Enhanced Virtual Networks (VPN+)

Stewart Bryant & Jie Dong (Huawei)

draft-bryant-rtgwg-enhanced-vpn-01

Overview

- VPNs have been widely deployed to support multi-tenancy in public operator networks.
- They are now expected to provide emerging new services/customers, with more stringent performance requirement (e.g. bandwidth, latency, jitter, etc.), on a shared network infrastructure
- An enhanced VPN is needed to:
    - Enable multiple customers with demanding services in a shared network.
    - Ensure high performance with reasonable cost and scalability.
    - Provide an underpin for 5G network slicing.

Requirements

- Isolation between VPNs
  - Greater isolation than routing table separation
- Guaranteed performance
  - Bounded packet loss, latency and jitter for critical services
- Integration
  - Between overlay and underlay
  - Network and service functions
- Customization
  - On-demand network topology and resource allocation
- Disruption-free Service Management
  - Need to add, modify and delete services without disrupting other services

Layered Architecture

- Overlay
  - Customized virtual network topology and service routing.
  - Normally compete for shared resource in underlay.
  - VPN+ needs to arbitrate such that the resources are allocated where needed.

- Underlay
  - Physical network with various network resources, which can be partitioned for different service needs.
  - Provide transport connectivity between overlay nodes.

- Overlay and underlay needs to be tightly integrated
  - For greater isolation and performance guarantee
  - VPN+ is more than simple connectivity

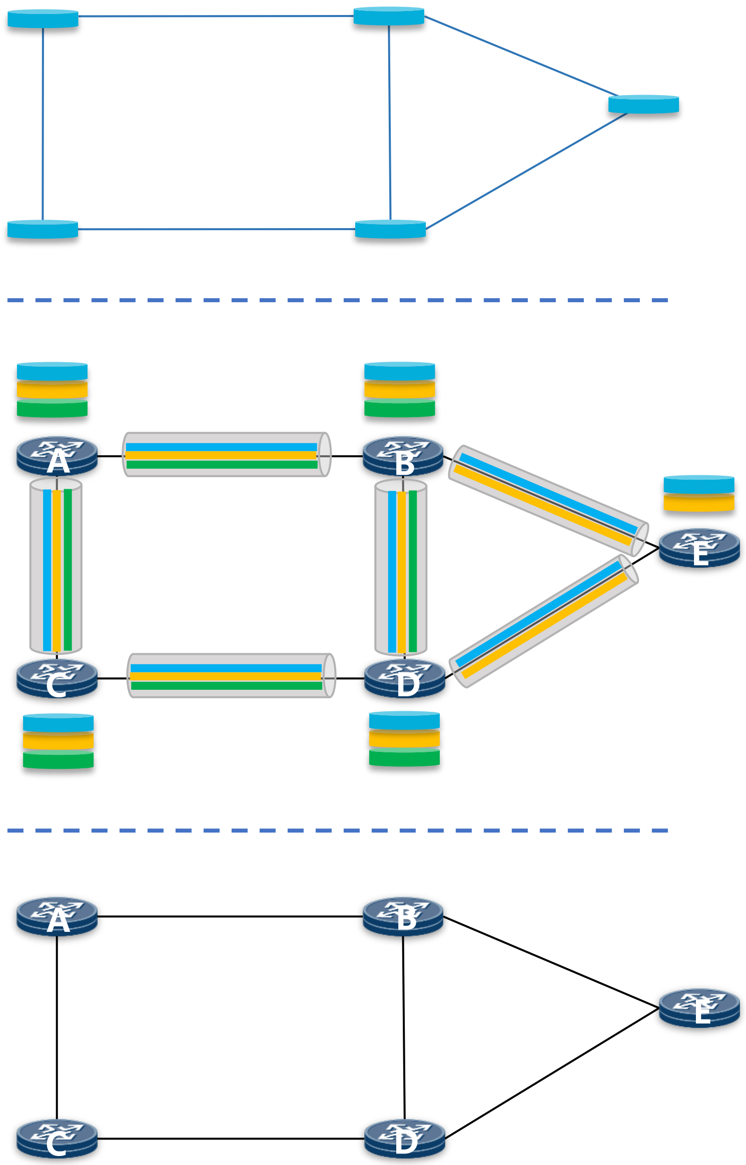

Traditional Overlay Mechanisms

- Overlay Virtual Networks
  - Each overlay network can have its own routing/forwarding table and separate address space
  - Overlays compete for network resources with each other, unless every connection in overlay is mapped to one dedicated TE-LSP for bandwidth reservation

- Problems
  - The increase of overlay tenants asking for guaranteed performance results in the increase of TE-LSPs, which ultimately leads to scalability problems
  - Bandwidth reservation is not enough to guarantee latency, jitter etc.

Multi-Topology Routing

- Multiple-Topology Routing (MTR)
  - Provides multiple customized network topologies for different services
  - Mapping of data packets to specific topology is not addressed by MTR design (Complex ACL structure normally needed)
  - MTR together with Segment Routing can solve the data plane mapping problem by steering packet through different paths

- Problems
  - MTR assumes best effort forwarding service.
  - Neither MTR or SR can provide resource reservation which is necessary for performance guarantee
  - Topology ID limits may be a problem

Proposed Mechanism

- Extend Segment Routing for resource reservation
  - Per-hop instead of per-path resource reservation
    - Follow the paradigm of SR, achieve resource reservation with much less state
    - Controller based resource reservation, no needs of signaling protocol
    - Dedicated SIDs for different partitions of link/node resources
    - Each logical network is constructed with a set of dedicated SIDs
    - Aggregate resource in network in low stress points, but disaggregate in the packet
      - A hybrid between aggregation and strict per-hop reservation
  - Flexible and fine-grained resource manipulation
    - SR SIDs can be used to represent various types of resources
  - VPN service maps to SR logical network efficiently
    - Reduce the provisioning overhead of per-tunnel binding to VPN

Proposed Mechanism

- Allocate dedicated SIDs for partitioned link/node resources
- Each SID associates with one particular virtual network

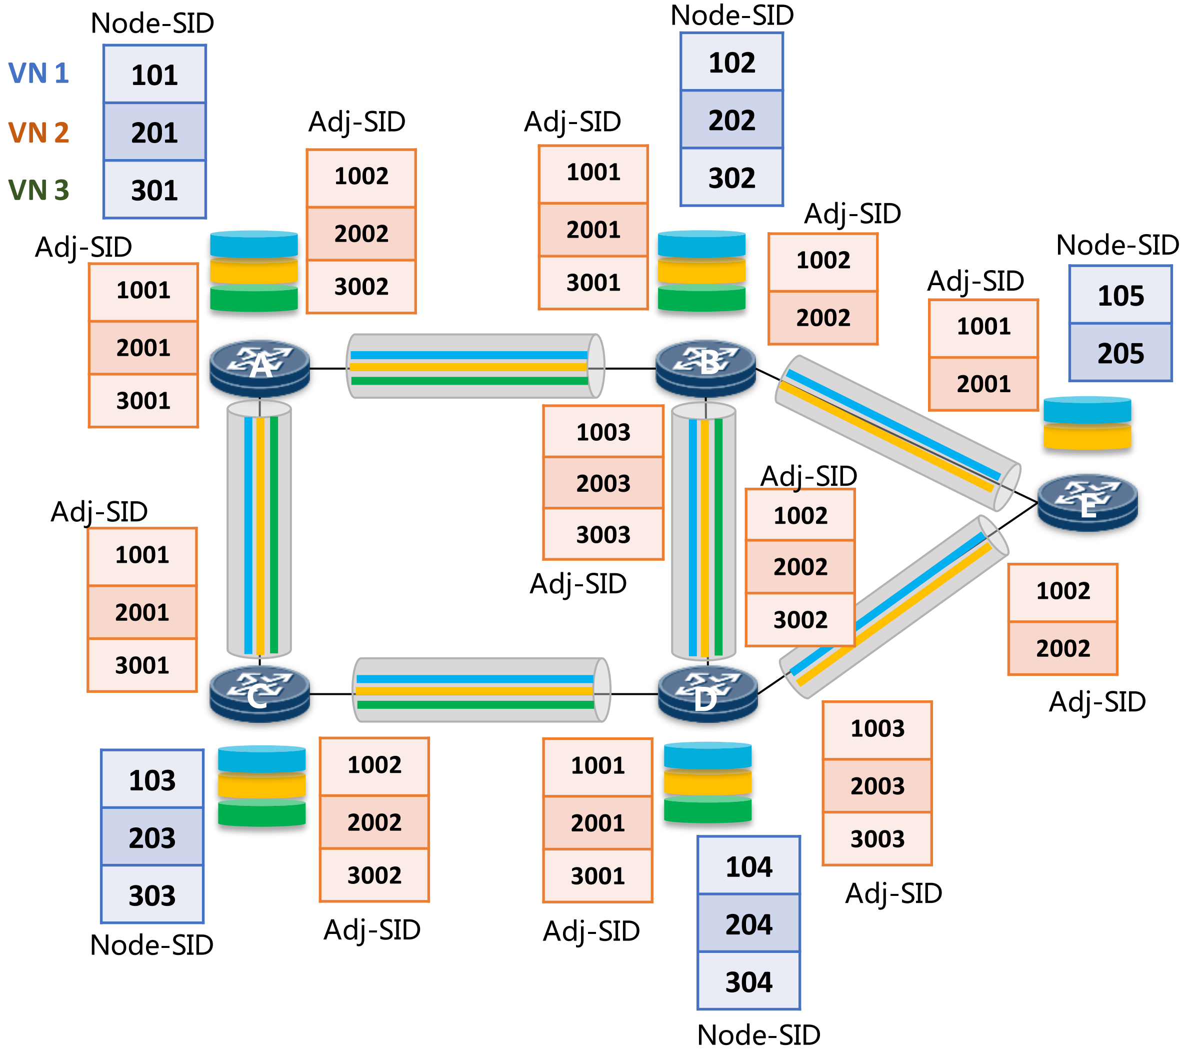

Customized SR virtual networks

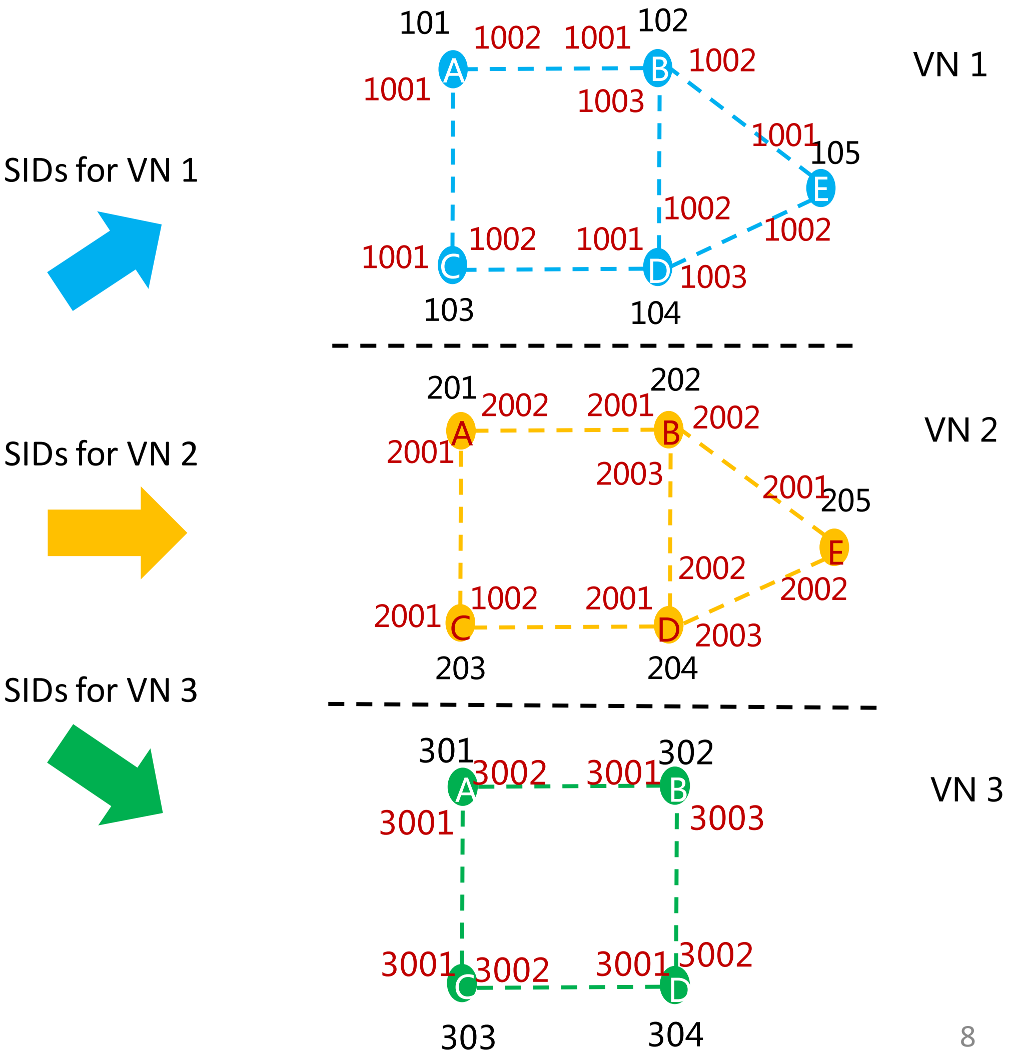

Proposed Mechanism: Forwarding Plane

● Strict path: A-B-D-E

- Packets in different virtual networks are encapsulated with different SID lists.
- On each hop, SID maps to reserved network resources.

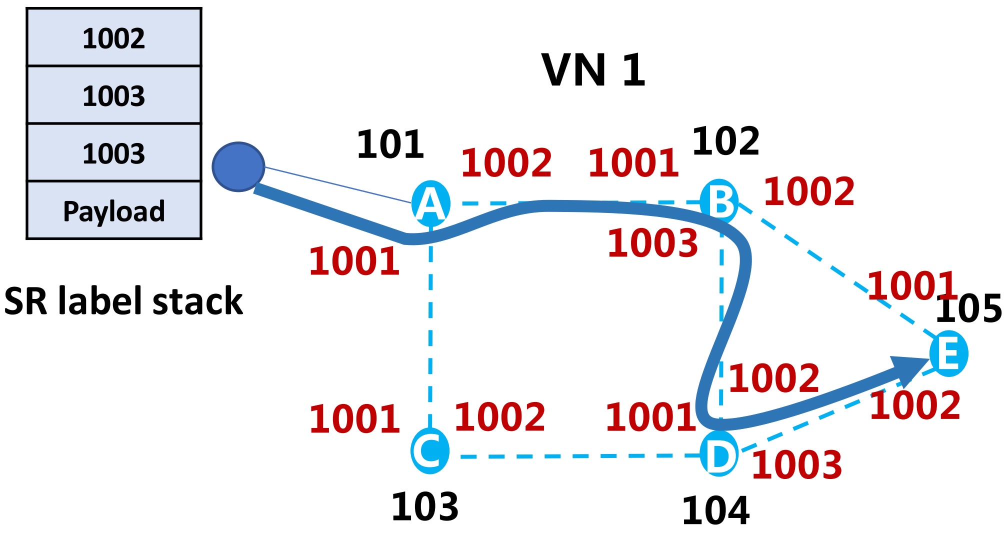

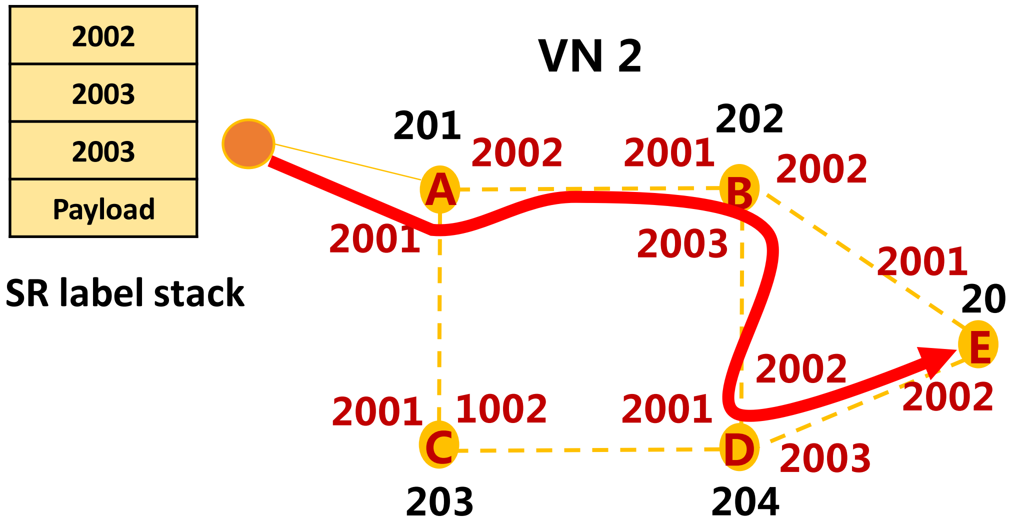

Proposed Mechanism: Forwarding Plane

Loose path: A-D-E

- Loose path forwarding is achieved with per-VPN node-SIDs.
- Loose path computation is constrained to specific virtual network.
- Resources strictly allocated to the VPN, but aggregated within the VPN.

104
105
Payload

SR label stack

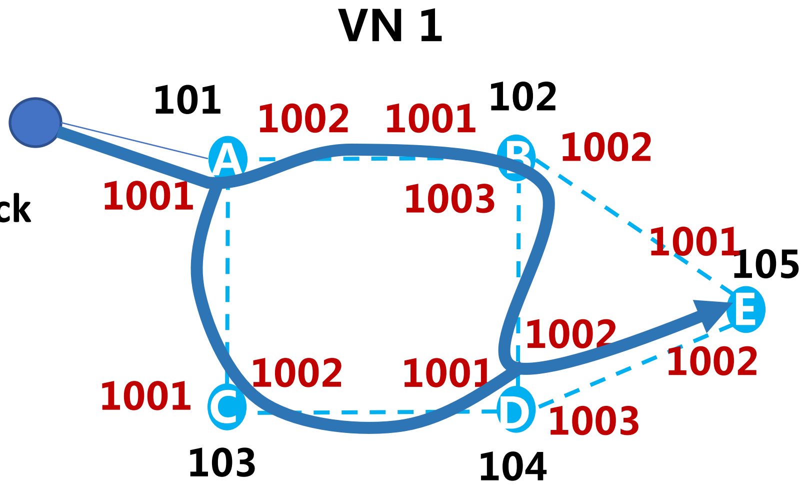

204
205
Payload

SR label stack

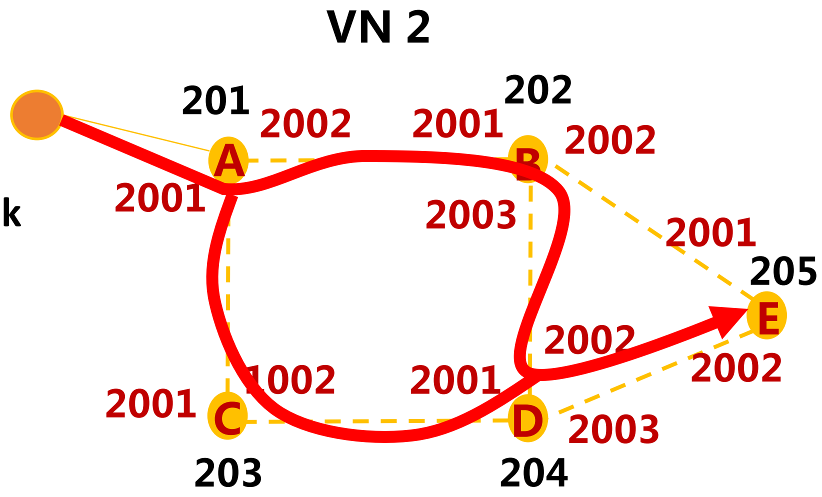

Relationship to DETNET

- The performance goals of DETNET are similar to VPN+:
  - No packet drop due to congestion
  - Low latency, but upper bound more important than minimum latency
- DETNET currently runs over a data plane that is unmodified apart from the replication-elimination process.
- There are proposals in DETNET to enhance the router queuing model. These need to be studied.
- Multi-tenancy is not currently part of the DETNET design.
- As VPN+ aims to provide multi-tenancy with performance guarantee, DETNET may be used within a VPN+ instance to provide these guarantees.
- To get the most out of the available underlay, VPN+ will need tighter integration with the underlay than currently planned for DETNET.
- Whether VPN+ extends DETNET, or whether it uses DETNET in an enhanced way, VPN+ needs to move some of the current DETNET boundaries.

Integration with SFC

- More functions within the network are being virtualized.
- Enhanced VPN tenants will require access to virtualized functions with similar performance and isolation characteristics to those needed for conventional network functions.
- This is consistent with the holistic network slicing view.
- This leads to a need to integrate service function chains with enhanced VPNs, and pushes us in the direction of a common technology.
- Segment routing is one candidate technology to achieve this integration. Other approaches are for further study.

Scalability Considerations

- When we move from best-effort to guaranteed performance we need to provide greater integration between the VPN and the underlay.
- We can
  - Introduce state into the packet (the SR approach)
  - Introduce state into the network (the RSVP-TE approach)
  - Provide a hybrid approach.
- Dynamic creation of VPN paths using SR requires less path state maintenance, but requires more latent state.
- This is an aspect of the problem that requires further study.

Disruption-Free Service Management

- One of the more challenging problems is the reconfiguration of a VPN+ instance without disrupting traffic in that instance and without disrupting other instances.
- SR handles path change well because there is always consistency between the intended path and the packet path identifier.
- SR loose paths are subject to looping unless convergence control technology is employed (IPFRR Loop Free Convergence).
- The hardest problem is disruption-free de-fragmentation, which is for further study.

Thank You
Further Questions?