

# MPLS Traffic Engineering across AS boundaries

CENSORED

Cristel Pelsser

*cpe@info.ucl.ac.be*

Université Catholique de Louvain  
Belgium



# Agenda

- ◆ Problem statement
- ◆ Constrained intra-AS path computation
- ◆ Current inter-AS routing
- ◆ Proposal for constrained inter-AS path computation
- ◆ Remaining issues

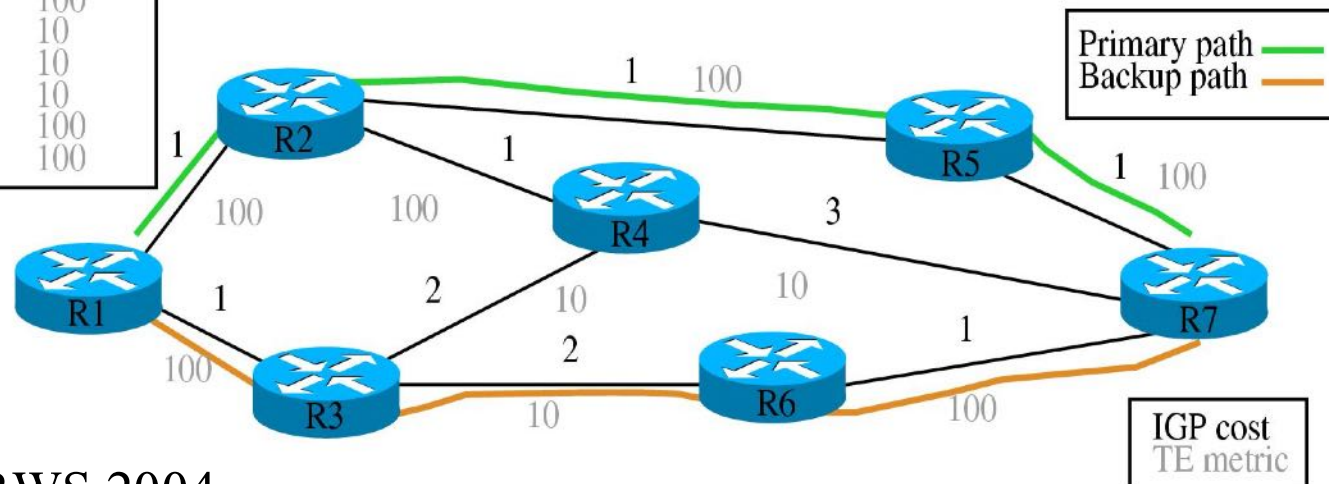
# Problem statement

- ◆ Use of MPLS across AS boundaries
  - ◆ VPNs
  - ◆ Faster recovery than with BGP
  - ◆ QoS
- ◆ Requirements are formulated at the IETF
  - ◆ ccamp (a lot of new drafts planned)
- ◆ Protocol extensions to RSVP -TE already proposed at the IETF
  - ◆ Establishment of inter-AS LSPs  
(draft-pelsser-rsvp-te-interdomain-lsp-00.txt)
  - ◆ Protection of inter-AS LSPs  
(draft-decnodder-mpls-interas-protection-01.txt)

# Constrained intra-AS path computation

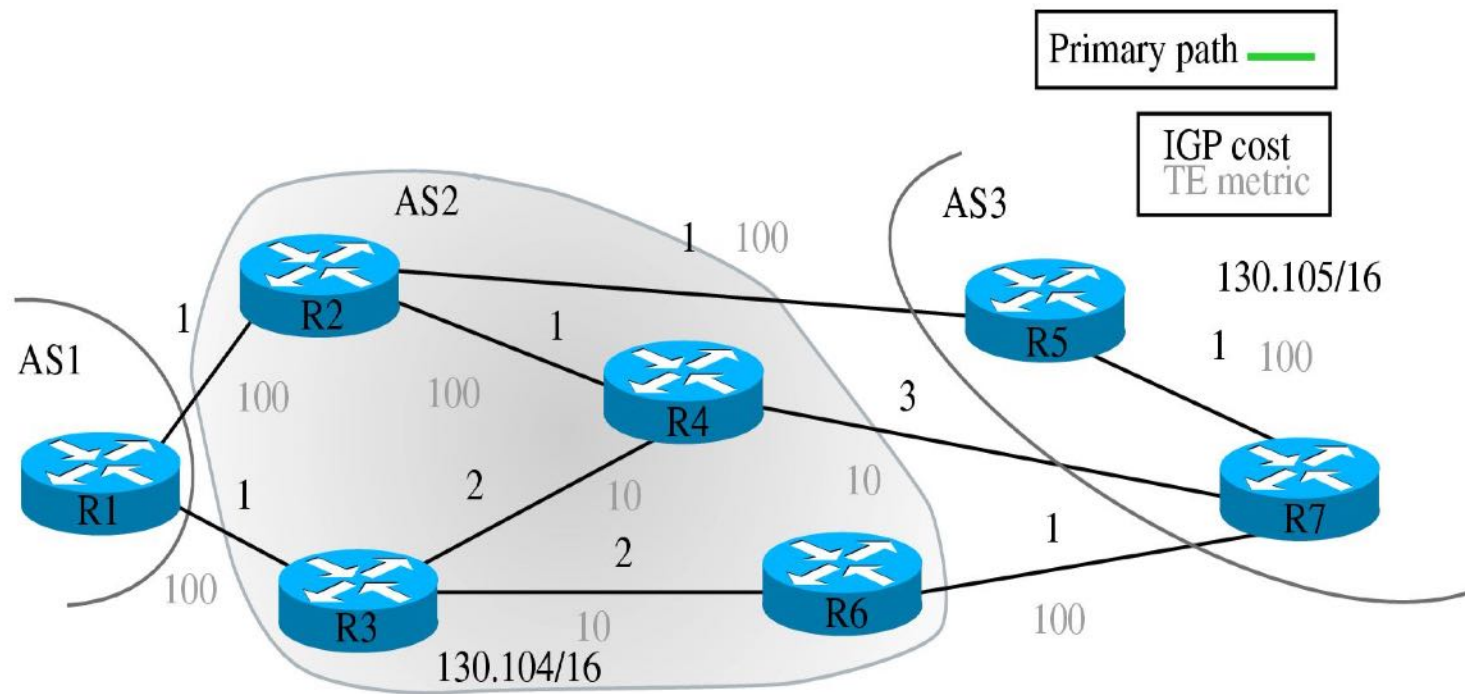
- ◆ **Each node** possesses the complete topology of its AS (No areas)
  - ◆ Link info:
    - ◆ IGP cost
    - ◆ TE info with OSPF-TE or IS-IS TE

Link	Cost	TE metric
R1-R2	1	100
R1-R3	1	100
R2-R4	1	100
R2-R5	1	100
R3-R4	2	10
R3-R6	2	10
R4-R7	3	10
R5-R7	1	100
R6-R7	1	100



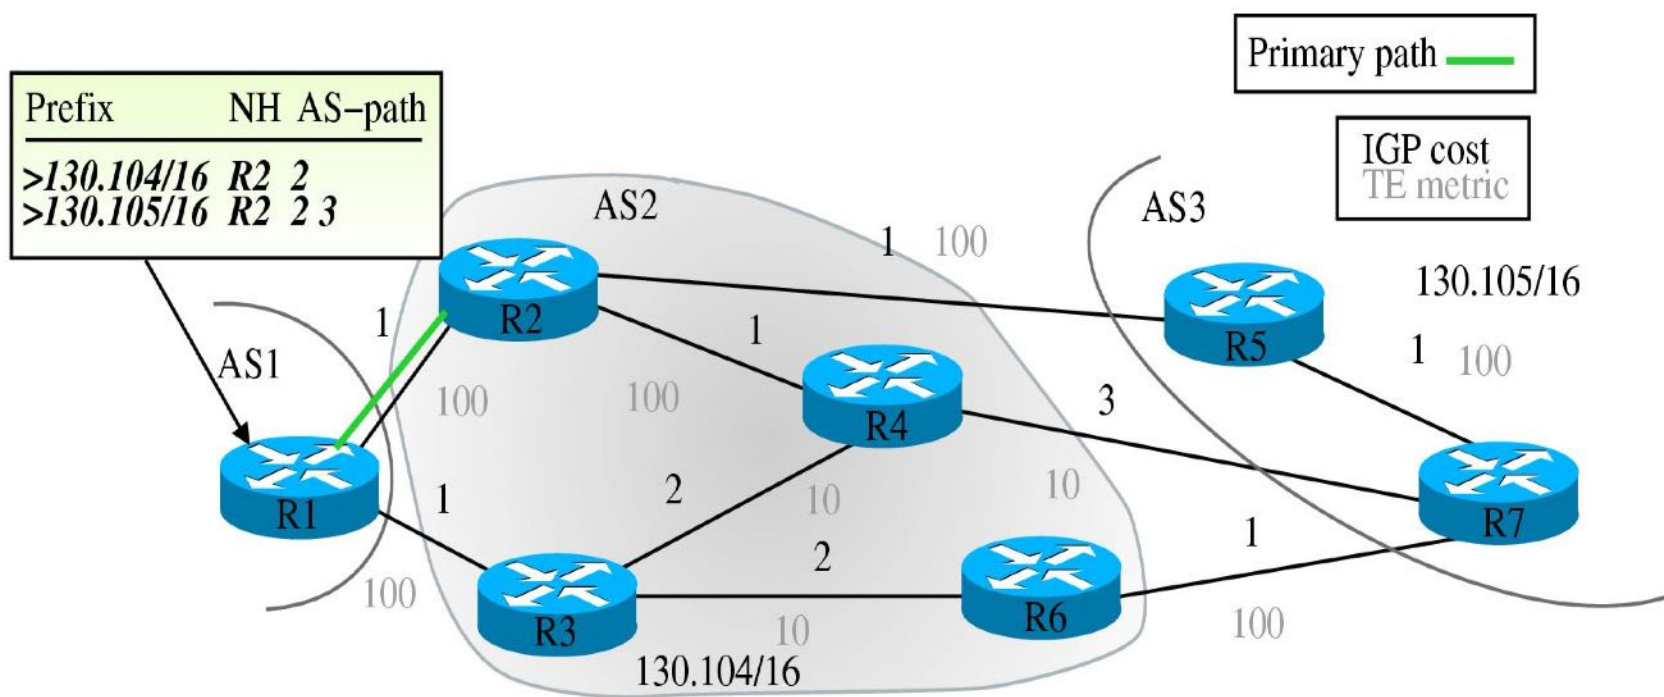
# Current inter-AS routing (BGP)

- ◆ The nodes only possess **reachability information** for prefixes outside the AS



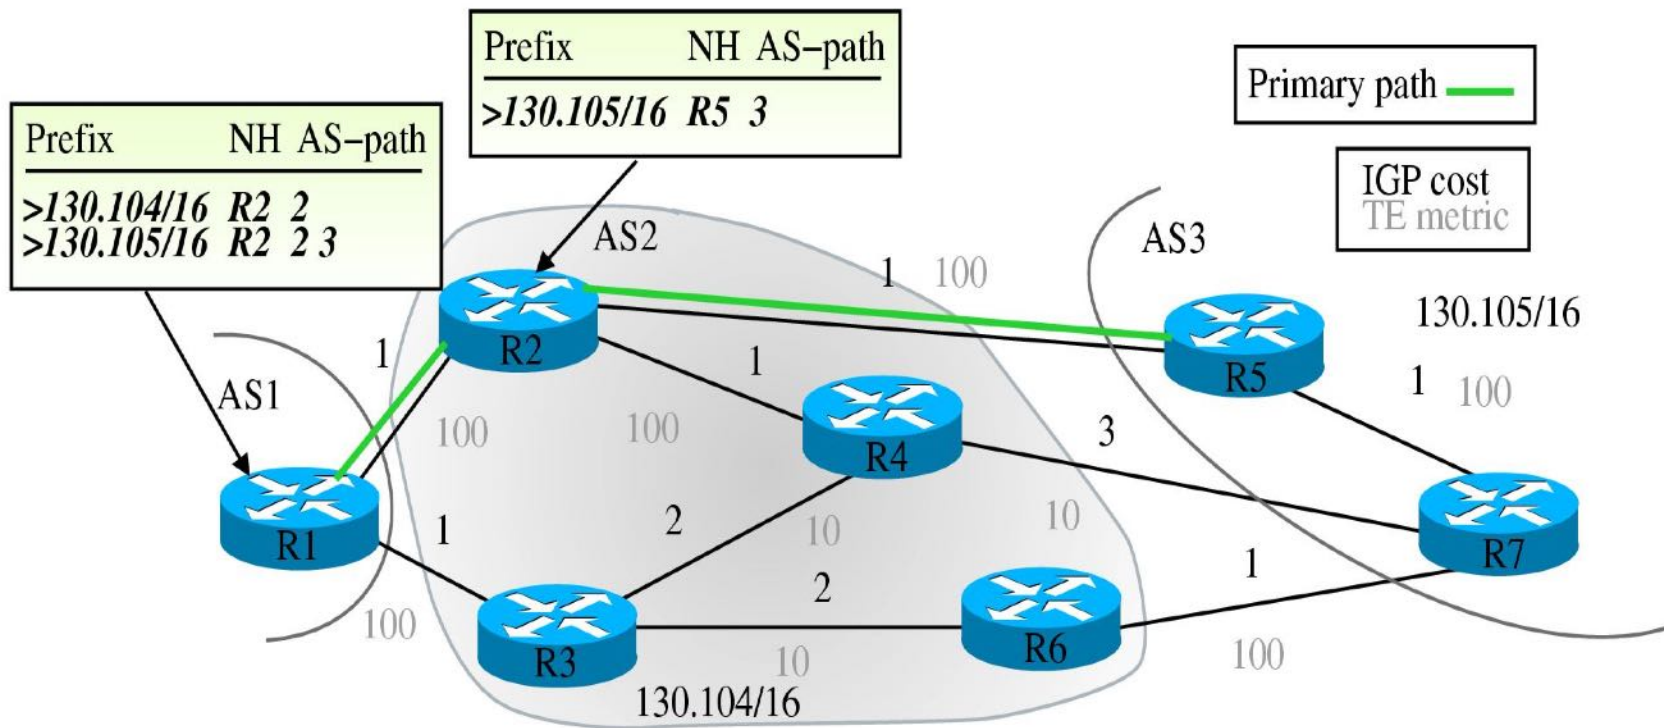
# Current inter-AS routing (BGP)

- ◆ The nodes only possess **reachability information** for prefixes outside the AS



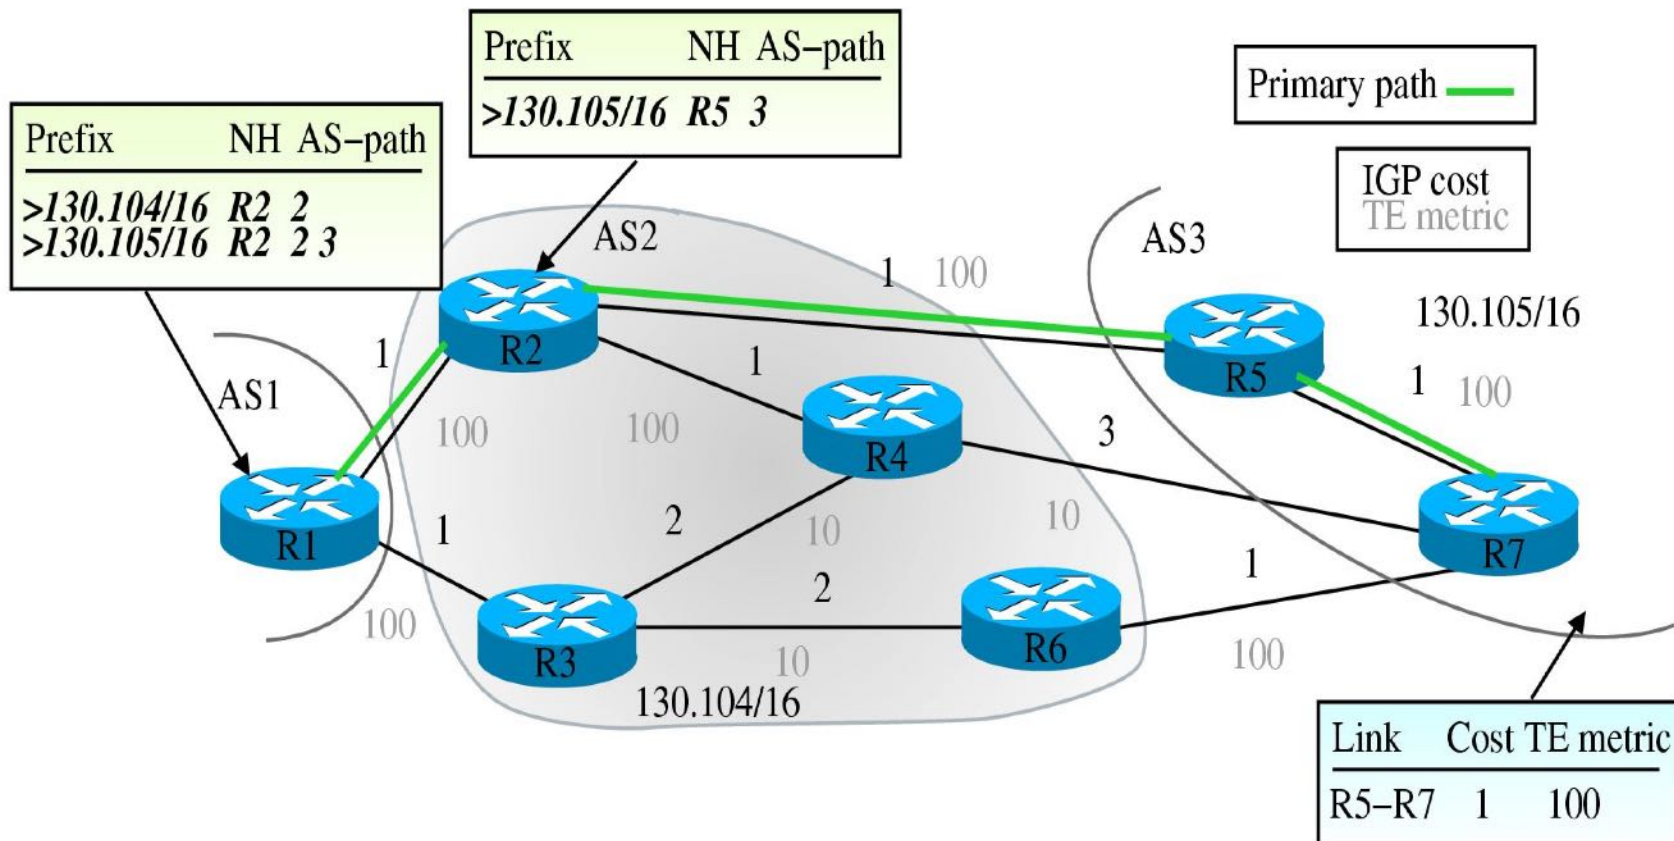
# Current inter-AS routing (BGP)

- ◆ The nodes only possess **reachability information** for prefixes outside the AS



# Current inter-AS routing (BGP)

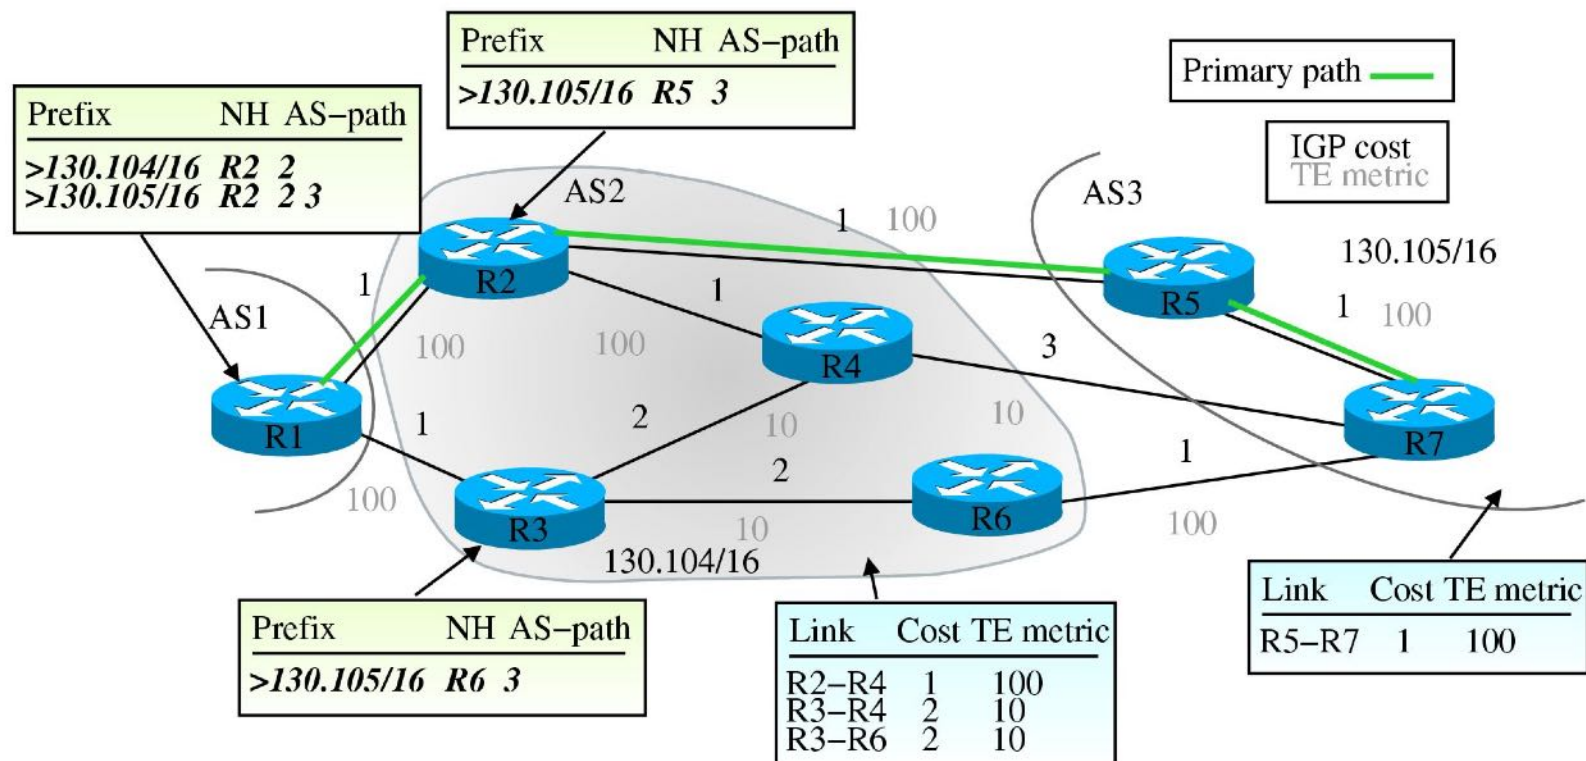
- The nodes only possess **reachability information** for prefixes outside the AS





# Current inter-AS routing (BGP)

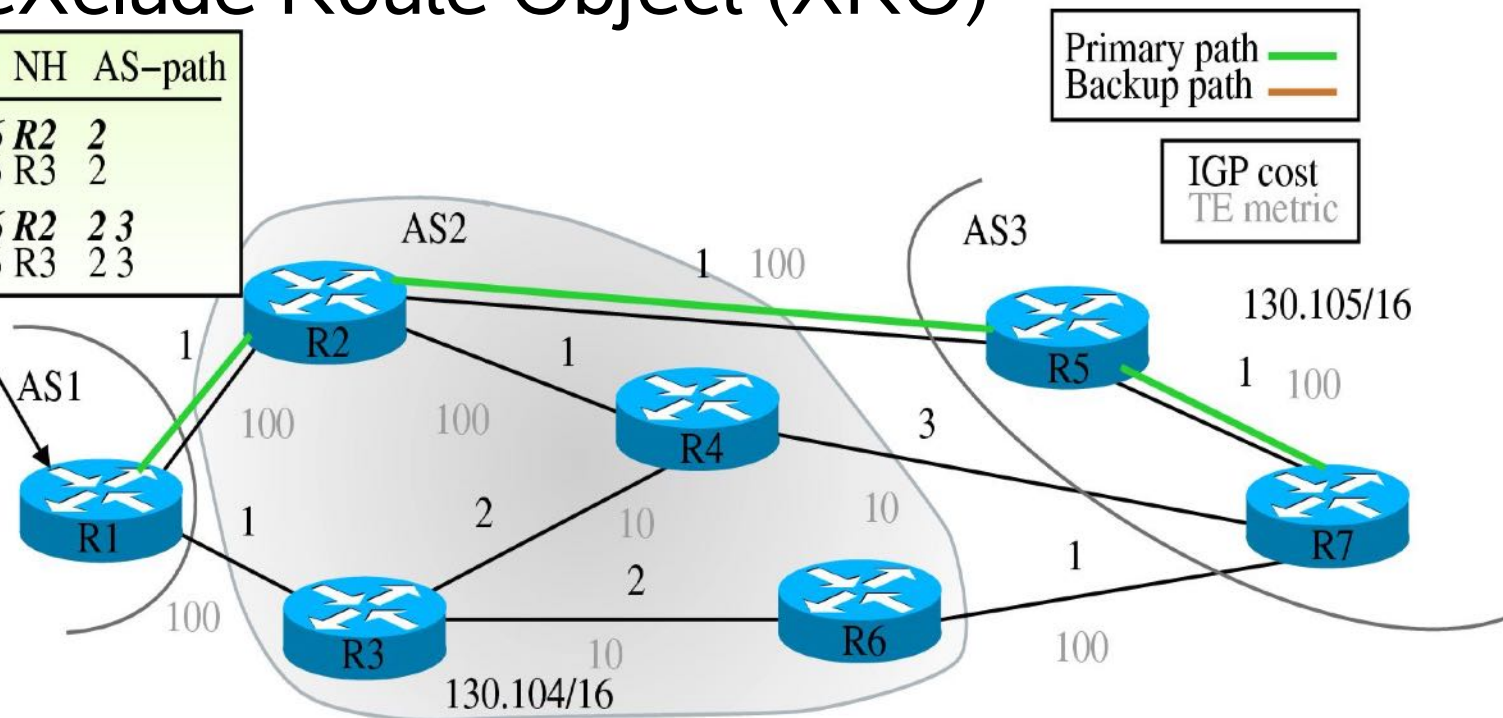
- ◆ The nodes only possess **reachability information** for prefixes outside the AS
- ◆ Alternate path through R3 never used



# Constrained inter-AS path computation : Proposal

- Compute disjoint path based on **local** Adj-RIB-In and eXclude Route Object (XRO)

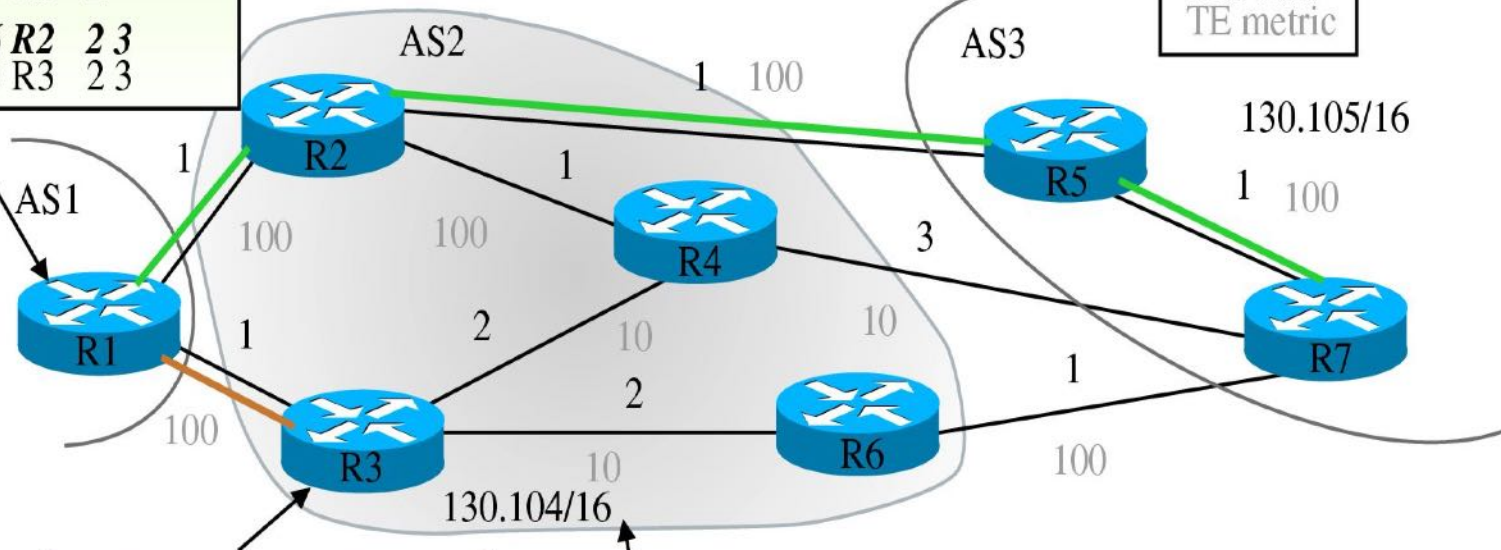
Prefix	NH	AS-path
>130.104/16	R2	2
130.104/16	R3	2
>130.105/16	R2	2 3
130.105/16	R3	2 3



# Constrained inter-AS path computation : Proposal

- Compute disjoint path based on **local** Adj-RIB-In and eXclude Route Object (XRO)

Prefix	NH	AS-path
>130.104/16	R2	2
130.104/16	R3	2
>130.105/16	R2	2 3
130.105/16	R3	2 3



Primary path —  
Backup path —

IGP cost  
TE metric

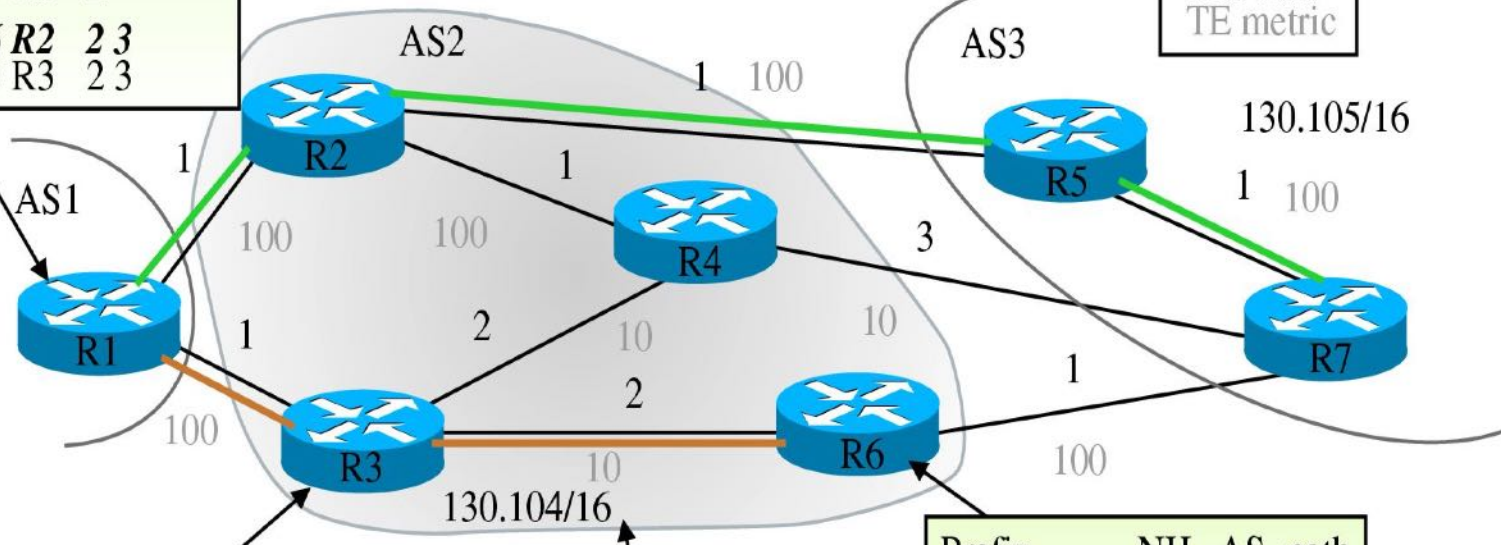
Prefix	NH	AS-path
>130.105/16	R6	3
130.105/16	R4	3
130.105/16	R2	3

Link	Cost	TE metric
R2-R4	1	100
R3-R4	2	10
R3-R6	2	10

# Constrained inter-AS path computation : Proposal

- Compute disjoint path based on **local** Adj-RIB-In and eXclude Route Object (XRO)

Prefix	NH	AS-path
>130.104/16	R2	2
130.104/16	R3	2
>130.105/16	R2	2 3
130.105/16	R3	2 3



Prefix	NH	AS-path
>130.105/16	R6	3
130.105/16	R4	3
130.105/16	R2	3

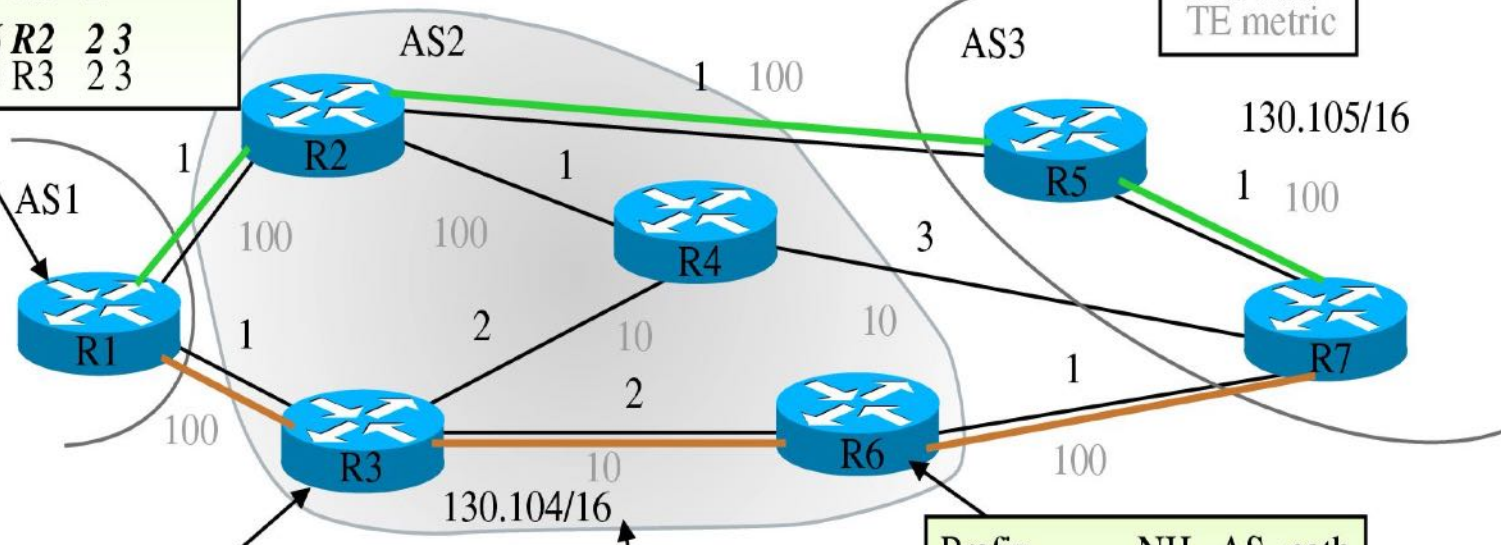
Link	Cost	TE metric
R2-R4	1	100
R3-R4	2	10
R3-R6	2	10

Prefix	NH	AS-path
>130.105/16	R7	3
130.105/16	R4	3
130.105/16	R2	3

# Constrained inter-AS path computation : Proposal

- Compute disjoint path based on **local** Adj-RIB-In and eXclude Route Object (XRO)

Prefix	NH	AS-path
>130.104/16	R2	2
130.104/16	R3	2
>130.105/16	R2	2 3
130.105/16	R3	2 3



Primary path —  
Backup path —

IGP cost  
TE metric

Prefix	NH	AS-path
>130.105/16	R6	3
130.105/16	R4	3
130.105/16	R2	3

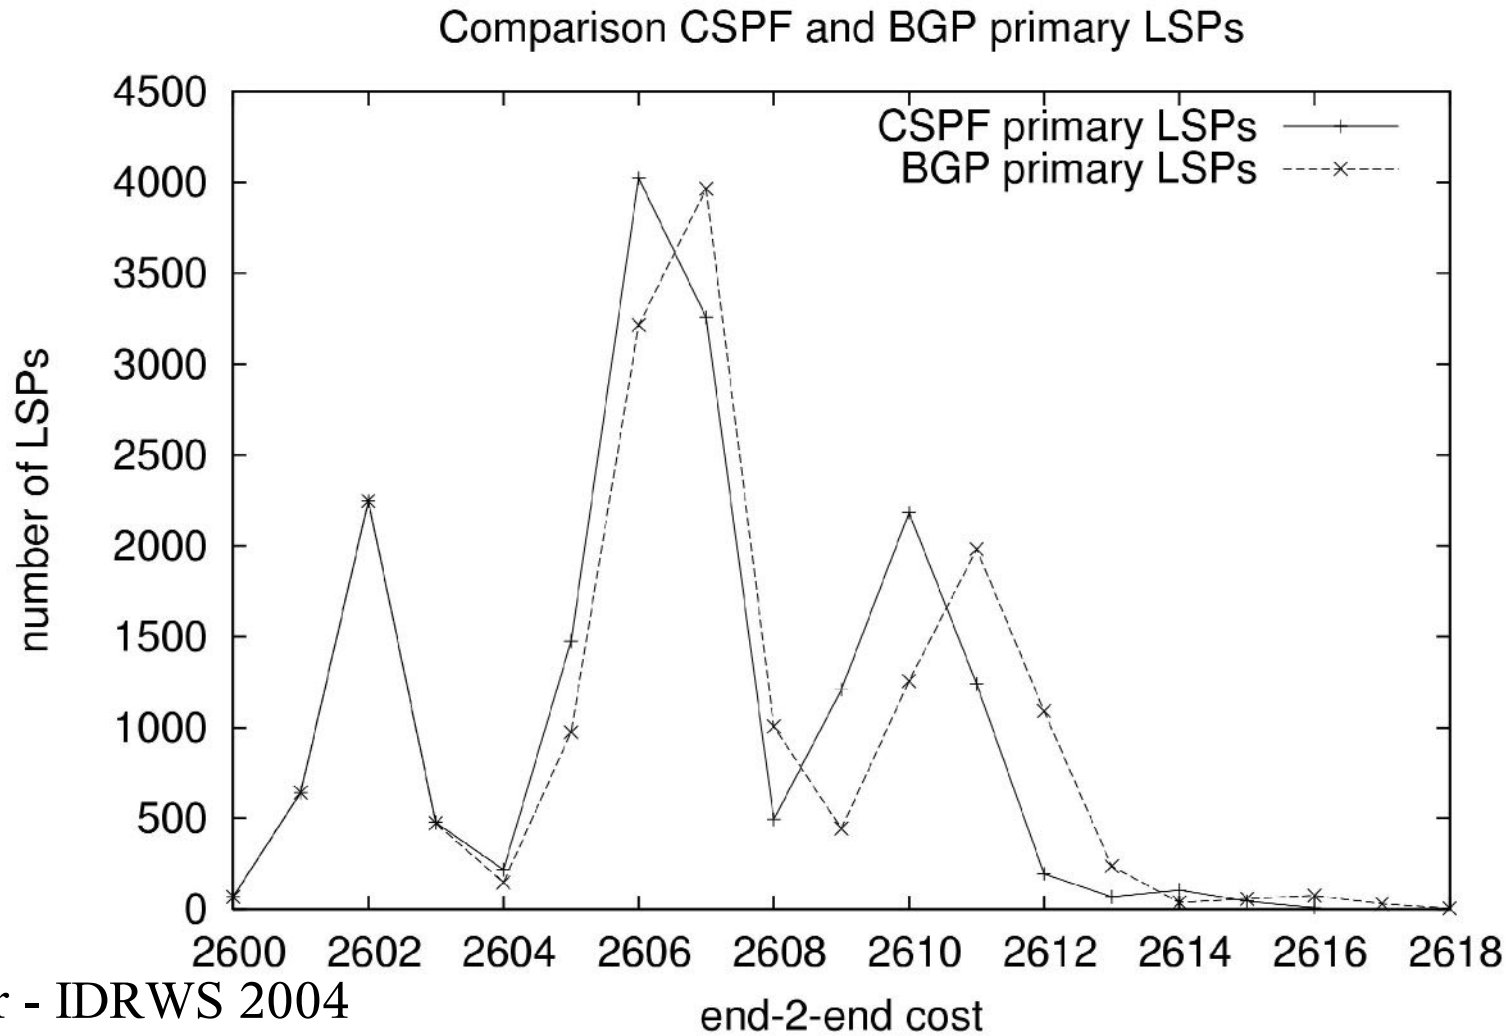
Link	Cost	TE metric
R2-R4	1	100
R3-R4	2	10
R3-R6	2	10

Prefix	NH	AS-path
>130.105/16	R7	3
130.105/16	R4	3
130.105/16	R2	3

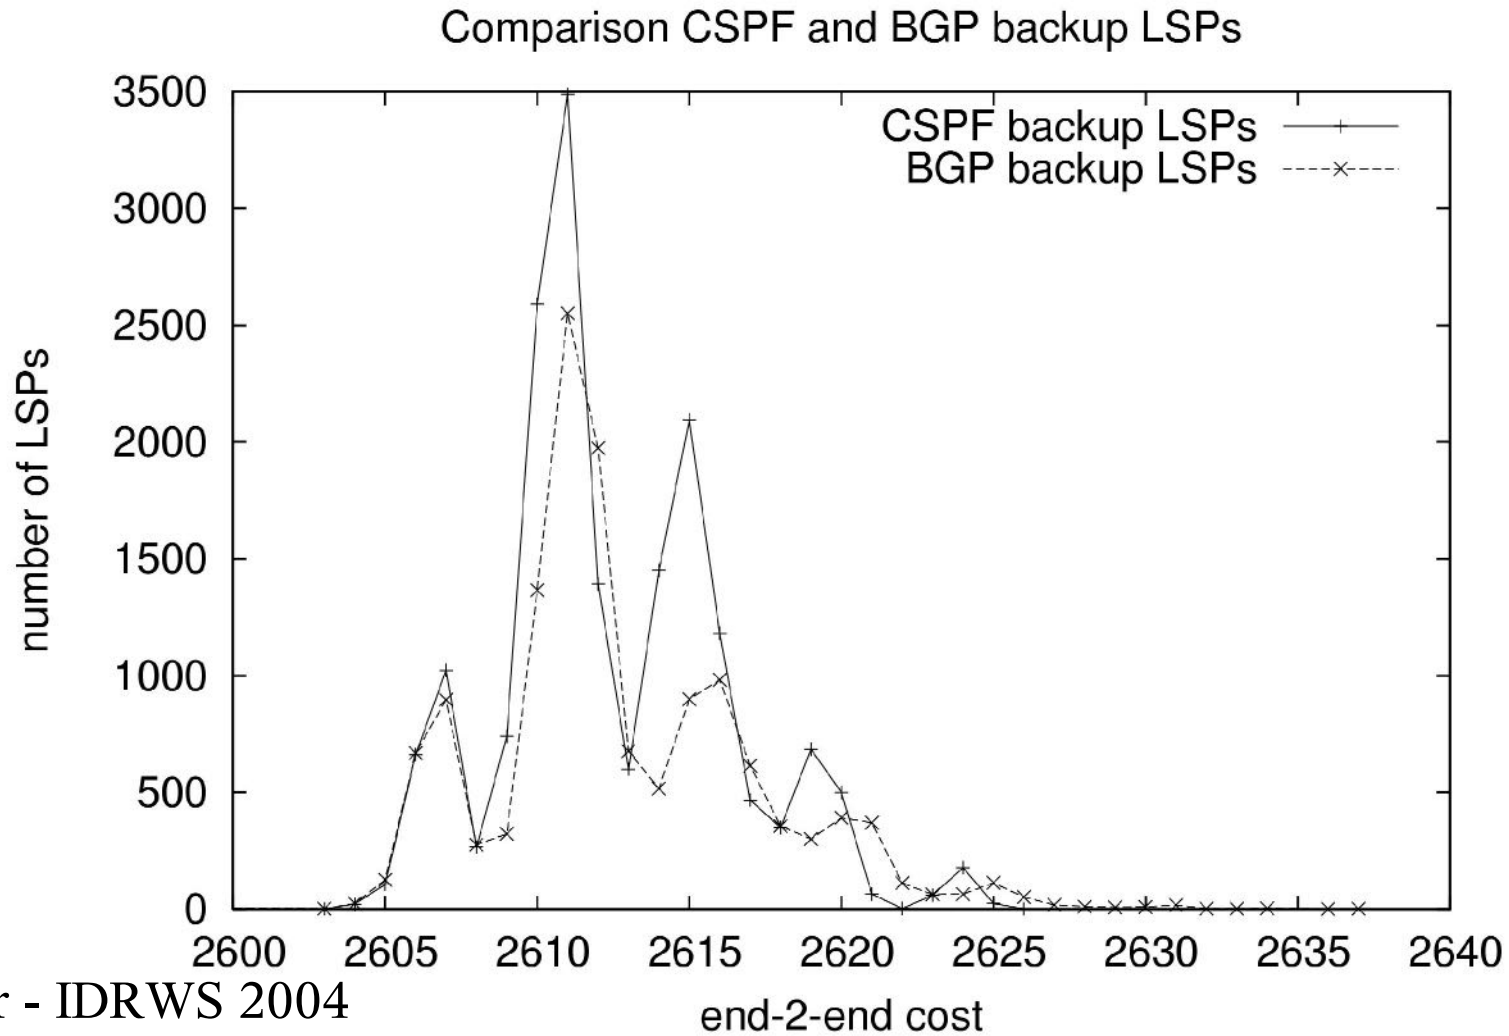
# Preliminary results

- ◆ Topology with :
  - ◆ 20 transits composed of 50 nodes
  - ◆ 190 stubs (all possible combinations of dual-homed stubs)
- ◆ Customer-provider policies between transit and stubs
- ◆ Constraint : node protection
- ◆ Optimise end-2-end cost (ex: delay)
- ◆ Backtracking (crackback) when no path available for the required constraint
- ◆ No incremental establishment of LSPs
  - ◆ available resources are not updated after each LSP establishment

# Preliminary results



# Preliminary results





# Remaining issues

- ◆ Work on heuristics for the choice of alternative next-hops (NH)
- ◆ All possible NH are not necessarily in the Adj-RIB-In of the local router
  - Full-mesh of iBGP session:
    - ◆ All routers only know the best route selected by the other routers in the iBGP mesh
  - Route-Reflectors (RR):
    - ◆ Clients only know the route selected by their RR
    - ◆ The RR should make the choice for its clients
- ◆ Work on link-state inter-AS routing protocols?

# Conclusion

- ◆ **Distributed** disjoint path computation possible based on
  - ◆ Adj-RIB-Ins
  - and
  - ◆ eXclude Route Object (XRO)  
(draft-ietf-ccamp-rsvp-te-exclude-route-01.txt)
- ◆ Applicable for
  - ◆ Protection, load-balancing and TE
    - ◆ The ISPs can choose the AS-path (difficult with BGP)
  - ◆ Establishment of **constrained** inter-AS primary LSP
    - ◆ bandwidth, delay, link affinities constraints