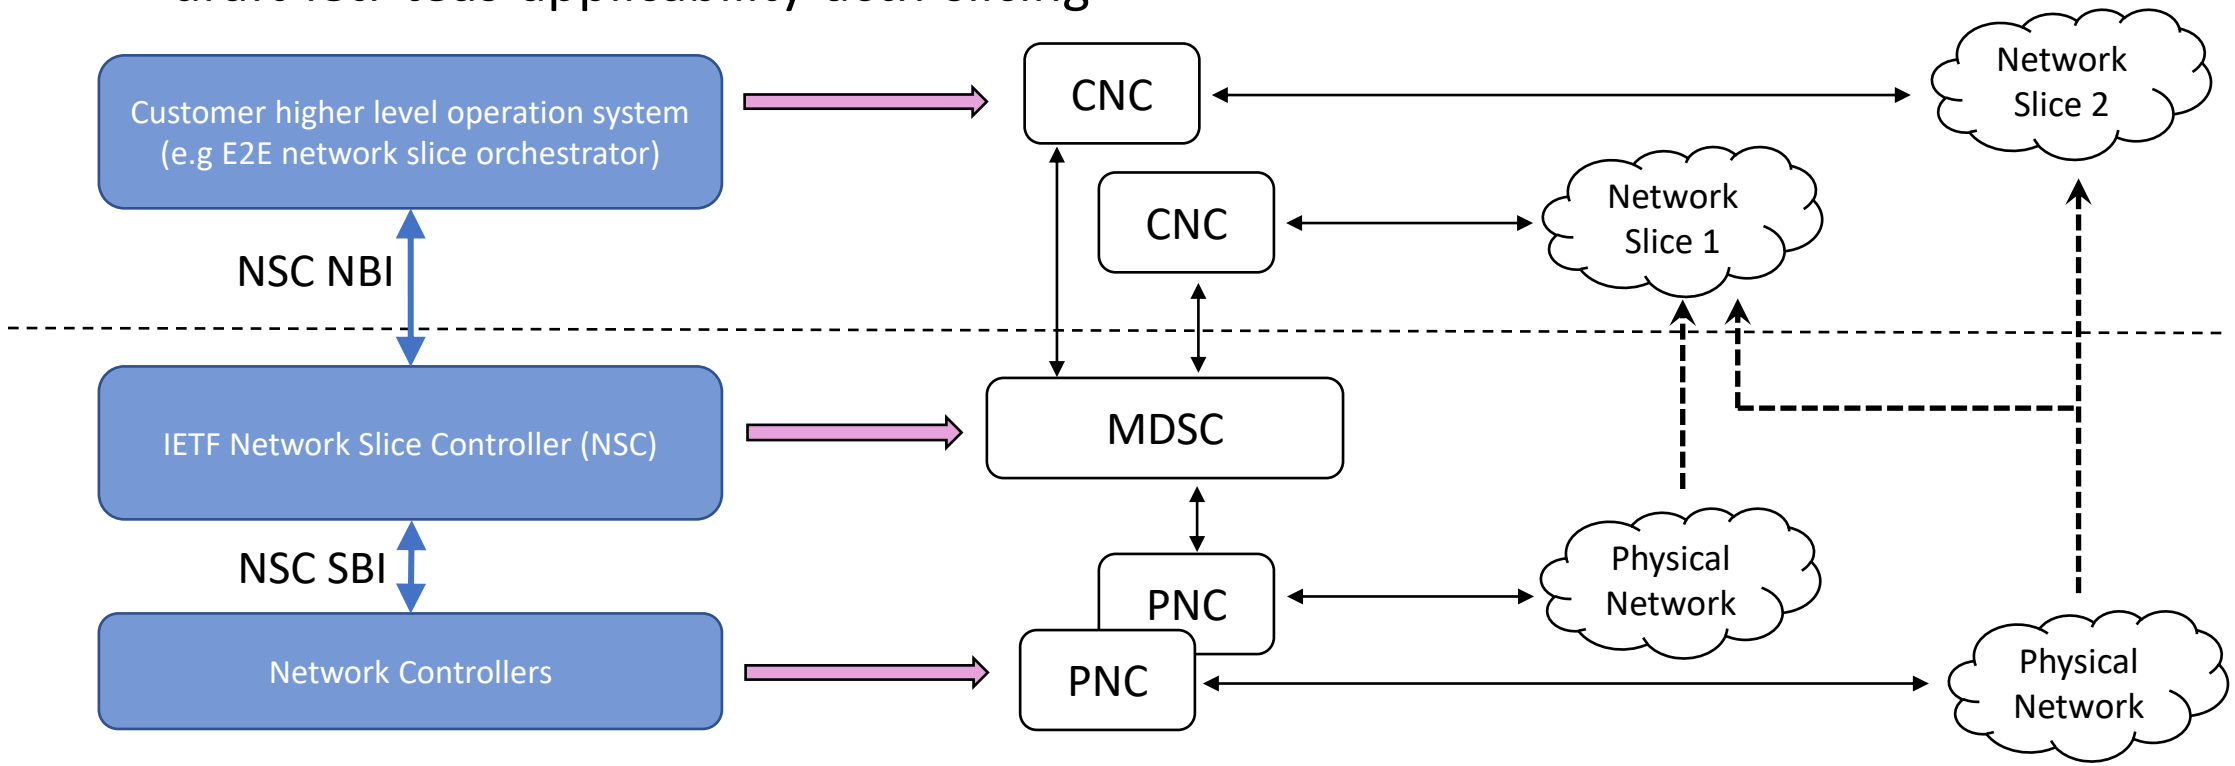


CNC - ACTN Customer Network Controller
MDSC - ACTN Multi-domain Service Coordinator
PNC - ACTN Provisioning Network Controller

# Enhanced VPNs

- Enhanced VPNs (VPN+) are defined in draft-ietf-teas-enhanced-vpn
  - A VPN service with additional requirements (such as SLOs)
- VPN+ can be used to deliver IETF Network Slices
  - A number of solution-oriented Internet-Drafts exist

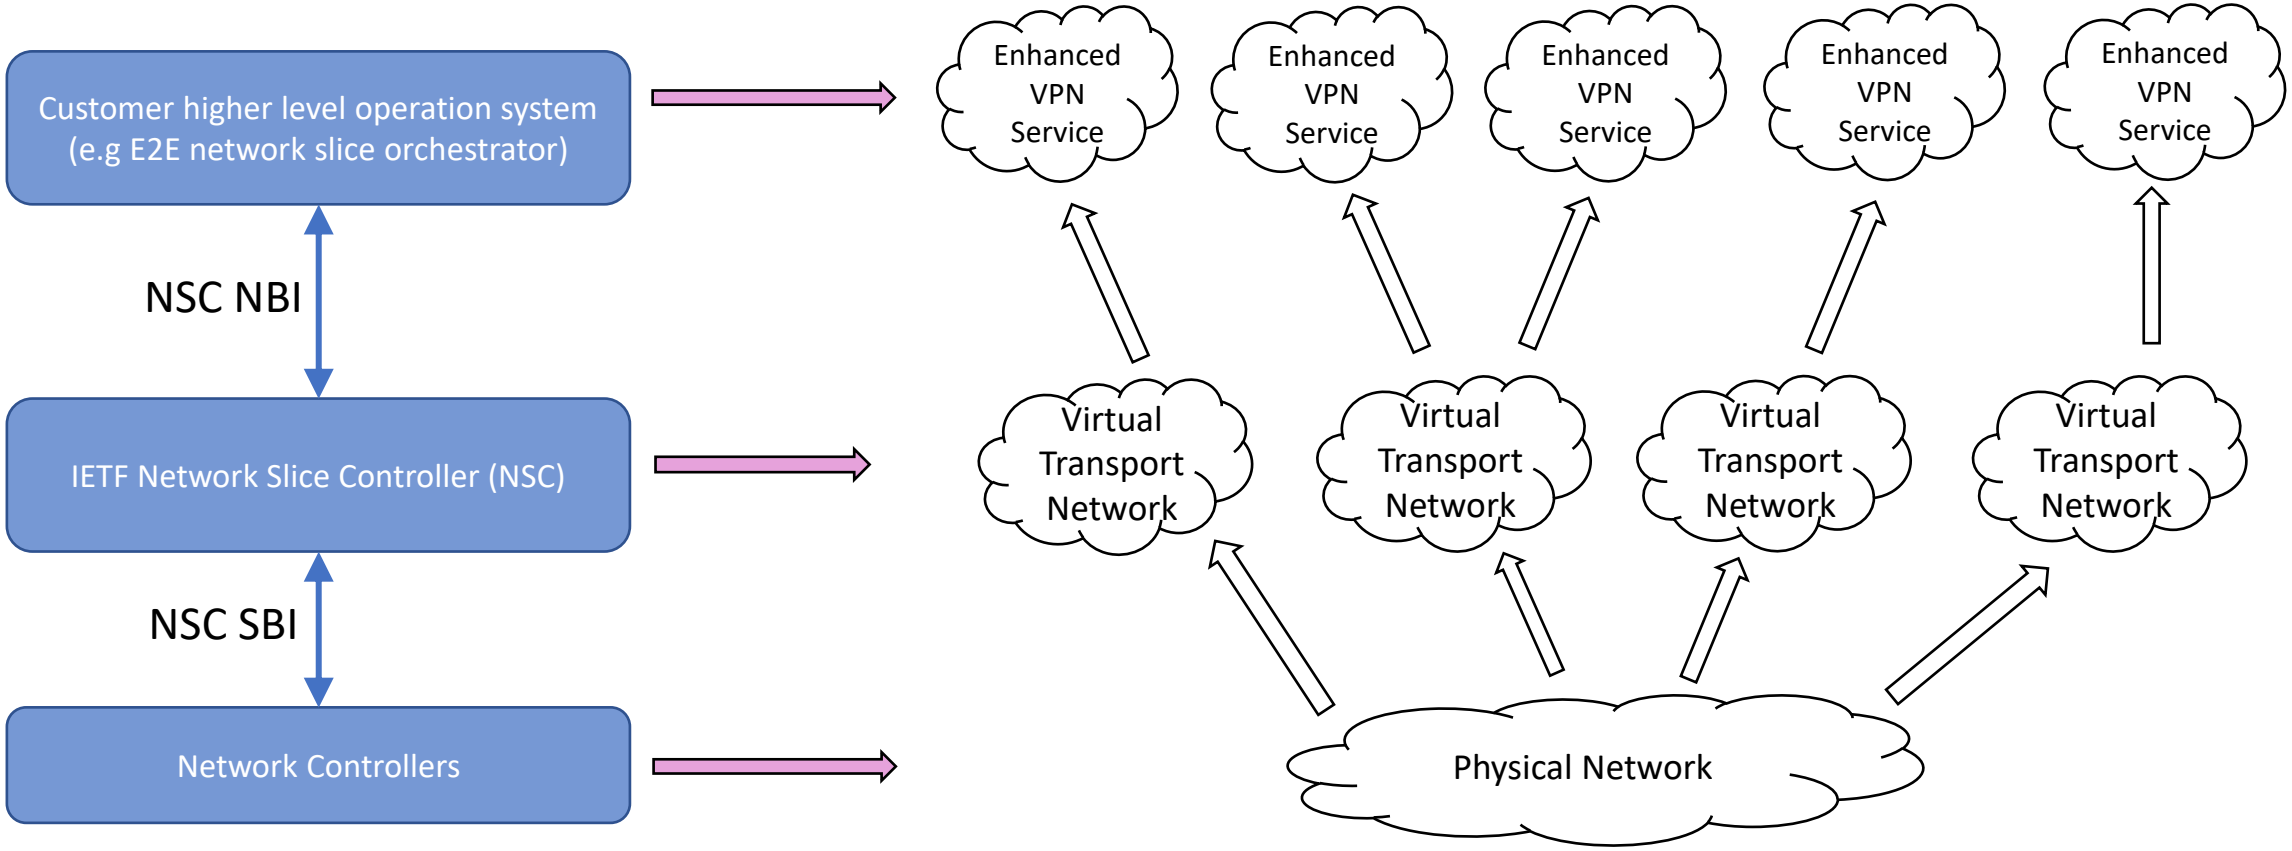

# Segment Routing Policy

- Segment Routing Policies (SR-TE) are a way of delivering IETF Network Slices in an SR network
  - Multiple solutions drafts exist
    - draft-ietf-spring-sr-for-enhanced-vpn
    - draft-bestbar-teas-ns-packet
    - draft-bestbar-teas-yang-slice-policy

- Approximate general process is…
  - Packets arrive at the ingress endpoint and are allocated to a network slice and connection matrix
    - May use port identifiers, VLAN tags, packet classification, etc.
  - Packets are marked (in some way) so that transit nodes can map them to SR policies
    - Possibly by using SIDs in then packets' SID stacks
  - Policies are programmed into the network
    - Network nodes handles packets differently according to how they are marked and the programmed policies

# Slicing All Manner of Networks

- IETF Network Slicing is an all-purpose tool
- Apply it to *any* network that uses IETF technology
  - IETF forwarding planes (IP, MPLS, SR)
  - IETF control protocols (MPLS, GMPLS, DetNet)
  - IETF management protocols (YANG)
- See also “Framework and Data Model for OTN Network Slicing
  - draft-zheng-ccamp-yang-otn-slicing

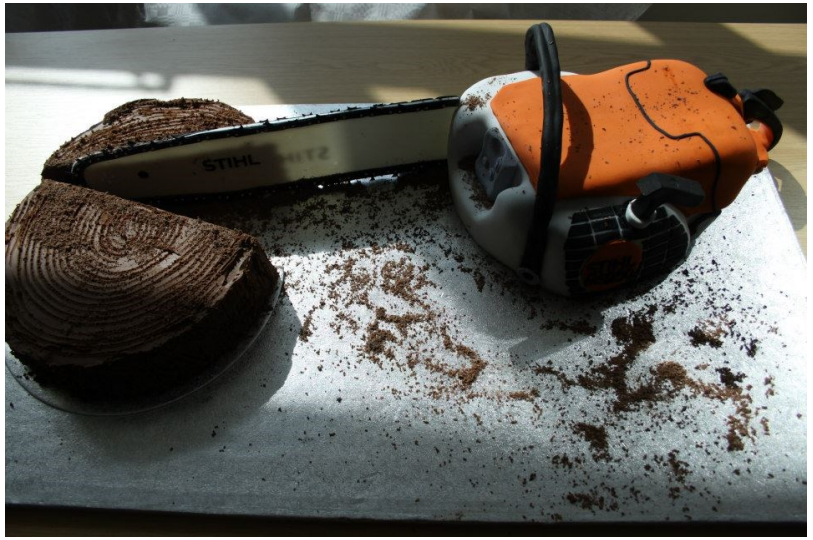

# What is TeraFlow?

- An EU-funded project to accelerate innovation in optical and IP network, supporting emerging high-capacity applications
- **Secured autonomic traffic management for a Tera of SDN flows**
  - 5G Integration with L3VPN/L2VPN up to the edge
  - 5G GTP Flow Definitions
  - Automated and Zero Touch Service Management for Transport Network Slices
  - Multi-vendor and multi-domain
- Output from the project will be "TeraFlow OS" for network applications
  - Cloud-native Network Operative System
  - Hierarchical SDN Orchestrator and Controllers, with connectivity services managed using blockchain distributed ledger
  - Open-Source Implementation with Apache License
- TeraFlow Project Site
  - www.teraflow-h2020.eu

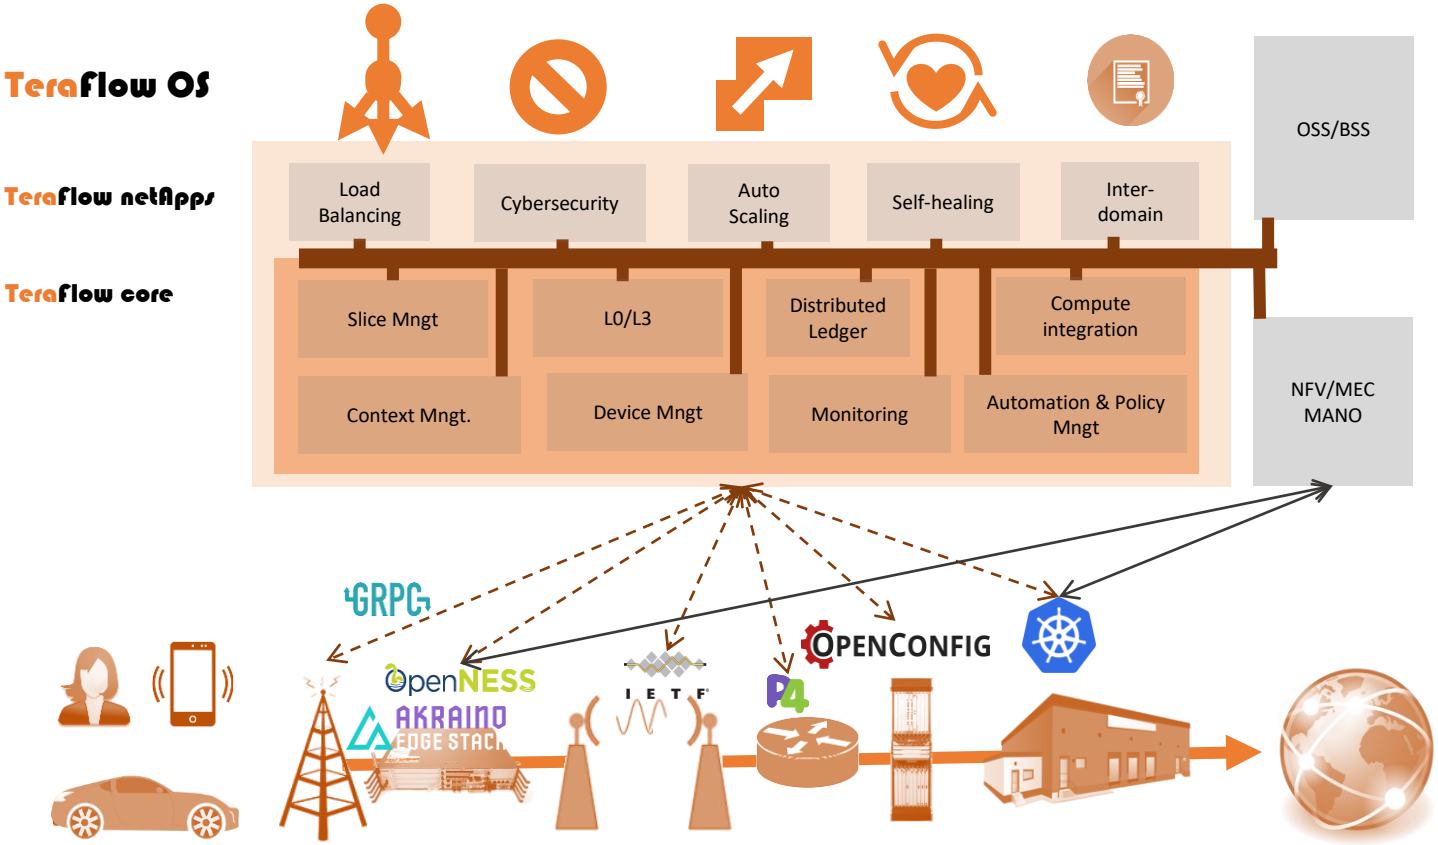

# TeraFlow Network Slicing

TeraFlow uses a combination generic service models and SDO device models, for network slicing deployment

- IETF, OpenConfig, OIF
- Hard and Soft slicing : Planned assignment of resources (VN) or Dynamic allocation of resources (VPN)
- Packet Optical Transport Network Slicing with Hard and Soft Isolation

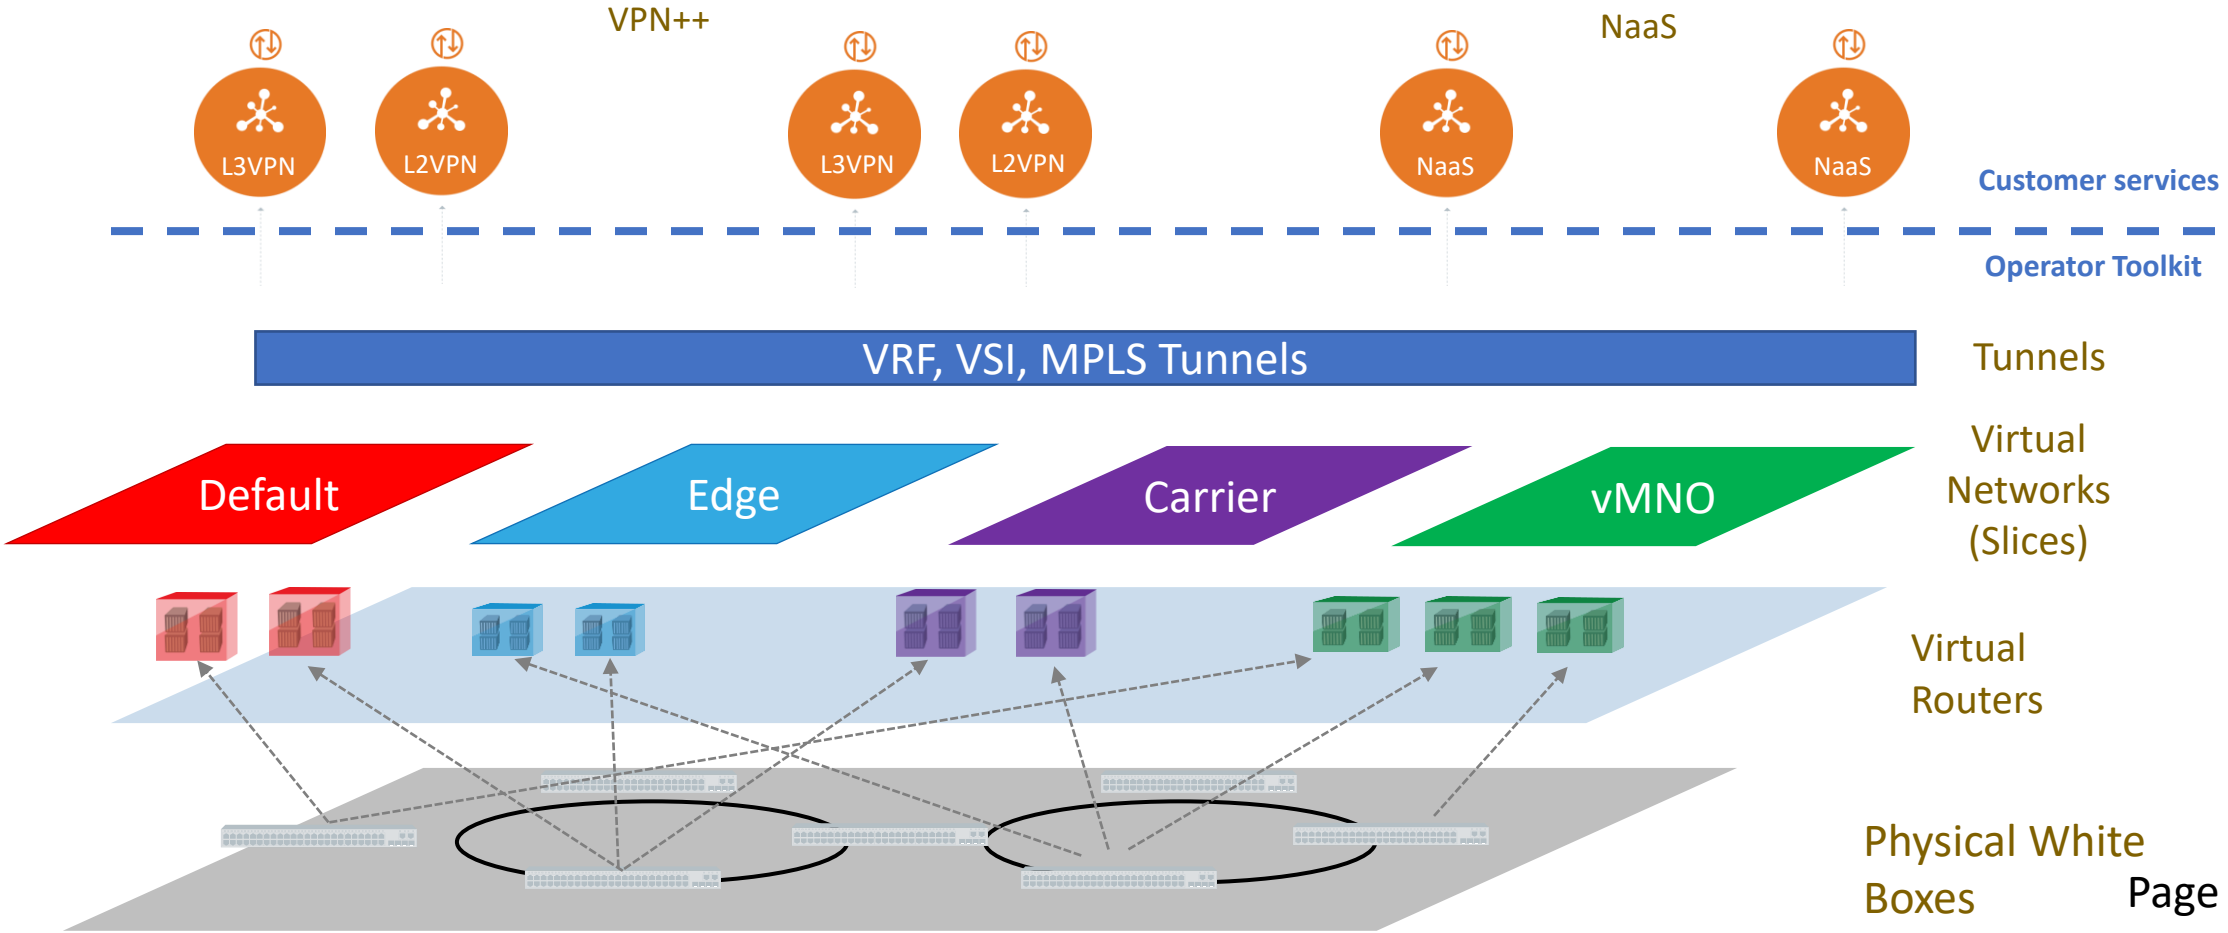

# Questions

adrian@olddog.co.uk