

# One Administrative Domain

James Uttaro (uttaro@att.com)

Saikat Ray (sairay@cisco.com)

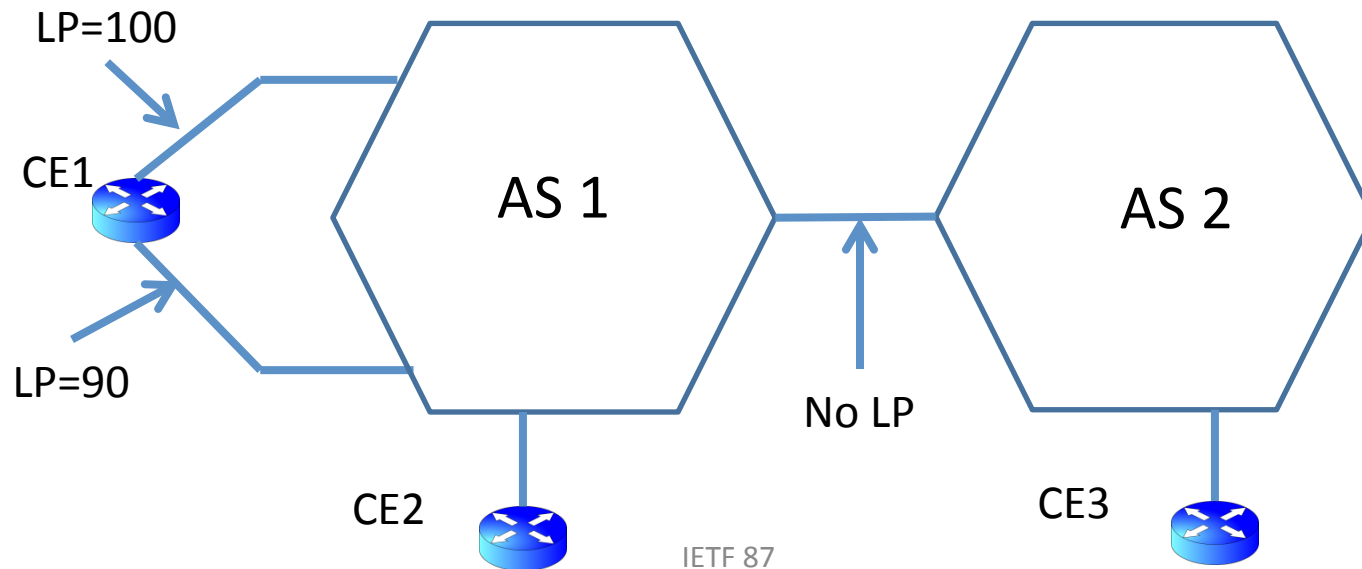
Pradosh Mohapatra  
(pmohapat@cumulusnetworks.com)

IETF 87

August 2013, Berlin

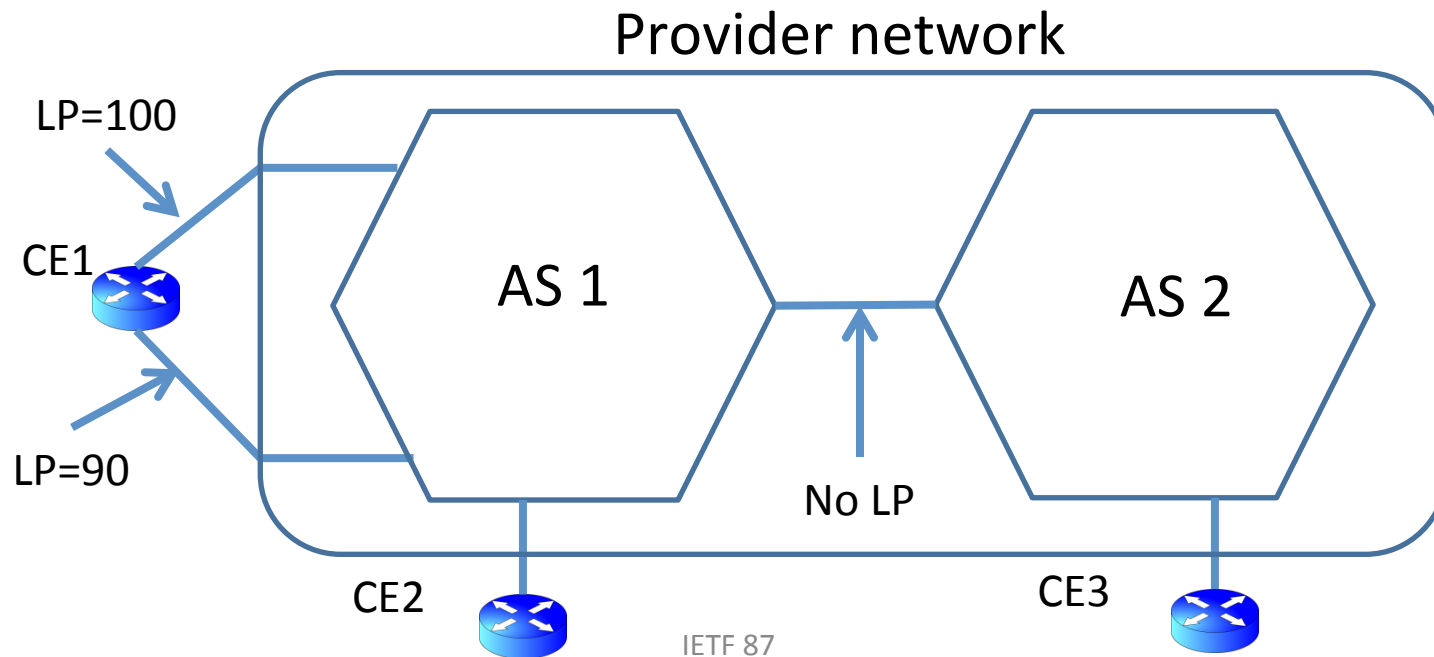
# OAD Overview

- Currently autonomous systems are equated with administrative domains
  - Different ASes assumed to have independent policies
    - Attributes like LOCAL\_PREF are not propagated across AS boundary
    - Traffic from CE2 honors LP; traffic from CE3 doesn't.



# OAD Overview

- Reality is, a large Provider may own multiple Ases
  - Renumbering would be nice, but may not be viable
  - The customers of the Provider expect all sites to behave the same way
    - Traffic from CE3 should use the primary link to reach CE1



# OAD Overview

- This draft proposes tunneling attributes across ASes to create “admin-domains” where uniform policy is used and enforced
  - An admin domain need not be contiguous
- Extend ATTR\_SET attribute for tunneling multiple sets of attributes
  - ATTR\_SET by itself cannot be used since it has an existing semantics for CE attributes

# ATTR\_SET\_STACK


- Optional, transitive attribute
- Creates a stack of ATTR\_SET attributes
  - Encodes multiple set of attributes
  - Encodes the sequence
- Rules for encoding two ATTR\_SET in this draft

```

+-----+
| Attr Flags (O|T) Code = TBD |
+-----+
| Length |
+-----+
| Attr Flags (O|T) Code = 128 | ^
+-----+
| Length (for the outer attrs) | |
+-----+ ATTR_SET 1
| Origin AS (provider network) | |
+-----+
. Path Attributes (variable) . v
+-----+
| Attr Flags (O|T) Code = 128 | ^
+-----+
| Length (for inner attributes) | |
+-----+ ATTR_SET 2
| Origin AS (customer network) | |
+-----+
. Path Attributes (variable) . v
+-----+
// //
// //
+-----+
| Attr Flags (O|T) Code = 128 | ^
+-----+
| Length (for inner attributes) | |
+-----+ ATTR_SET n
| Origin AS (customer network) | |
+-----+
. Path Attributes (variable) . v
+-----+

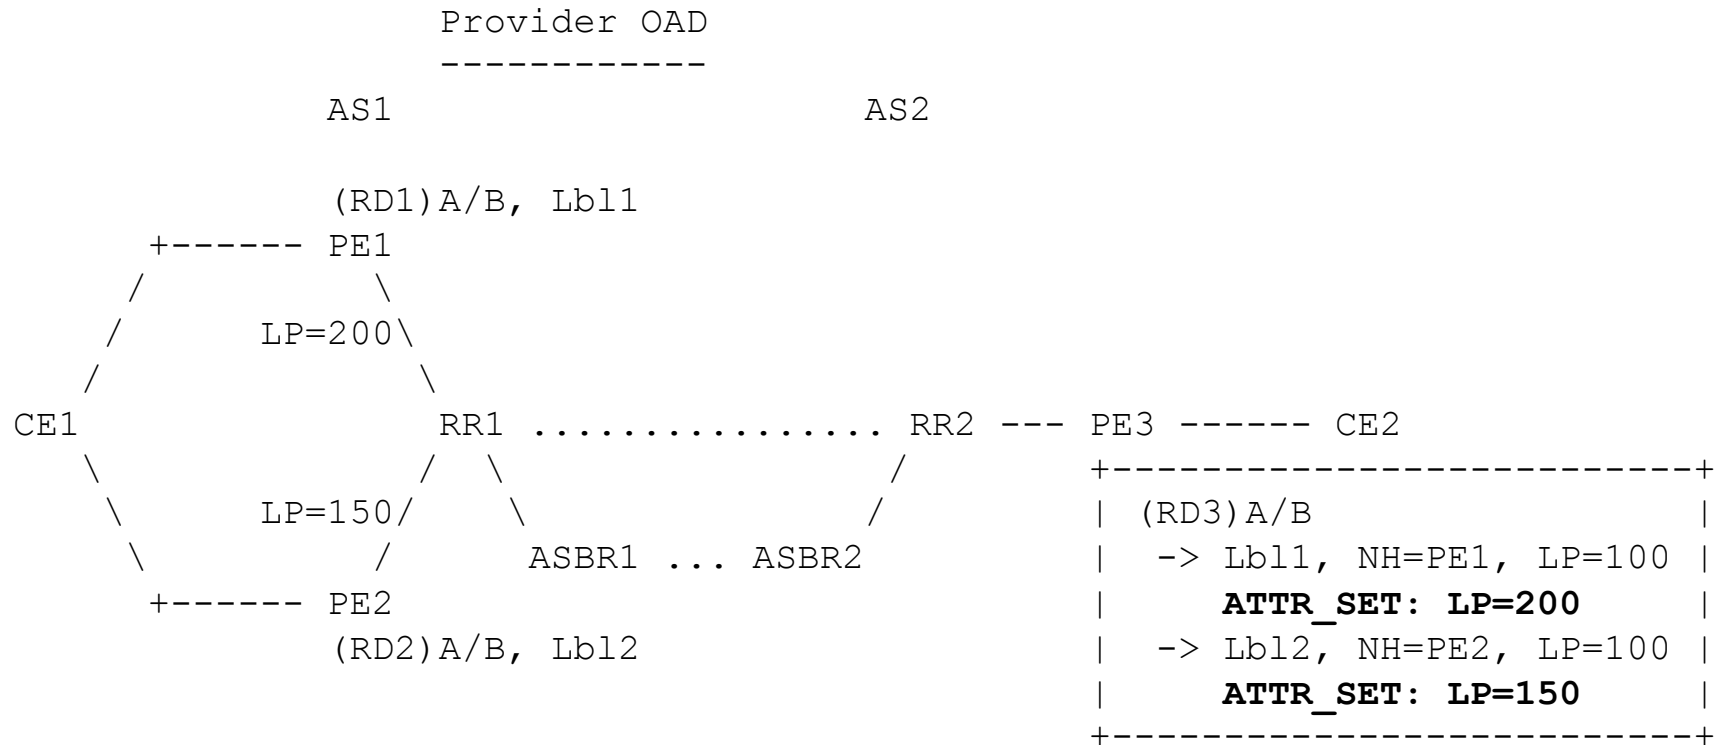
```

# Example: Option C



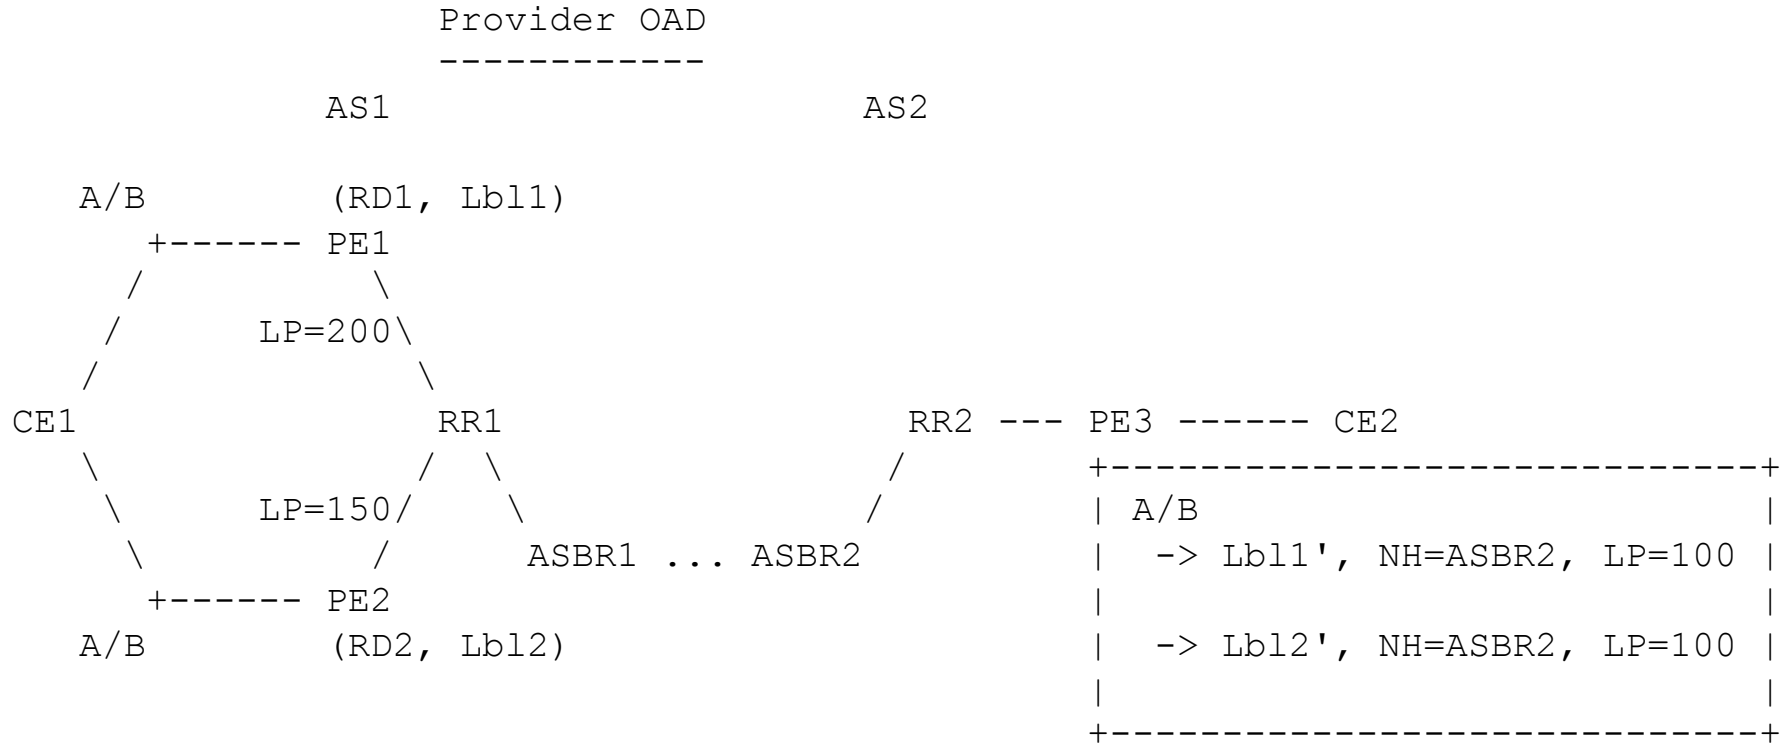
Existing behavior

# Example: Option C



OAD behavior

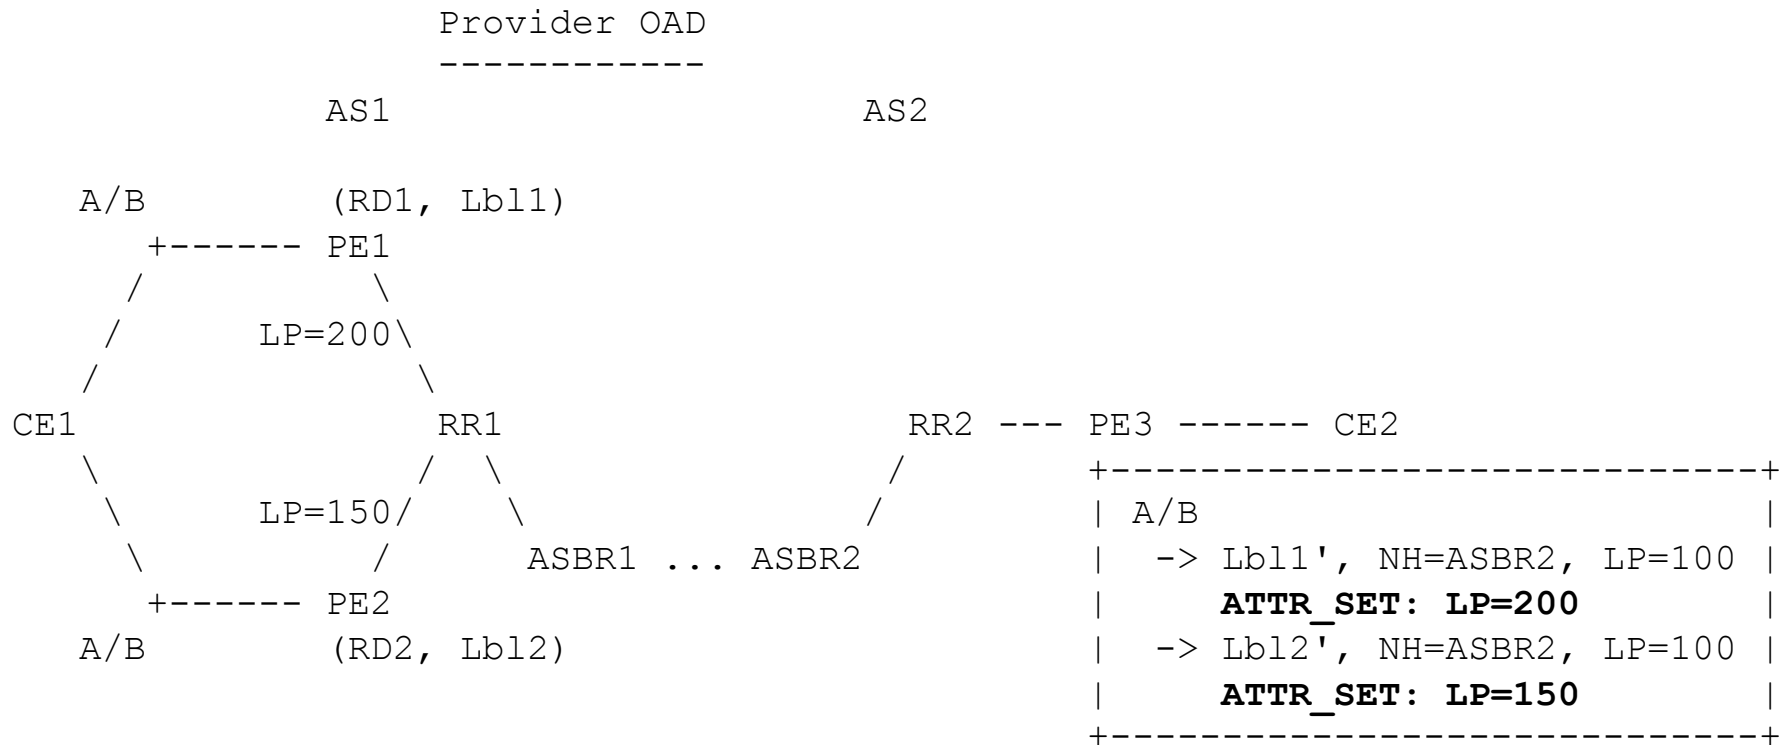
# Example: Option B



Existing behavior

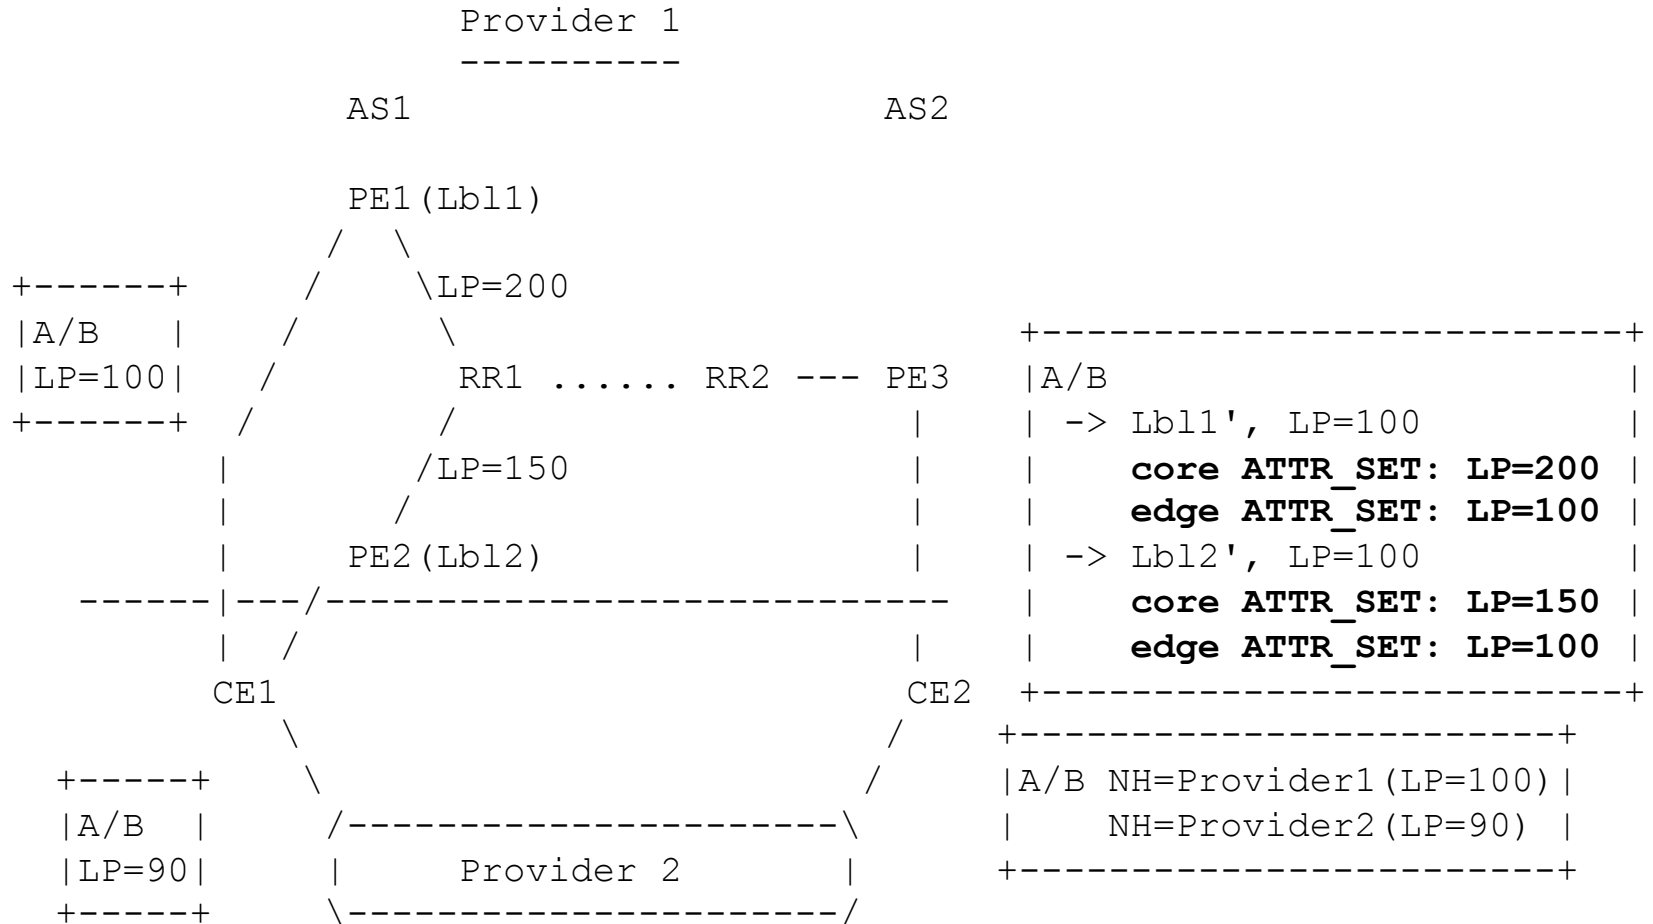


# Example: Option B



OAD behavior

# Example: Dual Provider case



OAD behavior

# OAD behavior

- ATTR\_SET\_STACK encodes both the edge attributes (iBGP PE-CE) and the core attributes (PE-PE/RR)
  - Needs the ATTR\_SET\_STACK wrapper even if only the core attributes are being sent.
- At the ingress PE, the core attributes (from the egress AS) as well as the local attributes are available
  - Local policy to make a choice
- Rules for encoding/use only at ingress and egress ASes; intermediate ASes do not use it
- Use policy to decide whether to send ATTR\_SET\_STACK attribute to an EBGP peer or not

# Next steps

- Make the draft a WG document
- Comments?