LABELED ARP

Kireeti Kompella
CTO, JDI

PROBLEM STATEMENT (DC)

Overlays are all the rage today in the data center

- but we've been doing overlays/underlays with MPLS since 1997

The DC overlays start at the host (server)

- which requires true "plug-and-play" operation

To have an MPLS underlay network, the host must be part of the underlay

- Here, we show how to make that easy and plug-and-play

Copyright © 2013 Juniper Networks, Inc.

PROBLEM STATEMENT (ACCESS)

Many have suggested that MPLS should start at the **access node** (DSLAM, OLT, cell-site gateway)

“Seamless MPLS” has proposed the use of LDP “Downstream on Demand” (DoD) for this

There haven’t been many implementations of LDP DoD from access node vendors

- Thus, a different approach and protocol for the same functionality seems advisable

Copyright © 2013 Juniper Networks, Inc.

OVERLAYS/UNDERLAYS

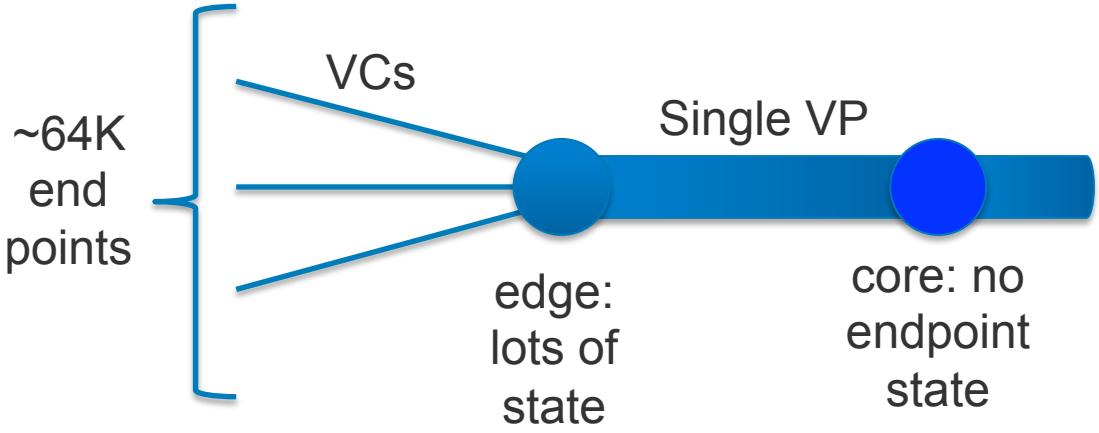

ATM overlay
(simple)

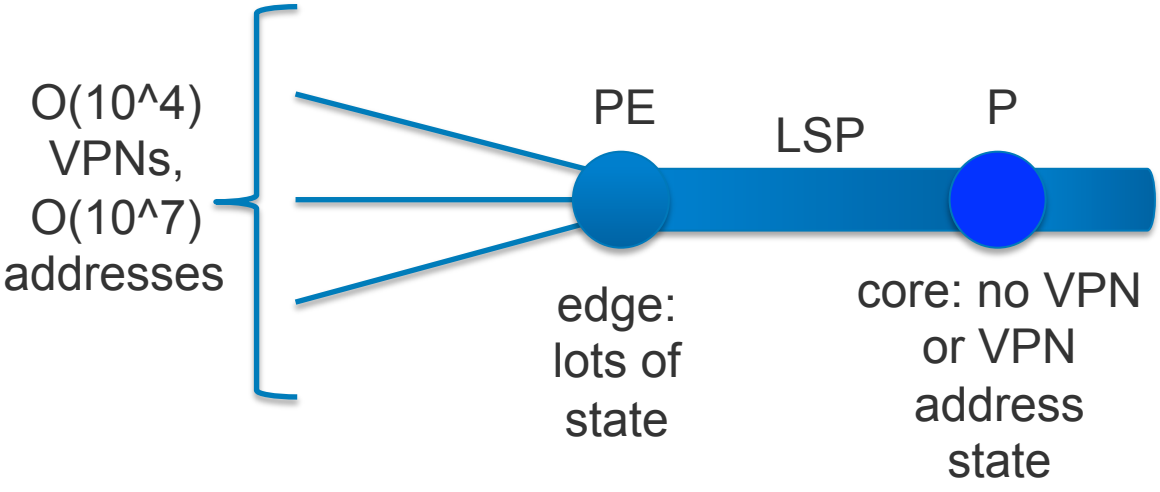

MPLS overlay
(sophisticated)

Copyright © 2013 Juniper Networks, Inc. www.juniper.net

OVERLAY/UNDERLAY CONTROL PLANES

MPLS has a very sophisticated, robust, scalable and interoperable control plane

- Various types of hierarchy are supported
- {BGP, Targeted LDP} [overlay]

  over {LDP, RSVP-TE, LDP/RSVP-TE} [underlay]

None of the new overlays encapsulations have well-specified, interoperable control planes for either the overlay or the underlay

- BGP for the overlay (EVPN/IPVPN over VXLAN) has just been proposed
- But as yet, there isn't a proposal for the underlay control plane

Copyright © 2013 Juniper Networks, Inc. www.juniper.net

CAN THE MPLS CONTROL PLANE BE *TOO* SOPHISTICATED?

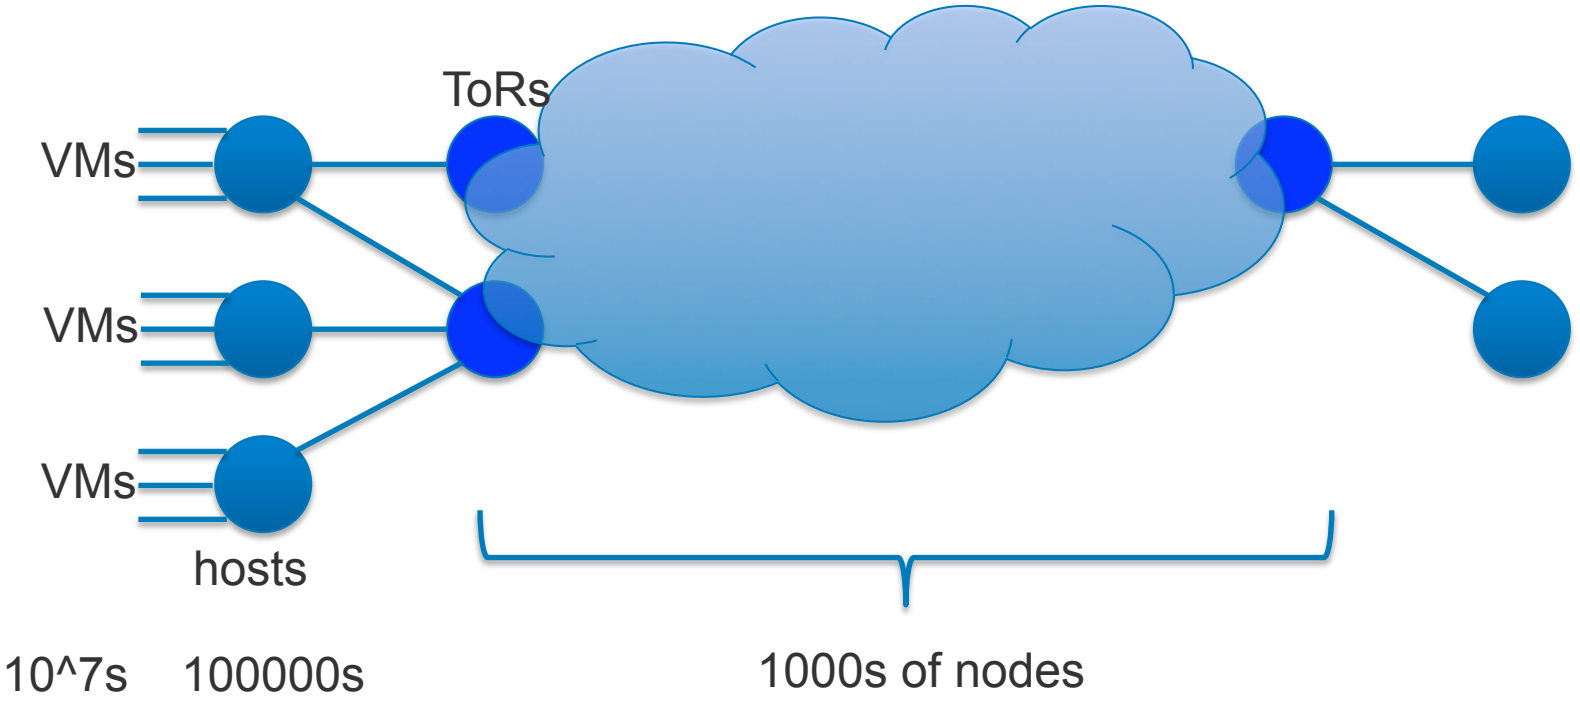

Can't have a flat IGP with so many hosts
LDP DoD with static routing is a possibility, but not ideal
Absolutely has to be plug-and-play
-- new hosts are added at a high rate

Copyright © 2013 Juniper Networks, Inc. www.juniper.net

PROXY ARP RECAP

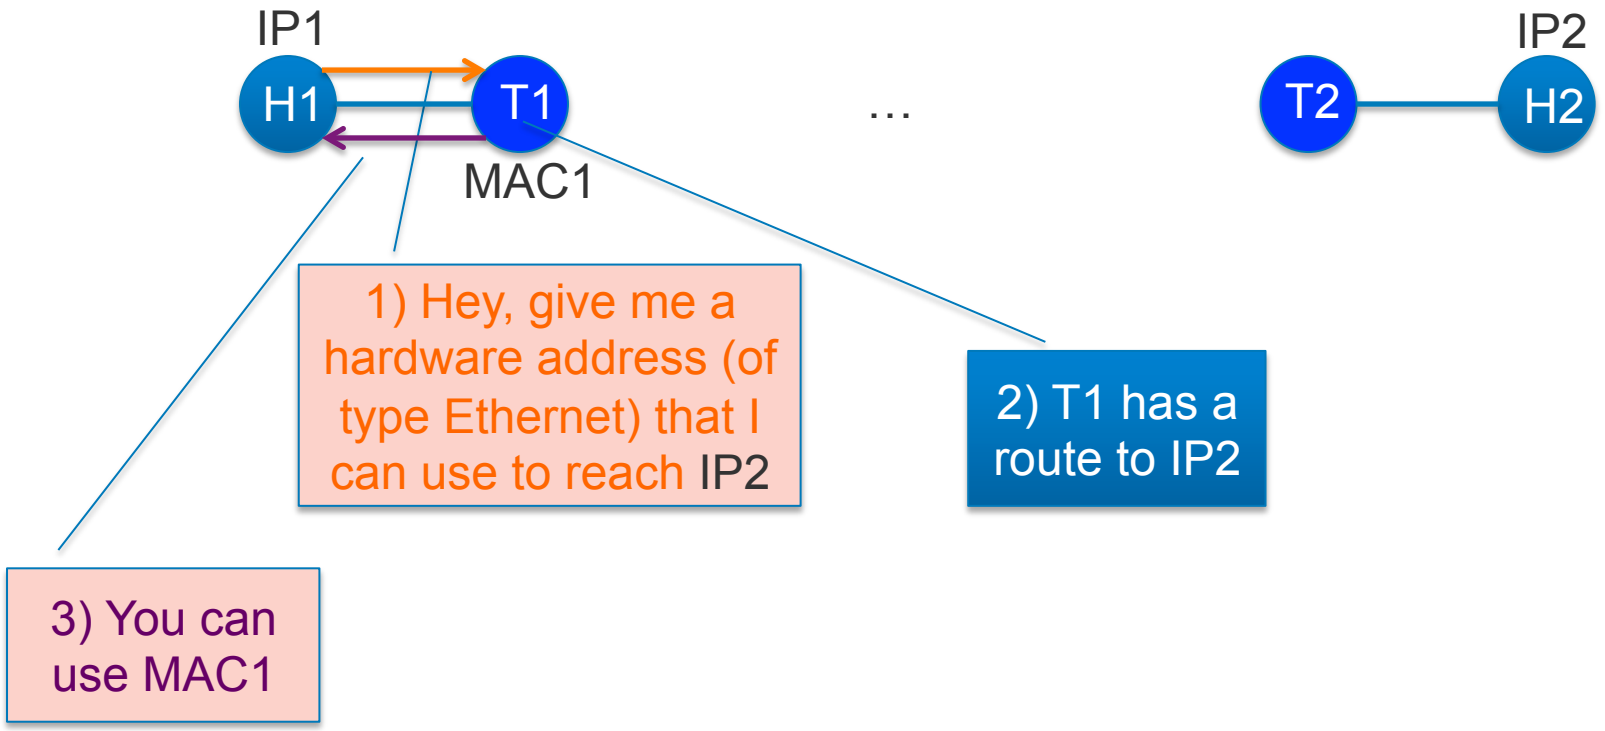

Copyright © 2013 Juniper Networks, Inc.

LABELED ARP: SET UP LDP AMONG ToRs

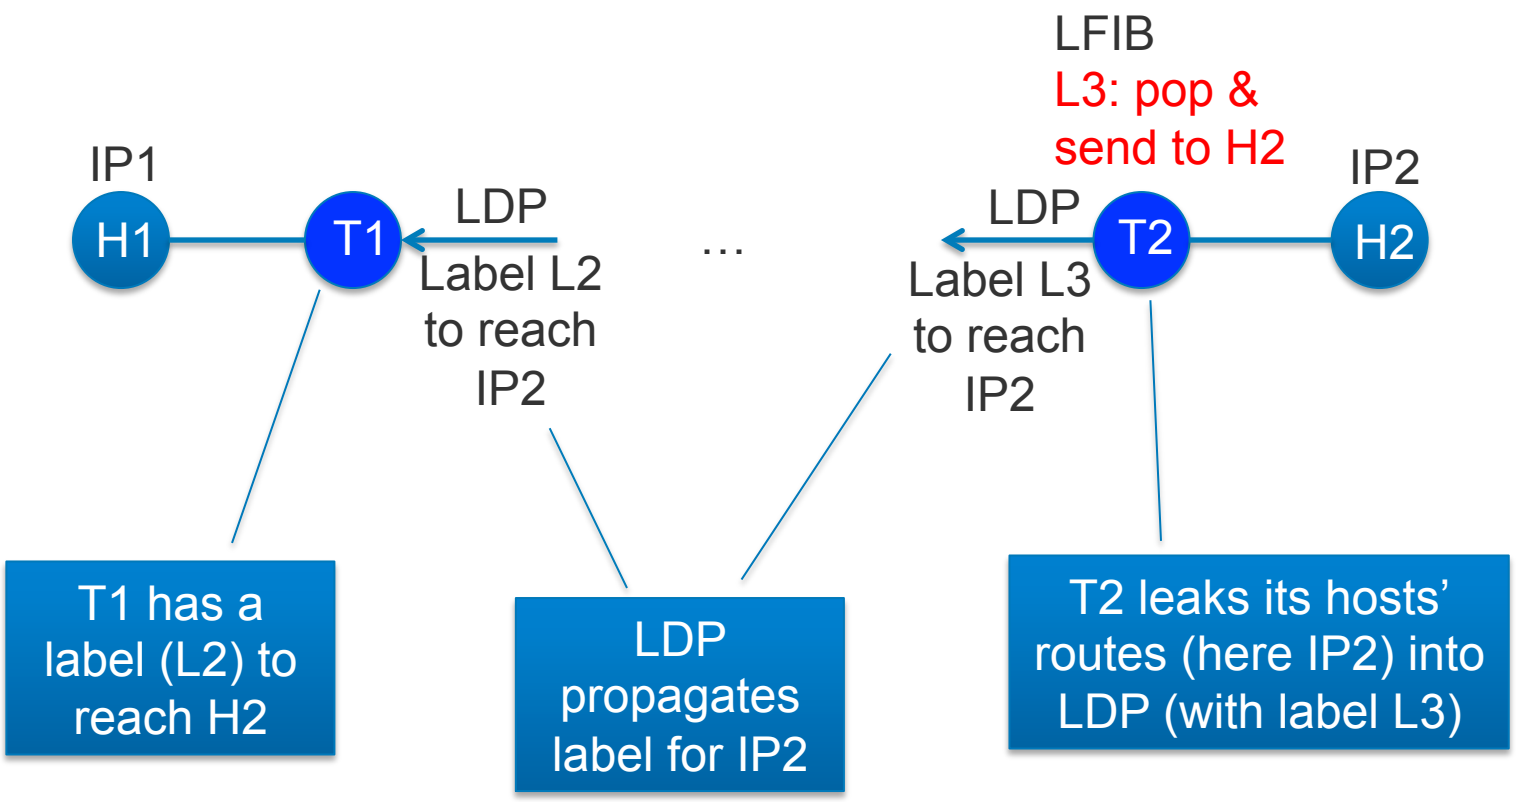

Copyright © 2013 Juniper Networks, Inc. www.juniper.net

LABELED ARP: NOW, H1 CAN REQUEST A LABEL

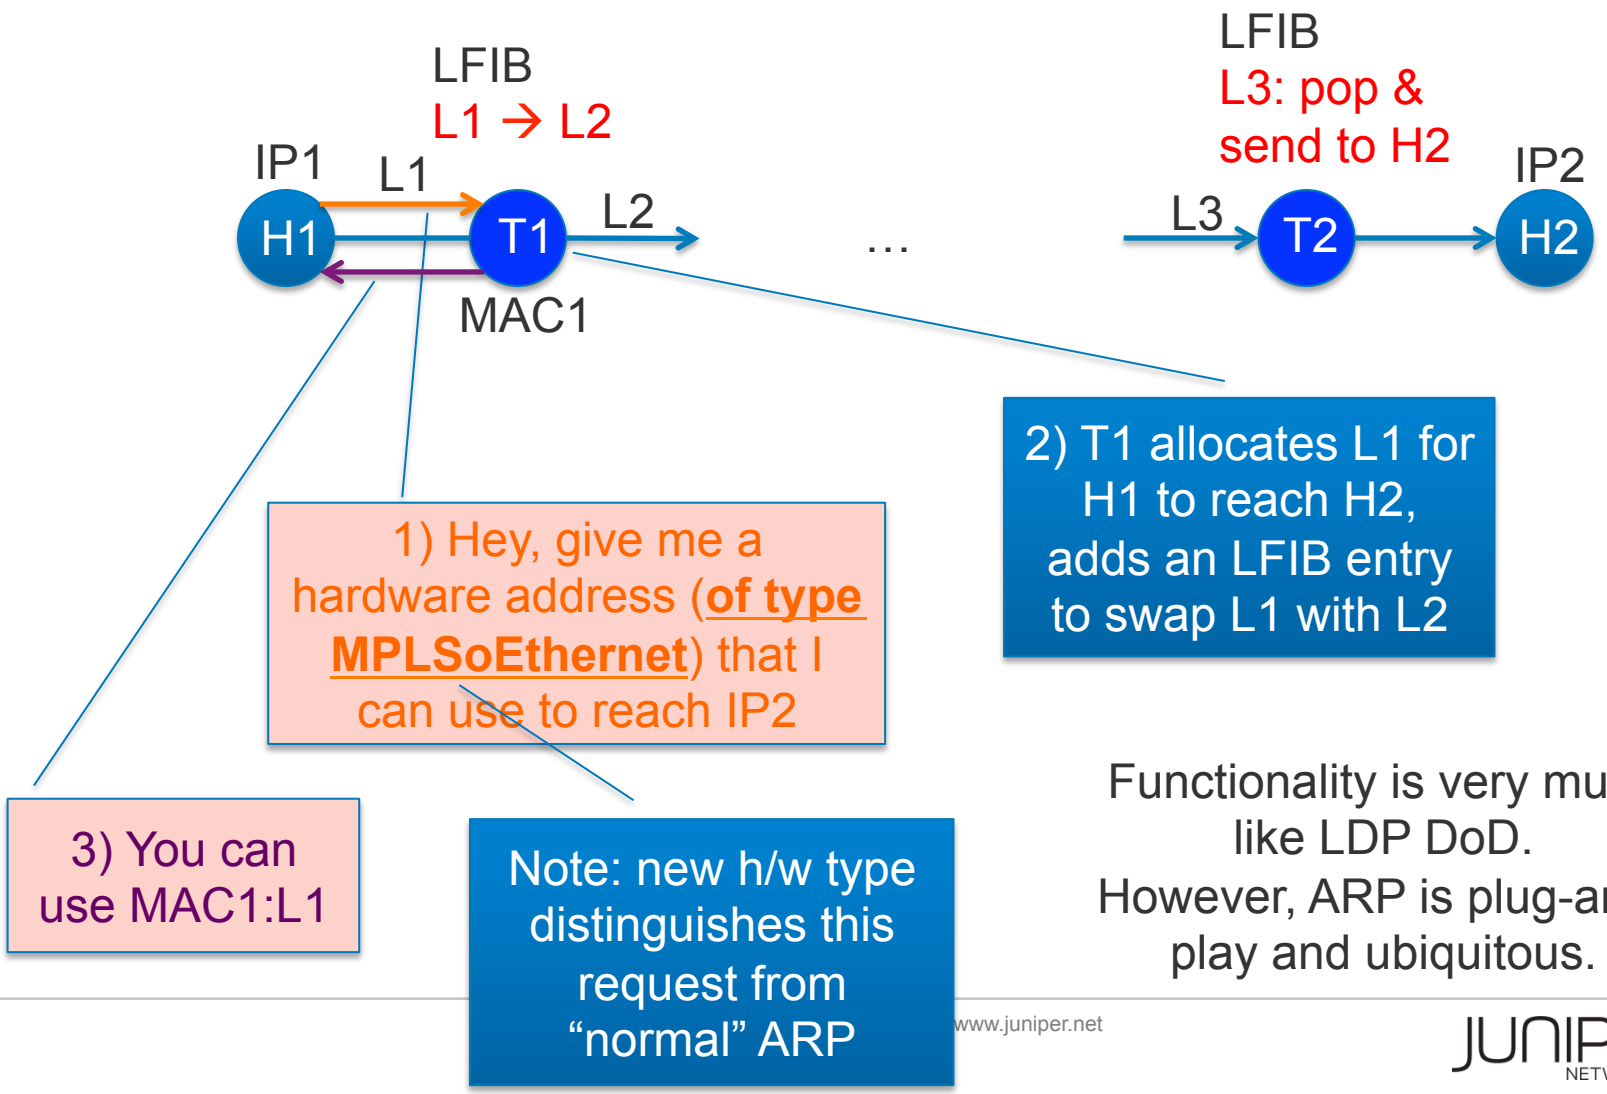

Functionality is very much like LDP DoD.
However, ARP is plug-and-play and ubiquitous.

USE CASE 1:
EGRESS PEERING TRAFFIC ENGINEERING

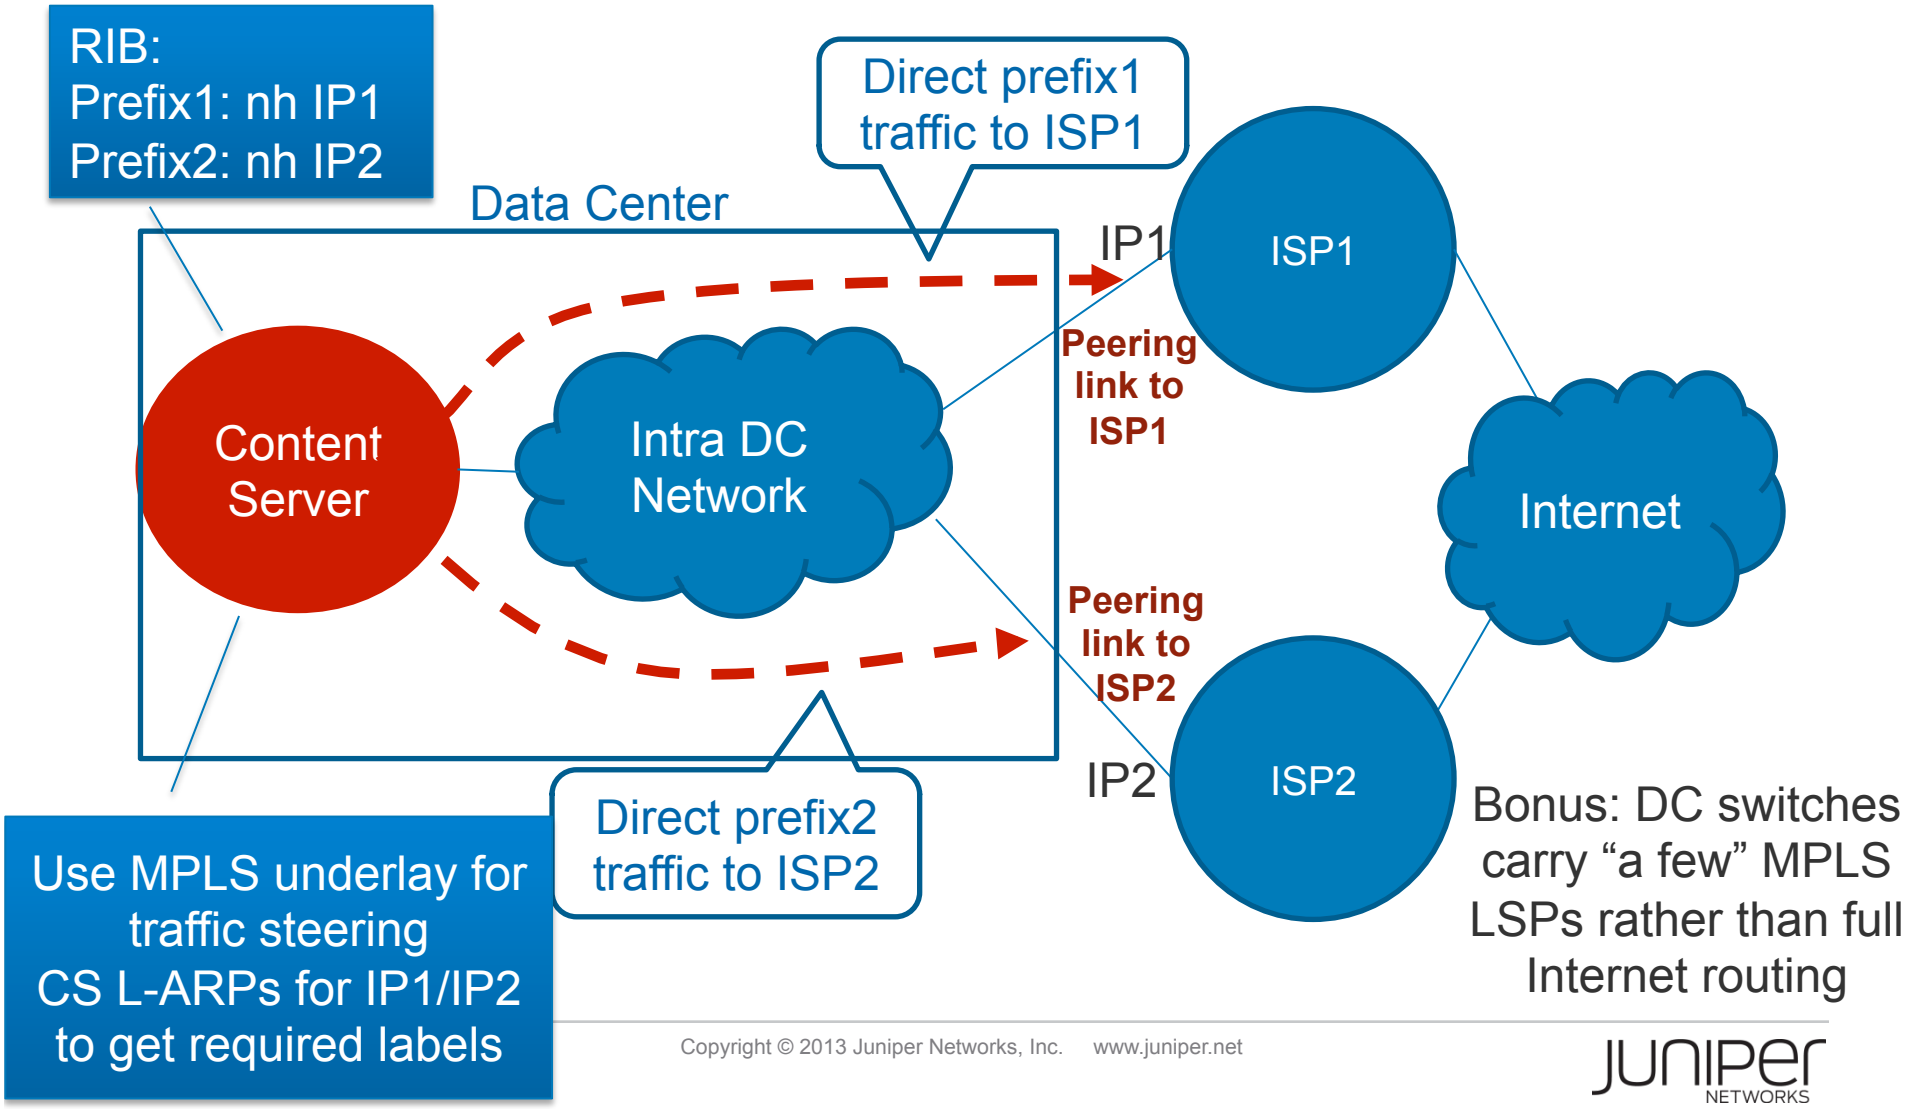

Copyright © 2013 Juniper Networks, Inc. www.juniper.net

USE CASE 2: MPLS UNDERLAY FOR DCs (WITH VRFs/E-VPNs FOR OVERLAY)

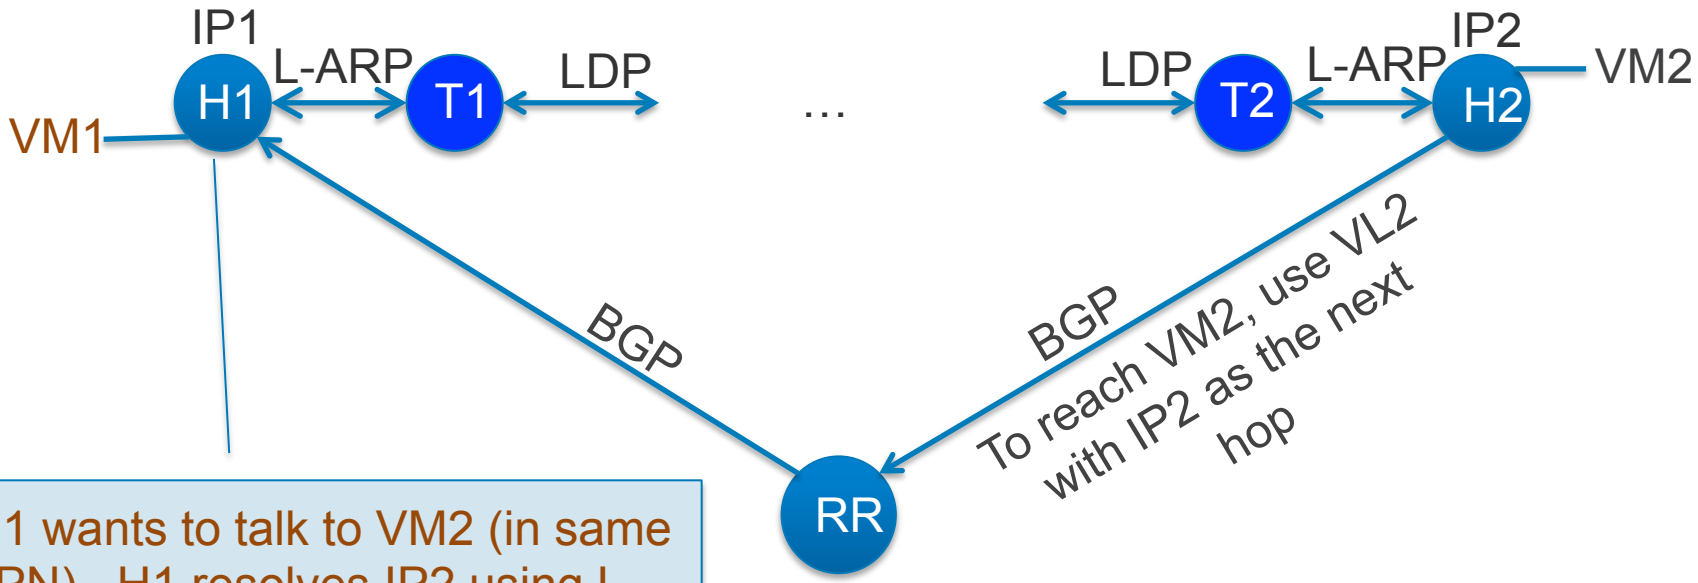

VM1 wants to talk to VM2 (in same VPN). H1 resolves IP2 using L-ARP. Then, packets from VM1 to VM2 are encapsulated with outer label L1 and inner label VL2

Copyright © 2013 Juniper Networks, Inc.

CONCLUSION

MPLS has been somewhat overlooked in consideration for data center use as it is deemed “too complex”

- in the DC (especially on hosts), protocols have to be plug-and-play

This proposal reuses a plug-and-play protocol, namely ARP, to allow the use of MPLS in the DC

- to stitch hosts into existing LSPs across the switching network

There are a few problems to resolve

- Main one: how to deal with label changes

We have prototype code (for Linux hosts)

- Can do both Labeled ARP and Ethernet ARP

Copyright © 2013 Juniper Networks, Inc. www.juniper.net