“Using BGP as an Auto-Discovery Mechanism for Network-based VPNs”

<draft-ouldbrahim-bgpvpn-auto-00.txt>

Hamid Ould-Brahim
Bryan Gleeson
Peter Ashwood-Smith
(Nortel Networks)
Eric C. Rosen
(Cisco Systems)
Yakov Rekhter
(Juniper Networks)

49th IETF, San Diego, December 2000

Motivations

- Define a common BGP based auto-discovery mechanism to distribute certain VPN related information between PEs (e.g., membership, reachability, topology, tunnel information, etc.).
- Allows a VPN to auto-discover its members, and to identify the set of PEs having a VPN in common.
- Used for both the virtual router and RFC2547bis architectures.
- Each scheme (VR or RFC2547bis) uses the mechanism to automatically discover the information needed by that particular scheme.
- Interworking scenarios between RFC2547bis and the virtual router models are addressed.

Network Reference Model

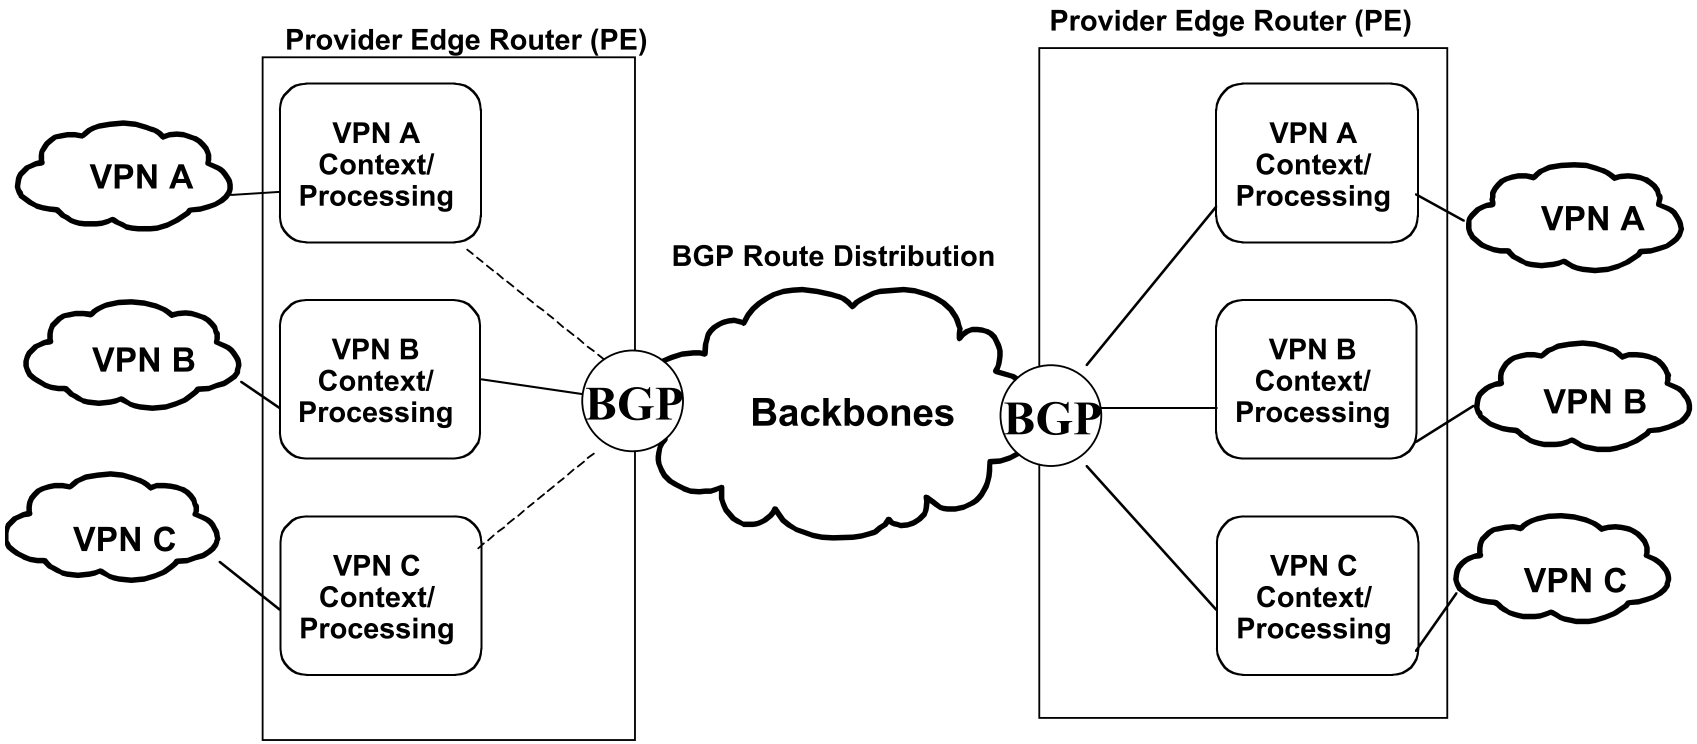

Each PE knows the type of architecture it is supporting.

What can be auto-discovered?

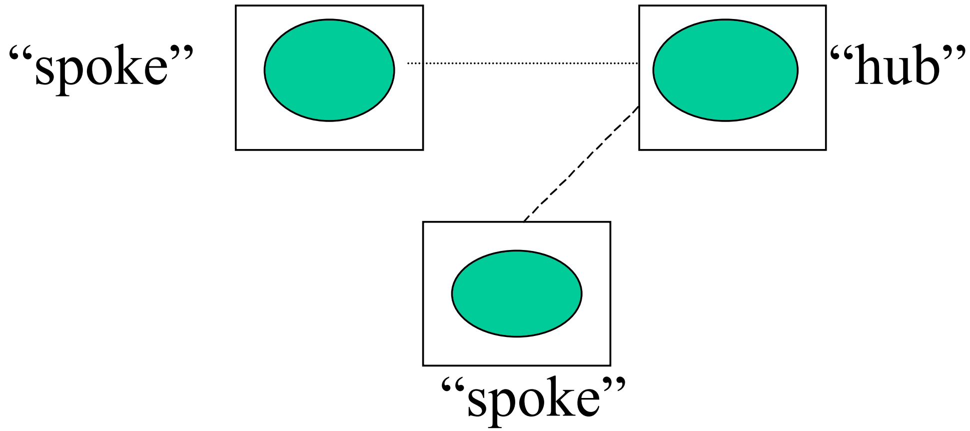

Information related to identifying the set of VPN Context/Processing (or VPN) intended to receive the vpn information? ➔ **Membership**

Information related to the nature of the interconnectivity (implicit or explicit) between theVPN contexts/Processing within the same VPN? ➔ **Topology**

Information related to the endpoints to use in order to achieve VPN/PE connectivity?
➔ **Tunnel Information (Endpoints)**

Information related to the set of VPN routes used to reach destinations in the VPN space?
➔ **Reachability**

What mechanism?

- BGP is used to implement the auto-discovery mechanism.
- BGP-4 multiprotocol extensions are used to carry VPN information for both VR and RFC2547 architectures.
- The NLRI is a VPN-IP address or a labeled VPN-IP address.
- VPN-specific information associated with the NLRI is encoded either as attributes of the NLRI, or as part of the NLRI itself, or both.
- The address prefix in the NLRI field is ALWAYS within the VPN address space, and therefore MUST be unique within the VPN.

- The address specified in the BGP next hop attribute is in the service provider addressing space.
- For the virtual router, the NLRI address prefix is an address of one of the virtual routers configured on the PE.
- For RFC2547bis the NLRI prefix represents a route to an arbitrary system or set of systems within the VPN.

RFC2547bis Scheme

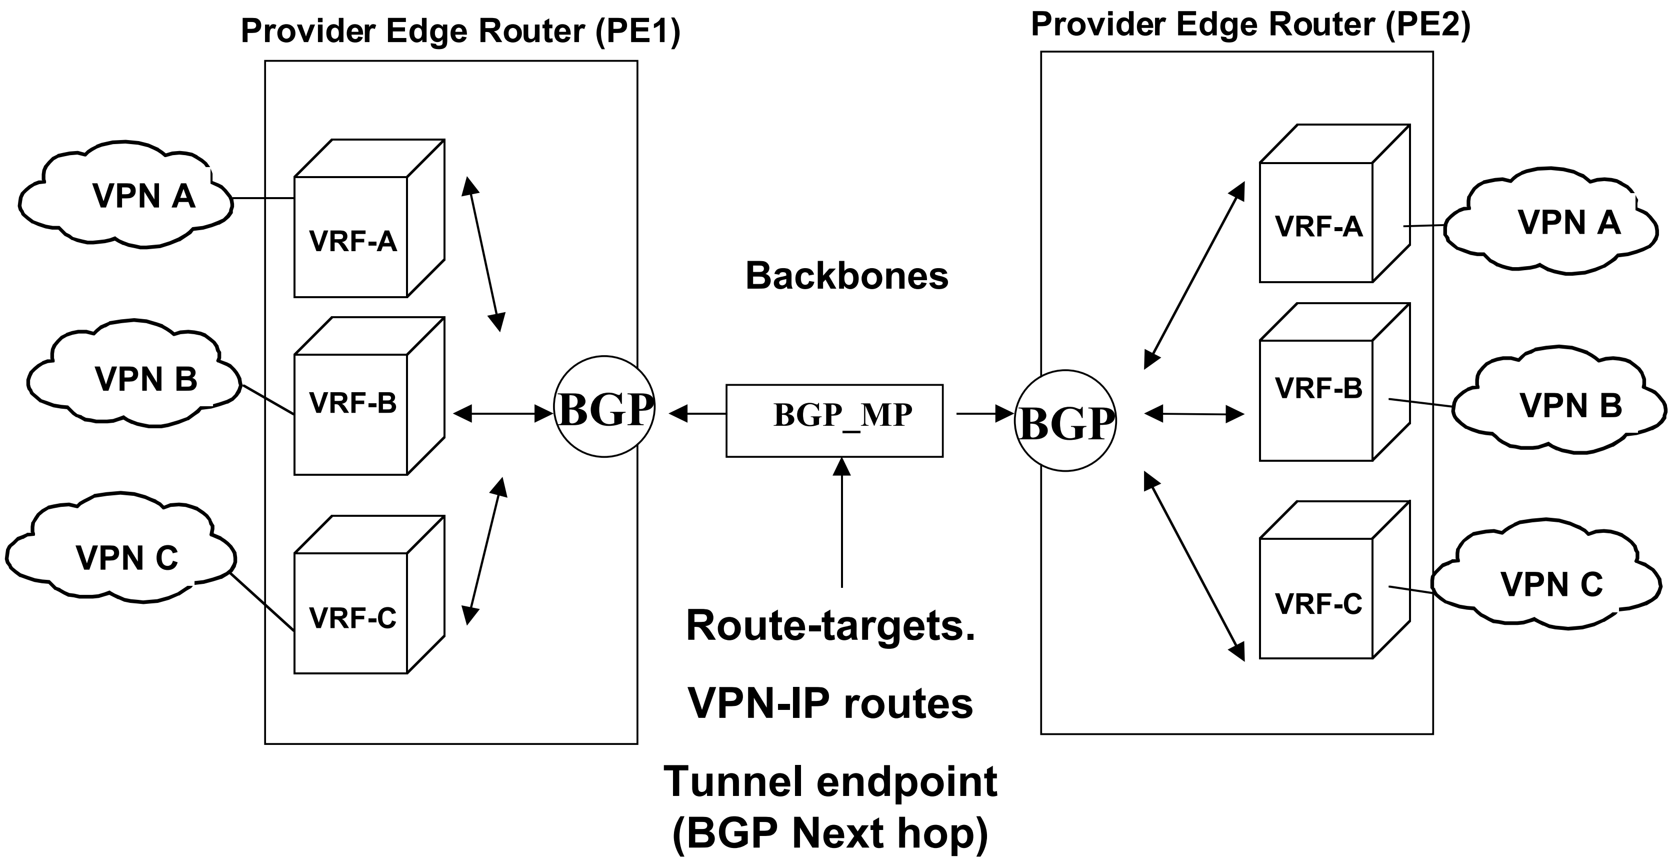

Virtual Router Scheme

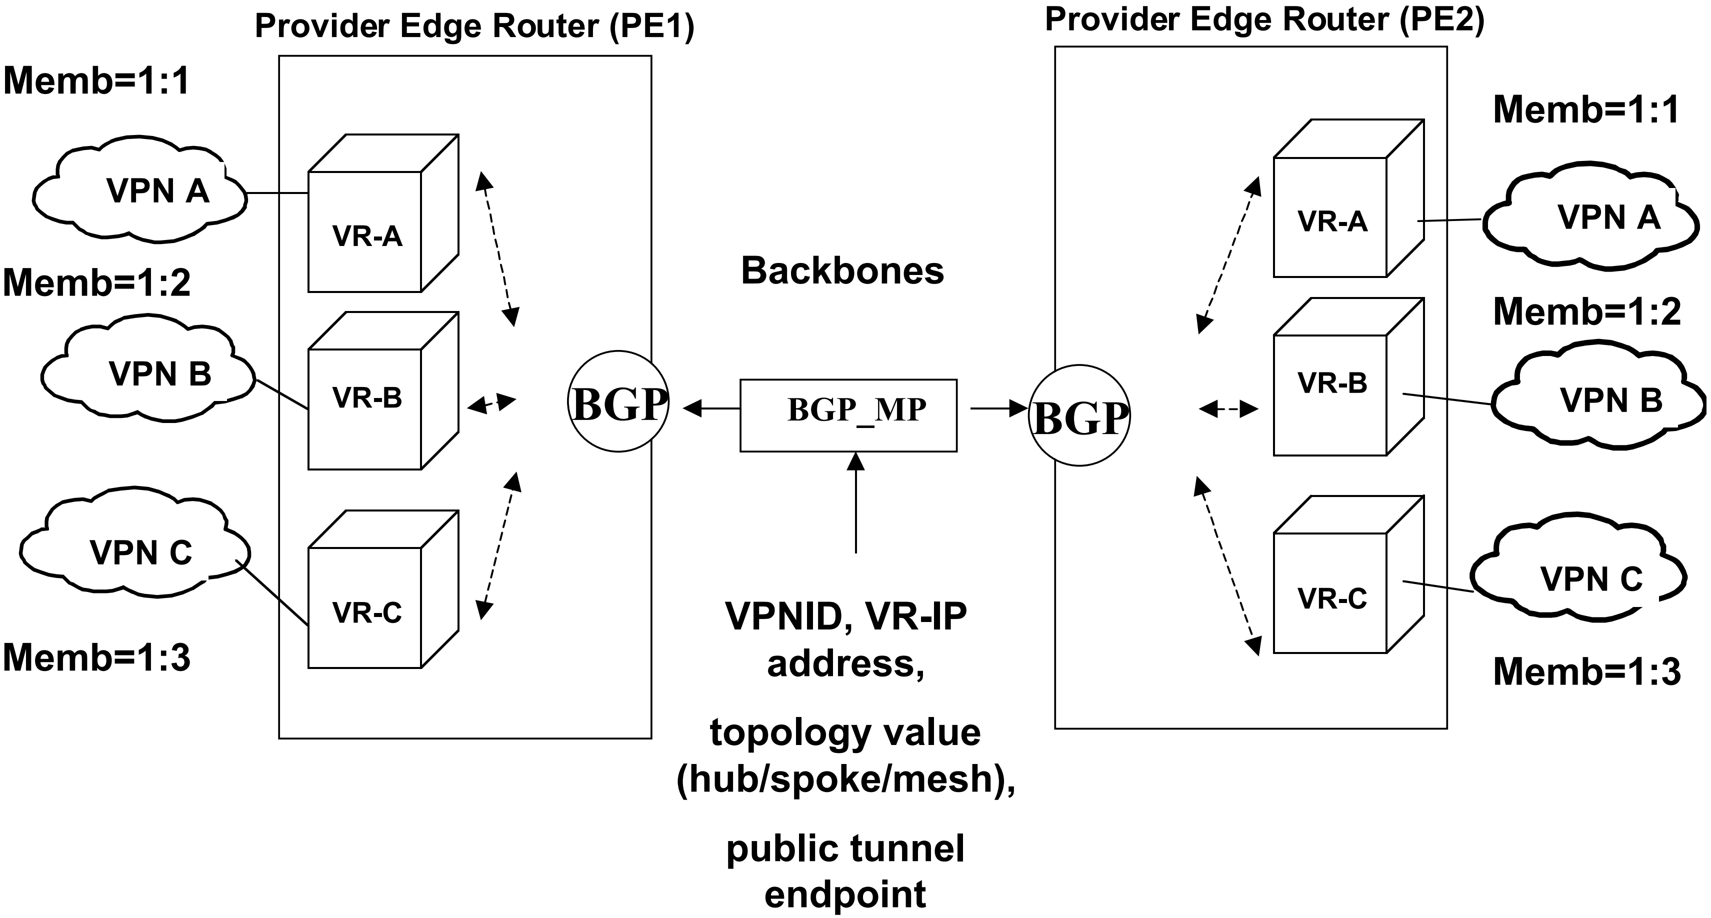

Interpretation of VPN Information in the RFC2547bis model

- The BGP attributes the Route Target Extended Community are used by the PE routers to assign the routes to particular VPN database/processing contexts, and hence implicitly determine the topology.
- The BGP Next Hop attribute specifies the remote end point of the tunnel to be used when sending packets whose destination addresses match the corresponding NLRI.

Interpretation of VPN Information in the Virtual Router Model

- A VPN-ID is carried in the NLRI in order to associate a particular VR address to the VPN.
- A value of 0x80 in the first byte of RD's type field indicates that the RD field is carrying the VPN-ID format. (range of 0x8000-0x80ff indicates the presence of VPN-ID format as defined in RFC2685).
- A new extended community attribute is used to carry the VPN-ID. A value of 0x20 indicates that the remaining 7 bytes following the first byte of the type field holds a VPN-ID value. (Range 0x2000-0x20ff will indicate the presence of the VPN-ID format).

Interpretation of VPN Information in the Virtual Router Model (Continued)

- A new extended community attribute is used to carry VPN topology values.
- 4 Bytes are allocated to hold the topology value. Values 1:Hub, 2:Spoke, 3:Mesh are reserved.
- Arbitrary topologies can also be constructed.
- The BGP next hop will carry the service provider tunnel endpoint address.

RFC2547bis and Virtual Router Interworking Scenarios

- Two scenarios are considered.
- Scenario 1: CE-PE relationship.
- Scenario 2: A single PE implementing both architectures.

CE-PE Relationship

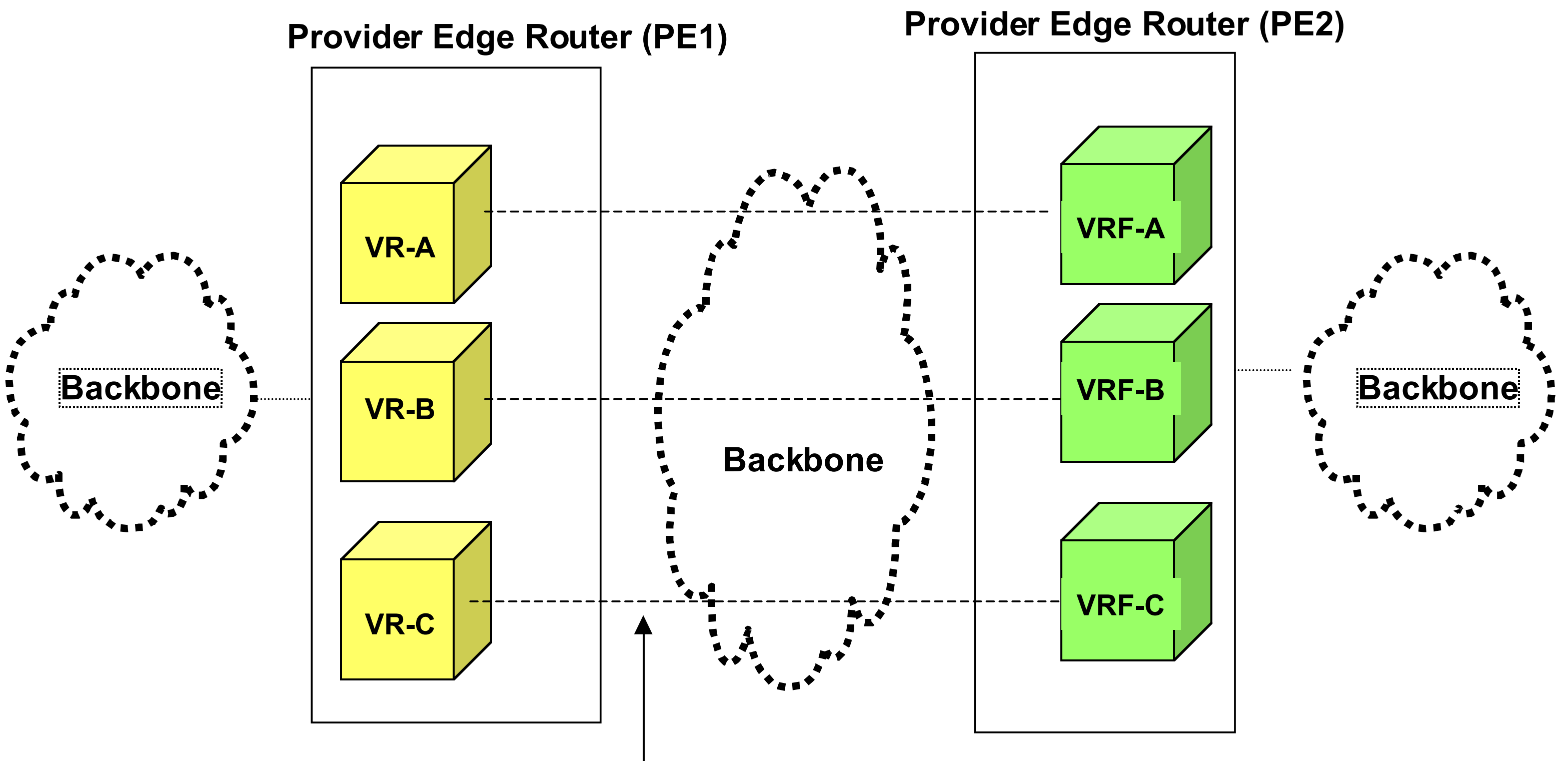

Link, Tunnel (e.g., IPSec)

PE with Both Architectures