Network Slicing and Enhanced VPNs

Adrian Farrel : Old Dog Consulting

<adrian@olddog.co.uk>

India Internet Engineering Society (IIESoc)

Connections : A pre-IETF India Forum, November 13th - November 14th, 2019

Agenda

- Virtual Connections and Virtual Networks
- Abstraction of TE Networks
- Network Slicing
- ACTN
- SDN and YANG Models
- Enhanced VPN (VPN+)
- References

Early Services Were Simple Connectivity

- Virtual Links, Private Lines, or Pseudowires
  - Connecting two sites over a shared infrastructure
- Sites consider themselves connected by a physical link
  - Service provided by the network meets a Service Level Agreement
  - Essentially a layer 2 service

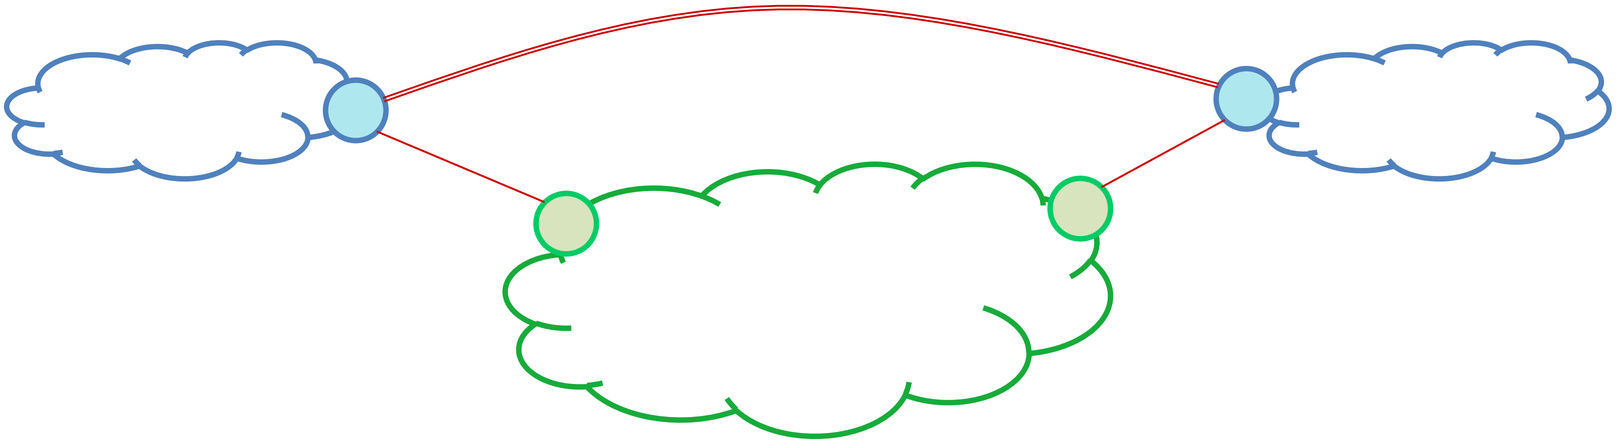

Connectivity Services Developed into Virtual LANs

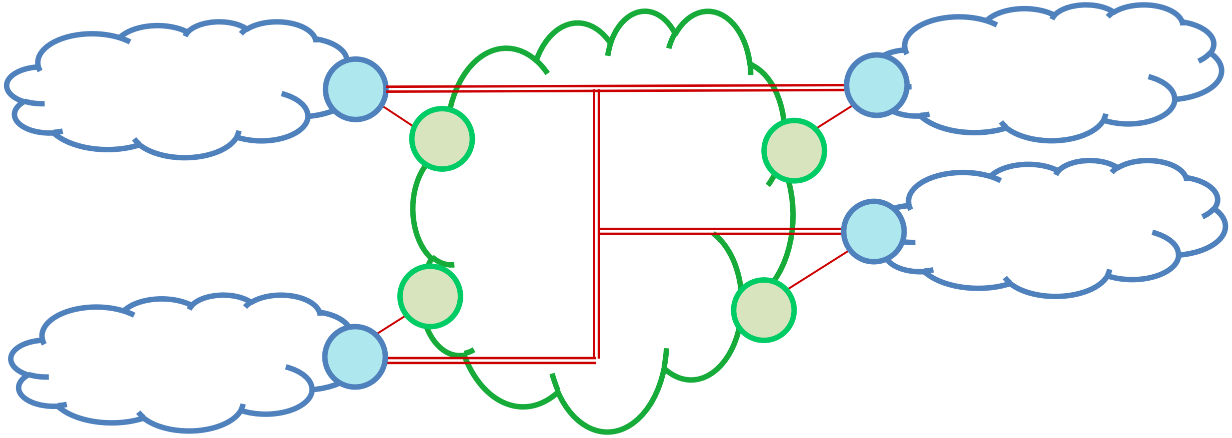

- Models basic LAN service
- Also a layer 2 service
  - Relatively simple SLA

Virtual Private Networks

- A generalisation of layer 2 connectivity services
- Also a very popular layer 3 service
  - Provides routed IP full or partial mesh
- VPN was the killer application for MPLS
- A VPN is virtual
  - It is not really a Network, but behaves somewhat like one
  - It is not really Private
    - Network resources are shared

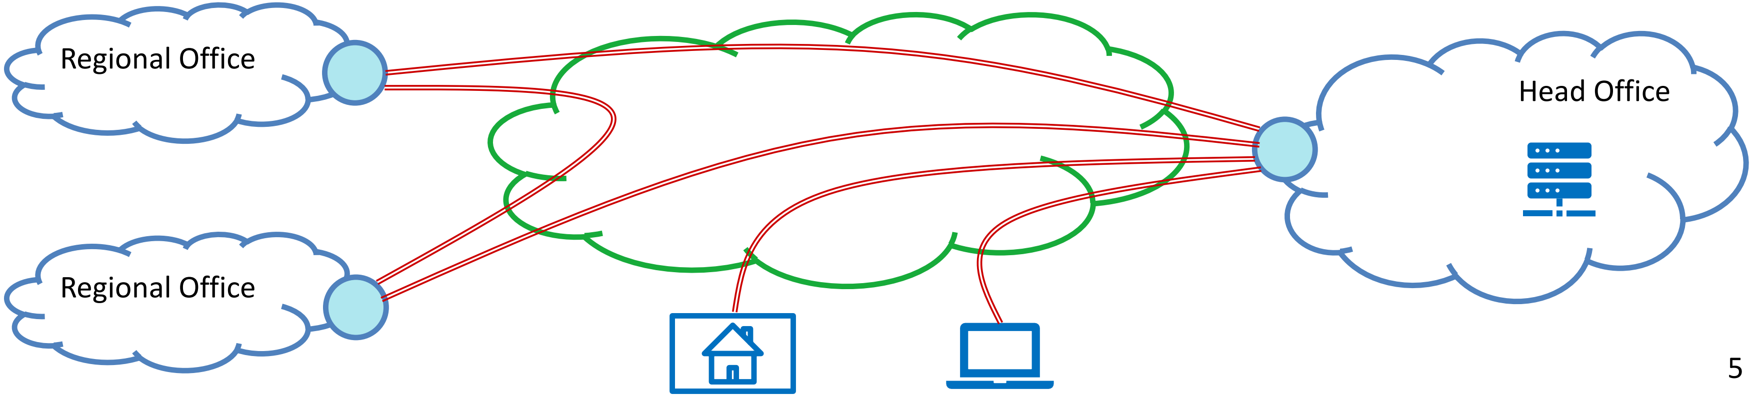

Topology Aggregation

- Abstraction Layer Network

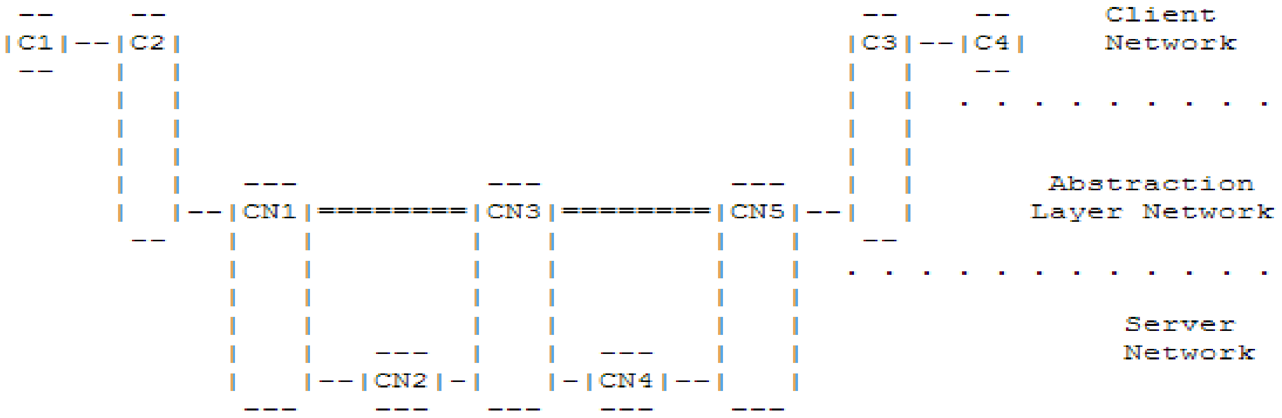

Client layer resources: C1, C2, C3, C4

Server layer resources: CN1, CN2, CN3, CN4, CN5

Abstraction layer resources:

Nodes: C2, CN1, CN3, CN5, C3

Physical links: C2-CN1, CN5-C3

Abstract links: CN1-CN3, CN3-CN5

Abstraction Leads to Virtualization

- Abstraction is about providing a summarised topology of potential connectivity
- Policy-based
  - Policies set by one network with knowledge of the other networks
  - Overcome issues of scaling, stability, confidentiality, and misinformation found in aggregation
    - Hint: virtual node representations may struggle
- Apply policy to the available TE information within a domain, to produce selective information that represents the potential ability to connect across the domain
  - Don't necessarily offer all possible connectivity options
  - Present a general view of potential connectivity
  - Consider commercial and operational realities
- Retain as much useful information as possible while removing the data that is not needed
- Can be further filtered to provide different views for different consumers

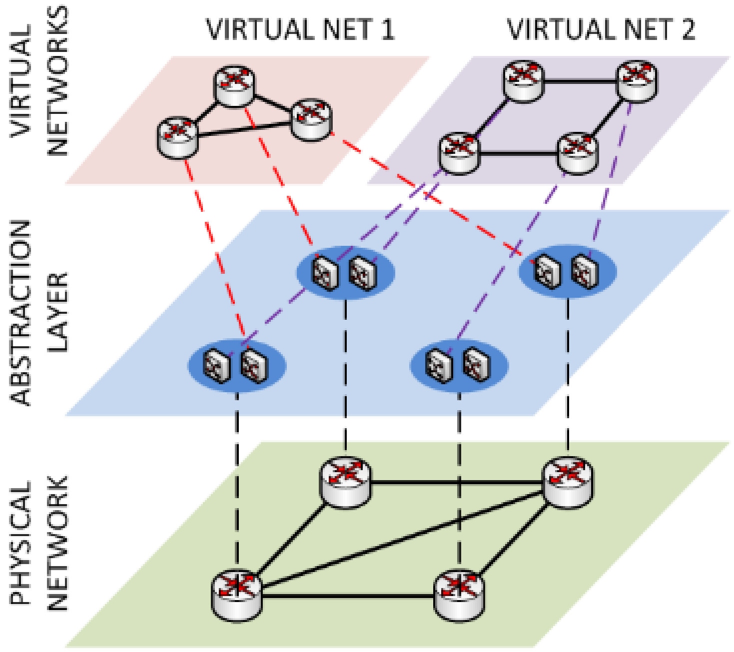

Virtual Networks (VNs)

- Network abstraction aggregates resources into
  - Virtual links (made from TE tunnels across links and nodes)
  - Abstract nodes (made from nodes and links)
- Describes edge-to-edge connectivity with certain qualities
- Available connectivity can be presented to the VN user (customer)
  - They can manipulate the VN as their own private network

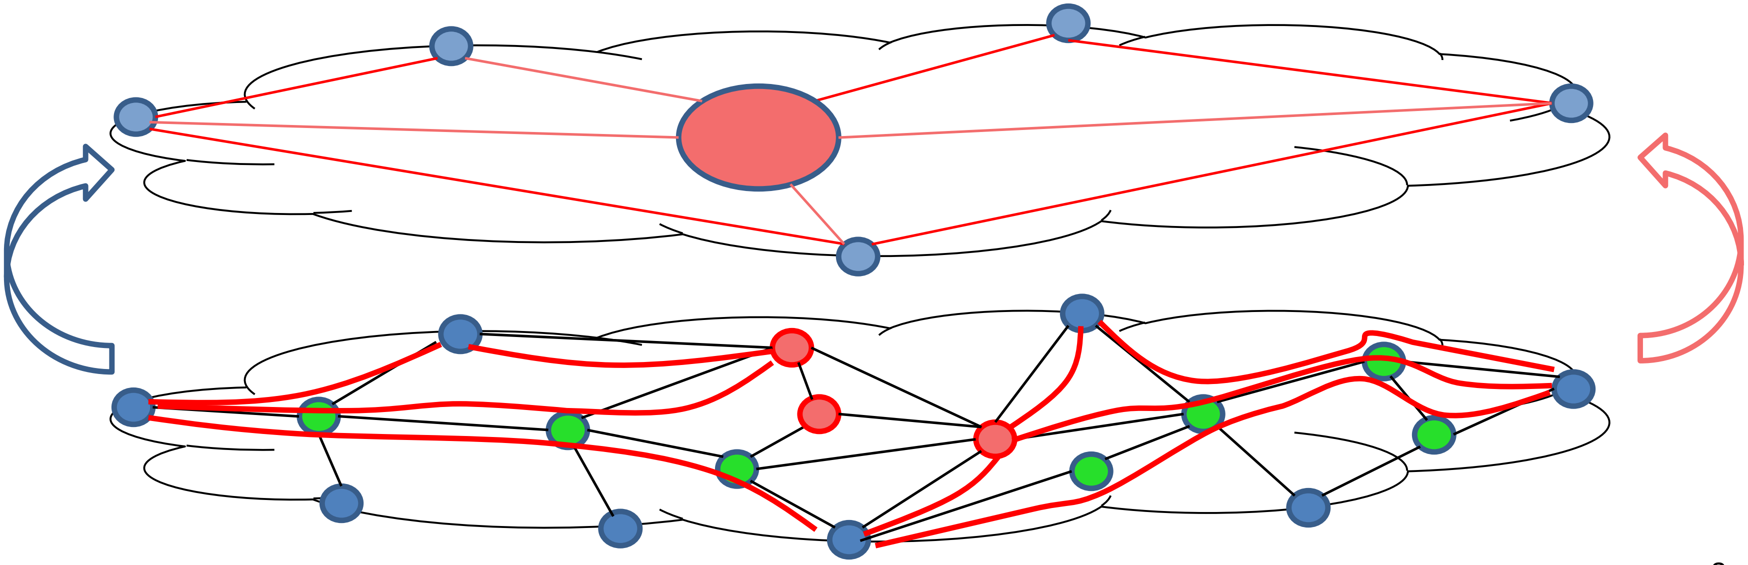

Application of VNs

- The VN concept covers a wide range of applications
  - Simple connectivity services (such as VPNs)
  - Enhanced connectivity services (such as VPNs with different per-site bandwidths)
  - Customer managed connectivity services (adding sites, connectivity, and bandwidth)
  - Customer-operated higher layer networks build from lower layer connectivity
    - Carrier’s carrier
    - IP division as a client of the transport division

The Traffic Engineering Database

- A collection of information about the network
- The topology of the controlled network
  - Nodes
  - Links
  - Nodes/Links connectivity

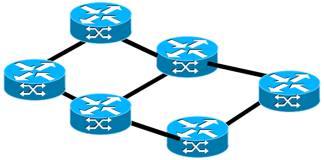

- The available resources and attributes
  - Available Link Bandwidth — 500 Mb/s
  - Link Metrics (e.g., costs) — 1
- The TED is an essential internal component of a Path Computation Element (PCE)
  - Provides the updated snapshot of the controlled network and its resources
  - PCE algorithms resort to TED as primary information source input

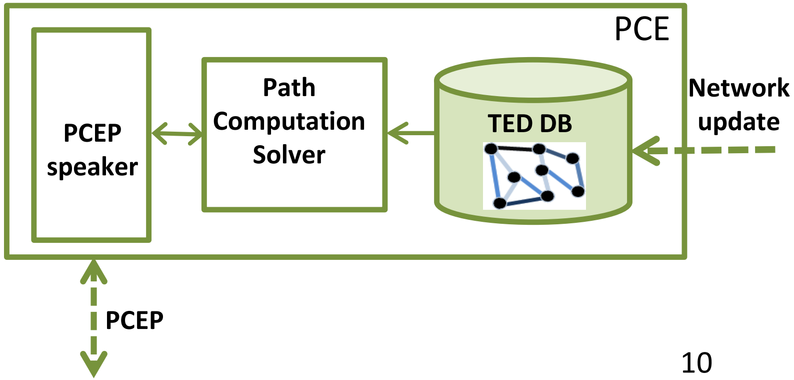

The Traffic Engineering Database

- Traffic Engineering Database (TED) is essential internal component of a PCE
  - provides the updated snapshot of the controlled network and its resources
  - PCE algorithms resort to TED as primary information source input

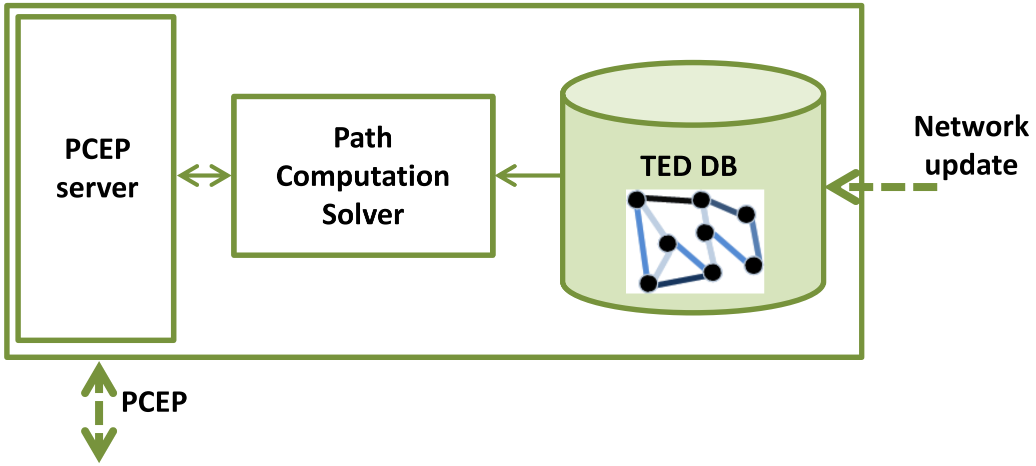

Building the TED from the Network

- This is the process of building a model of the network
  - Different mechanisms may be used
  - The functional architecture doesn’t care how the TED is built
- Information can come from different sources
  - From the network
  - From management systems
  - Through policy
- All kinds of ways to get information from the network
  - Passive peering with OSPF-TE or IS-IS-TE
  - Through Link State BGP (BGP-LS)
  - Reading from the network devices (e.g., YANG)
  - PCEP Notifications
- Abstraction can be done by:
  - The receiver of the information applying policy (e.g., OSPF plus policy)
  - The exporter of the information applying policy (e.g., BGP-LS plus policy)

Application-Based Network Operations (ABNO)

- First attempt at describing a system for network virtualisation
  - Pull together many components already described by the IETF
  - RFC 7491
- Path Computation and Traffic Engineering
  - Network Topology (LSP-DB, TED, Inventory Management)
  - PCE, PCC (RFC 4655)
  - Online & Offline (RFC 7399)
  - Stateful & Stateless (RFC 8231)
- Service Coordination
  - Application-Layer Traffic Optimization (ALTO) (RFC 5693)
- Multi-layer Coordination
  - Virtual Network Topology Manager (RFC 5623)
- Network Signalling & Programming
  - RSVP-TE (RFC 3209)
  - OpenFlow
  - Interface to the Routing System (RFC 7921)
- Additional components
  - ABNO Controller (Orchestrator)
  - Policy Agent
  - OAM Handler
  - Provisioning Manager

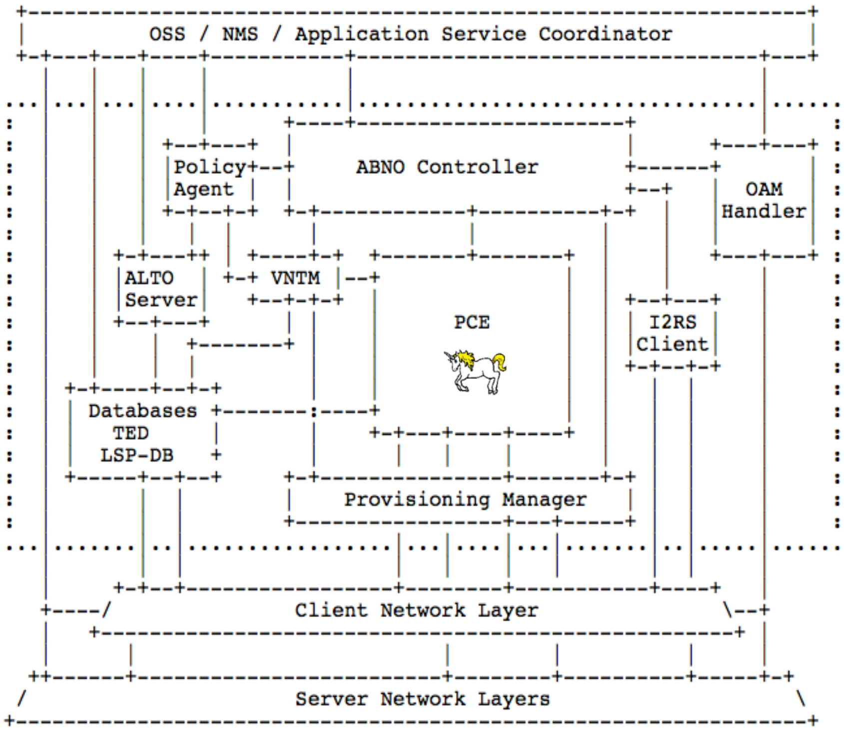

Figure 1: Generic ABNO Architecture

Slicing the Network

- Consider network resource separation
  - Partitioning the resources of a network for specific uses
  - Not a new thing:
    - VPNs, virtual networks, overlay networks
    - RSVP-TE, queuing/buffering schemes
  - Current interest is driven by 5G
    - But take care!
      - Slicing in the 5G radio network is not slicing in the Aggregation or Transport network
      - The granularity is completely different
- Aim to guarantee a level of service delivery without impacting or being impacted by other services
  - Service level can be:
    - Throughput
    - Latency
    - Jitter
- Reserving resources in a network for a customer or service
  - “Resources” may be:
    - Bandwidth on links
    - Compute and storage
    - Service functions

Network Slicing in More Detail

- Provide connectivity and function to serve customers with a wide variety of service needs
  - Low latency, reliability, capacity, and service function specific capabilities
- Requirements for Network Slicing
  - **Resources**: Partition the available network resources and provide specific Service Functions with correct chaining logic
  - **Network & Function Virtualization:** Virtualise physical resources and support recursive virtualisation
  - **Isolation:** Operate concurrent network slices across a common shared underlying infrastructure
    - **Performance**: Behaviour of one slice doesn't (can't) cause changes in behaviour of another slice
    - **Security**: Attacks or faults occurring in one slice must not have an impact on other slices. Traffic in a slice must be kept safe
    - **Management**: Each slice can be independently viewed, utilised, and managed as a separate network
  - **Control and Orchestration:** Orchestration is the overriding control method for network slicing
    - **End-to-end Orchestration**: End-to-end service delivery requires concatenation of networks
    - **Multi-domain Orchestration**: Services can be managed across multiple administrative domains

Why Standards for Slicing?

- Standards are about ensuring interoperability
  - Protocol standardisation is well-known
  - Data models form an increasing part of standardisation
- Network slicing in the IETF is:
  - Use of existing tools to manage networks
    - Routing protocols can advertise link information
    - Signalling protocols can collect path information
    - BGP-LS can share network abstractions and PCE can compute overlays
    - Management protocols can partition and configure networks
  - Three foundational pieces of work in progress
    - Software Defined Networking
    - Abstraction and Control of Transport Networks (ACTN)
    - Enhanced VPNs (VPN+)

SDN Is Key

- SDN is the buzzword of the decade
- Software control of distributed resources

- Facilitates network management and enables programmatically efficient network configuration
- Based on a shared architecture of orchestrators and controllers
- Provided through software APIs and common data models

SDN with a Control Plane

- A common misunderstanding…
  - "SDN implies node-by-node programming of the network"
- SDN is about centralised view and control
  - How to convert into network state is an open issue
- One option is node-by-node programming
  - Such as OpenFlow from a "controller"
- Or we can choose a hybrid approach
  - Central control leveraging an active control plane
  - Keeps autonomy and smarts in the network
  - Adds central, programmable control
  - Allows migration to SDN
    - Supports existing deployment models
- Key component is the TED

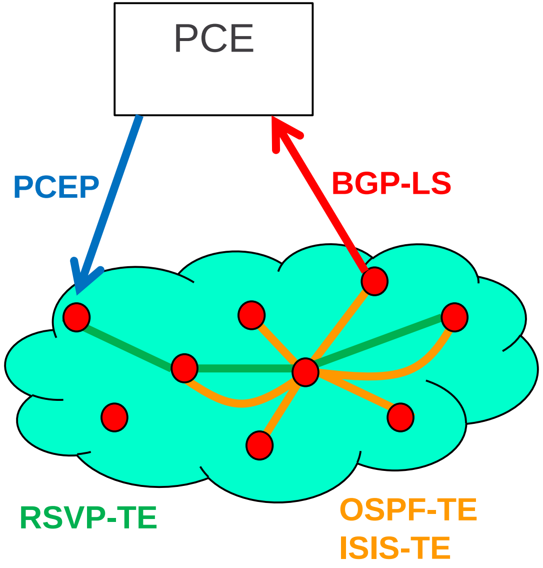

YANG Models Are Everywhere

- Data models are an essential tool for SDN
- A model describes a system
  - Allows it to be modelled, observed, and controlled
- YANG is today's modelling language of choice
  - Replaced MIBs in the IETF
  - Used widely in Open Source
- Hundreds of YANG models have been written
  - Sometimes multiple models for the same thing
- Gradual increase in standardization
  - Enables interworking of components from different vendors and Open Source projects

Why Build a Standard Topology Data Model?

- Data models let us represent information in a well-known way
- Useful for moving it between implementations
  - Export from the network
    - From a single network node talking about its local resources
    - From a network node that collects and aggregates it from the network
  - Share between servers
    - Exchange between PCEs that synchronize state
  - Store, test, and experiment
    - Archive the network at a point in time
    - Conduct offline tests and experiments on stored topologies
    - Debug networks and software
    - Share topologies between researchers or with suppliers

Abstraction and Control of TE Networks (ACTN)

- Abstraction is a way of representing connectivity across a TE network
- This allows a server network to present connectivity options to a client
- ACTN is an architecture for requesting and managing abstractions
- A customer (a client) requests connectivity from an operator
  - Delivered as a VN or a TE topology
- ACTN components map the customer requests to network resources
  - Orchestration can select and instruct networks
  - Controllers can program the network devices
  - TE links (tunnels), abstract nodes, and virtual networks are constructed
  - Services are mapped to the TE resources

Base ACTN Architecture

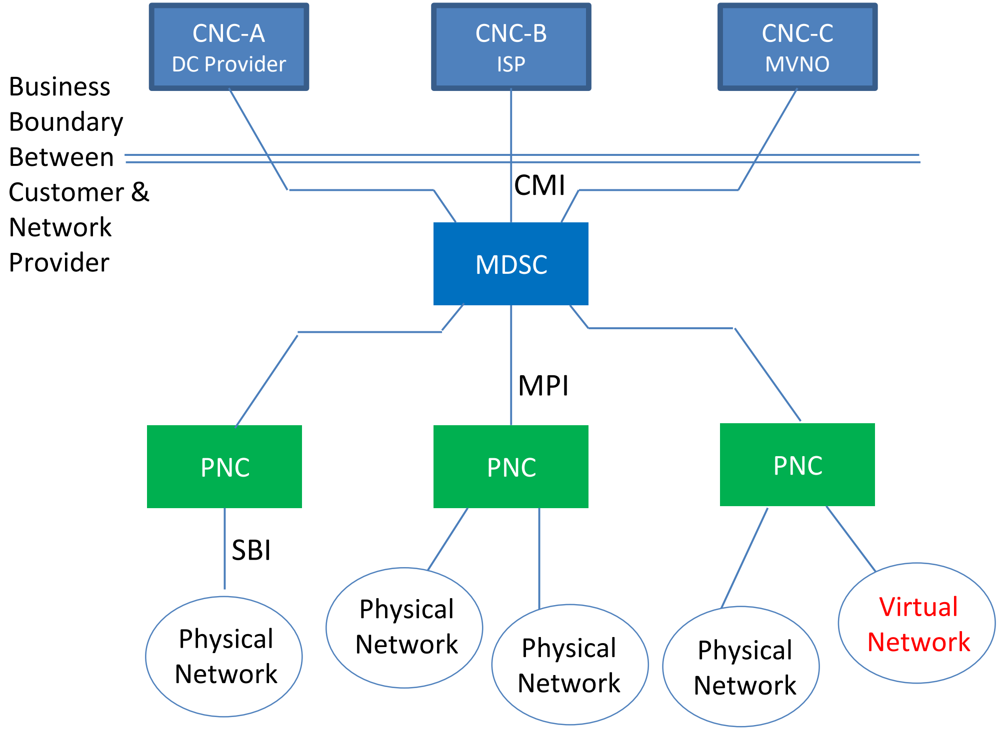

- Three components
  - Customer Network Controller
    - Formulates requests for clients/customers
  - Multi-Domain Service Coordinator
    - Maps service requests to one or more underlying network
  - Provisioning Network Controller
    - Classic SDN controller
    - With or without control plane
- Three interfaces
  - CNC-MDSC Interface (CMI)
  - MDSC-PNC Interface (MPI)
  - Southbound Interface (SBI)
- Note separation of Customer and Network Provider
- Note recursive nature for carrier's-carrier

Functional Split of MDSC Functions in Orchestrators

- SDN architecture can be mapped to ACTN
- Key features are:
  - Service orchestration
  - Network orchestration
  - Domain control
- MDSC function can be split between orchestrators
- Additional functions may be provided alongside
- YANG models serve as the interfaces
  - Categorized per RFC 8309

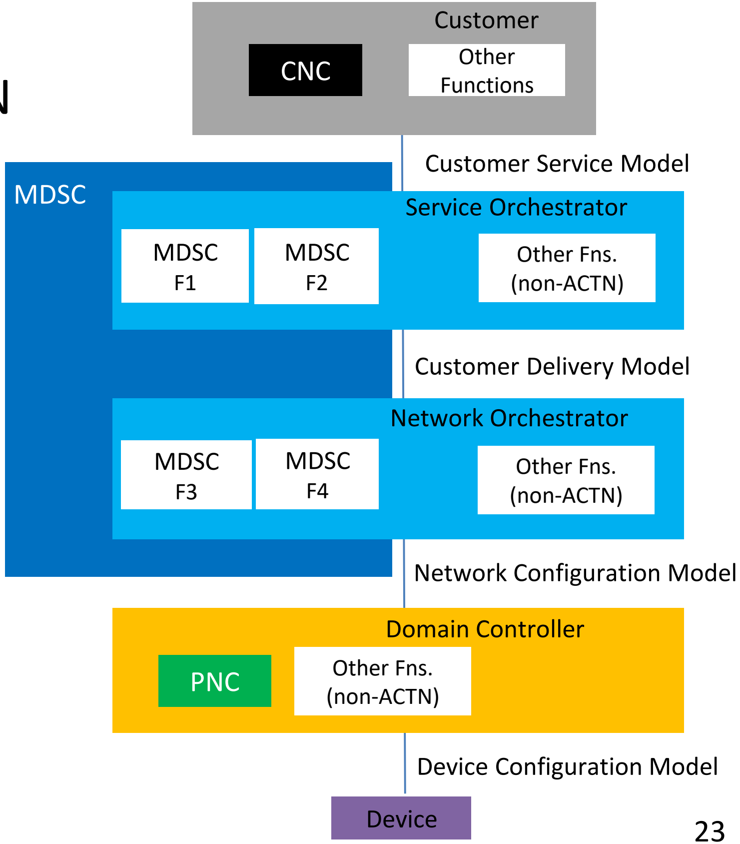

Network Slicing – The Key to 5G

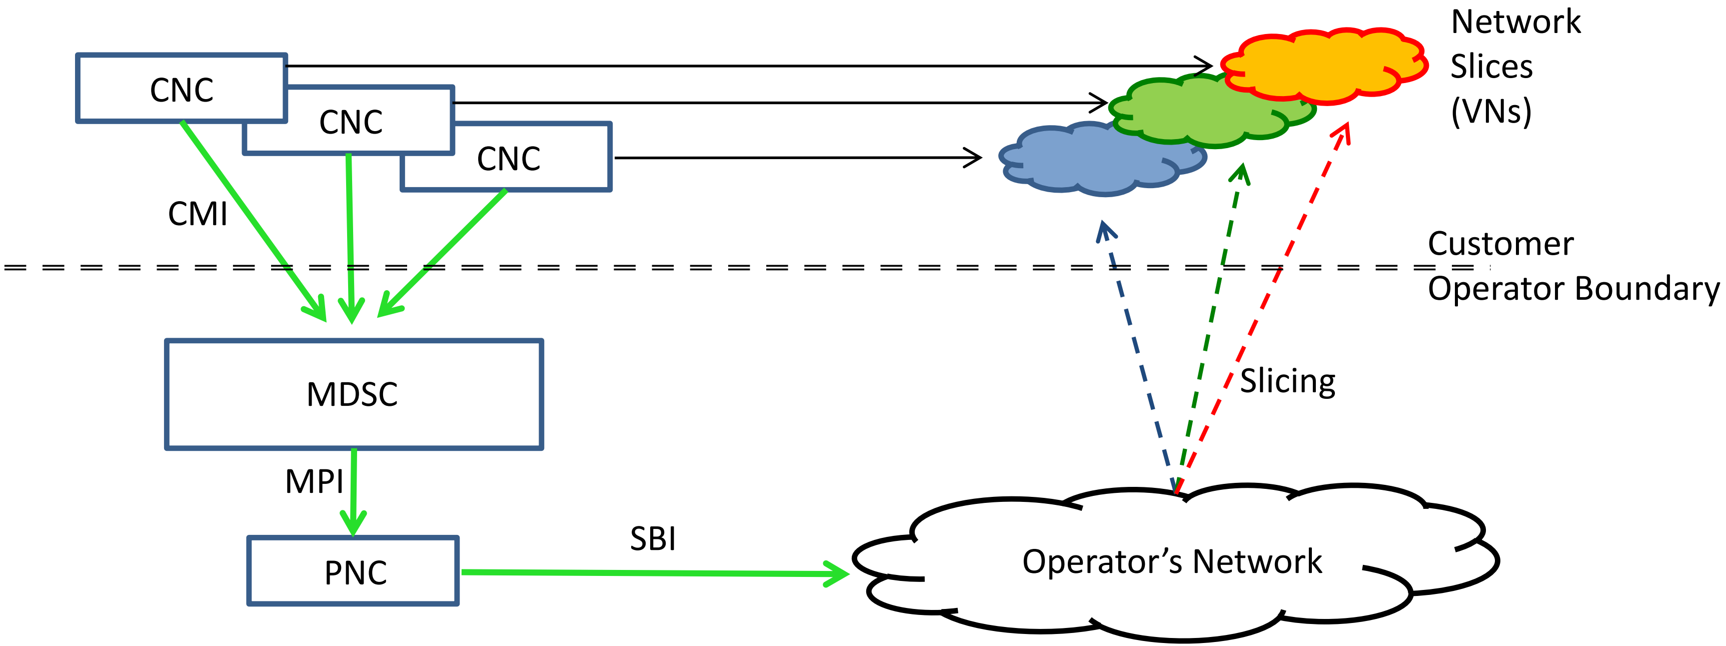

- Virtual Networks (VNs) are slices of the Operator's Network
- They are "private" slices of the nodes and links
- Help to guarantee specific service types

ACTN Progress in the IETF

- Demonstrates the time-line for developing significant pieces of work

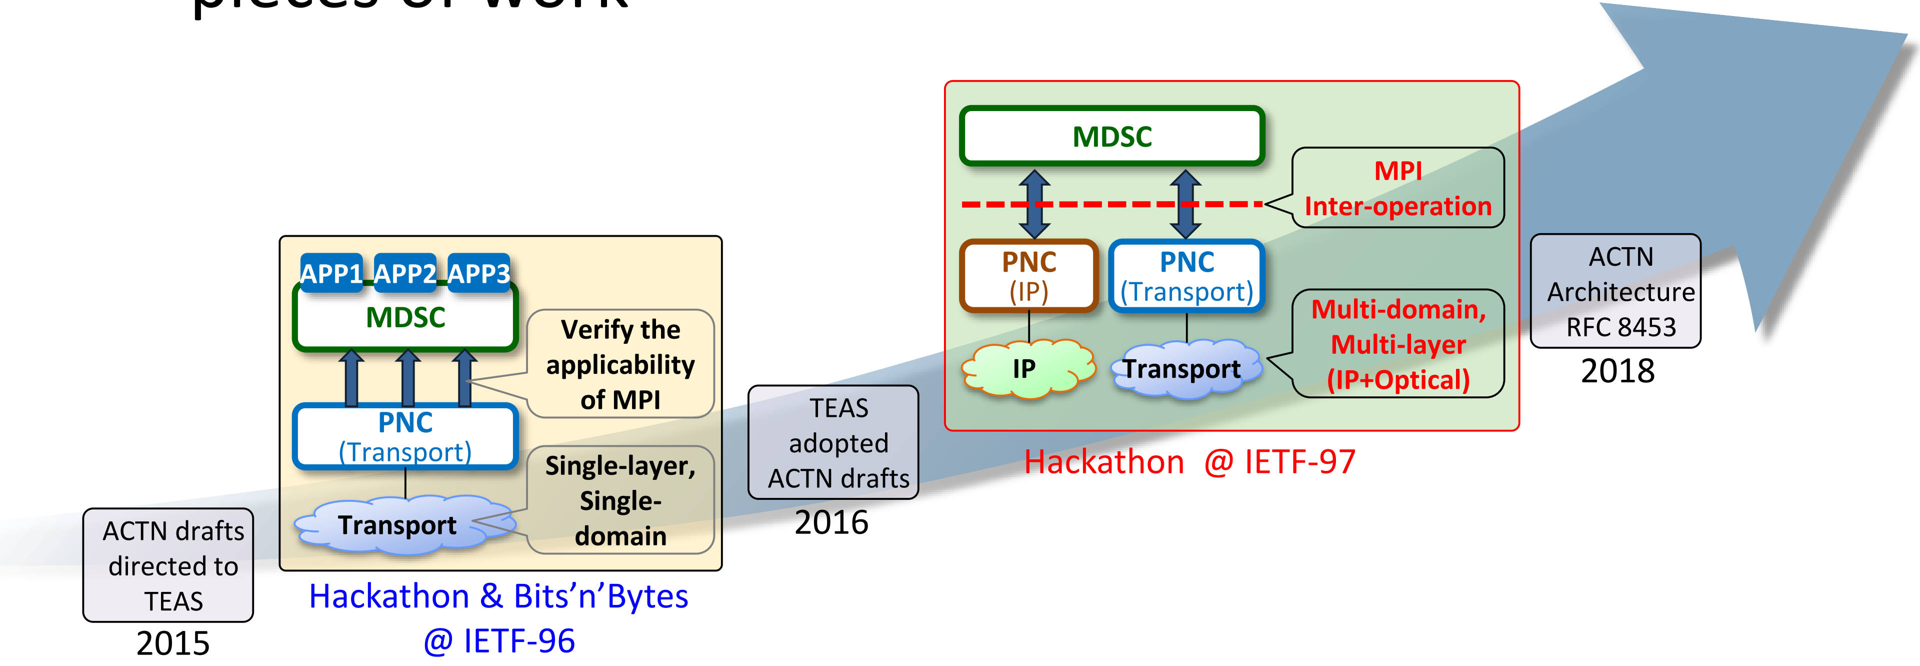

YANG Models for ACTN and TE

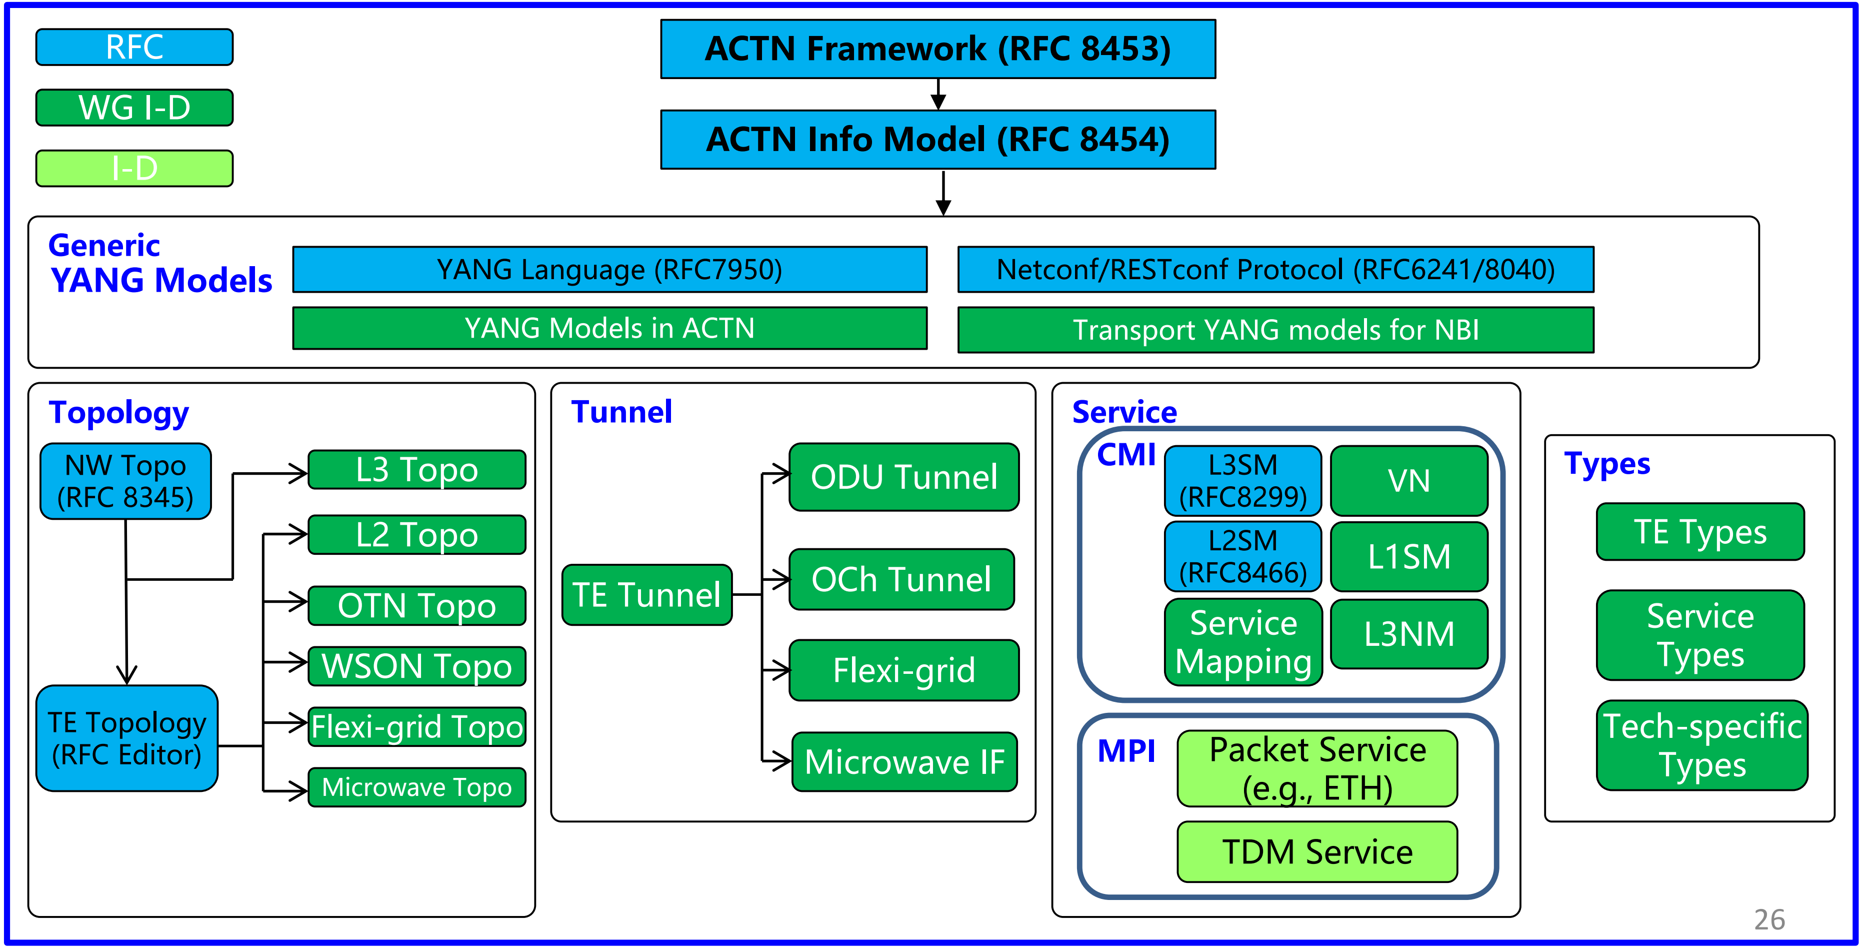

Enhanced VPN Why?

- New applications
  - Particularly applications associated with 5G services
  - “Enhanced overlay services”
- New requirements
  - Isolation from other services
    - Changes in network load or events in other services have no effect on the throughput or latency of this service
    - Drives some form of “partitioning” of the network – Network Slicing
  - Performance guarantees
    - Bandwidth, latency limits, jitter bounds
  - Some level of client control of underlay resources
- Existing technologies
  - VPNs have served the industry well
    - Provides groups of users with logically isolated access to a common network
  - Re-using existing tools, techniques, and experience is very effective
  - Look for an approach based on existing VPN technologies
    - Add features that specific services require over and above traditional VPNs – Enhanced VPN (VPN+)

Architecture of Enhanced VPN

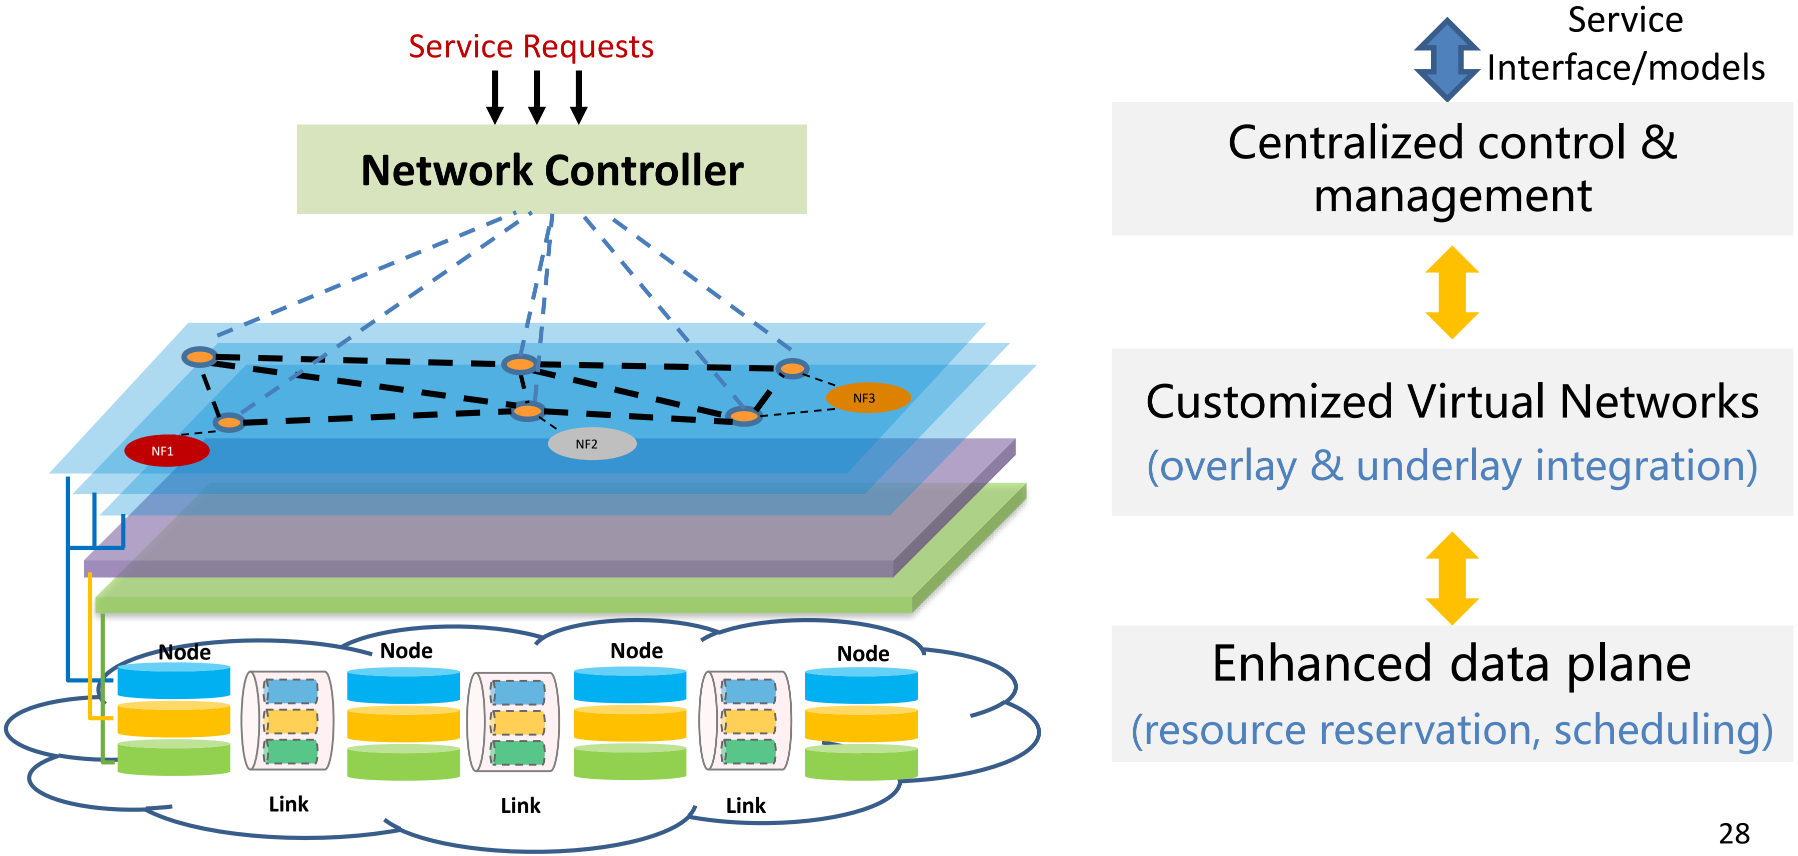

Scope of VPN+ Work

- Enhanced data plane
  - Different levels of isolation (from soft isolation to hard isolation)
  - Determinism of packet loss and delay
  - Identification of network slice and the associated network resources
- Control protocols
  - Both centralized and distributed
  - Information distribution, collection and computation to build the required virtual networks
  - Scalability considerations: the amount of state introduced
- Management plane
  - Life-cycle management: planning, creation, modification and deletion
- OAM, protection, inter-domain/inter-layer considerations

Candidate Technologies for VPN+

Layer 2 Underlay Data Plane	• Flexible Ethernet (FlexE) • Dedicated queues • Time Sensitive Networking (TSN) • …
Layer 3 Underlay Data Plane	• MPLS-TE • SR-MPLS/SRv6 • Detnet • …
Control Plane	• Distributed: RSVP-TE, IGP, BGP… • Centralized: PCEP, BGP-LS…
Management Plane	• ACTN architecture and data models • Service models: L3SM, L2SM, etc.

Enhanced Data Plane for VPN+

- The foundation of service performance assurance
- Many new work in progress to solve the requirement of low/bounded latency, jitter and packet loss, scalability, etc.
- Need to figure out which and how to integrate into VPN+ architecture
- Further discussion about soft and hard isolation

VPN+ Challenges

- Existing VPN sites are connected by RSVP-TE tunnels
  - So what's new? Why not just do that?
- Scaling is a challenge
  - VPNs are typically aggregated over tunnels
    - Resources are shared and only concerns are capacity and routing
  - But network slices need to be isolated at every hop
    - How many slices will there be?
- Alternatives exist with new technologies
  - A combination of central planning and Segment Routing
    - Central planning is able to determine optimal paths
    - Central planning can keep track of bandwidth usage
    - SR can steer packets onto appropriate paths without (much) state in the network
    - But network nodes still need to be programmed with information about slices

VPN+ Work In Progress at the IETF

- draft-ietf-teas-enhanced-vpn
  - A Framework for Enhanced Virtual Private Networks (VPN+) Service
    - Overview of functions and requirements for VPN+
- draft-dong-spring-sr-for-enhanced-vpn
  - Segment Routing for Enhanced VPN Service
    - Overview of how to use SR to achieve VPN+
- draft-dong-teas-enhanced-vpn-control-plane
  - Control Plane Considerations for Enhanced VPN
    - Control plane requirements, functions, and considerations for VPN+
- draft-dong-lsr-sr-enhanced-vpn
  - IGP Extensions for Segment Routing based Enhanced VPN
    - Floods Multi-Topology or SR Flex-Algorithm information
      - Scaling concern depends on number of enhanced VPNs
    - Packets are tagged
- draft-drake-bess-enhanced-vpn
  - BGP-LS Filters : A Framework for Network Slicing and Enhanced VPNs
    - Targeted programming (rather than flooding)
      - Better scaling, but a second protocol
    - Packets use DSCP or SR
- draft-zhuang-bess-enhanced-vpn-auto-discovery
  - BGP Extensions for Enhanced VPN Auto Discovery
    - Builds on L3VPN auto-discovery

Resources

- Most relevant working groups
  - TEAS
  - PCE
  - BESS
- VPN
  - RFC 4364 : BGP/MPLS IP Virtual Private Networks (VPNs)
- PCE
  - RFC 4655 : A Path Computation Element (PCE)-Based Architecture
- ABNO
  - RFC 7491 : A PCE-Based Architecture for Application-Based Network Operations
- BGP-LS
  - RFC 7752 : North-Bound Distribution of Link-State and Traffic Engineering (TE) Information Using BGP
- Virtual Networking
  - RFC 7926 : Problem Statement and Architecture for Information Exchange between Interconnected Traffic-Engineered Networks
- SDN
  - RFC 8283 : An Architecture for Use of PCE and the PCE Communication Protocol (PCEP) in a Network with Central Control
- ACTN
  - RFC 8453 : Framework for Abstraction and Control of TE Networks (ACTN)
- VPN+
  - draft-ietf-teas-enhanced-vpn : A Framework for Enhanced Virtual Private Networks (VPN+) Service

Questions and Follow-up

adrian@olddog.co.uk