Introduction to MPLS

Basic Terminology, Functions, and Applications

Colorado Springs Cisco User's Group

9 January, 2008

Jeff Doyle
Principal
Jeff Doyle and Associates, Inc.

A Little Background

- What do we mean by multiprotocol?
  - Generic encapsulation
- What do we mean by label switching?
  - Fixed address switching similar to ATM or FR
- Original motivations for MPLS
  - ATM-like switching speeds
- Modern motivations for MPLS
  - Virtual circuits over IP networks
  - Service separation and virtualization
  - Multiservice networks
    - Network consolidation

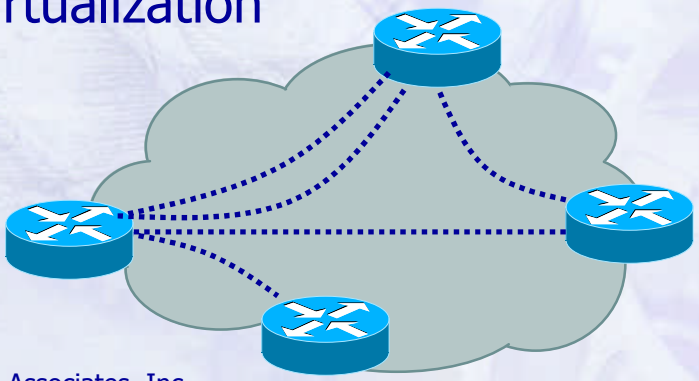

MPLS Applications

- Link and node failure protection
- Traffic engineering
- Virtual point-to-point connections
- Virtual point-to-multipoint connections
- ATM/FR interworking
- Virtual Private Networks (VPNs)
  - Layer 3 VPNs
  - Layer 2 VPNs
  - Virtual Private LAN Service (VPLS)

Basic Concepts and Terms

- Label
  - A fixed-length (20-bit) address
  - Local significance (link scope)
- Label Switched Paths (LSPs)
  - An MPLS virtual circuit
  - LSPs are unidirectional
- Label Switching Routers (LSRs)
  - Any router supporting MPLS
- Forwarding Equivalence Classes (FECs)
  - All packets:
    - To be forwarded out the same interface
    - With the same forwarding treatment (CoS)
    - To the same next hop
  - This is not a new concept!

Basic Concepts and Terms

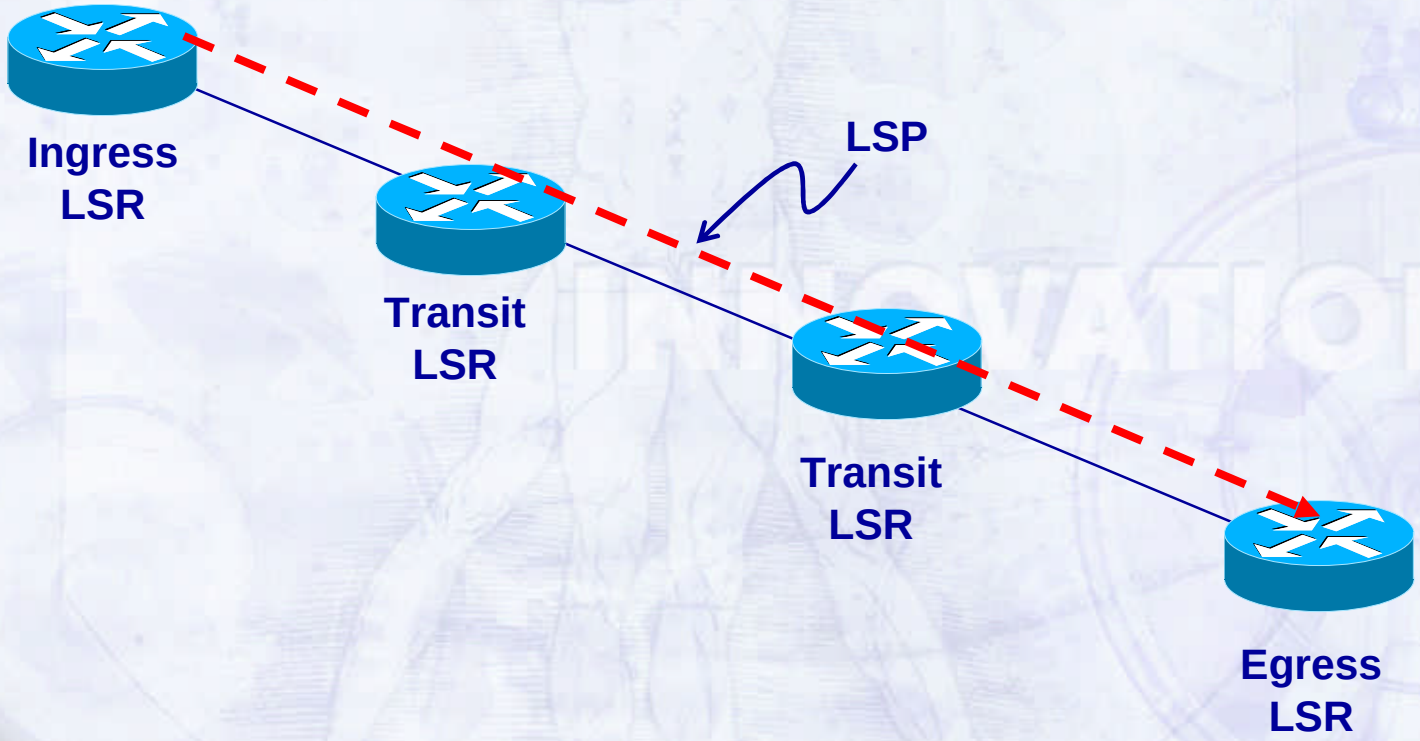

- LSPs are unidirectional
- Ingress, transit, and egress are relative to a given LSP
- A given router can be ingress, egress, and transit for different LSPs

The MPLS Switching Table

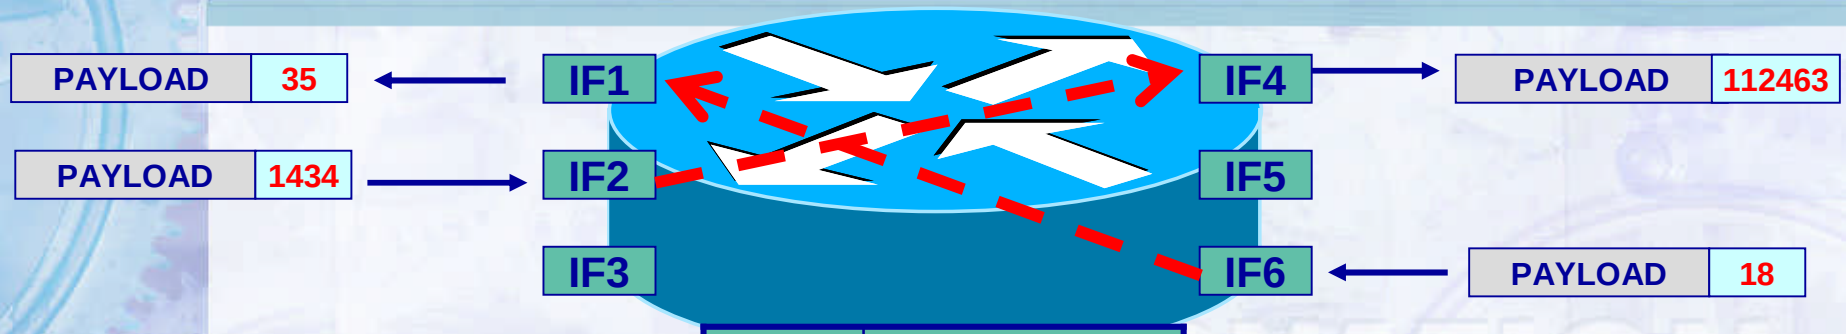

IN	OUT	
Label	Label	IF
18	35	1
22	5175	2
105	16	3
1434	112463	4
9295	17	5
26312	17	5
100034	16	6

- Labels have *local significance*
- *Incoming* labels are assigned by local router
  - Label distribution

Pushing, Swapping, and Popping

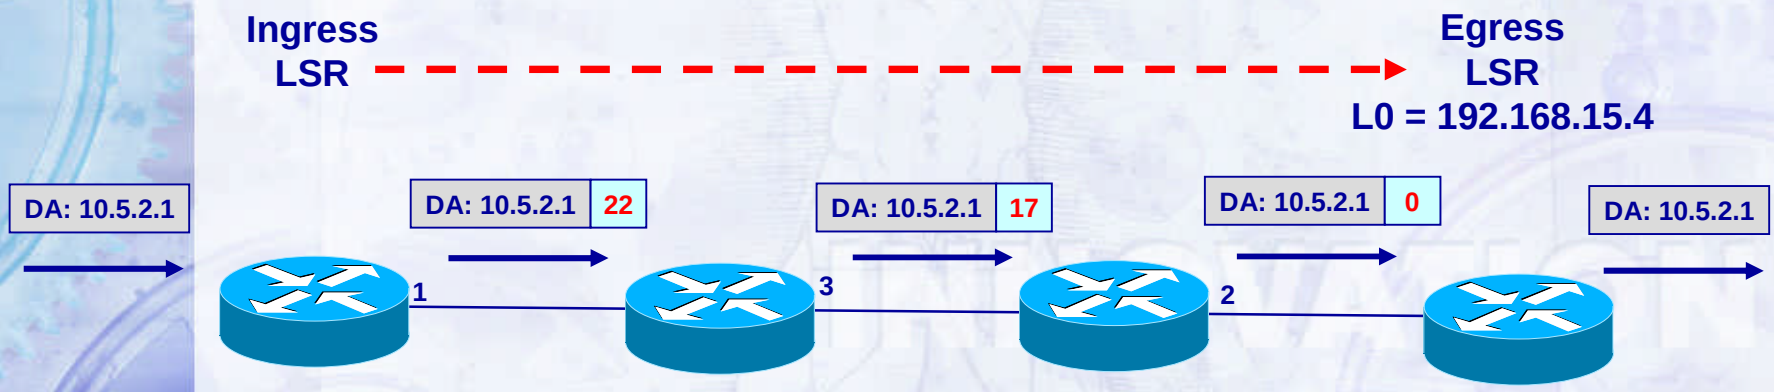

Routing Table

Prefix	Next Hop
10.5.0.0/16	192.168.15.4 **PUSH** 22, IF 1

MPLS Switching Table

IN	OUT
22	**SWAP** 17, IF 3

MPLS Switching Table

IN	OUT
17	**SWAP** 0, IF 2

MPLS Switching Table

IN	OUT
0	**POP**

Routing Table

Prefix	Next Hop
10.5.0.0/16	10.1.16.3

- Label 0 = *Explicit Null*

Penultimate Hop Popping

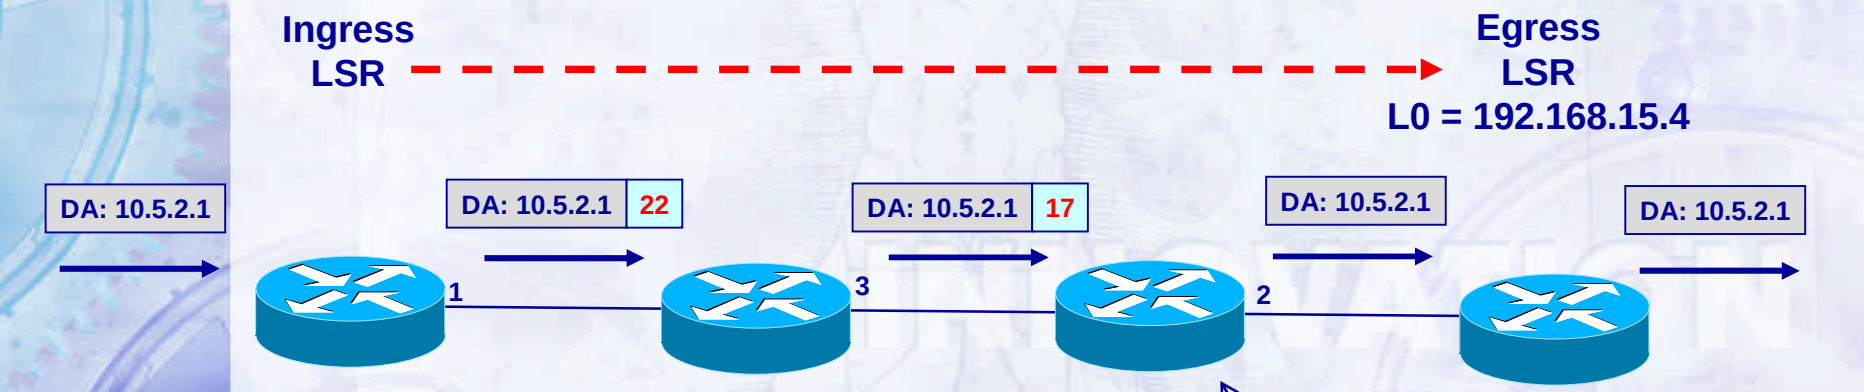

Routing Table

Prefix	Next Hop
10.5.0.0/16	192.168.15.4 **PUSH** 22, IF 1

MPLS Switching Table

IN	OUT
22	**SWAP** 17, IF 3

MPLS Switching Table

IN	OUT
17	3, **POP**

Routing Table

Prefix	Next Hop
10.5.0.0/16	10.1.16.3

- Label 3 = *Implicit Null*

***Penultimate LSR*:**
Last transit LSR before egress

Label Values

0 - 15 Reserved

LABEL	DESIGNATION
0	IPv4 Explicit Null
1	Router Alert
2	IPv6 Explicit Null
3	Implicit Null
4-14	Reserved for Future Use
15	OAM
16 - 2^{20-1}	Production Use

Label Stacking

Label Stacking allows LSPs to be tunneled in other LSPs

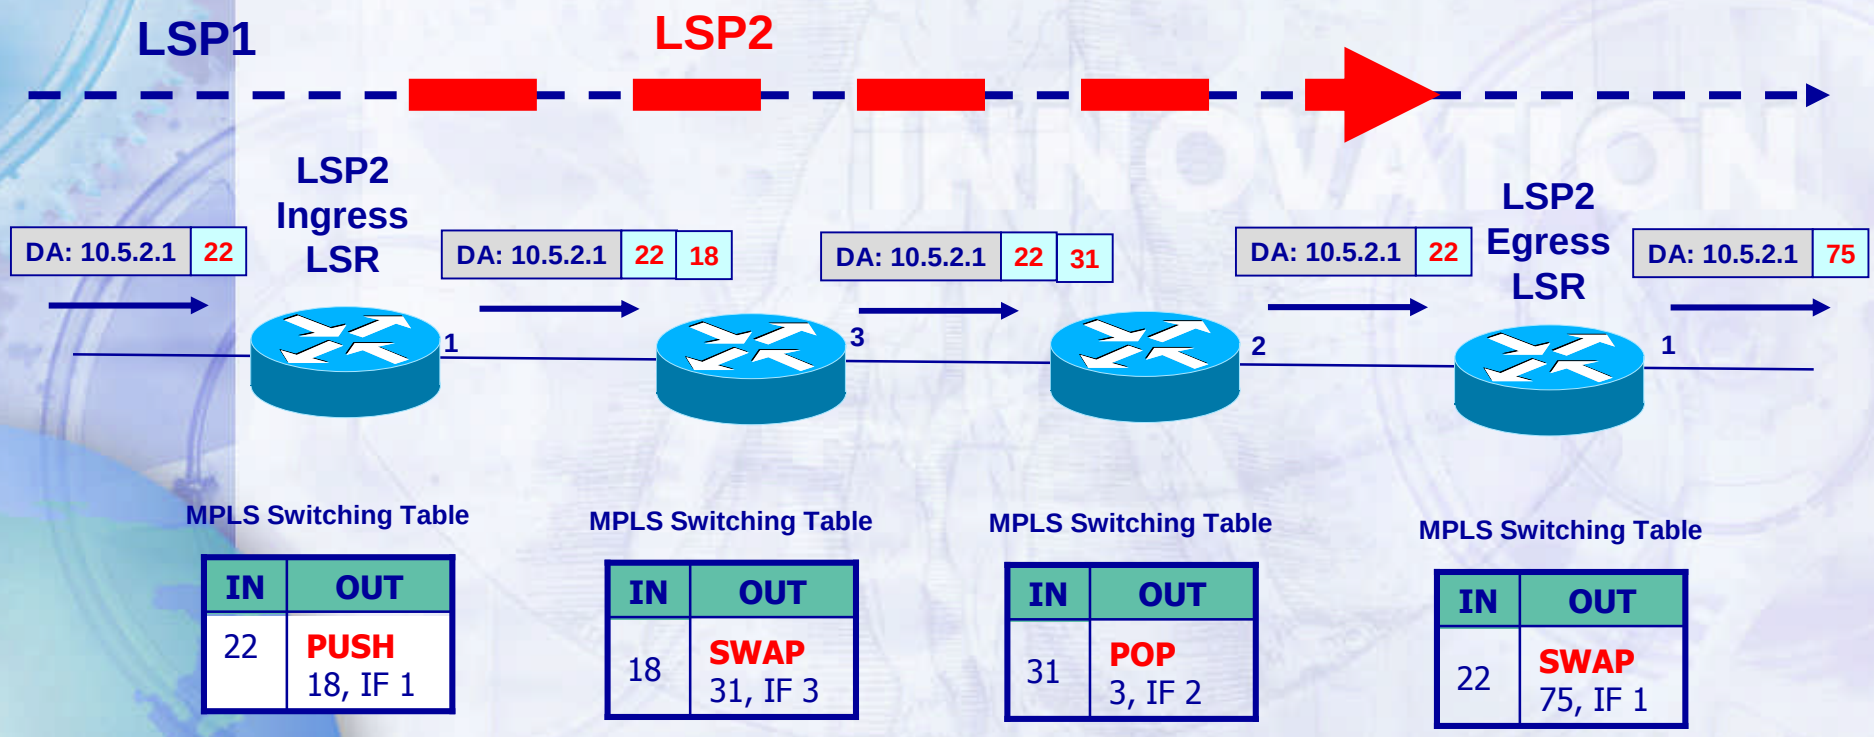

The MPLS Header

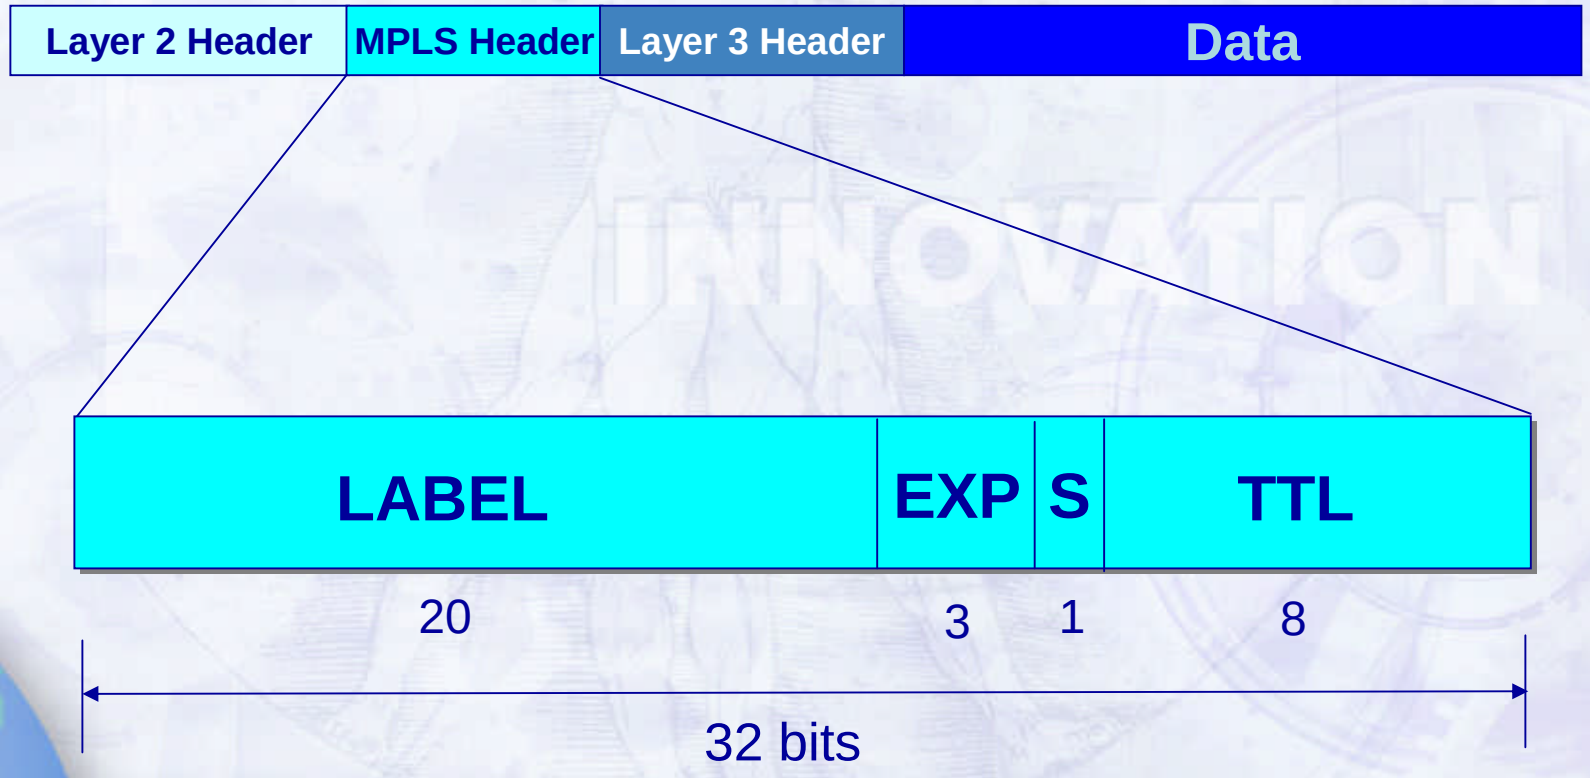

QUESTIONS?

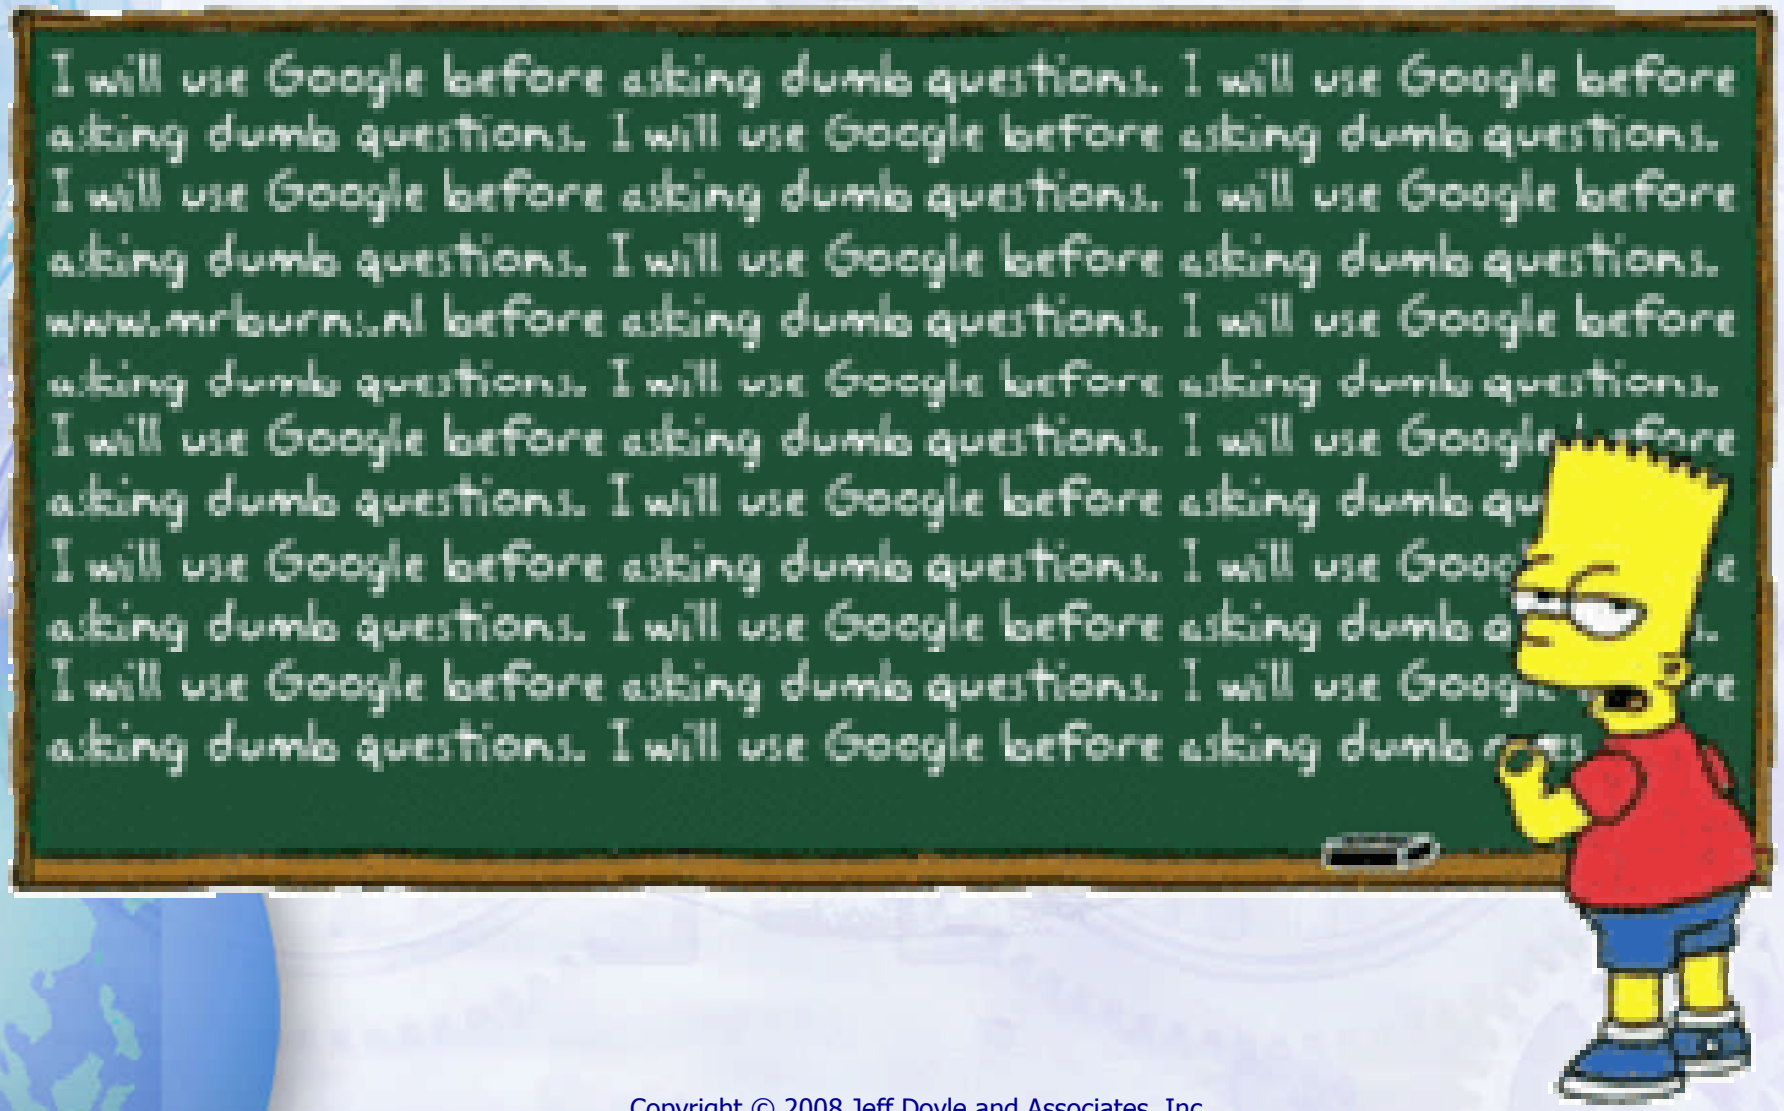

Label Distribution

- Requests for labels flow downstream
  - Ingress ==> Egress
  - Because ingress is the LSR that established the LSP
- Assignment of labels (label binding) flows upstream
  - Egress ==> Ingress
  - Because LSRs need to map *incoming* labels to some action (Push, Swap, Pop)

Label Distribution Protocols

- Label Distribution Protocol (LDP)
  - Hop-by-hop label distribution
  - Follows IGP best path
    - No traffic engineering capabilities
  - Highly scalable
    - Best suited for apps using thousands of LSPs (VPNs)
- Resource Reservation Protocol with Traffic Engineering Extensions (RSVP-TE)
  - End-to-end LSP signaling
  - Enables specification of path constraints
  - Less scalable -- LSRs maintain soft state
    - Best suited for traffic engineering in the core
- Constraint-Based Routed LDP (CR-LDP)
  - TE-capable LDP
  - Never widely deployed
- MP-BGP
  - Can distribute labeled-unicast address families
  - Best suited for inter-AS VPNs

Label Distribution: LDP

- Hop-by-hop label distribution
- Always follows IGP best path
- IP addresses are locally bound to labels
- Bindings are stored in Label Information Base (LIB)
- All bindings advertised to all peers
  - No split horizon

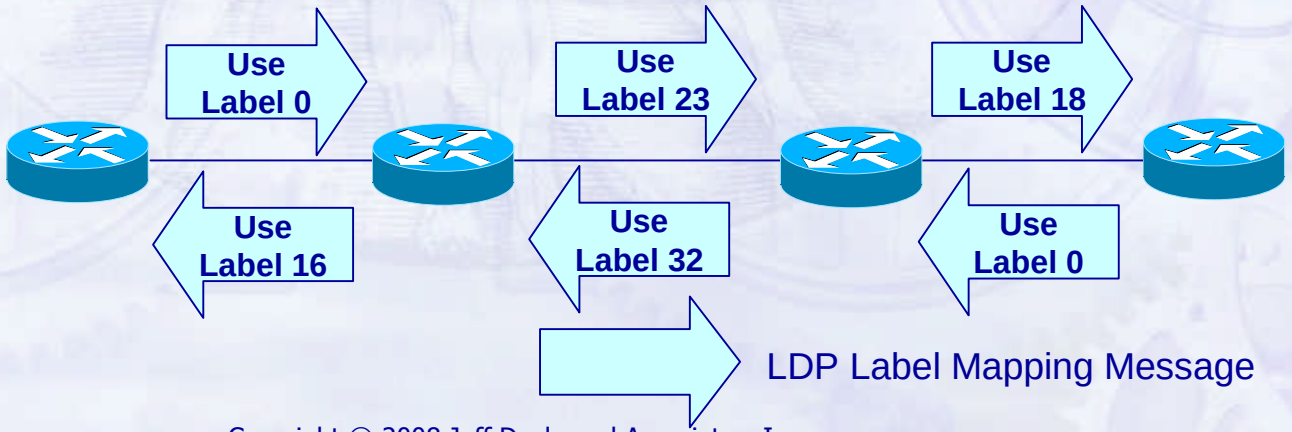

Label Distribution: RSVP

- End-to-end *constrained* path signaling
- Enabled by OSPF or IS-IS with TE extensions
  - Extended IGPs flood TE interface parameters:
    - Maximum Bandwidth
    - Maximum Reservable Bandwidth
    - Unreserved Bandwidth
    - TE Metric
    - Administrative Group (aka Link Affinity or "Link Coloring")
- Interface parameters used to build *Traffic Engineering Database* (TED)
- *Constrained Shortest Path First* (CSPF)
  - Calculates best path based on specified constraints
- *Explicit Route Object* (ERO) passed to RSVP

CSPF Calculation

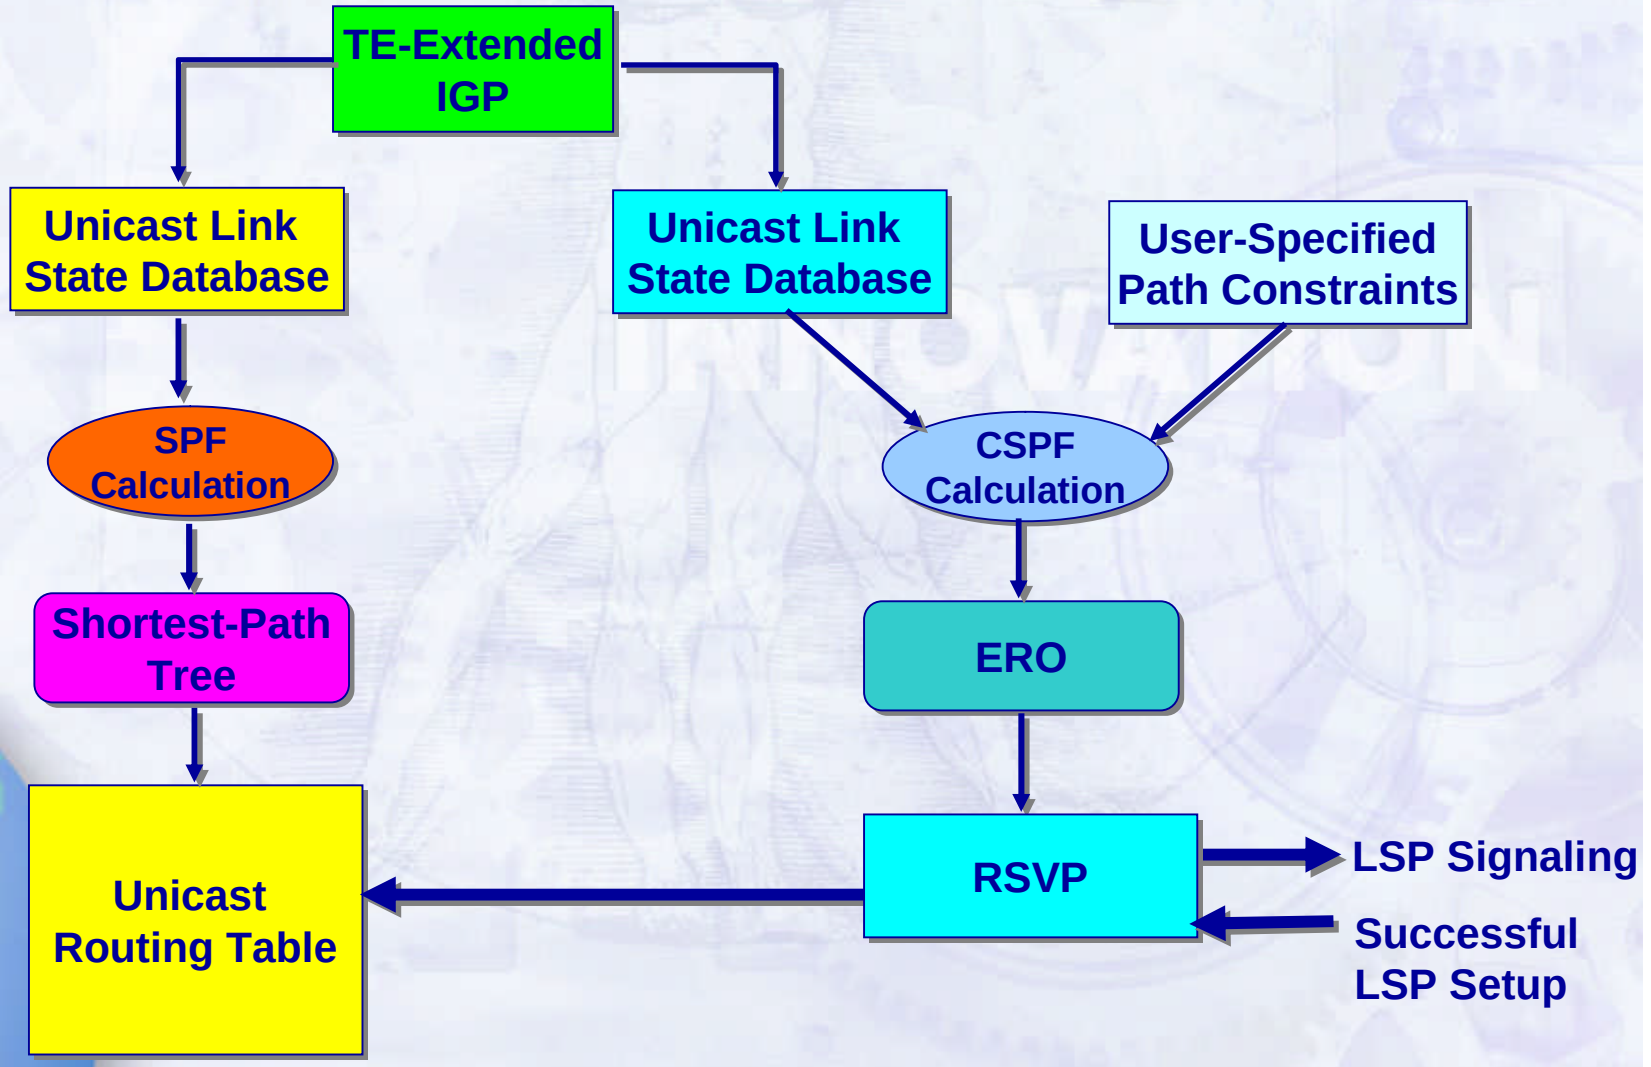

RSVP LSP Signaling

ERO

B Strict;
E Loose;
G Strict;
H Strict

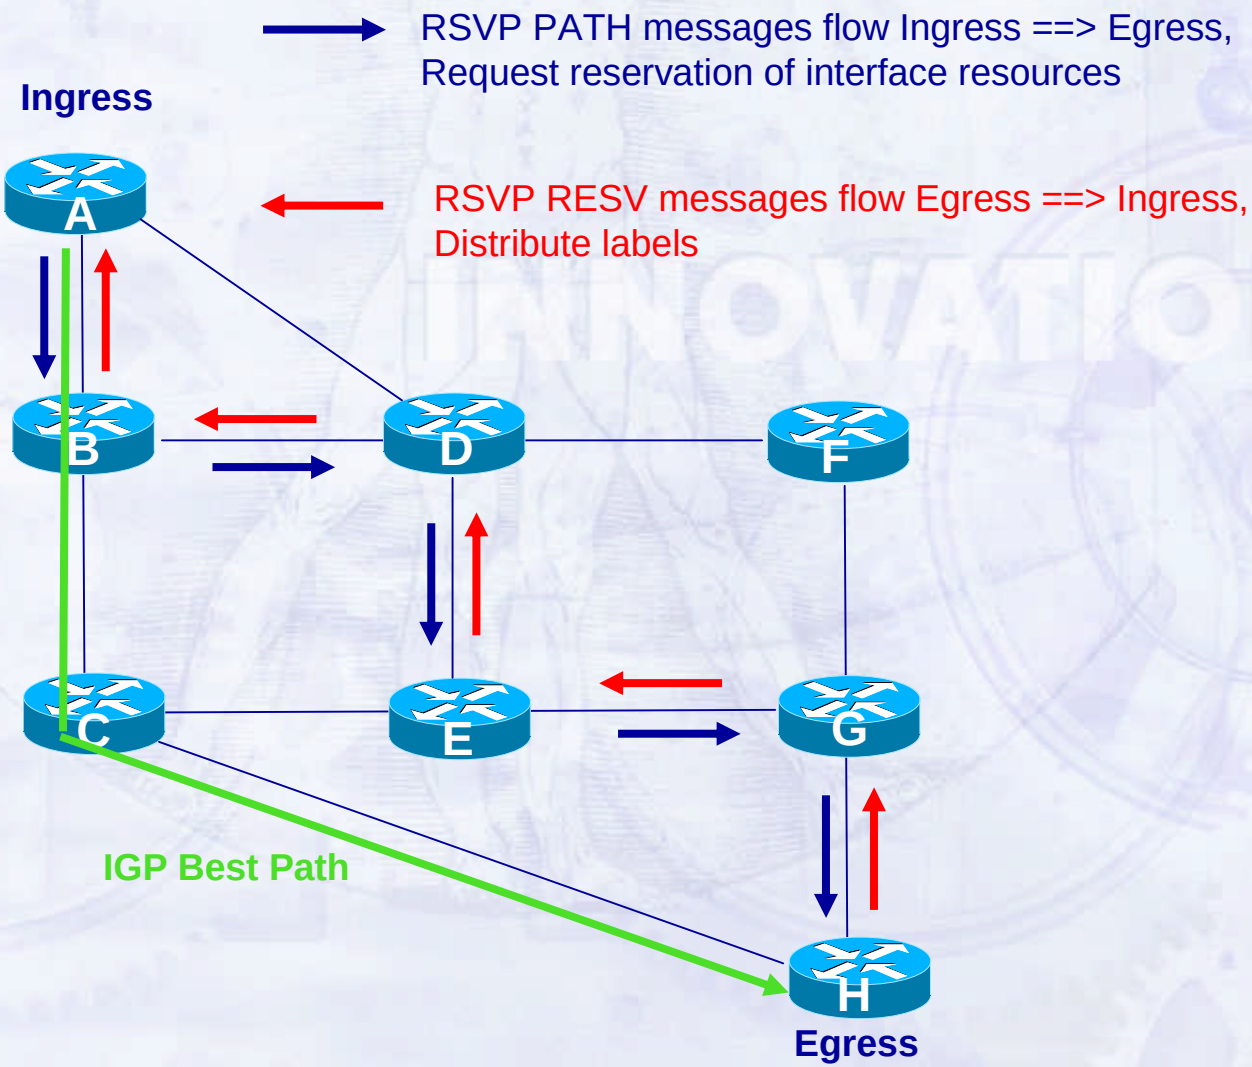

QUESTIONS?

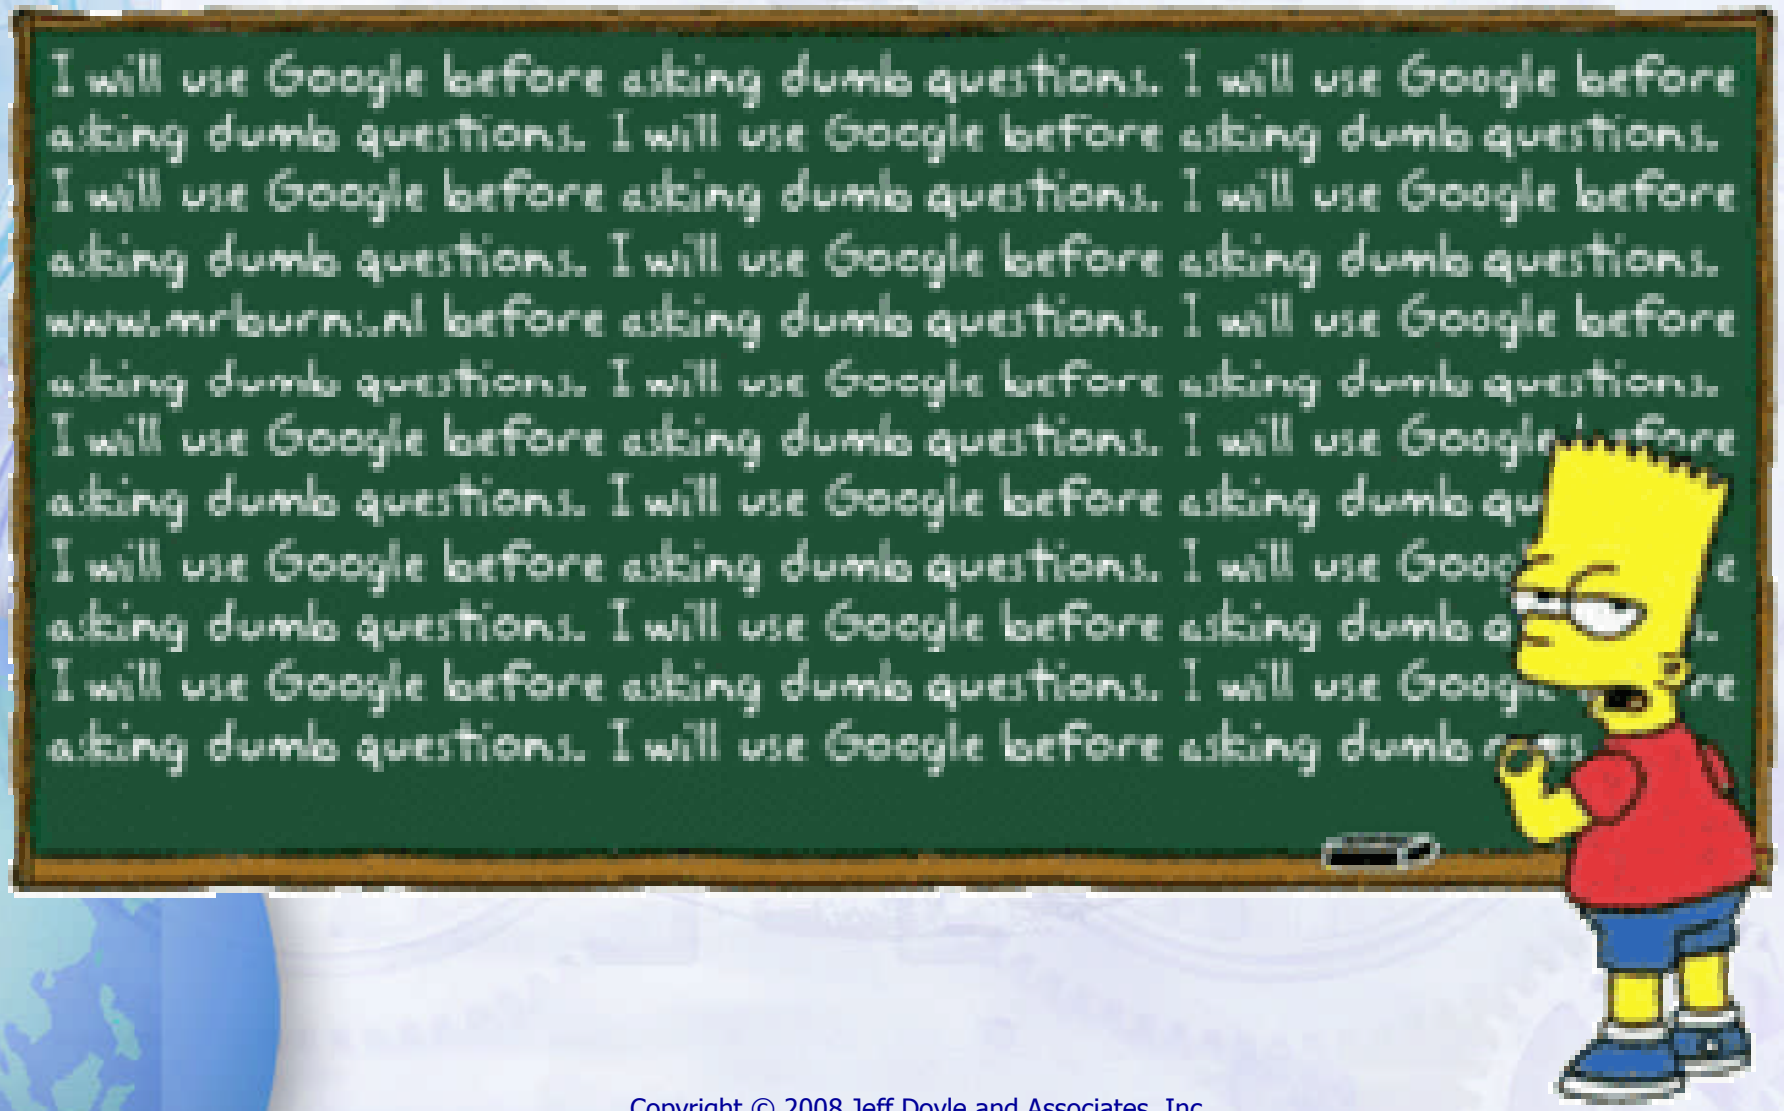

RSVP and LDP Applications

Typical PoP architecture:

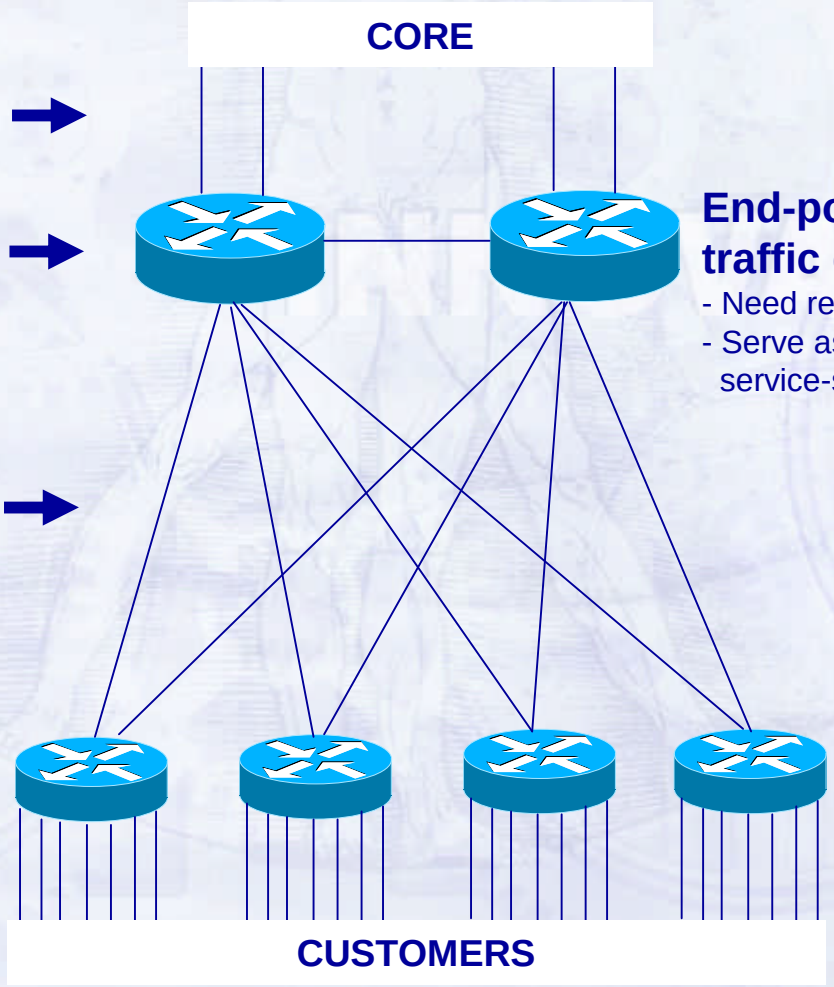

High-bandwidth core uplinks

Core routers
Primary requirement is high performance

Aggregation routers might or might not exist here

Edge routers
Primary requirements are:
•Service intelligence
•Customer-facing interface density

End-point for RSVP traffic engineering core LSPs
- Need relatively few of these
- Serve as PoP-to-PoP tunnels for service-specific edge LSPs

End-point for LDP service-specific LSPs
Might be hundreds or thousands of these

Using RSVP and LDP LSPs Together

LDP-based customer (and/or service) specific LSPs at edge are tunneled through core in RSVP-TE LSPs

Both LDP scalability and RSVP TE capabilities are leveraged

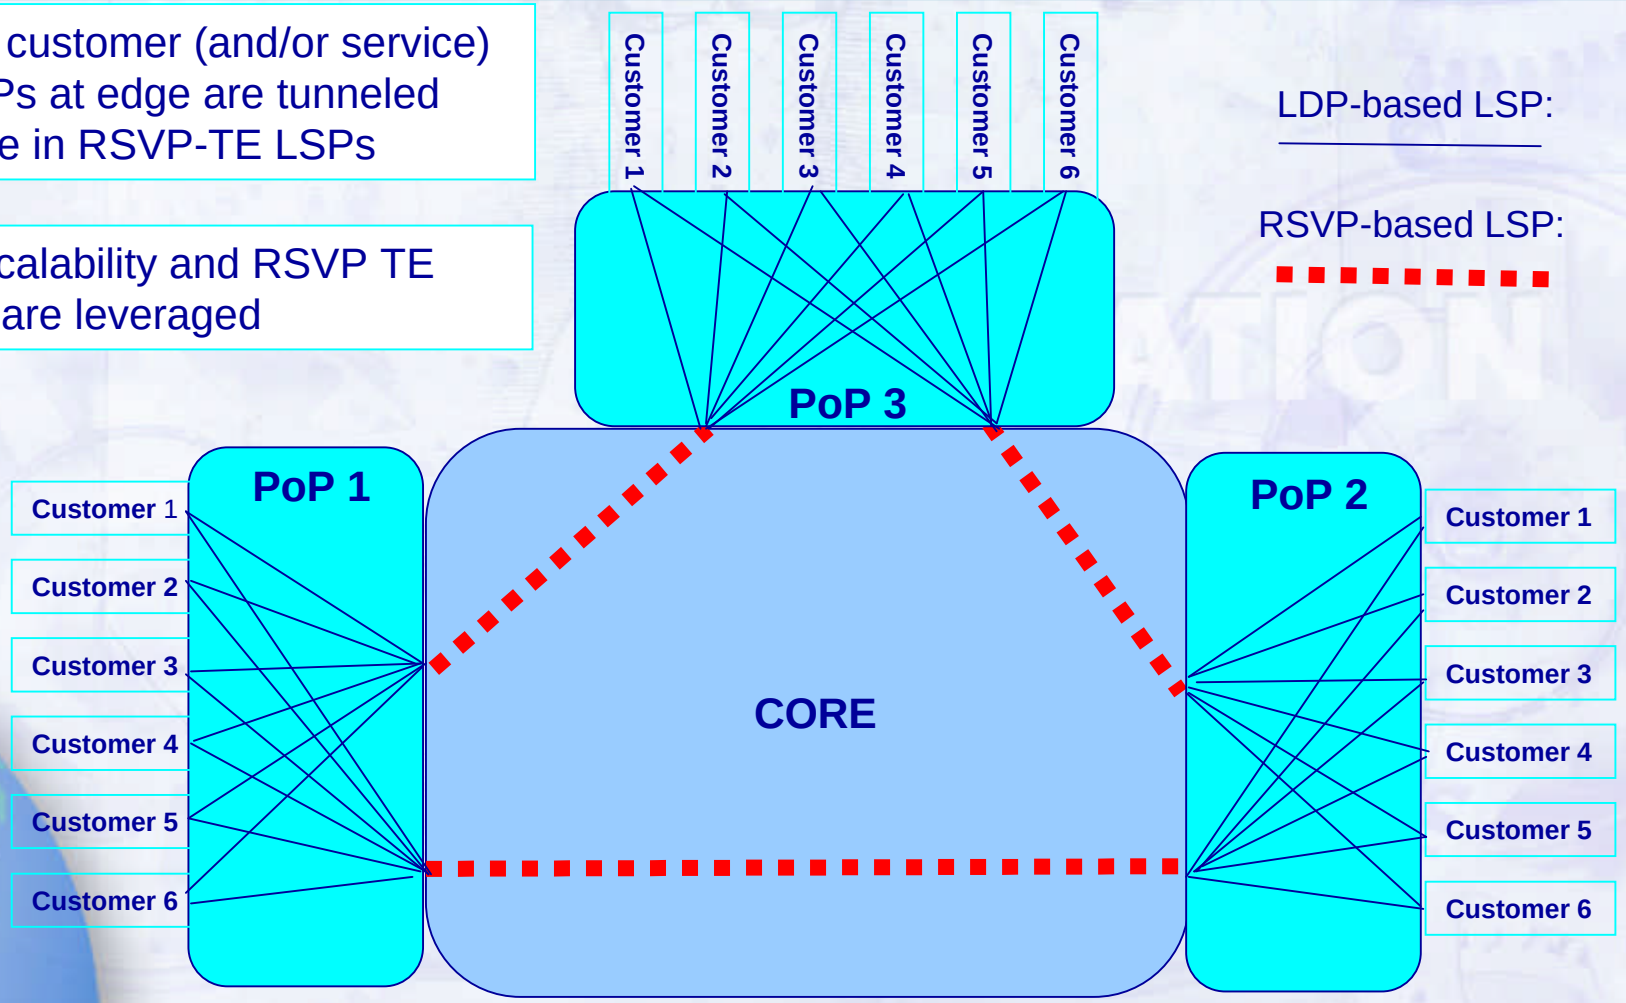

A Quick Glance at MPLS VPNs

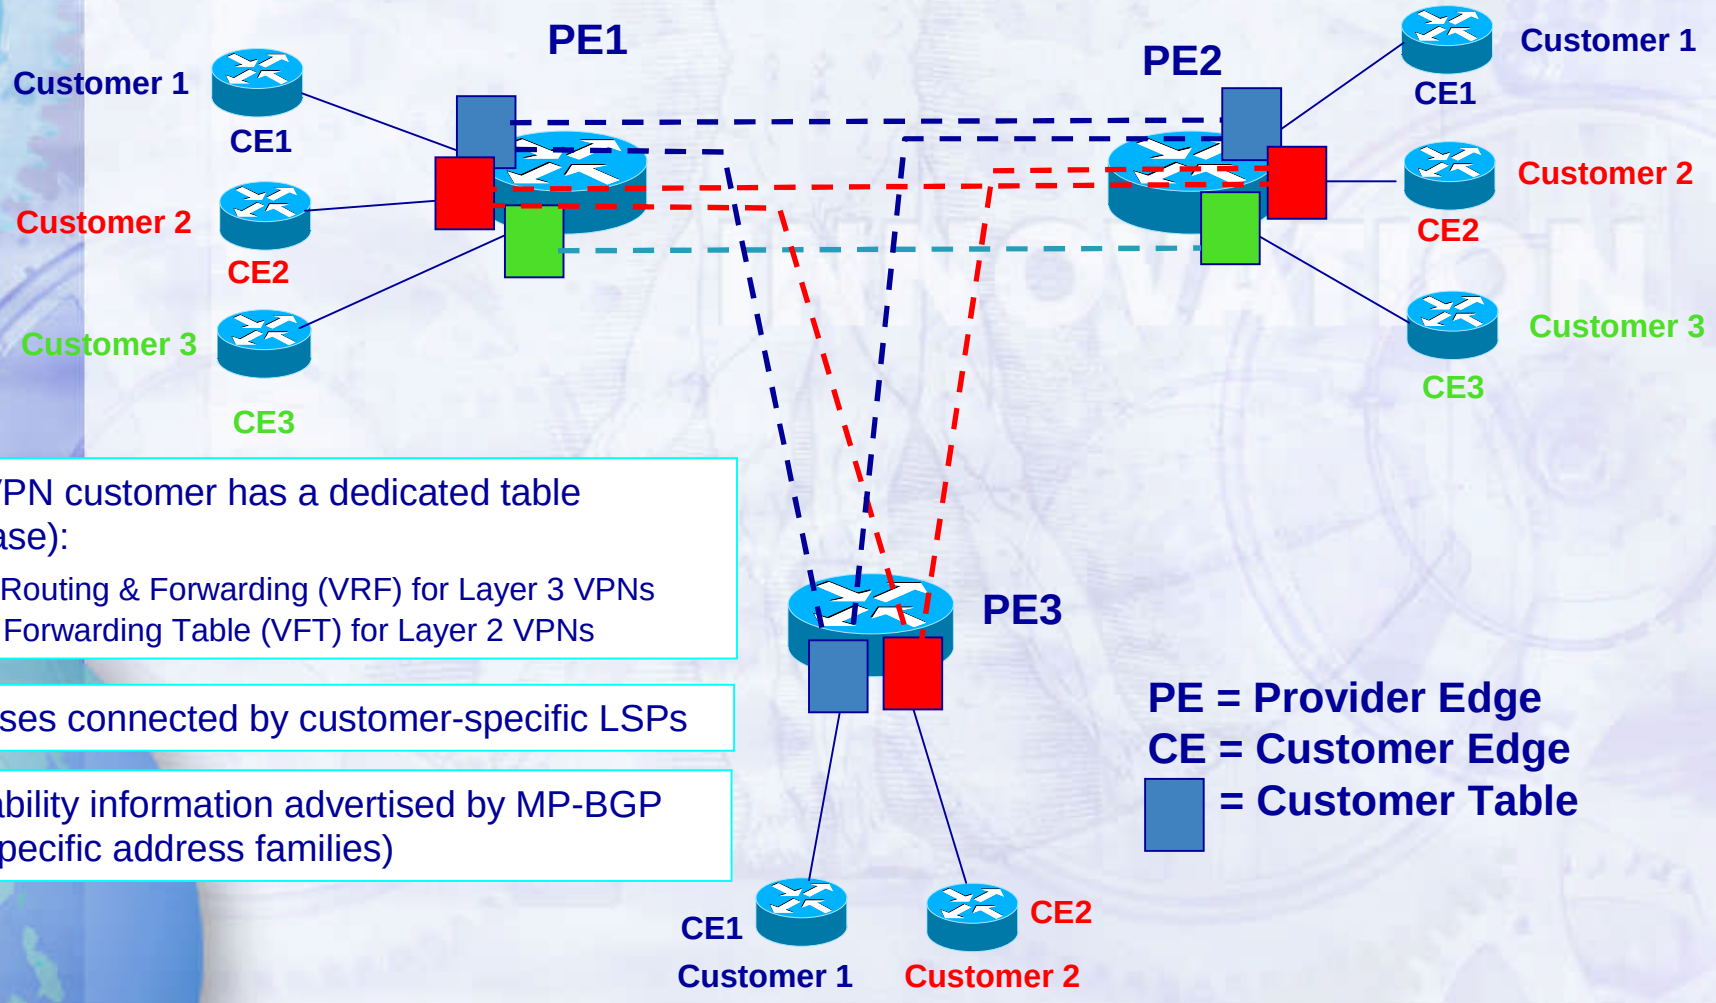

Each VPN customer has a dedicated table (database):

- VPN Routing & Forwarding (VRF) for Layer 3 VPNs
- VPN Forwarding Table (VFT) for Layer 2 VPNs

Databases connected by customer-specific LSPs

Reachability information advertised by MP-BGP (VPN-specific address families)

PE = Provider Edge
CE = Customer Edge
= Customer Table

QUESTIONS?

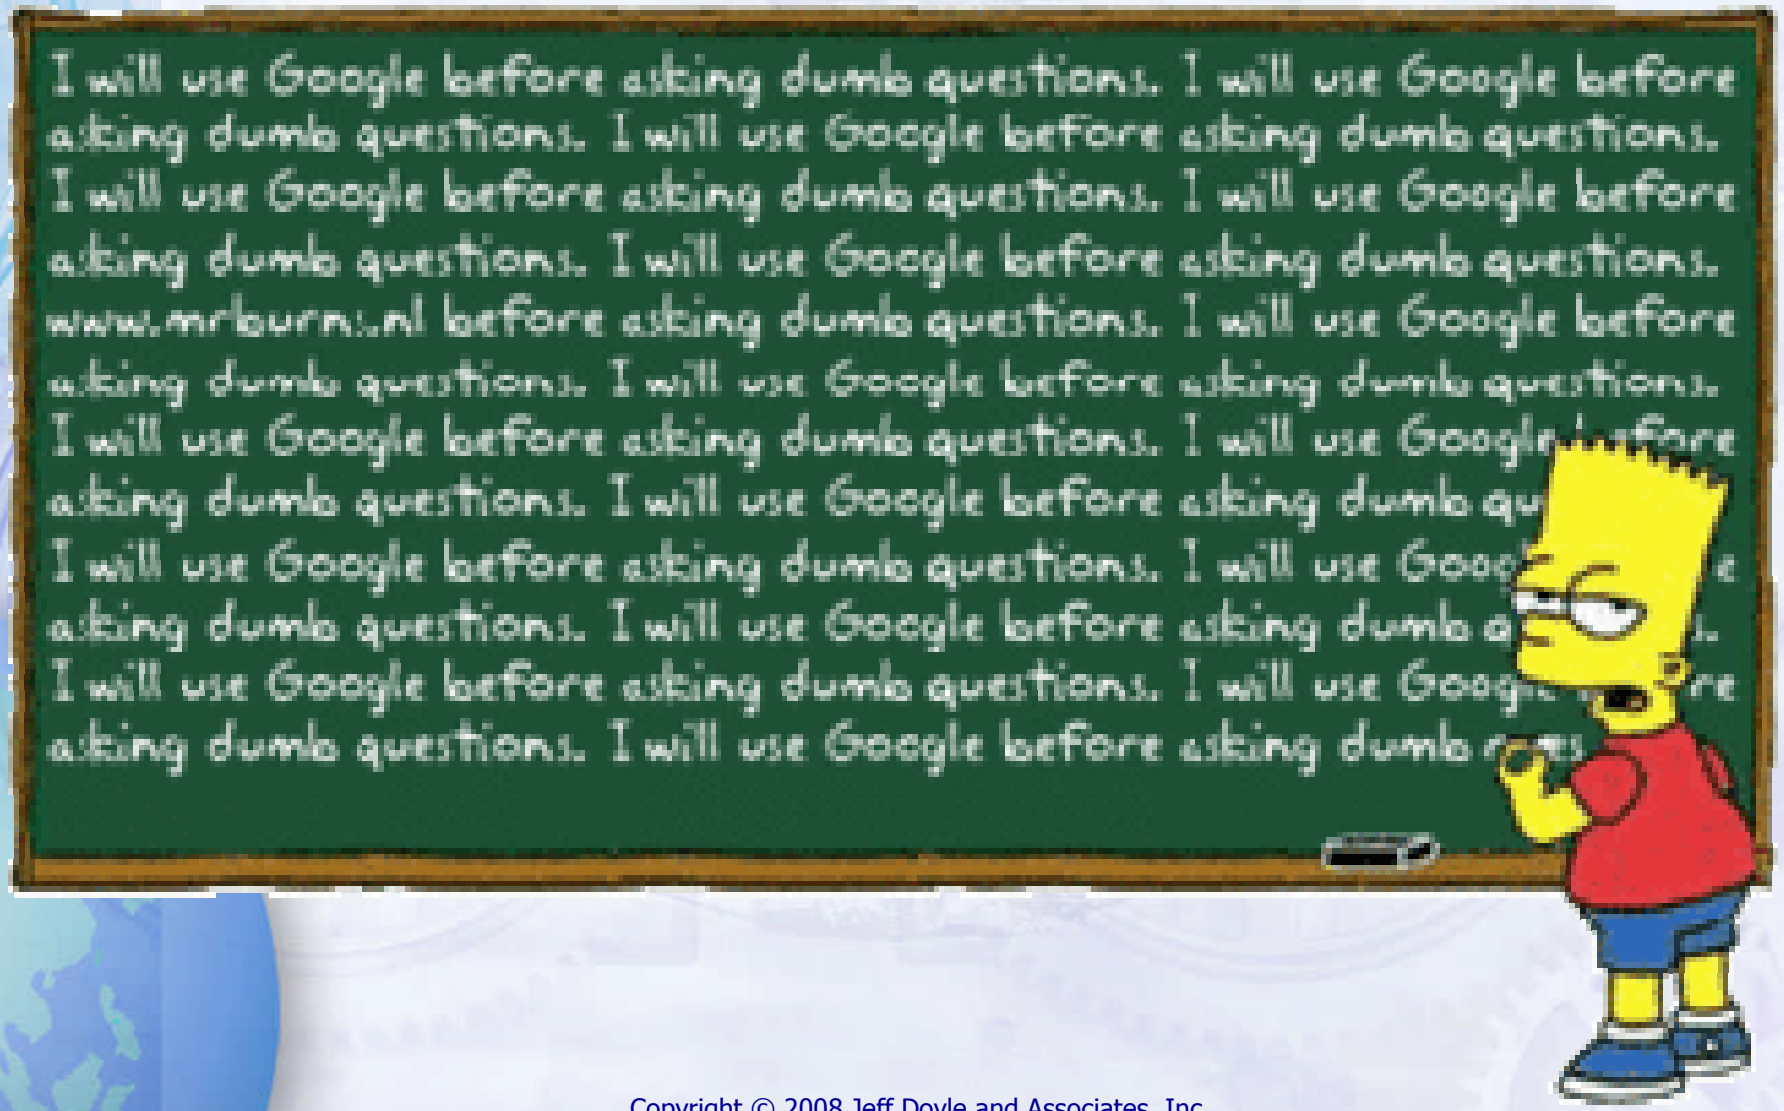

Thank You!

jdoyle@doyleassociates.net

(303) 428-4680

www.doyleassociates.net