Scalability Analysis of L2 and L3 VPN Technologies

Service Provider Perspective

Alex ZININ
CTO IPD APAC, Alcatel-Lucent

What is Scalability Analysis?

Different approaches have different complexity (overhead)

- Computational (CPU utilization)
- Data (amount of state)
- Signaling (required BW)
- Management & provisioning (OPEX)

Goal of scalability analysis:

- Understand growth of overhead as complexity input parameters increase

Method--asymptotic complexity analysis

- Asymptotic behavior of *f(n)* == growth of *f(n)* as *n* gets larger
- Typically ignore small values of *n*
- Slower asymptotic growth – better
- *O*-notation for "order" of growth, e.g.:

$$f(n) = n^2 + 25*n - 10 = O(n^2)$$

$$f(n) = 100*n^2 = O(n^2)$$

Example: CPU overhead functions

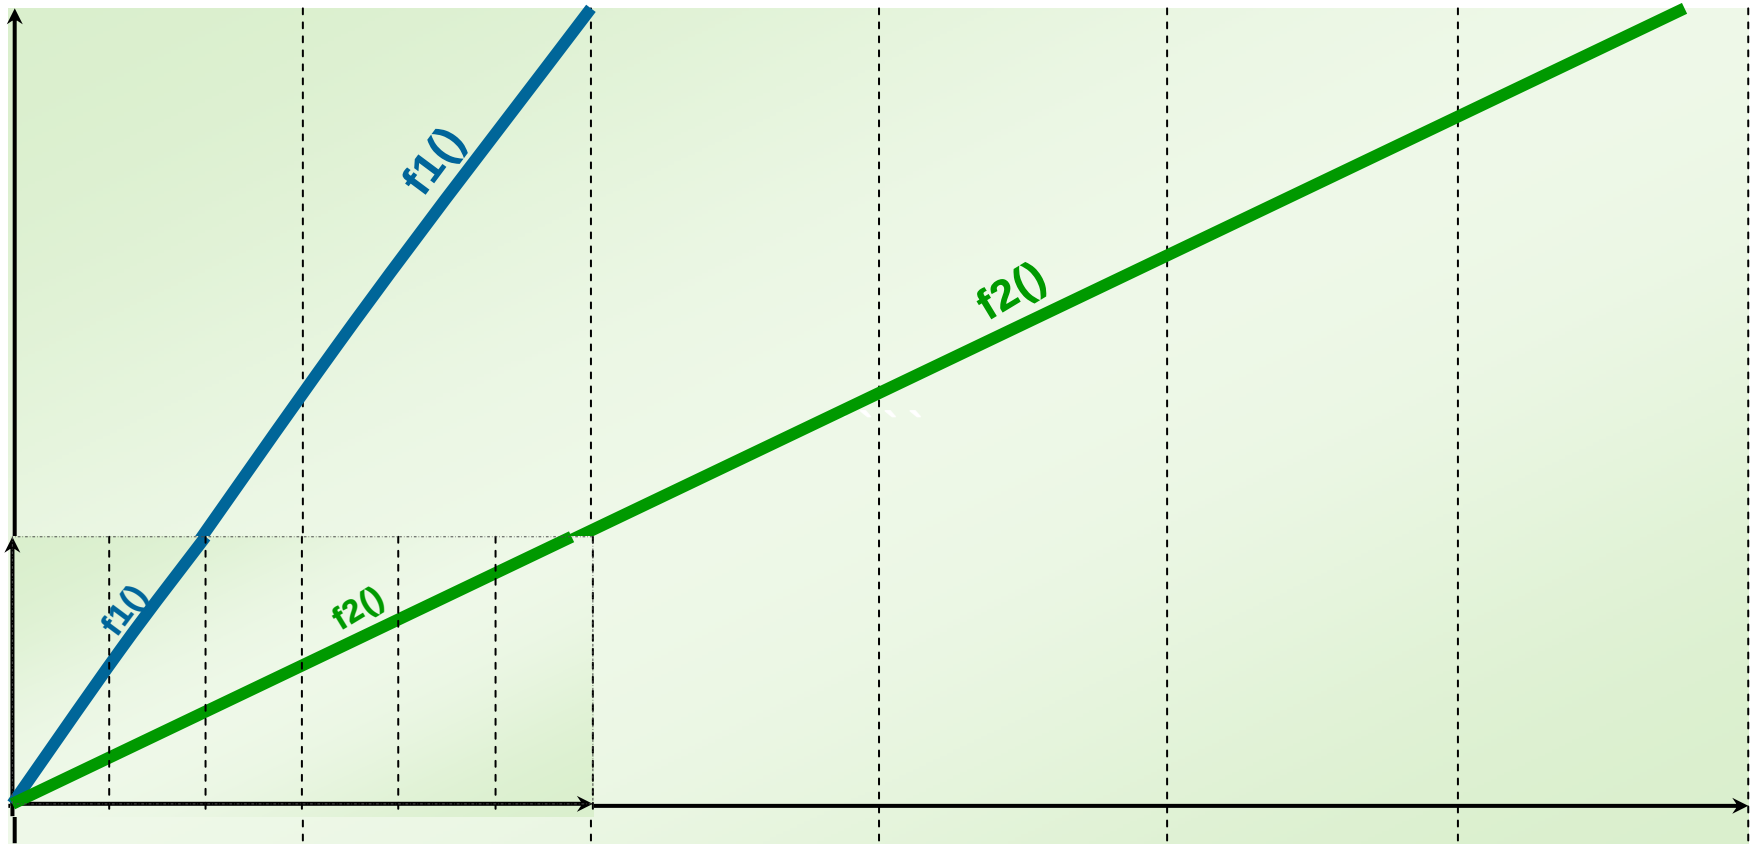

Compare overhead functions for two algorithms f1() and f2()

Both appear linear f2() better than f1()

Example: CPU overhead functions

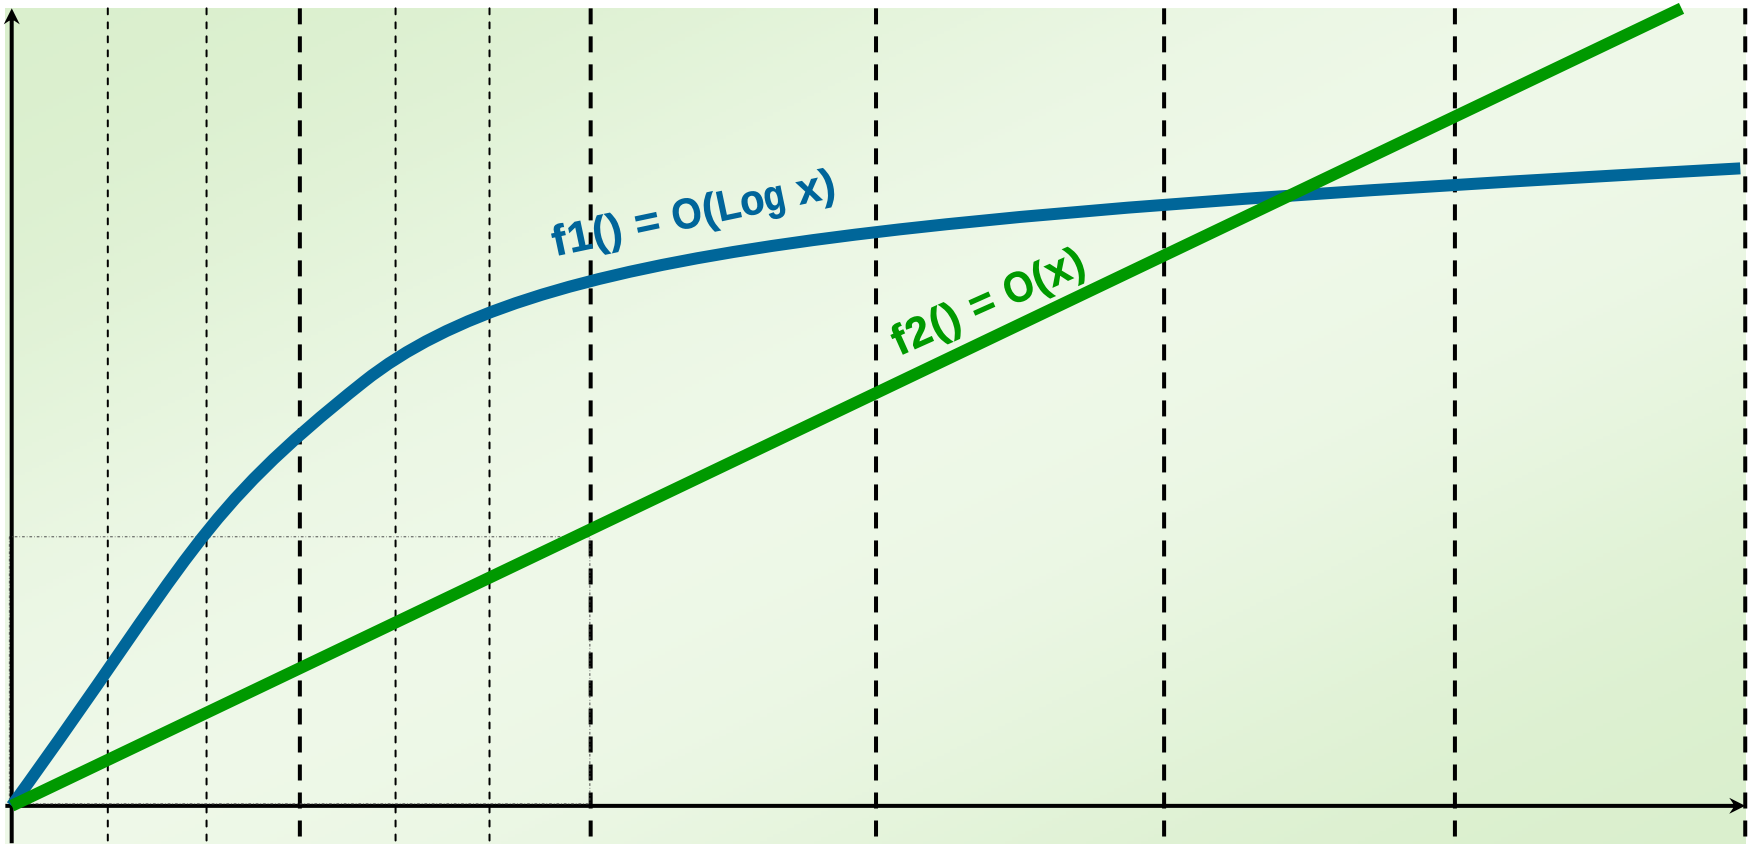

On a different scale
f1() better than f2()

Example: CPU overhead functions

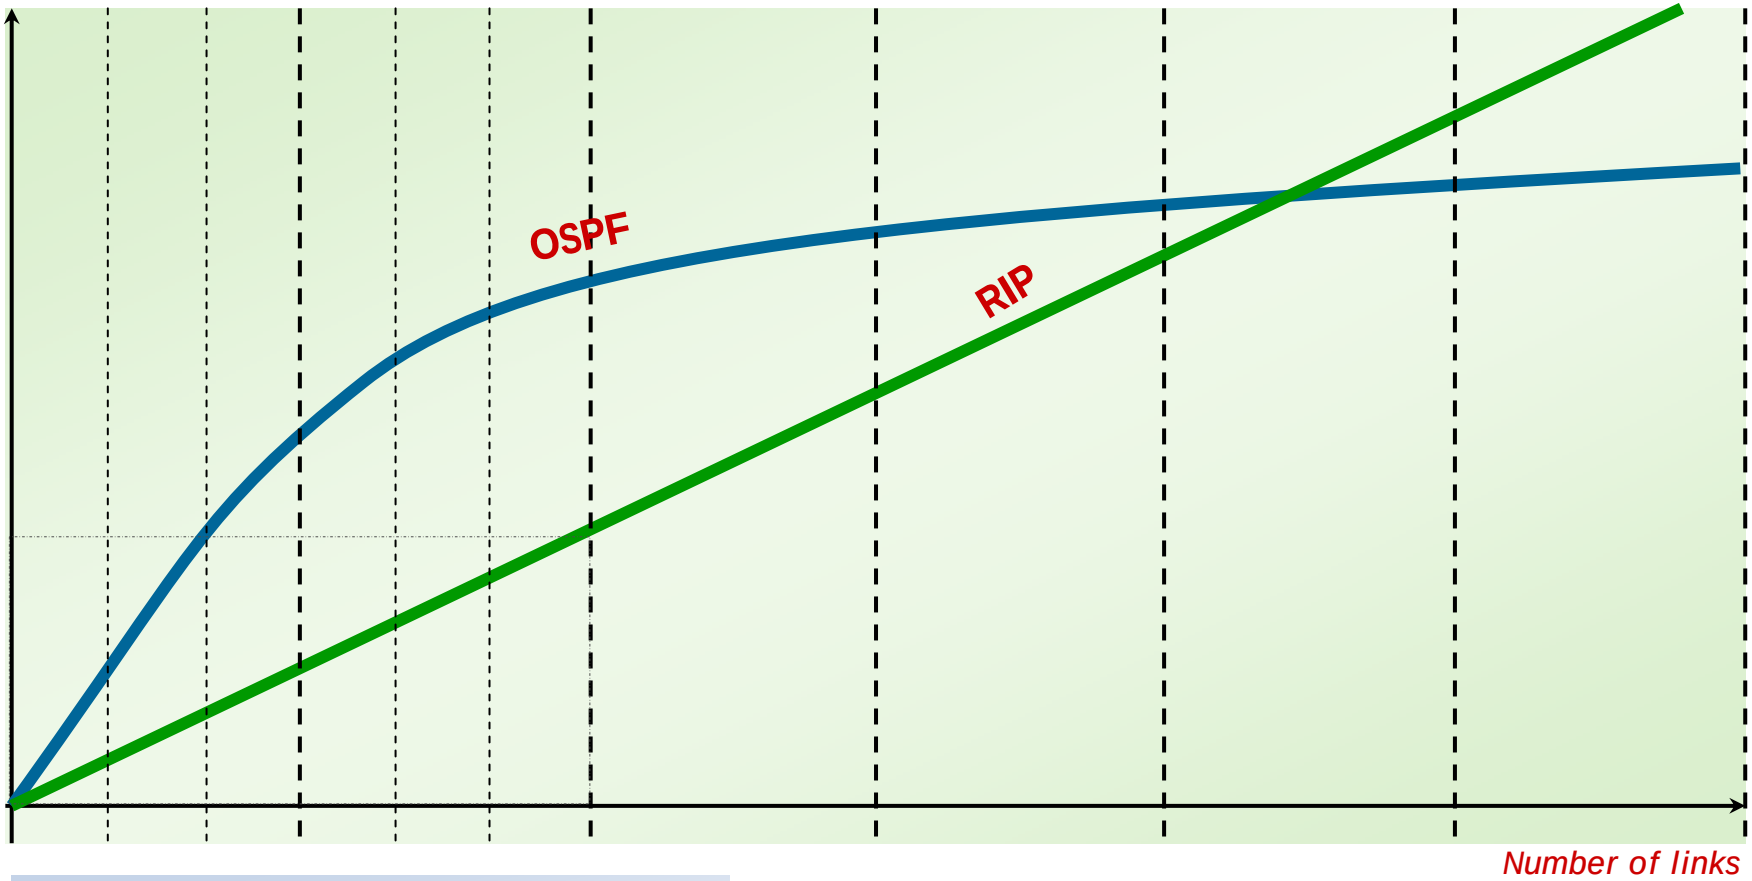

Choice of better solution depends on problem scope

BUT

Generally inferior solution can be better for small problems

Network Protocol Scaling Analysis

Control Plane

- Number of protocol sessions (adjacencies, TCP connections, etc.)
- Amount of state (LSAs, routes, etc.)
- Computational overhead (CPU load)
- Signaling overhead (required BW)

Data Plane

- Amount of state (routes, MAC addresses, etc.)
- Computational overhead (forwarding, flooding, replication, etc.)

Management Plane

- Provisioning overhead
- Troubleshooting complexity

Reason for Scaling Analysis

Understand technology aspects

- Maximum number of nodes, users, services
- Potential bottlenecks
- Sensitive implementation areas

Understand cost functions

- Cost of initial deployment
- Cost growth curve as function of number of subs/svcs

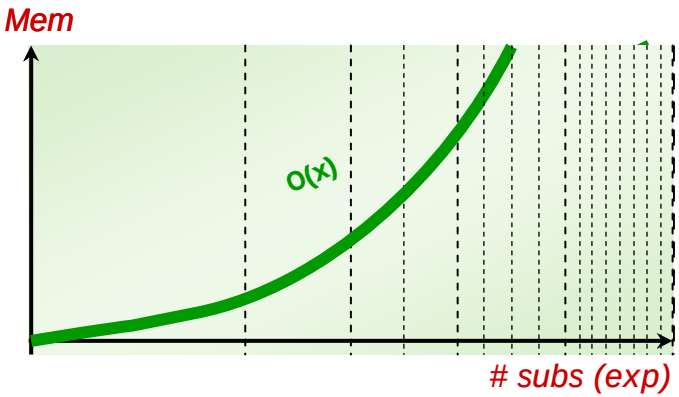

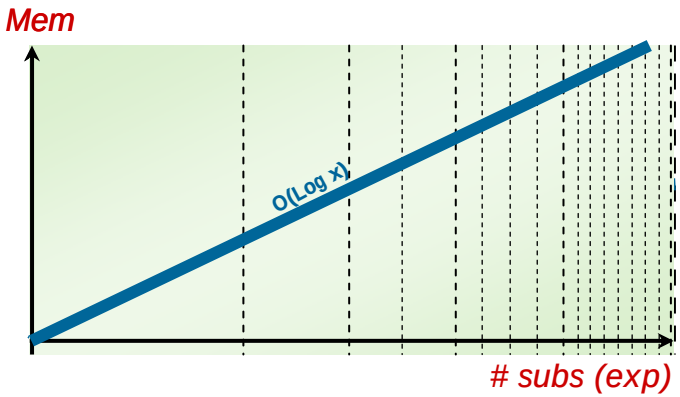

VPN Scaling Analysis—General Aspects

What is the scaling factor?

- Number of CPE devices?
- Number of end users?

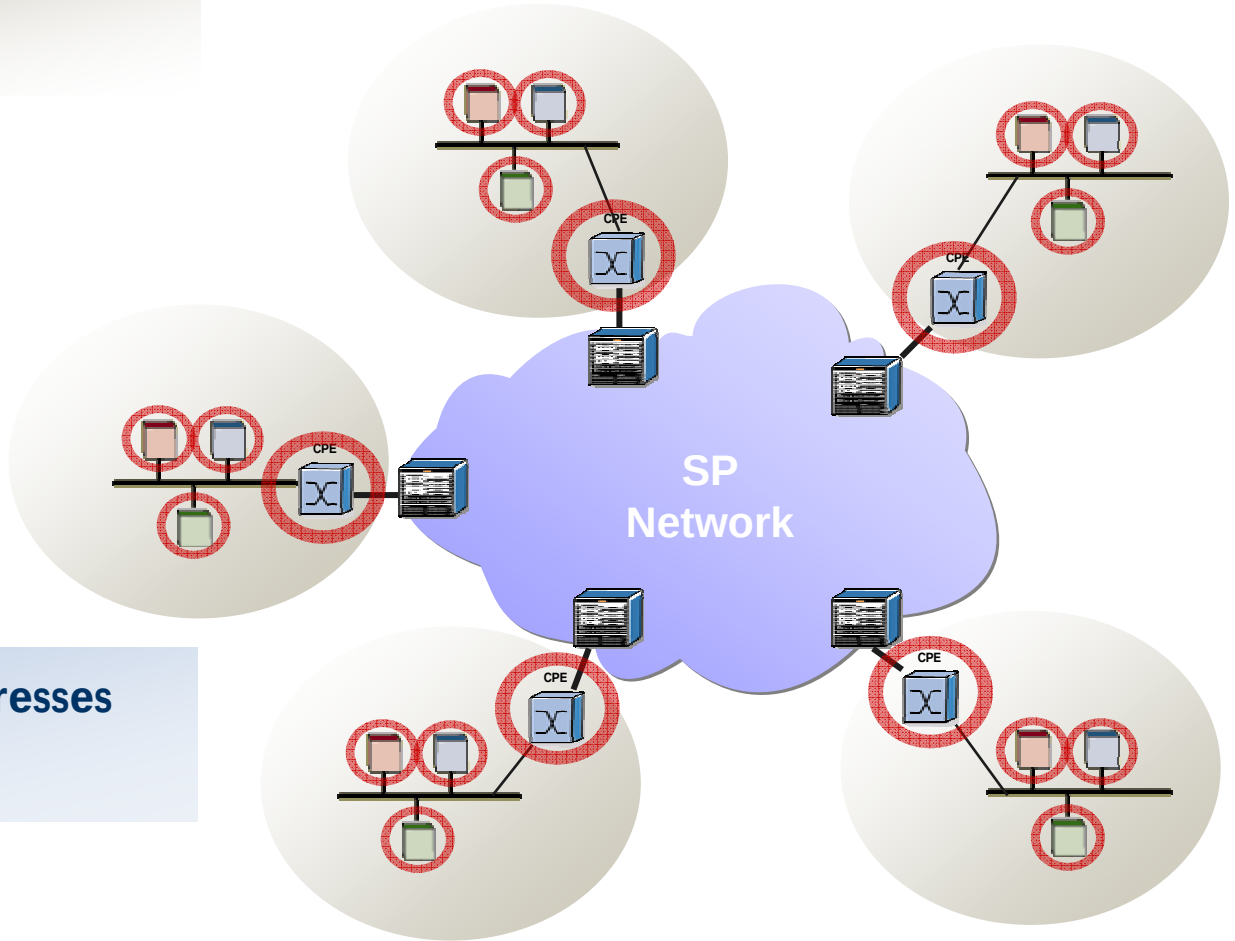

L2VPNs: sites + MAC addresses

L3VPNs: sites + IP routes

VPN Scaling Parameters

L2 VPNs: Control Plane

- Sessions: T-LDP
- State: PWs, MAC addresses
- CPU: LDP, MAC table
- Signaling: PW setup, status, MAC withdraw

L3 VPNs: Control Plane

- Sessions: iBGP
- State: VPN Routes
- CPU: BGP, PE-CE
- Signaling: route updates, withdraw

L2 VPNs: Data Plane

- State: MAC table
- Computation: MAC lookup, flooding, mcast replication

L3VPNs: Data Place

- State: VPN Routes
- CPU: FIB lookup, mcast replication

Basic Configuration

L2 VPNs: flat VPLS

- Full mesh T-LDP
- No multicast optimizations

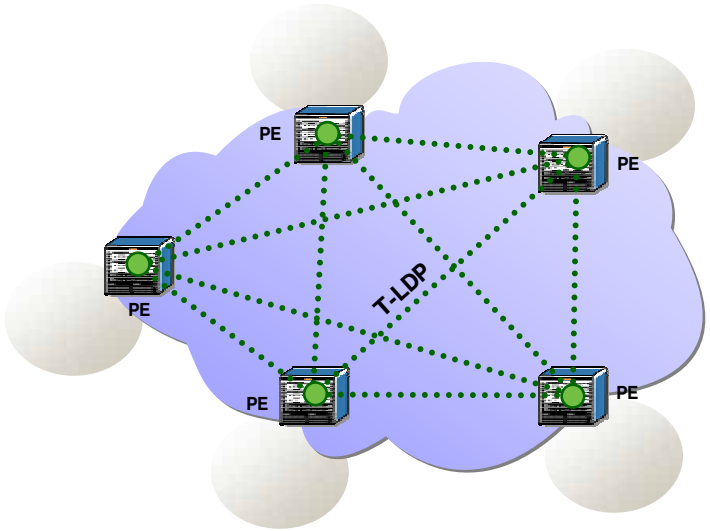

S: number of sites
N: number of PE nodes ~O(Log S)
M: number of MAC addresses
V: number of VPN instances

L3 VPNs: basic MPLS/BGP

- Full mesh iBGP
- No multicast optimizations

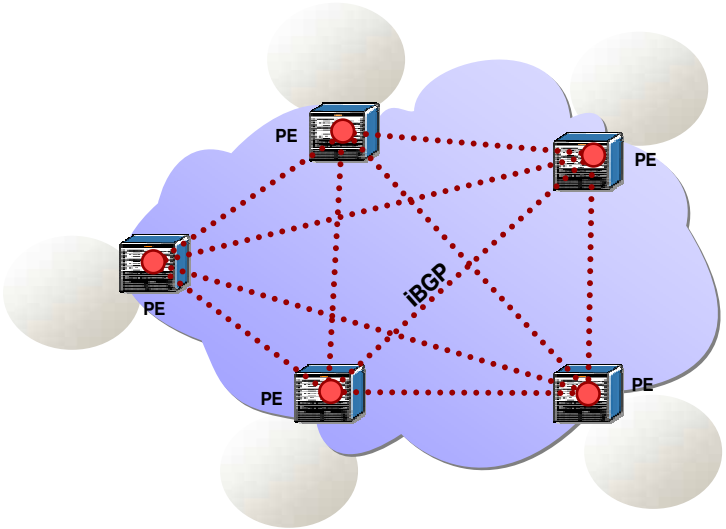

S: number of sites
N: number of PE nodes ~O(Log S)
R: number of VPN Routes
V: number of VPN instances

VPN Scaling: Basic config

L2VPN

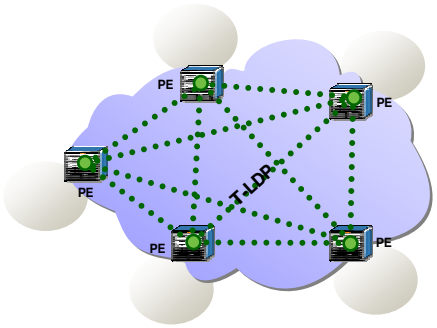

L3VPN

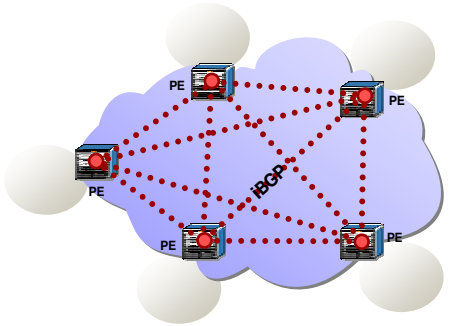

CP: sessions	**O(N) per PE** **O(N^2) total**	**O(N) per PE** **O(N^2) total**
CP: state	**O(M*V)**	**O(R*V)**
CP: CPU	**LDP complexity**	**BGP complexity**
CP: Signaling	**O(N) PW signaling**	**O(R*V) route propagation**
DP: state	**O(M*V)**	**O(R*V)**
DP: computational	**O(Log M) lookup** **O(N) per replicated packet**	**O(Log R) lookup**

Basic setup: scaling aspects

Number of PEs and protocols sessions to configure and bring up

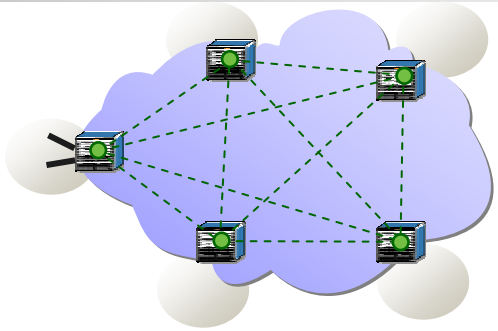

Number of PW/Route labels to signal

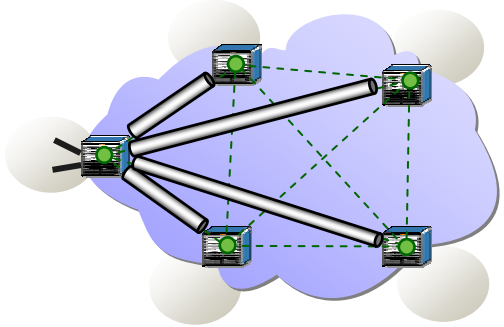

Replication overhead

MAC address table/RIB size

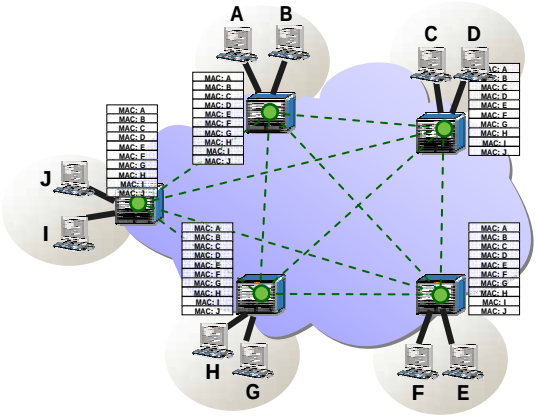

Hierarchical Configuration

L2 VPNs: H-VPLS

- Hub-n-spoke T-LDP
- No multicast optimizations

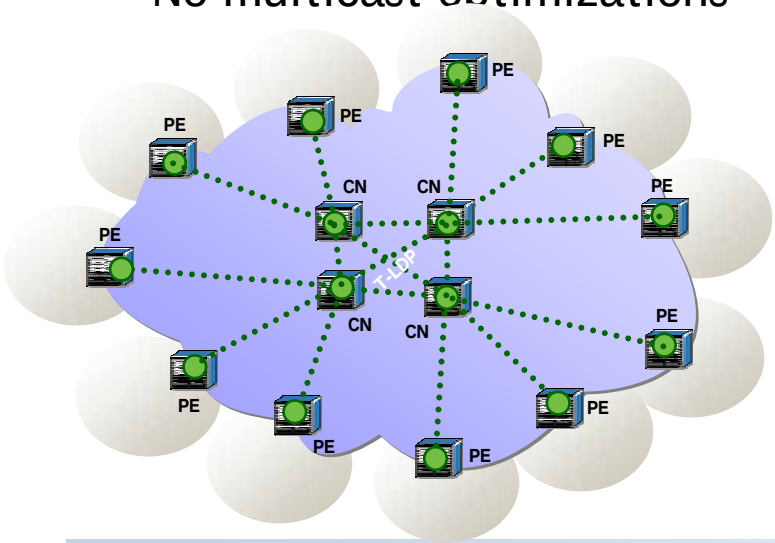

S: number of sites
N: number of PE nodes ~O(Log S)
C: number of H-VPLS core nodes
M: number of MAC addresses
V: total number of VPN instances
v: VPN instances at PE

L3 VPNs: Hierarchical MPLS/BGP

- RRs for iBGP
- No multicast optimizations

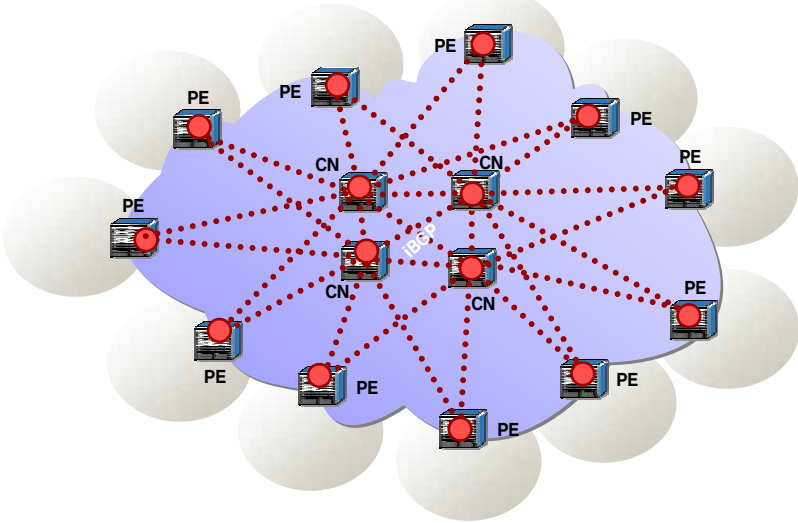

S: number of sites
N: number of PE nodes ~O(Log S)
C: number of RRs
R: number of VPN Routes
V: number of VPN instances
v: VPN instances at PE

VPN Scaling: hierarchy

L2VPN

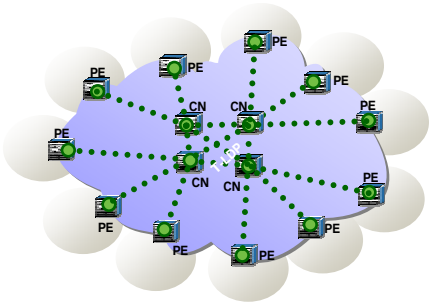

L3VPN

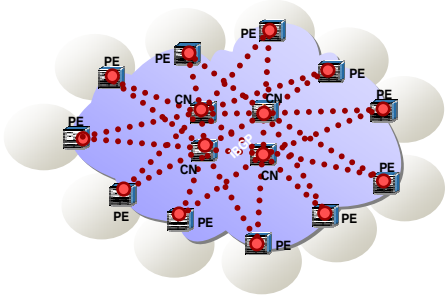

CP: sessions	**PE:** *O(1)* **CN:** *O(C)*	**PE:** *O(1)* **CN:** *O(C)*
CP: state	**PE:** *O(M*v)* **CN:** *O(M*V)*	**PE:** *O(R*v)* **CN:** *O(R*V)*
CP: CPU	**PE:** *O(LDP)* **CN:** *O(LDP * C)*	**PE:** *O(BGP)* **CN:** *O(BGP * C)*
CP: Signaling	**PE:** *O(1 PW)* **CN:** *O(N PW)*	**PE:** *O(R*v) route propagation* **CN:** *O(R*V)*
DP: state	**PE:** *O(M*v)* **CN:** *O(M*V)*	**PE:** *O(R*v)* **CN:** *O(R*V)*
DP: computational	*O(Log M) lookup* **PE:** *O(1) per replicated packet* **CN:** *O(N) per replicated packet*	*O(Log R)* **lookup**

Scaling Multicast: Default Operation

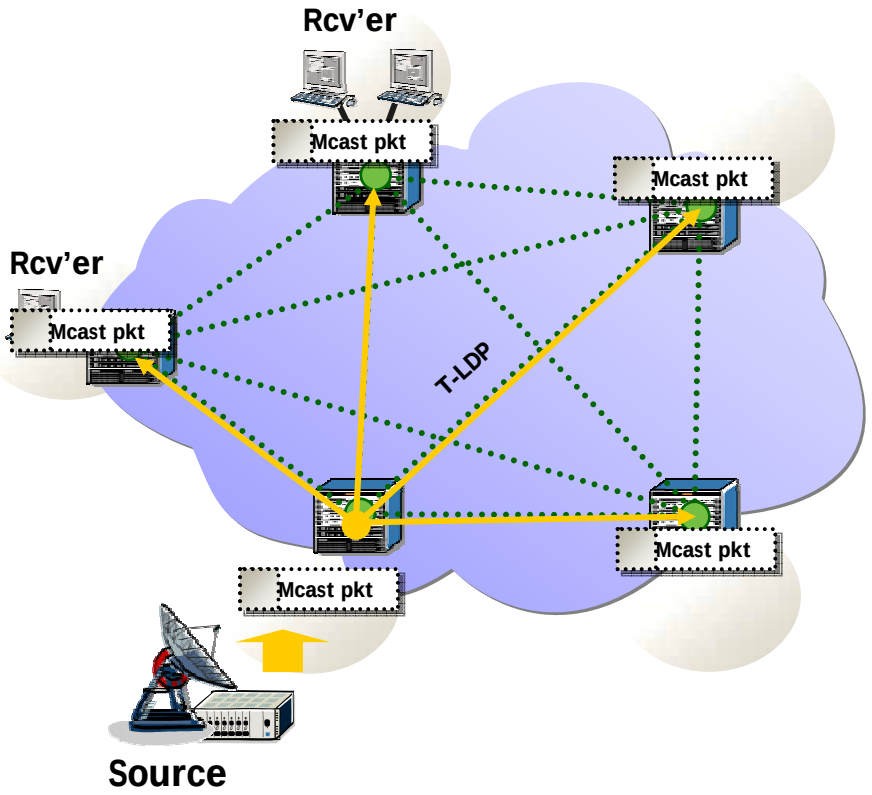

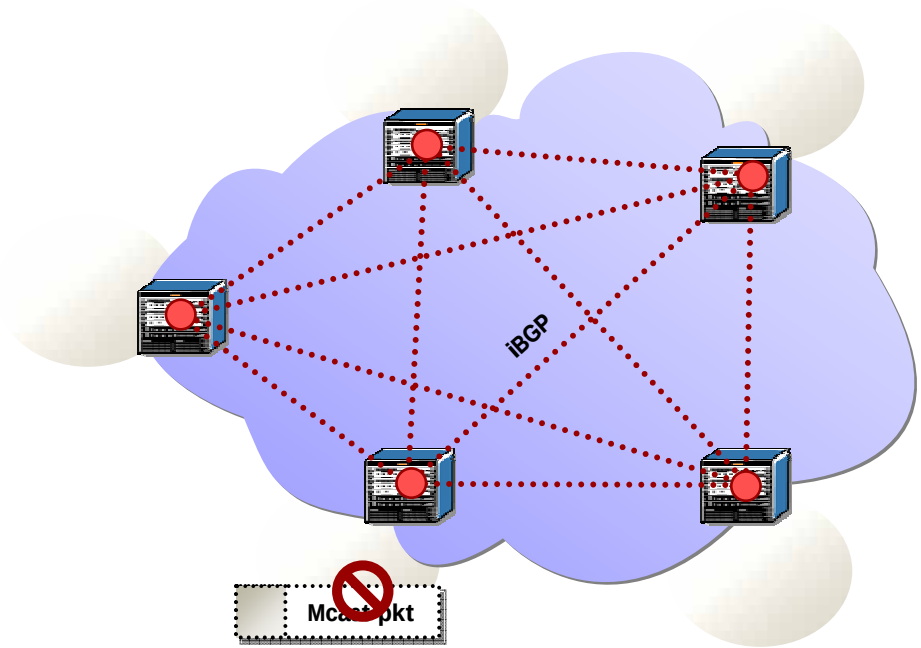

Default multicast replication in VPLS is flooding:

- **O(N) per replicated packet**
- **BW inefficiency**

No multicast in MPLS/BGP VPN basic operation

Efficient Multicast Operation

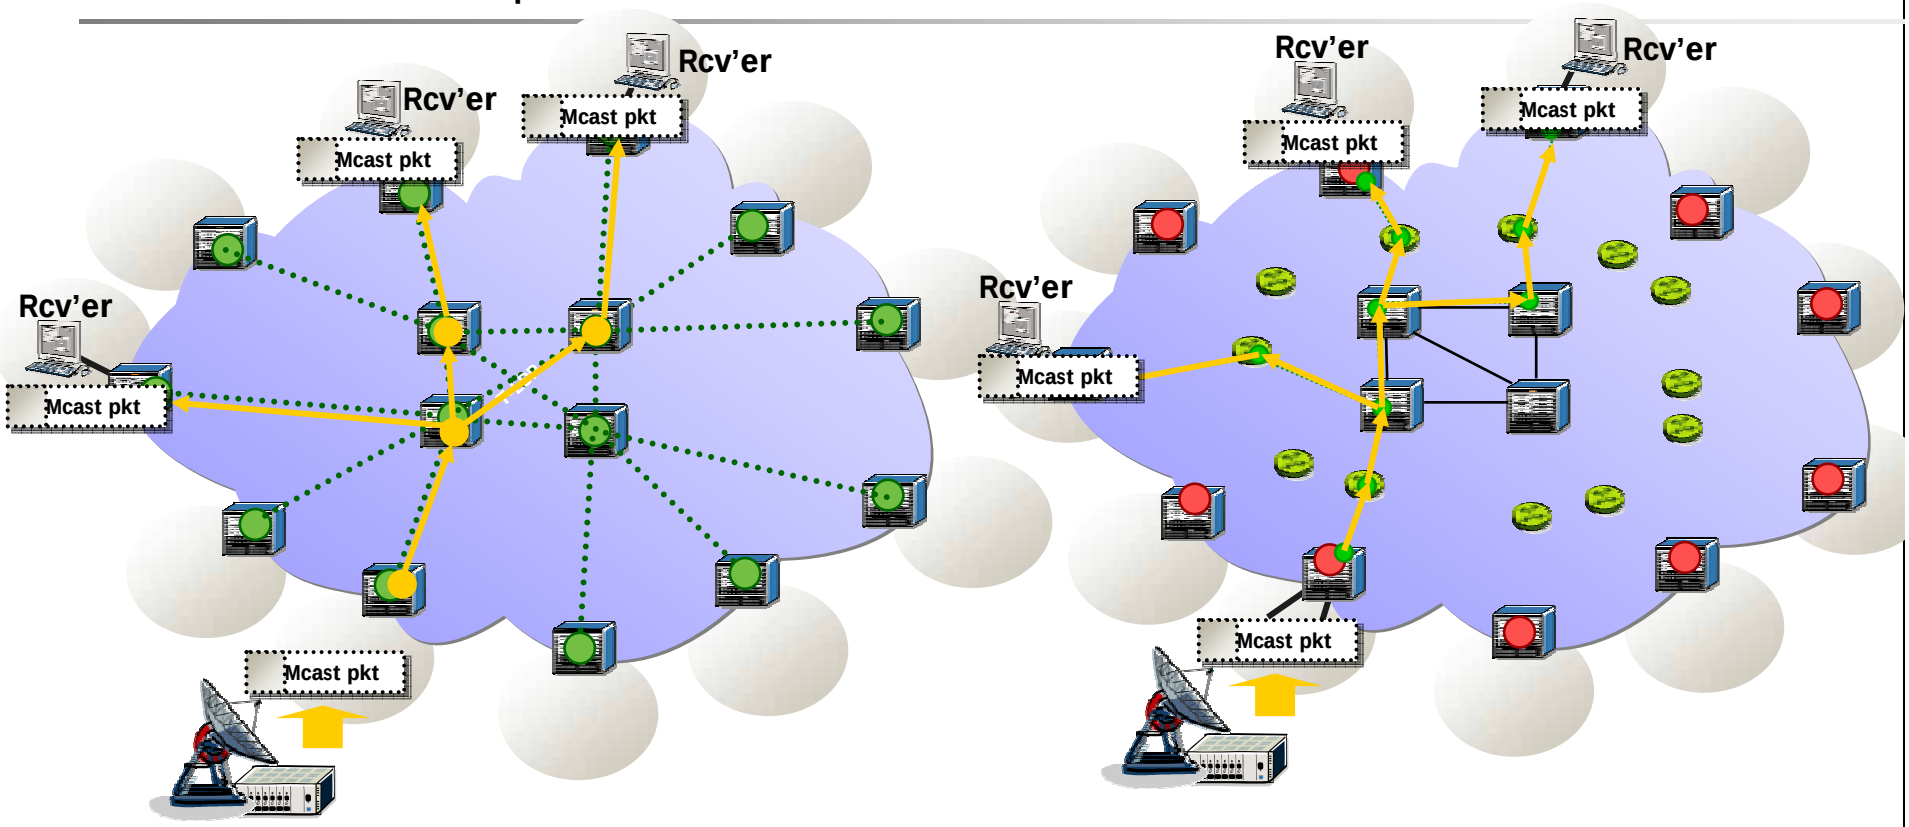

VPLS multicast optimizations

- IGMP snooping
- H-VPLS replication
- BW suboptimalities still possible

MPLS/BGP multicast support:

- Pre-built mcast tree (PIM-SSM)
- Mcast packets encapsulated in GRE

Scaling VPN services: controlling amount of network state

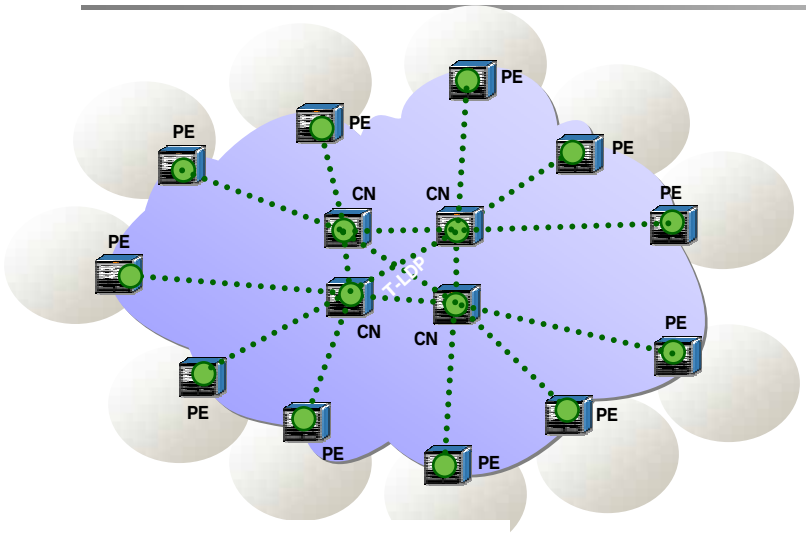

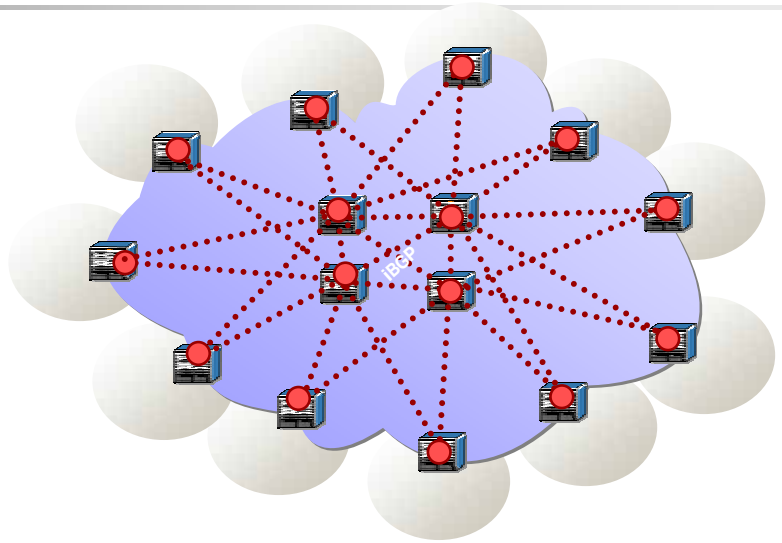

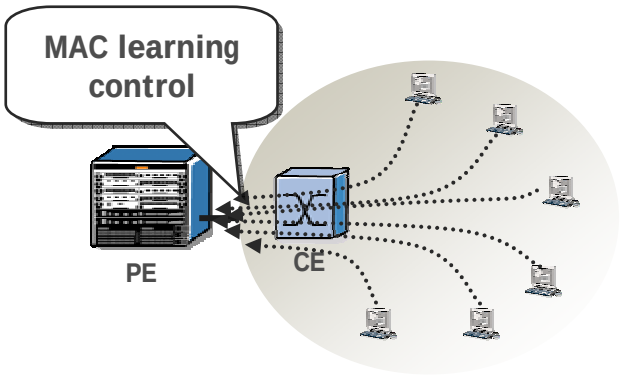

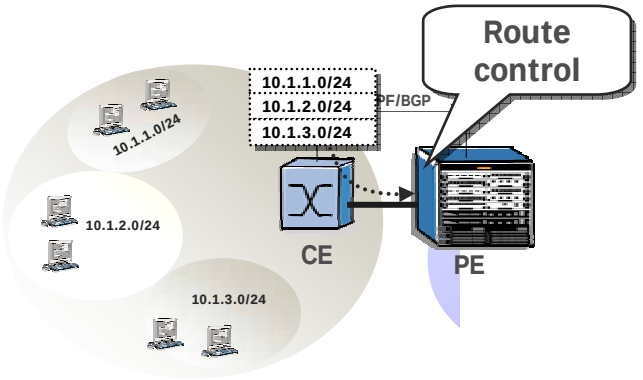

Scaling VPLS services further: MAC address challenge

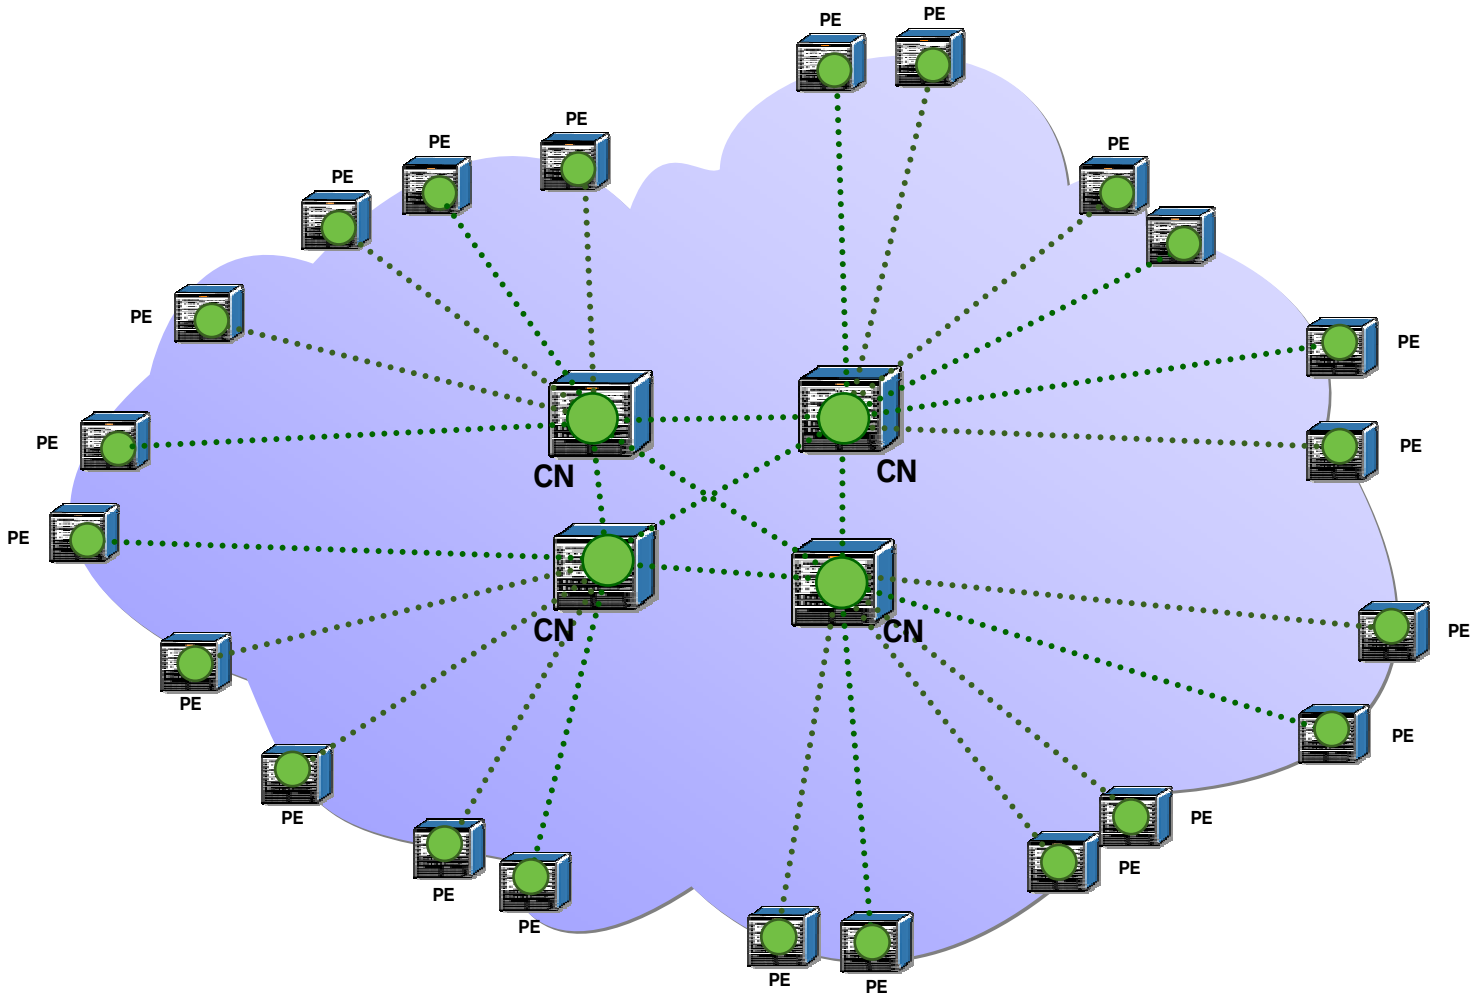

CN MAC Table scaling: *O(M*V)*

Scaling VPLS services further: MAC-in-MAC for VPLS

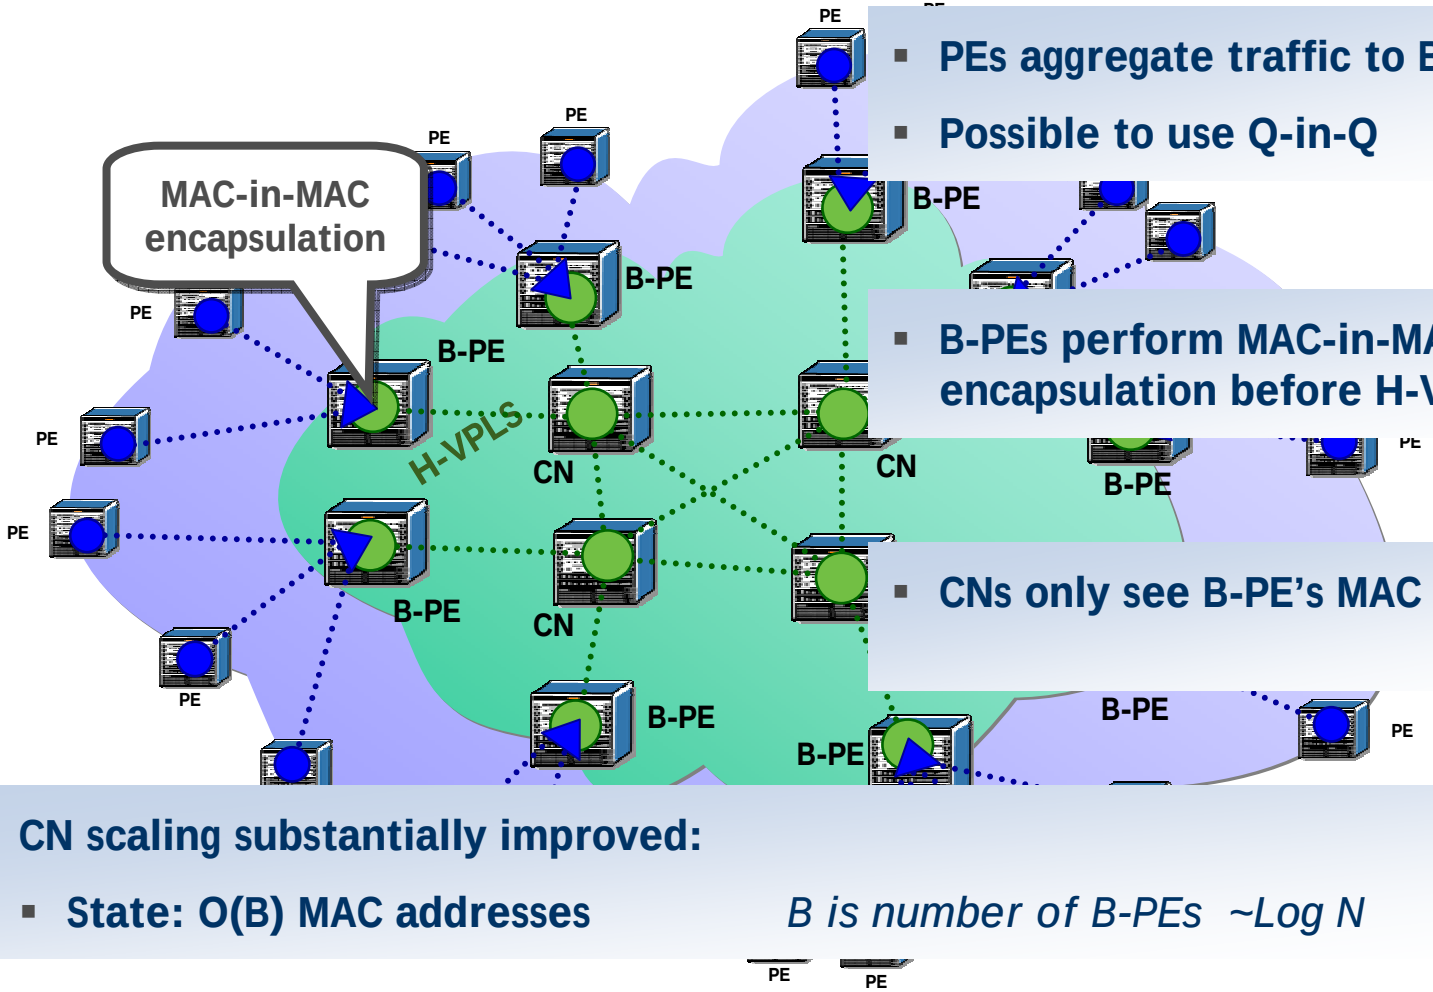

- PEs aggregate traffic to B-PEs
- Possible to use Q-in-Q
- B-PEs perform MAC-in-MAC encapsulation before H-VPLS
- CNs only see B-PE's MAC addresses

CN scaling substantially improved:

- **State: O(B) MAC addresses** *B is number of B-PEs ~Log N*

Management & Operations: scaling aspects

L2 VPN

- T-LDP: needs provisioning
- MACs: need to be controlled
- OAM: VPN-aware tools needed

L3 VPN

- MP-BGP: needs provisioning
- VPN Routes: need to be controlled
- OAM: VPN-aware tools needed

Scalable provisioning platform is a must

Implementations need fine control over MAC table population

OAM tools to look for:

- **VPLS MAC ping**
- **VPLS MAC traceroute**
- **VPLS MAC populate/purge**

Implementations need control over learned and propagated VPN routes

OAM tools to look for:

- **VRF ping**
- **VRF traceroute**

Summary

VPLS and MPLS/BGP VPN technologies:

- Different scaling aspects (MAC addresses vs IP routes)
- Different scaling improvement methods
- Yet similar operational impact and approach

www.alcatel-lucent.com