I2RS Topology Use Cases

Nabil Bitar, Samita Chakrabarti, Luyuan Fang, Wassim Haddad, Giles Heron, Ram Krishna, Himanshu Shah

General Description

- Mechanism to learn and to provision:

1. customer topology
2. service topology
3. service chaining

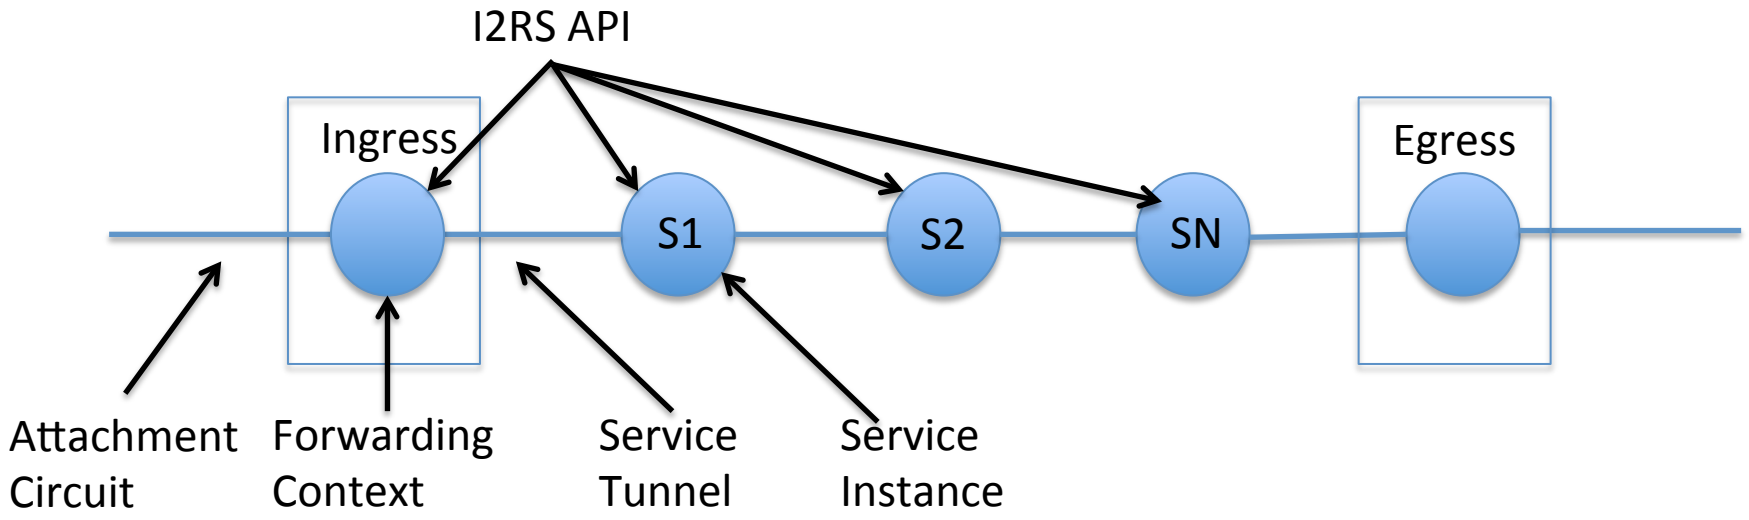

Customer Topology

- Interfaces
- Network Service and Attributes

Service Topology

- Service node per service
  - Liveliness
  - Performance (monitoring and stats)
- Service instantiation
  - Pre-built, spin up on demand, n:1 vs 1:1
- Service node capacity
  - CPU, B/W, policies, ACLs etc.

Service Chaining

- Service chaining / forwarding
- Classification of packets
- Service profile
- 2 models:
  1. Create whole chain at ingress (requires metadata in packet header?)
  2. Recursive model where provision each link in the chain
- Need to support service multiplexing/demux
  - Mux example – single firewall for multiple users
  - Demux example - stateful load balance etc.

Scope and Scale

- Applicable to SP edge network and to NFV in the data-center.
- Applicable to service-aware network elements (PE routers, NFV elements)
- 3 types of operations:
  - Discovery/publishing – when resources change. Can be learned dynamically or stored in centralized database.
  - Configuration – when setting up a service
  - Monitoring/statistics – very high frequency (should these be passed over I2RS API?)

What's Different for I2RS

- Via CLI and/or AAA systems (RADIUS or DIAMETER)
- Data/information model standardization
- Standardized APIs with rich semantics
- I2RS solution requires:
  - Data Model / Information Model
  - APIs
  - Provisioning / De-provisioning mechanics

Use-Case In Context

- Network application here is assumed to be customer service management (enabled FCAPS for network services).
- Network services can be L2/L3VPN, L4-L7 apps (firewalls, DDoS etc.).
- Assume application is layered above an SDN controller (or plugged into the controller)
- Will now focus on L3VPN use-case

L3VPN Service Provisioning with I2RS

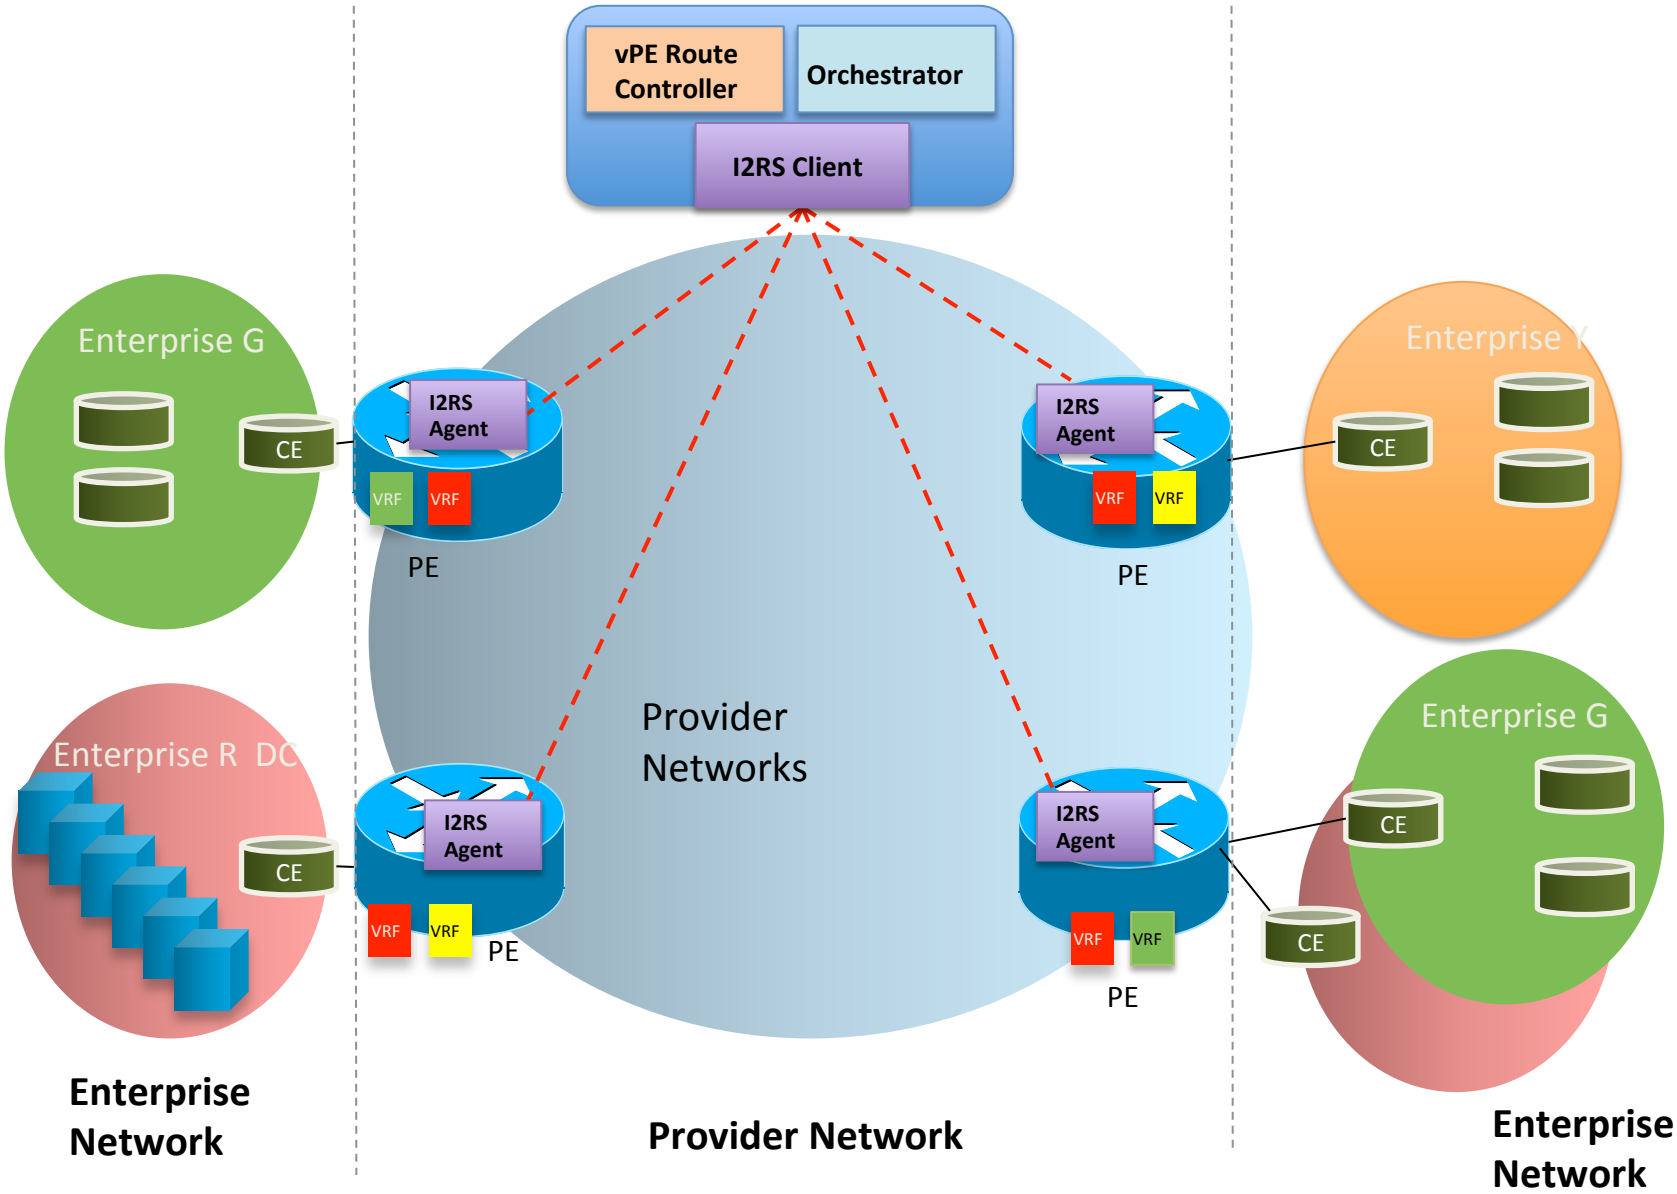

L3VPN Service Chaining with I2RS

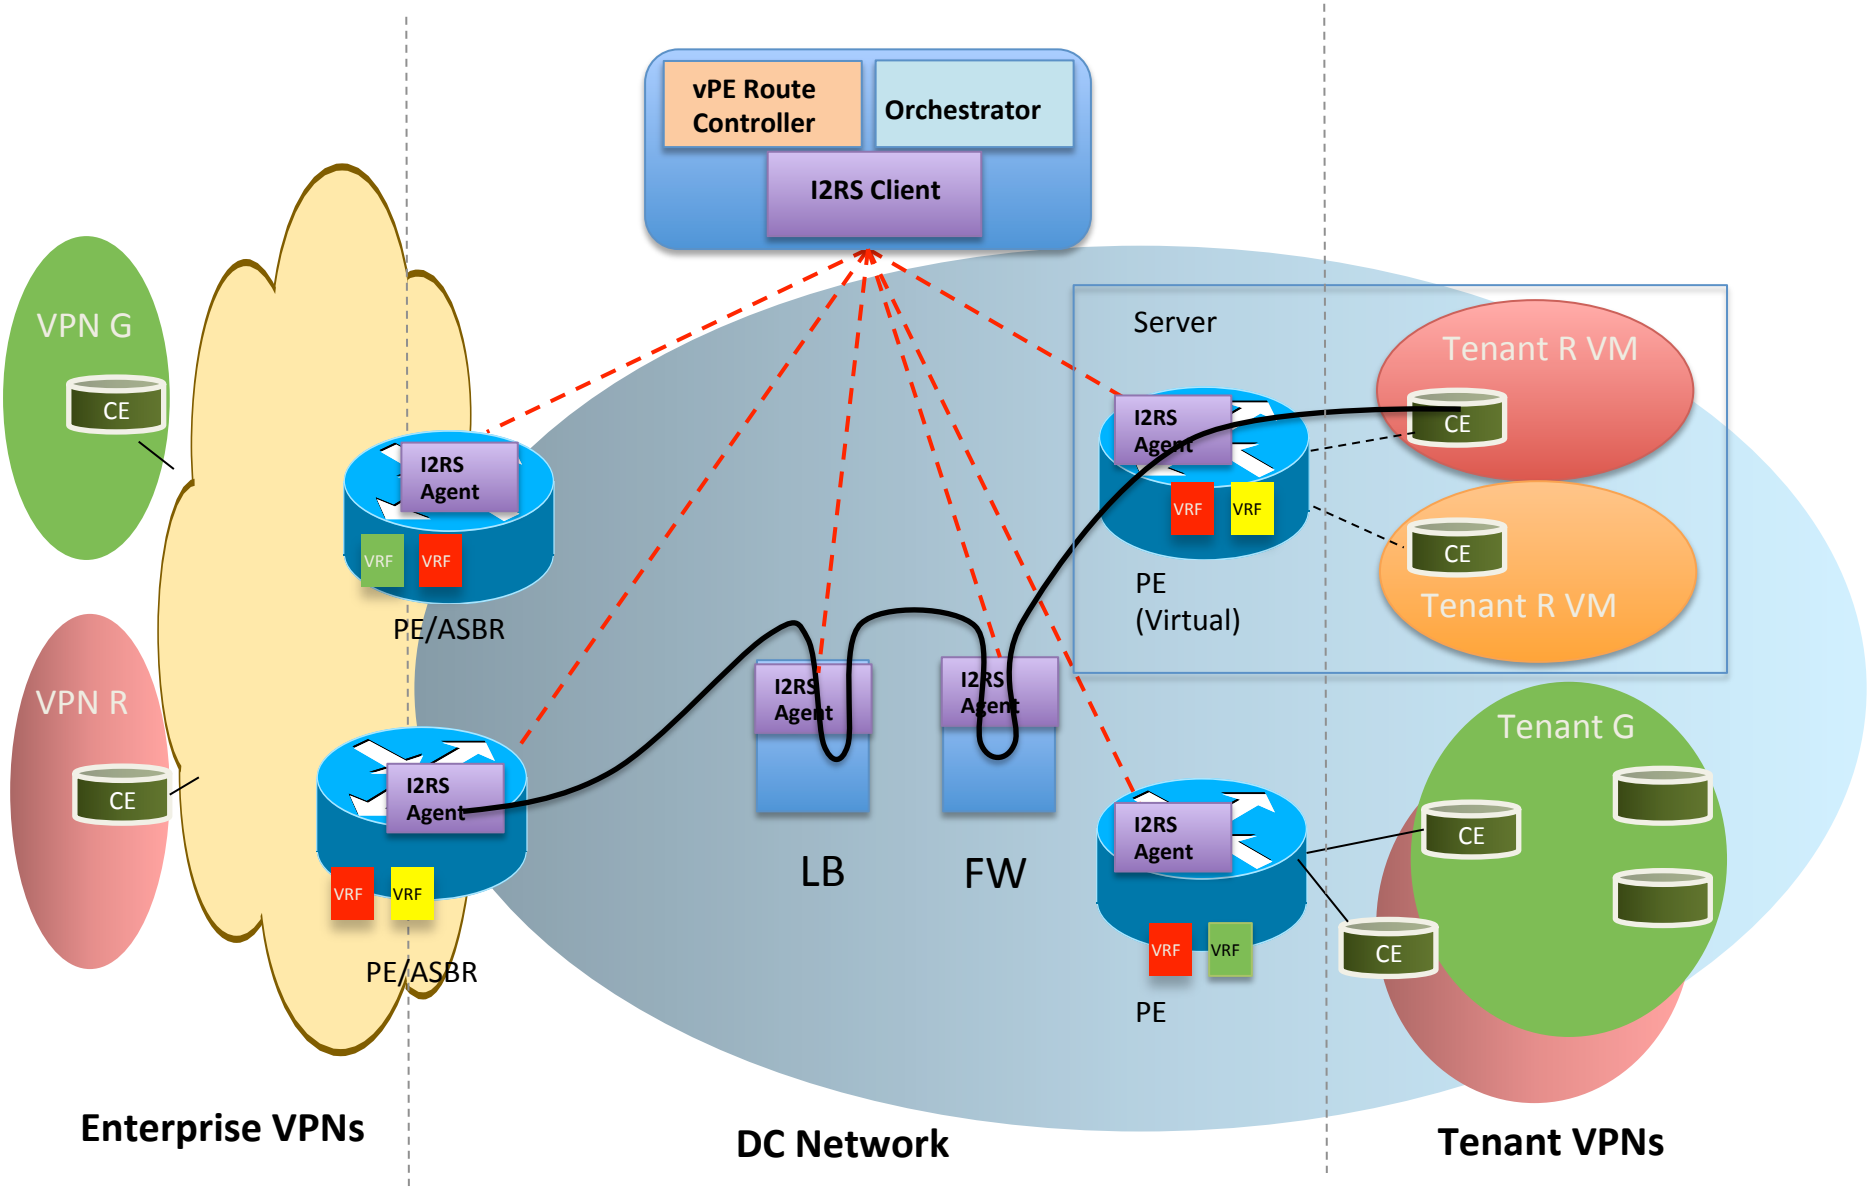

L3VPN use case context

- Assumptions
  - Service Provider network already supports L3VPN services
  - Hybrid model: I2RS + traditional service provisioning/monitoring
- Key Attributes of the services
  - Multi-tenant Layer 3 connectivity
  - Multi-point managed services

What events are needed

- PE/CE interface up/down
- PE/CE routing protocol up/down
- Control plane thresholds (e.g. number of routes learned exceeds limit)
- Data plane thresholds (e.g. bandwidth exceeds contracted rate per CoS)
- Others?

Data – Client pull from Agent

- List of candidate IP interfaces
- Data plane resource availability per candidate interface (Policers / Shapers / ACLs)
- Control plane resource availability (route memory, BGP sessions, VRFs, CPU etc.)
- Traffic statistics (per CoS etc.)

Data – Client push to Agent

- VRF creation
- RD/RT assignment to VRF
- Interface mapping to VRF
- PE/CE routing protocol configuration (protocol selection, parameters, route limits etc.)
- Interface OAM
- QOS: policers, shapers
- ACLs

Data – published by agent

- Data plane operational state (interface up/ down, OAM status etc.)
- Control plane operational state (routing protocol status, number of routes etc.)
- Traffic statistics (per CoS etc.)

What needs to be written/modified?

- Customer interface and VRF configuration needs to be written (see data pushed to agent on earlier slide)
- What does “state bulking” mean?

Client-Agent Interaction

- Client pushes data down to agent on a per-agent basis.
  - No requirement for transactions across multiple NEs
- Hybrid model:
  - Resources set aside for I2RS provision (other services may be configured using traditional CLI/AAA models)
- NE publishes events (see earlier slides)
- Bulk stats published by agent or pulled by client