

The Benefits of Locator/ID Separation



AGAVE: A liGhtweight Approach for Viable End-to-end IP-based QoS Services

Joint Work:

L. Iannone, O. Bonaventure, B. Quoitin, D. Saucez, C. De Launois


Roadmap

- **Motivation**
- **Loc/ID Separation**
 - How does it work
- **Benefits**
 - Improved QoS
 - FIB Reduction
- **Costs**


Roadmap

- **Motivation**
- **Loc/ID Separation**
 - How does it work
- **Benefits**
 - Improved QoS
 - FIB Reduction
- **Costs**

Multihoming => Path diversity




Multihoming => Path diversity

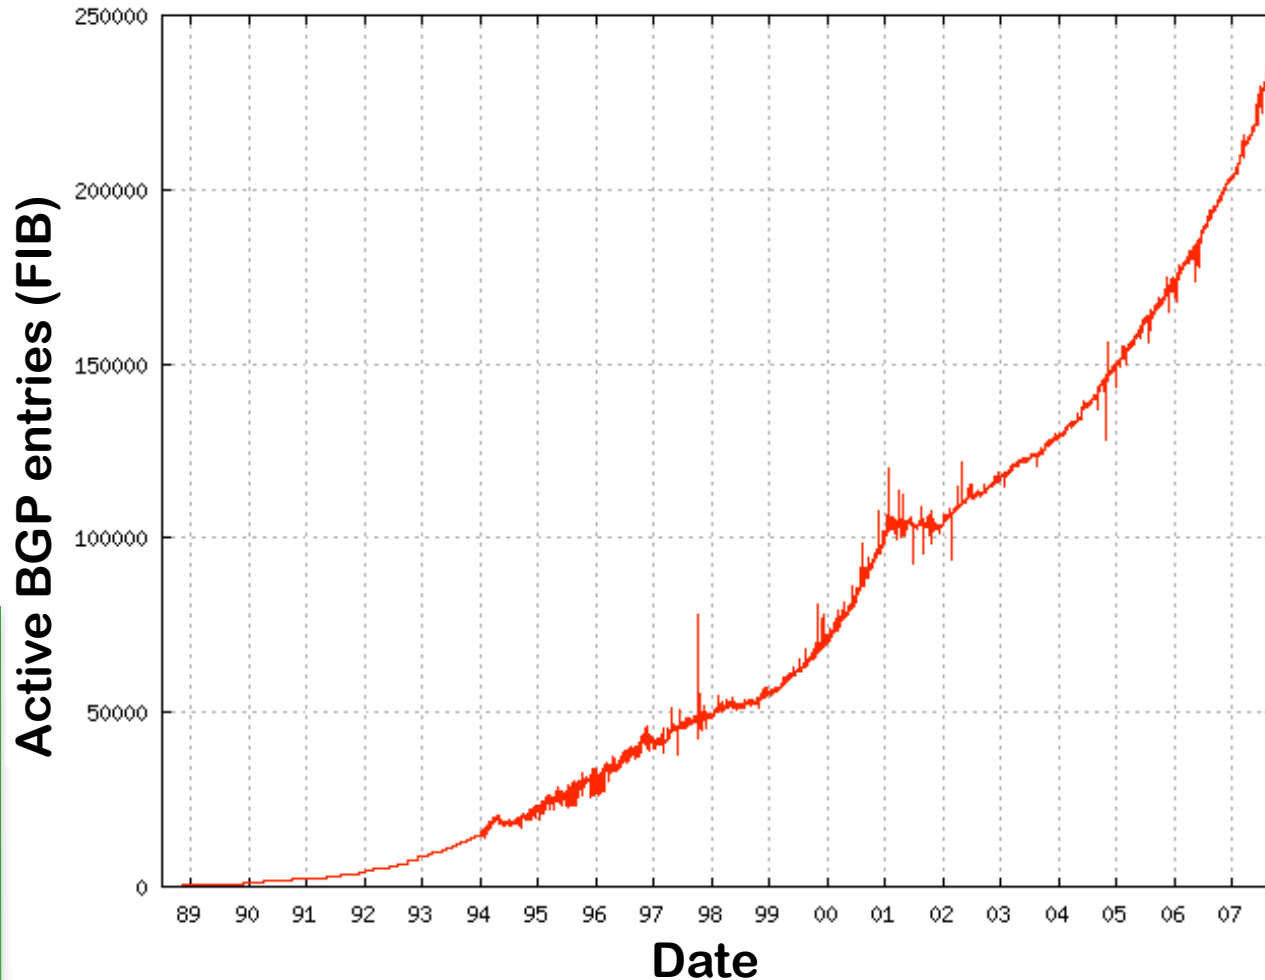


Multihoming => Path diversity

Path-Vector nature of BGP does not allow full usage of Path Diversity



BGP table growth



Main causes:


- PI (Provider Independent) prefix assignment
- Multi-homing
- Traffic-Engineering
- ...

01-Jul-89 to 21-Sep-07 Source: <http://www.cidr-report.org>

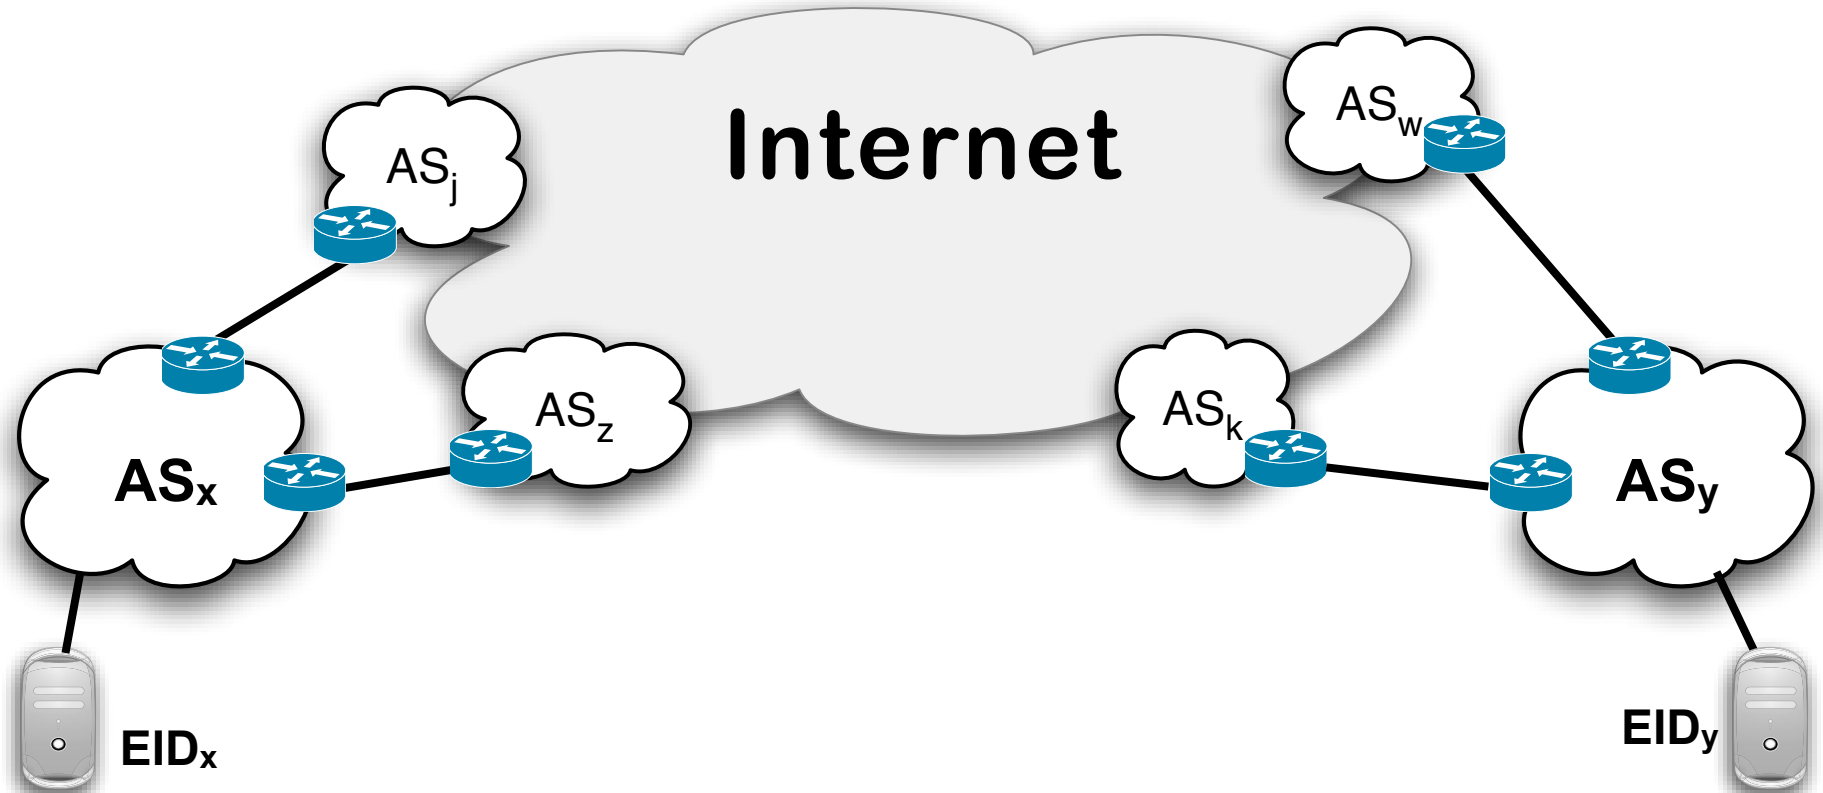
Roadmap

- Motivation
- **Loc/ID Separation**
 - How does it work
- Benefits
 - Improved QoS
 - FIB Reduction
- Costs

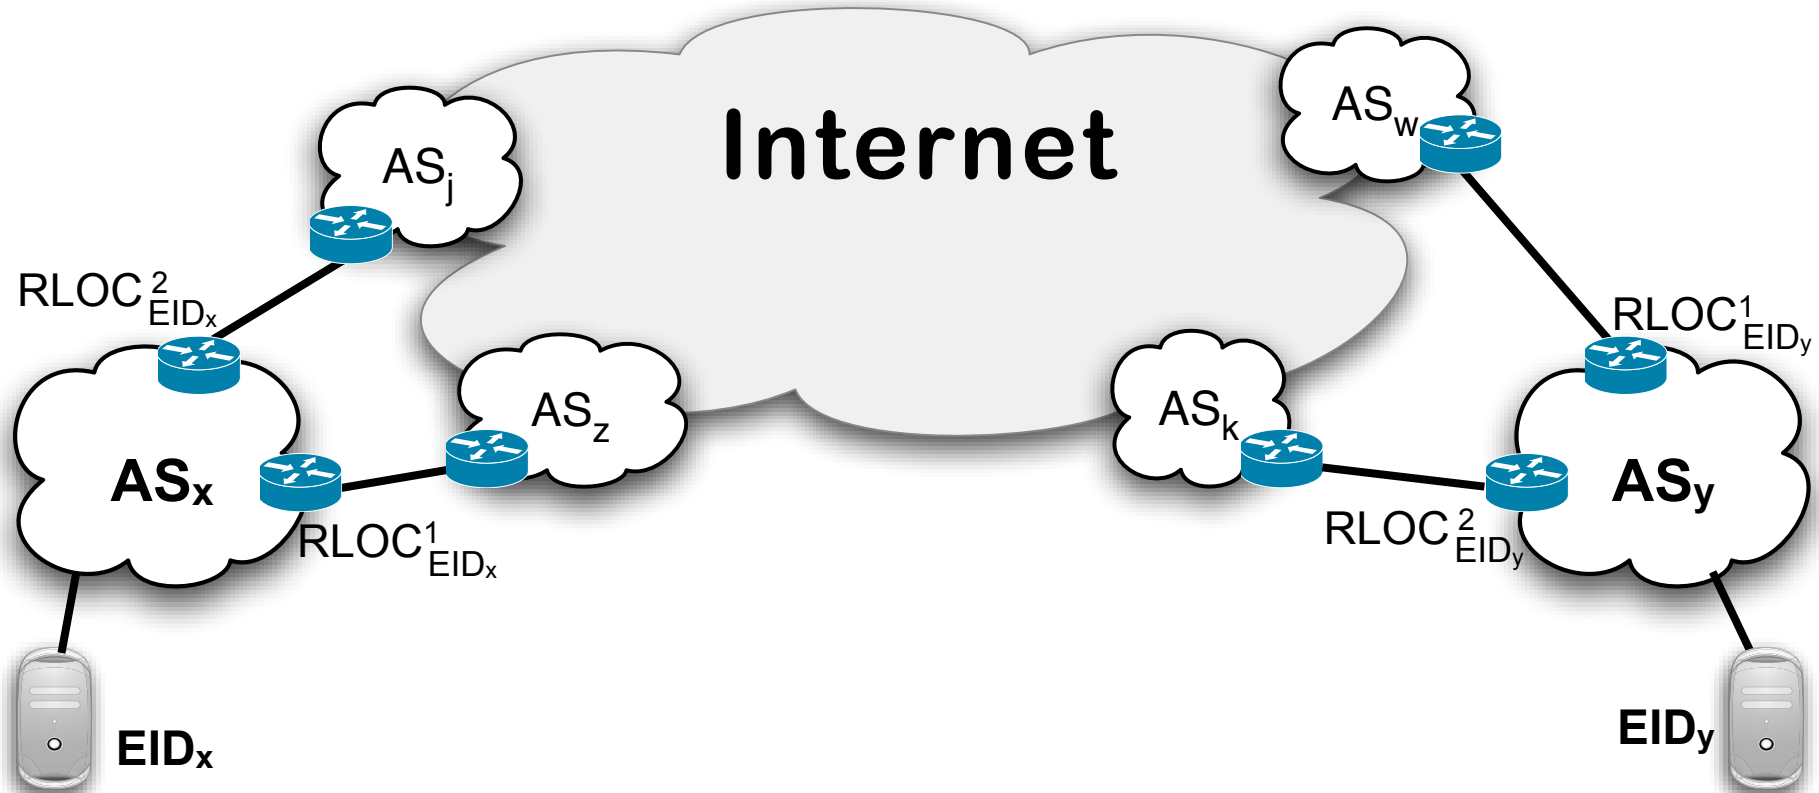
Separating Locators and IDs



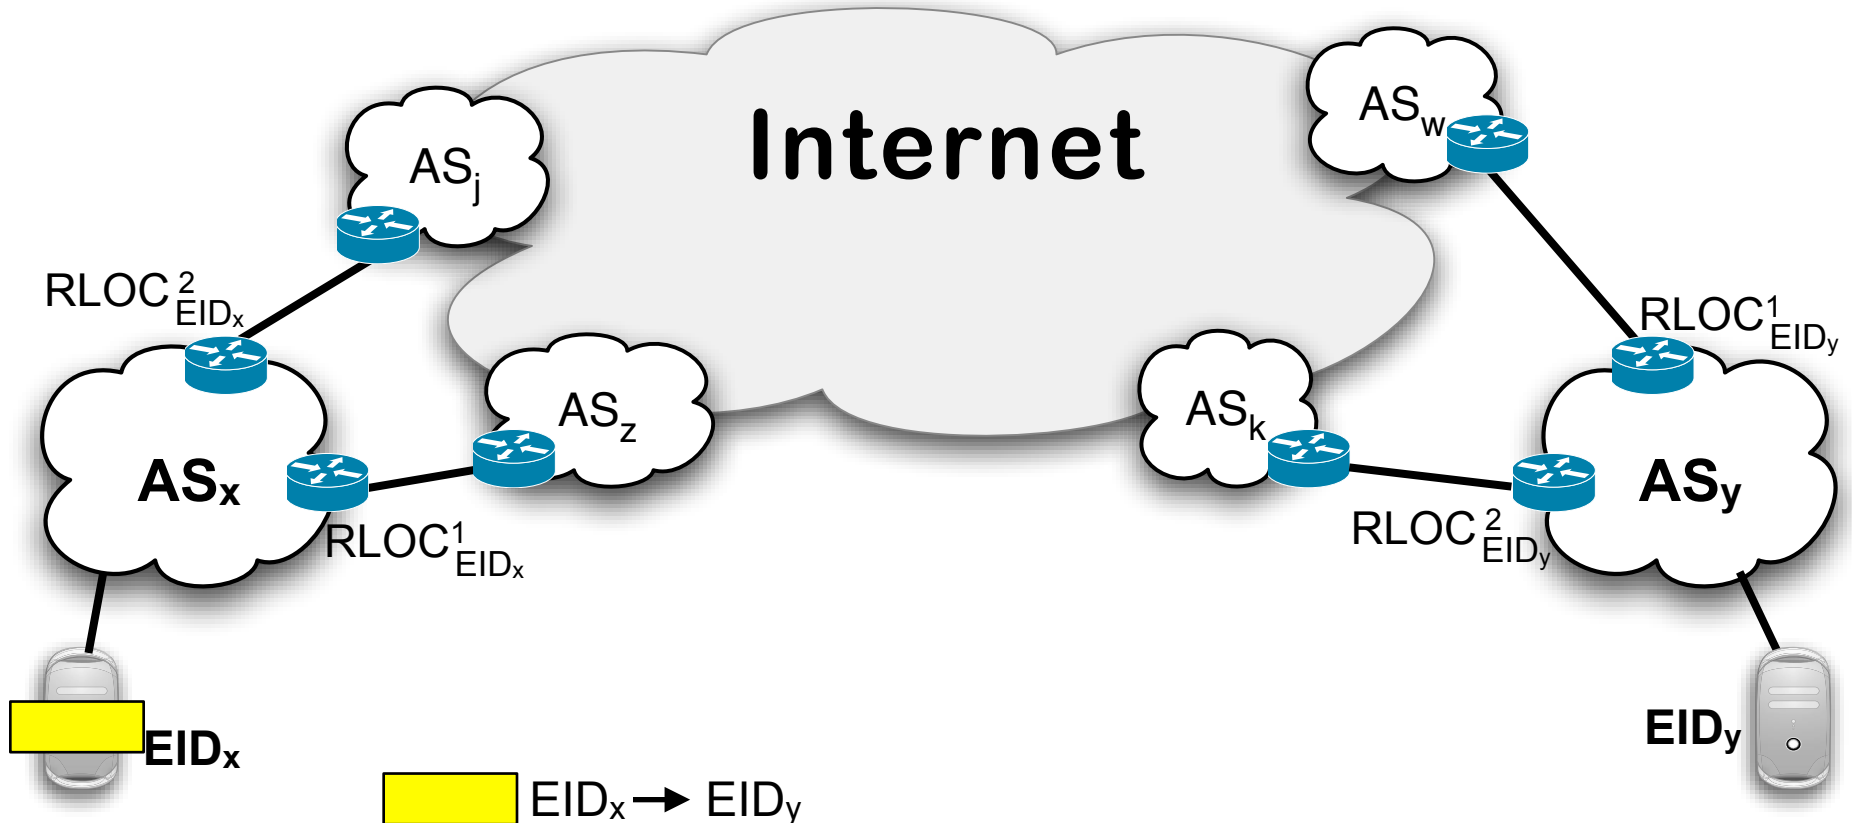
Separating Locators and IDs



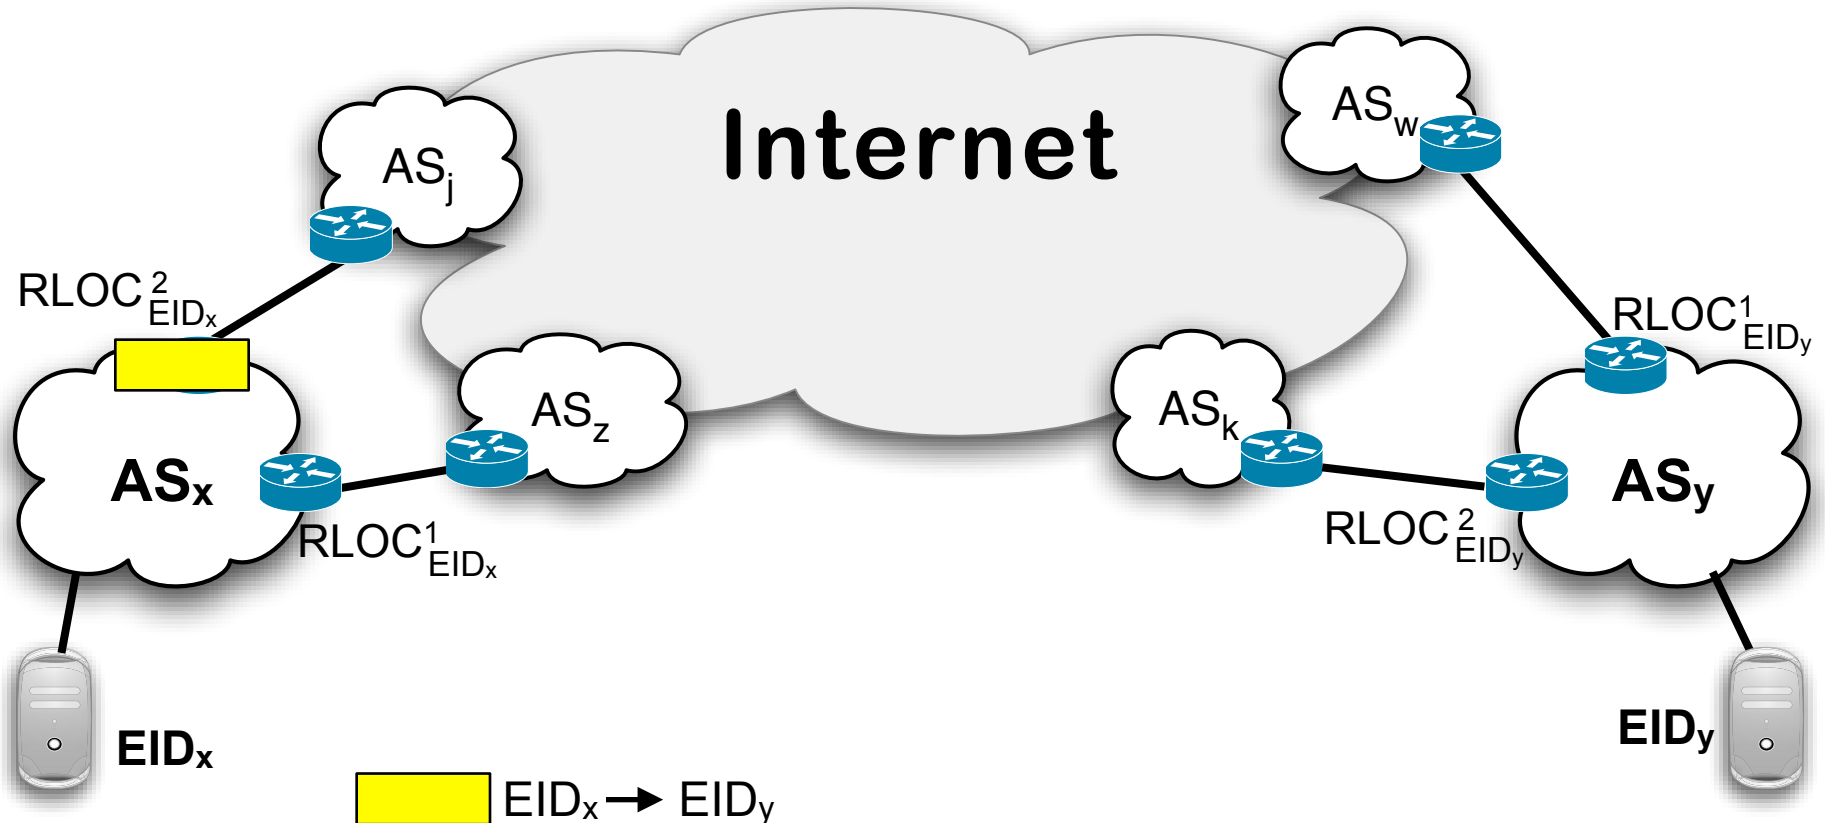
Separating Locators and IDs



Separating Locators and IDs

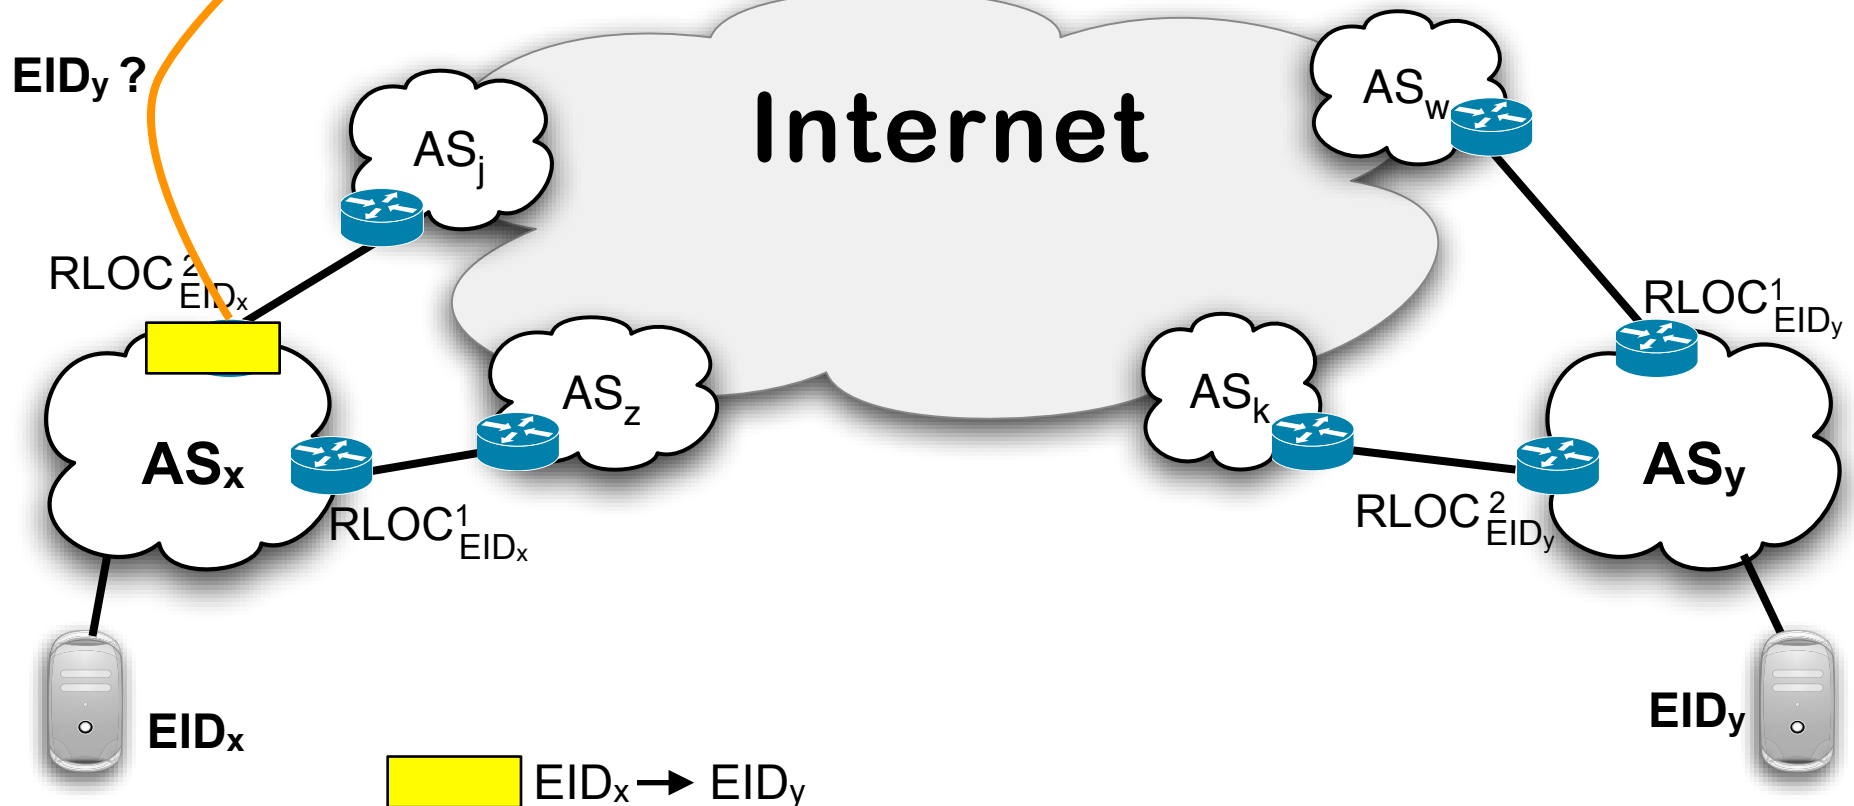


Separating Locators and IDs



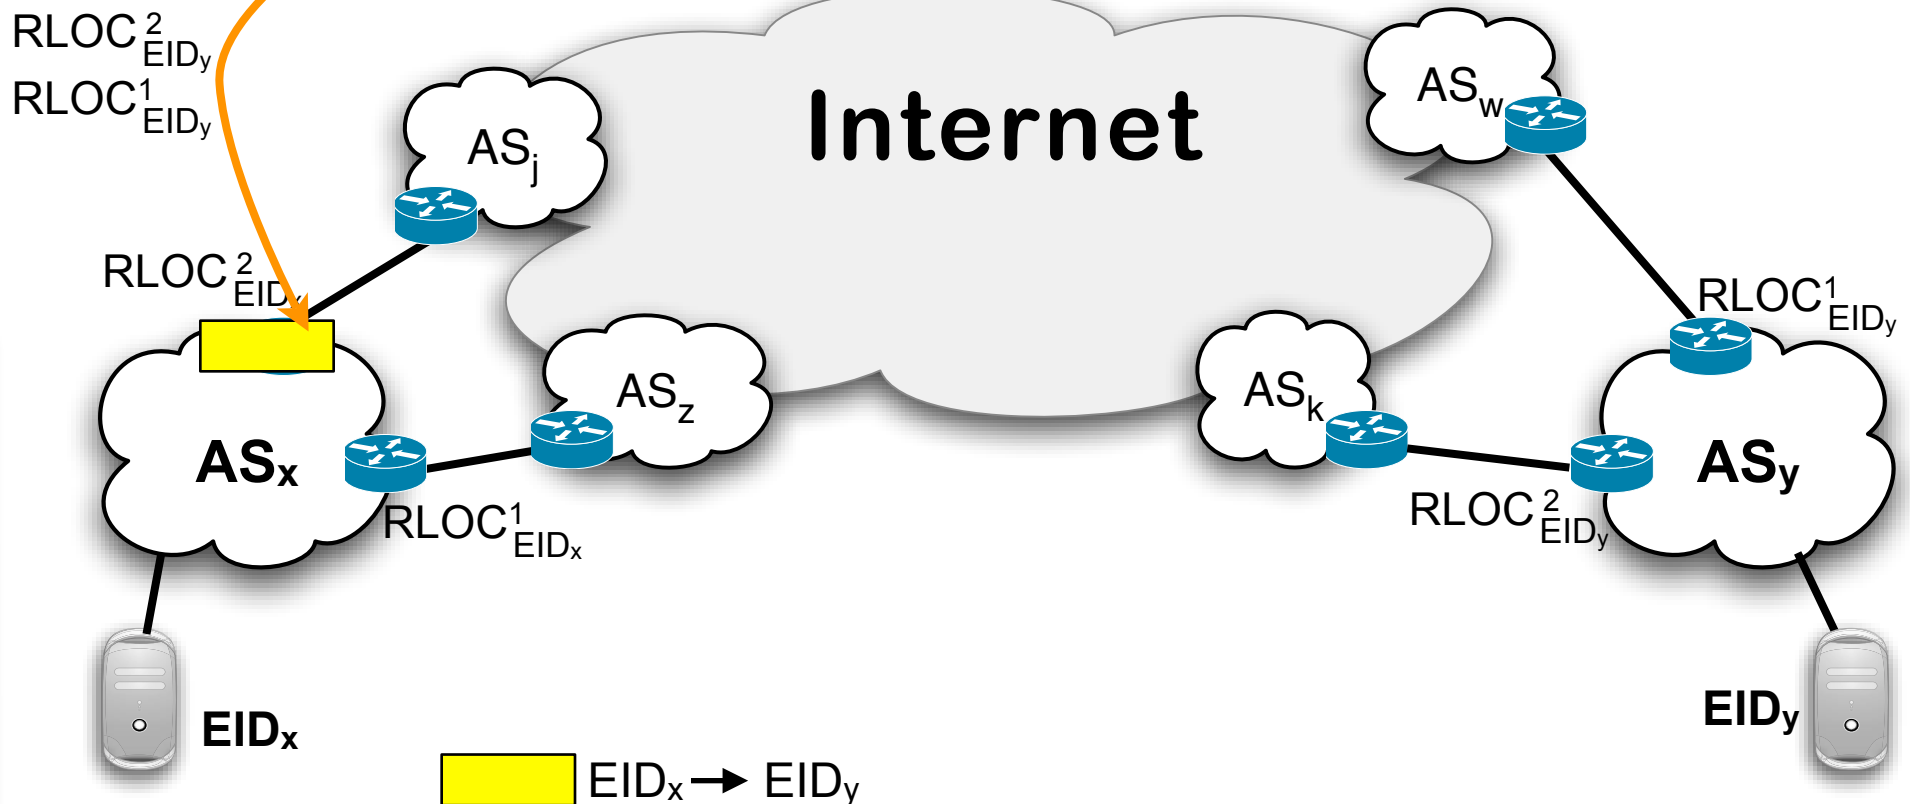
Separating Locators and IDs

Mapping Distribution System

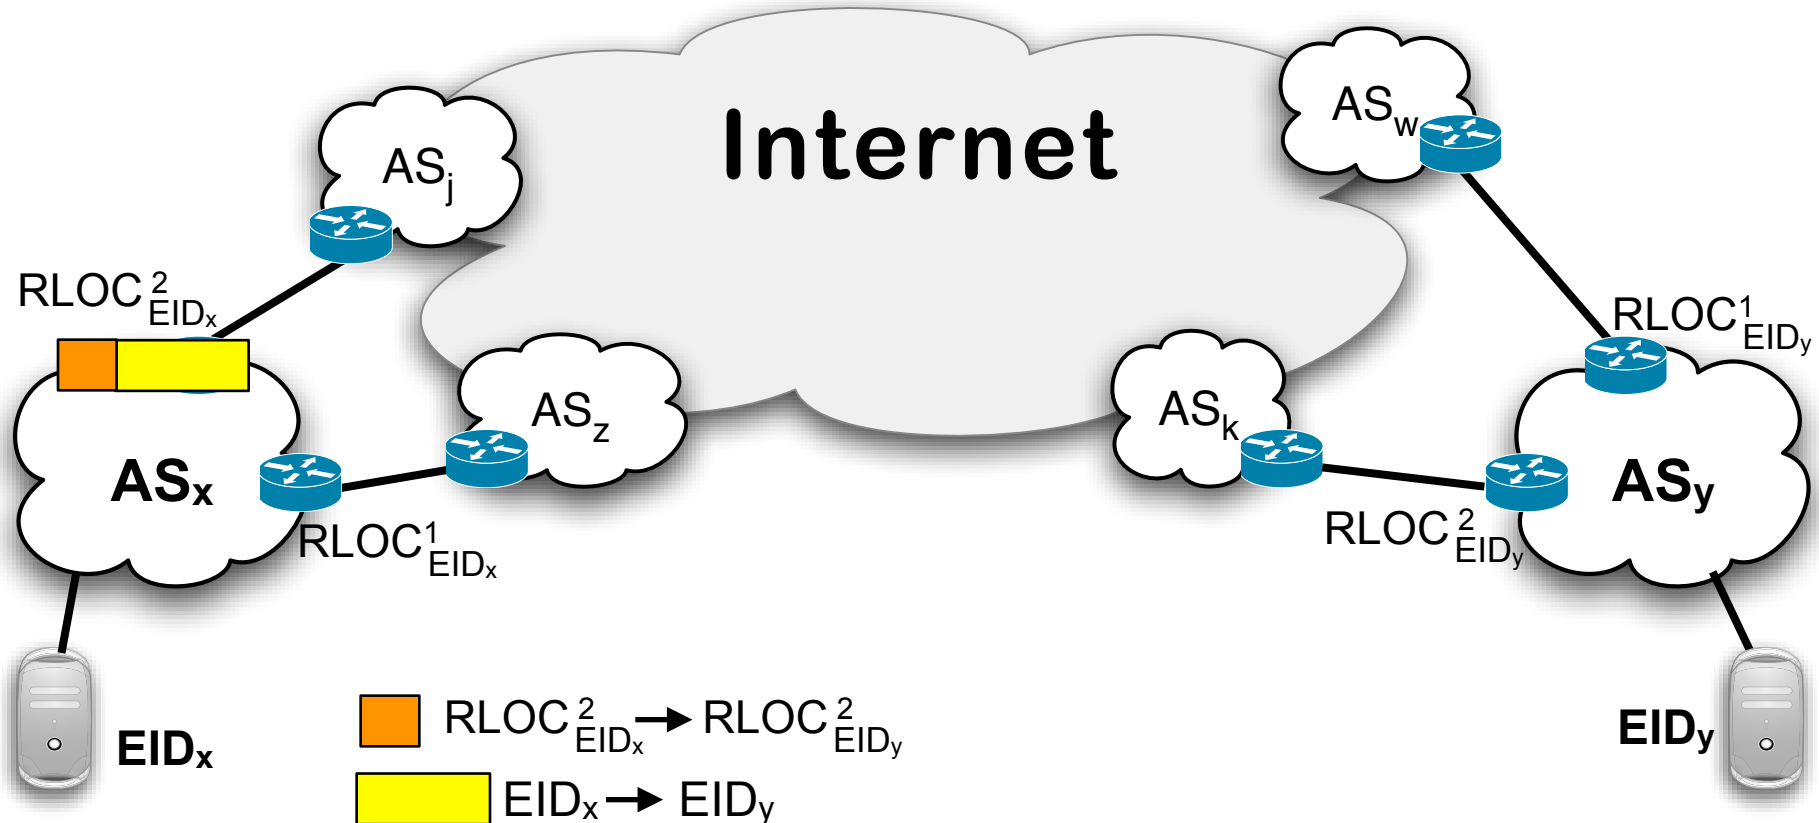


Separating Locators and IDs

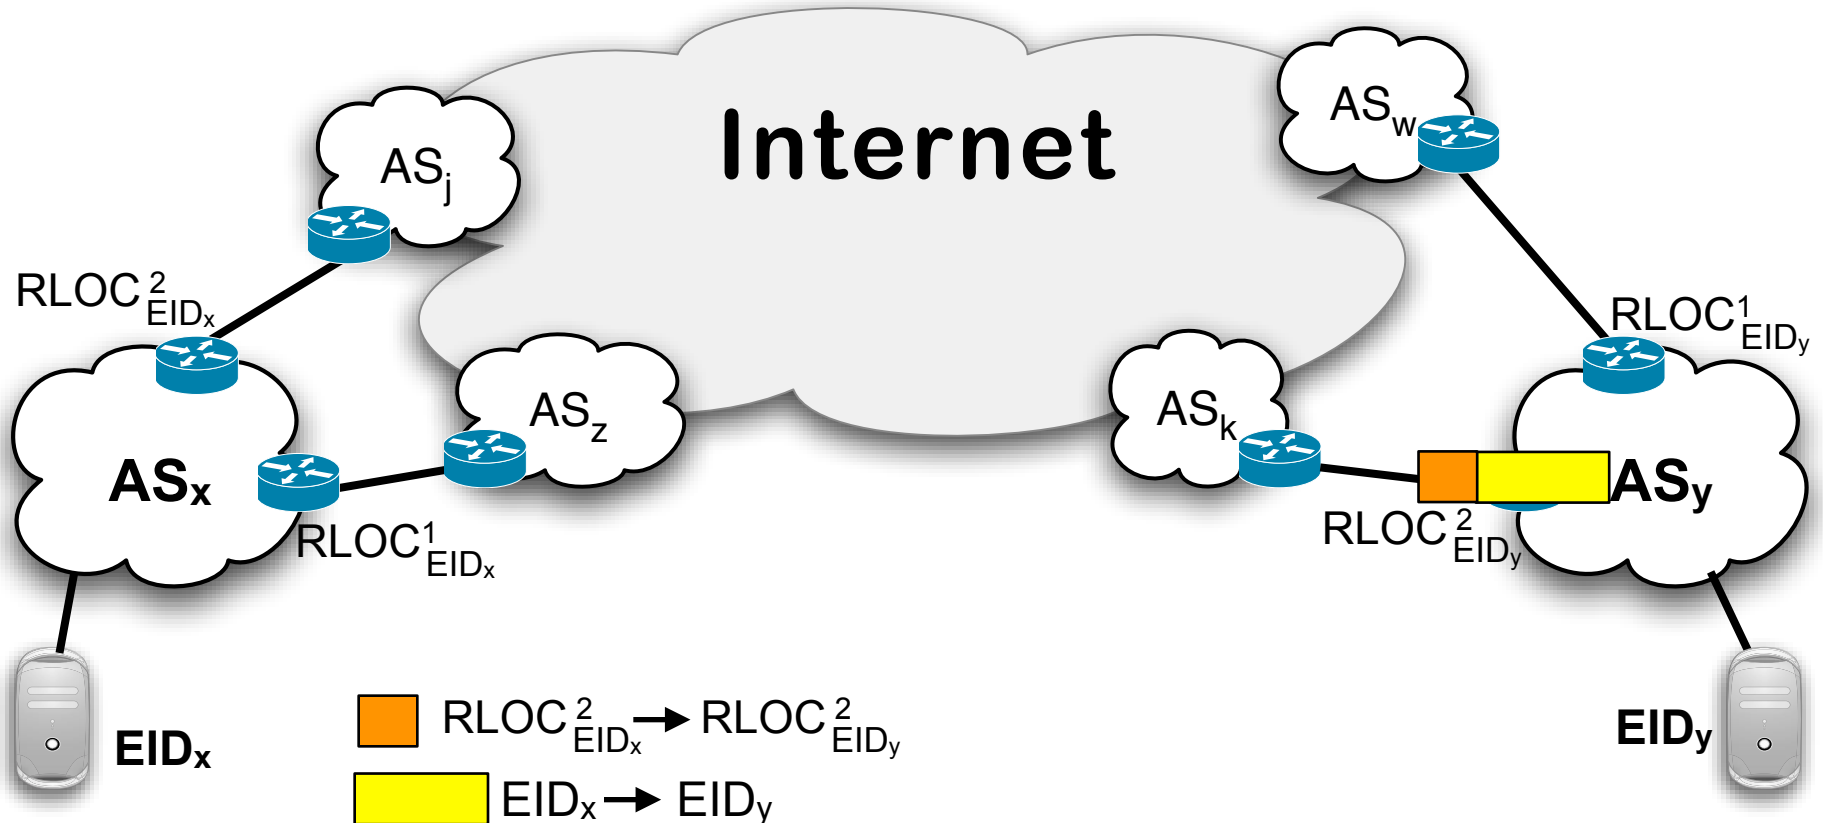
Mapping Distribution System



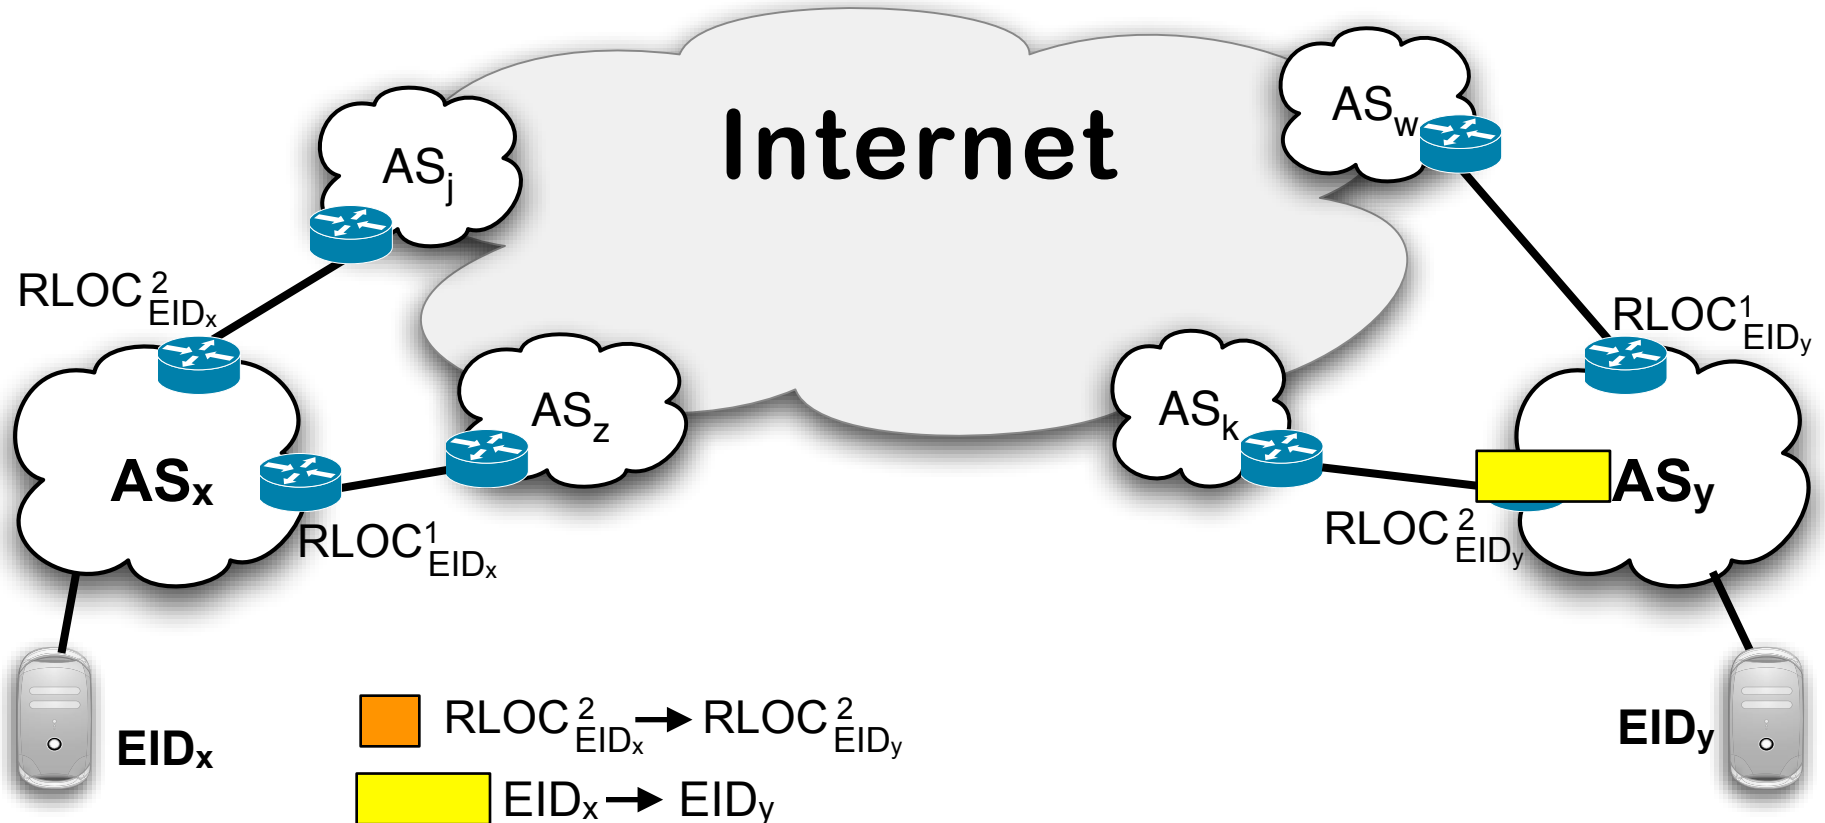
Separating Locators and IDs



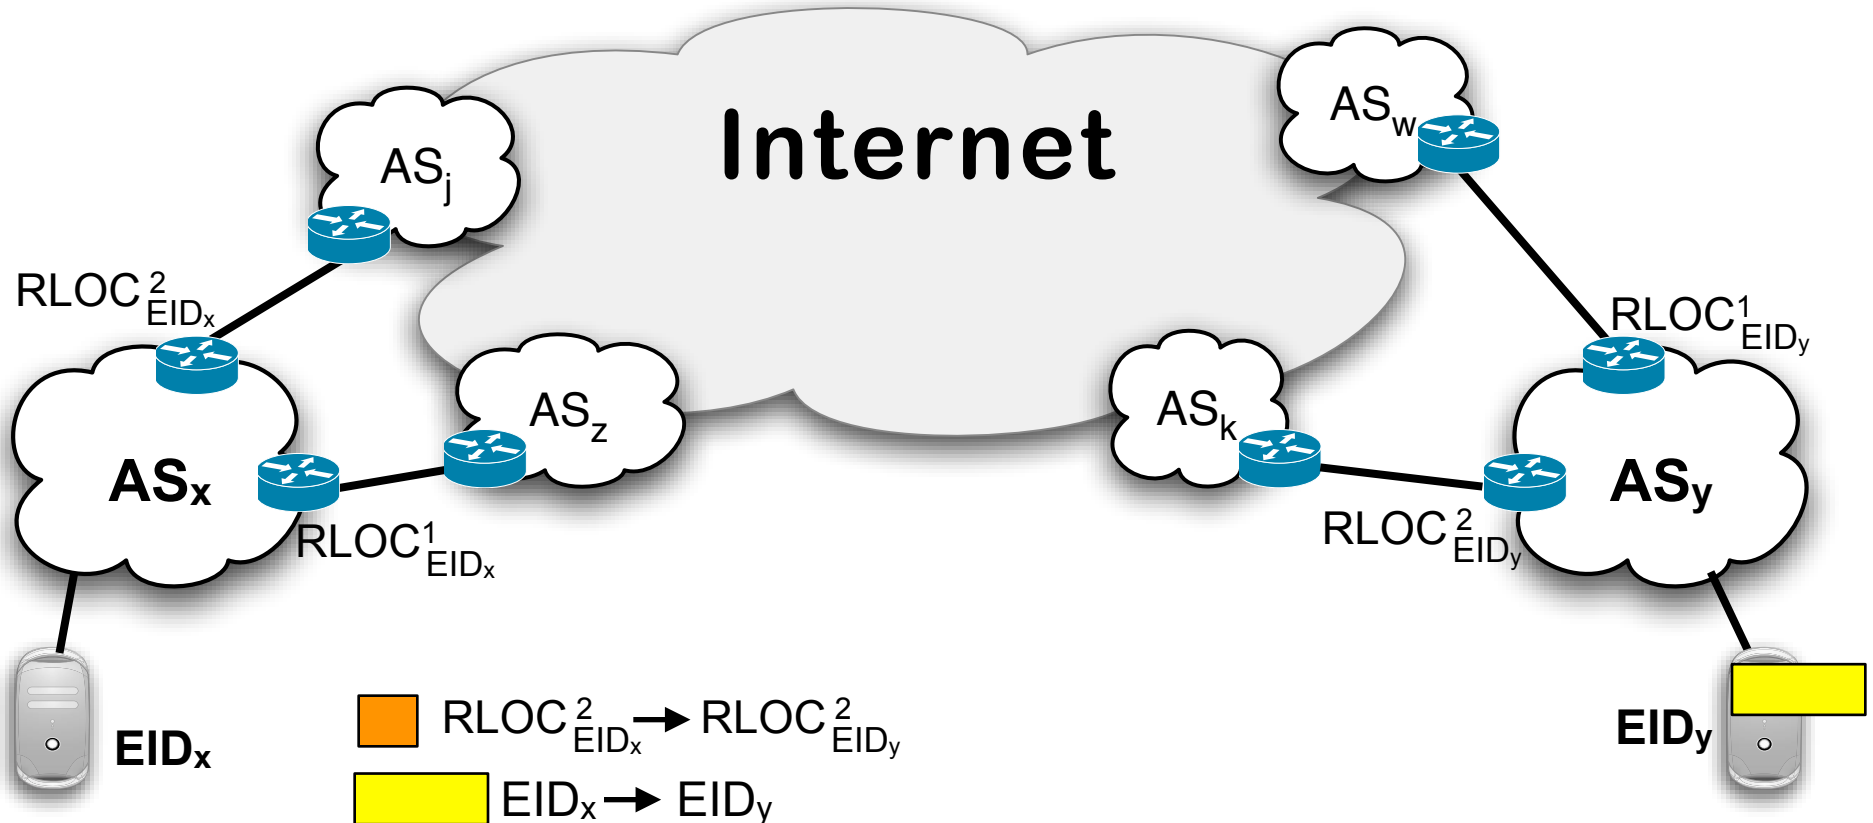
Separating Locators and IDs



Separating Locators and IDs



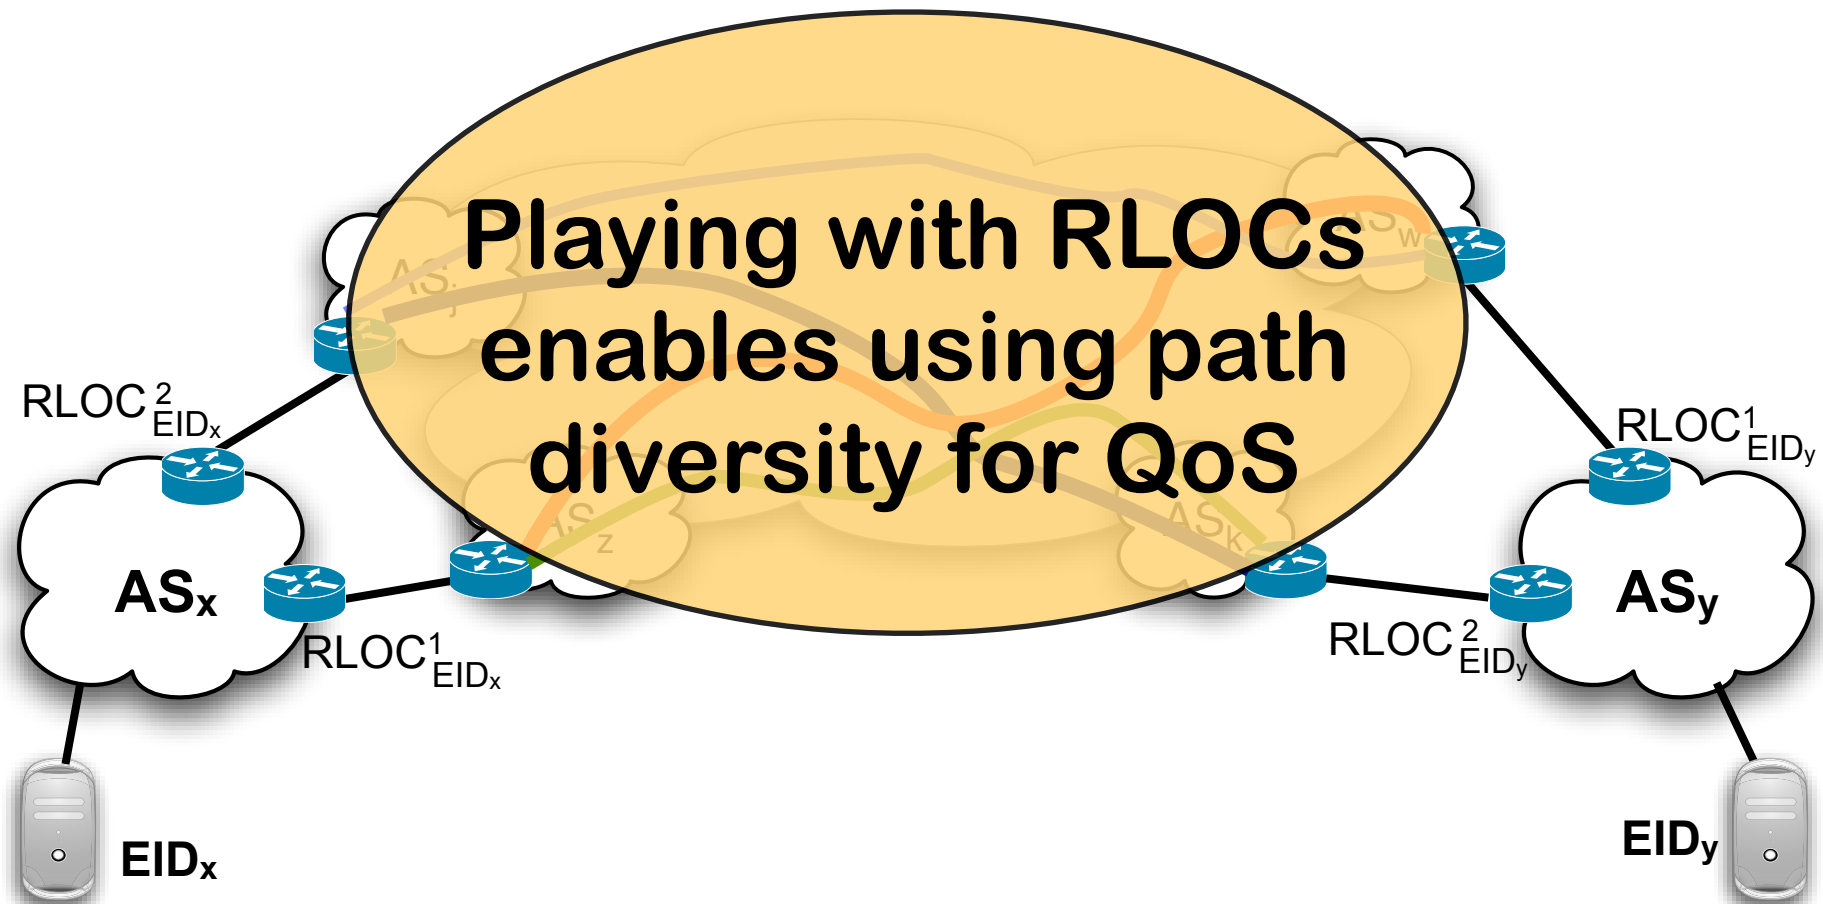
Separating Locators and IDs




Roadmap

- Motivation
- Loc/ID Separation
 - How does it work
- **Benefits**
 - Improved QoS
 - FIB Reduction
- Costs

Advantage 2: Path diversity




Internet Path-diversity



Simulations of Multihomed ASes

Dataset	Routes	Peers	Pairs	M-h Stubs	M-h Prefixes
Routeviews	5,750,380	32	496	6,402	29,575
RIPE NCC00	1,641,618	11	55	6,247	29,934

Exploiting Path-diversity



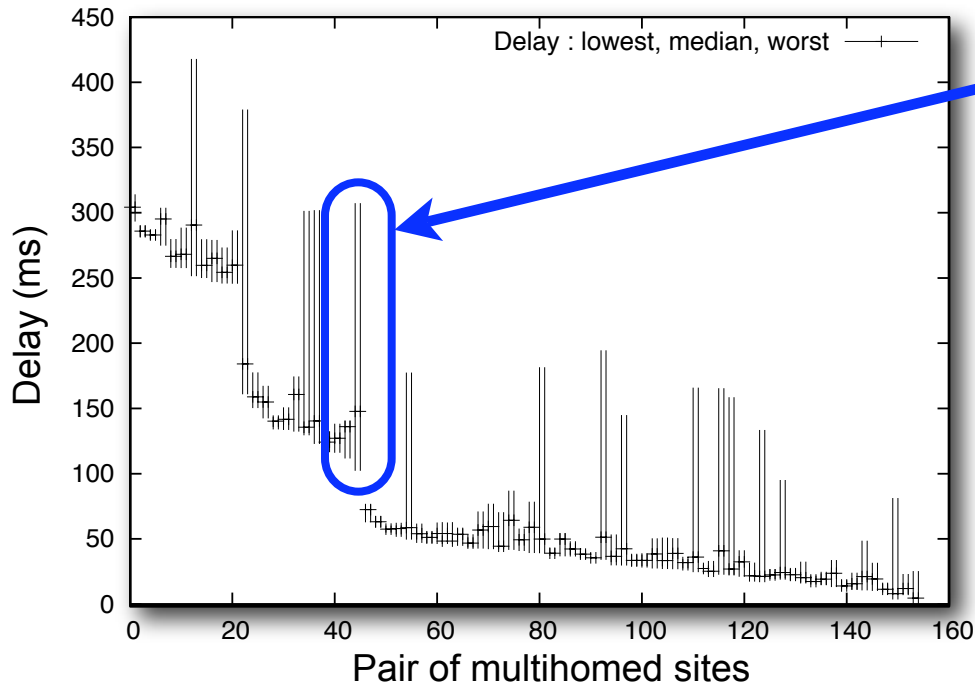
Simulations of 13 multi-homed ASes (156 pairs)

- Based on RIPE NCC Traffic Measurements Service
- Delays averaged using the Vivaldi algorithm

How to take advantage of path diversity:

- TSC (Tunnel Service Controller) to select end-points

Exploiting Path-diversity



Pair 45:

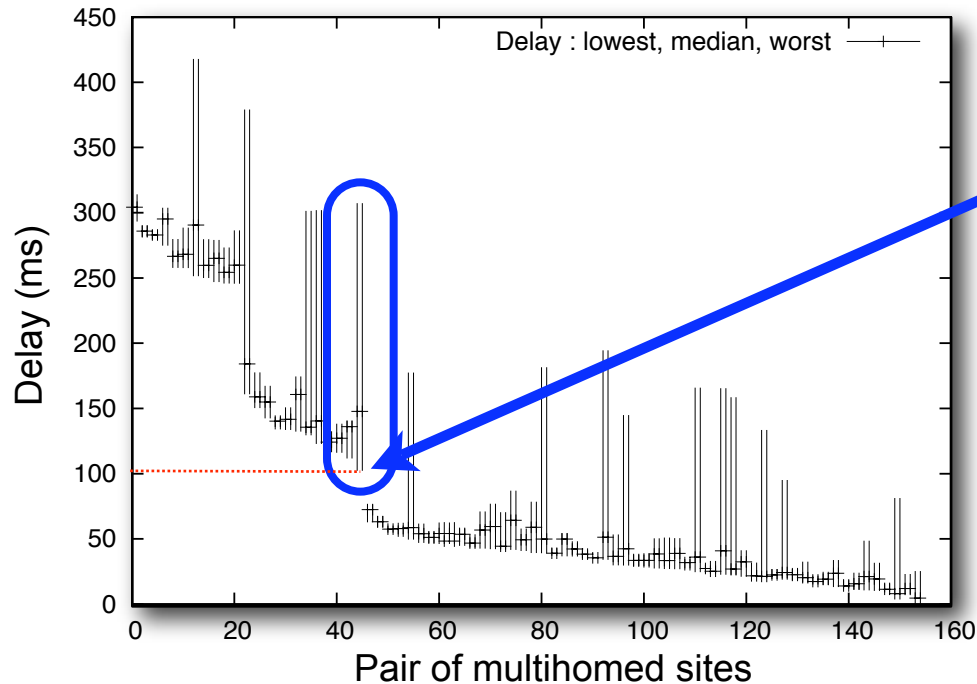
Simulations of 13 multi-homed ASes (156 pairs)

- Based on RIPE NCC Traffic Measurements Service
- Delays averaged using the Vivaldi algorithm

How to take advantage of path diversity:

- TSC (Tunnel Service Controller) to select end-points

Exploiting Path-diversity



Pair 45:

Best Path: ~100 ms


Simulations of 13 multi-homed ASes (156 pairs)

- Based on RIPE NCC Traffic Measurements Service
- Delays averaged using the Vivaldi algorithm

How to take advantage of path diversity:

- TSC (Tunnel Service Controller) to select end-points

Exploiting Path-diversity



Pair 45:

Best Path: ~100 ms

Worst Path: ~300 ms


Simulations of 13 multi-homed ASes (156 pairs)

- Based on RIPE NCC Traffic Measurements Service
- Delays averaged using the Vivaldi algorithm

How to take advantage of path diversity:

- TSC (Tunnel Service Controller) to select end-points

Exploiting Path-diversity



Pair 45:

Best Path: ~100 ms

Worst Path: ~300 ms

Median: ~ 150 ms

Simulations of 13 multi-homed ASes (156 pairs)

- Based on RIPE NCC Traffic Measurements Service
- Delays averaged using the Vivaldi algorithm

How to take advantage of path diversity:

- TSC (Tunnel Service Controller) to select end-points

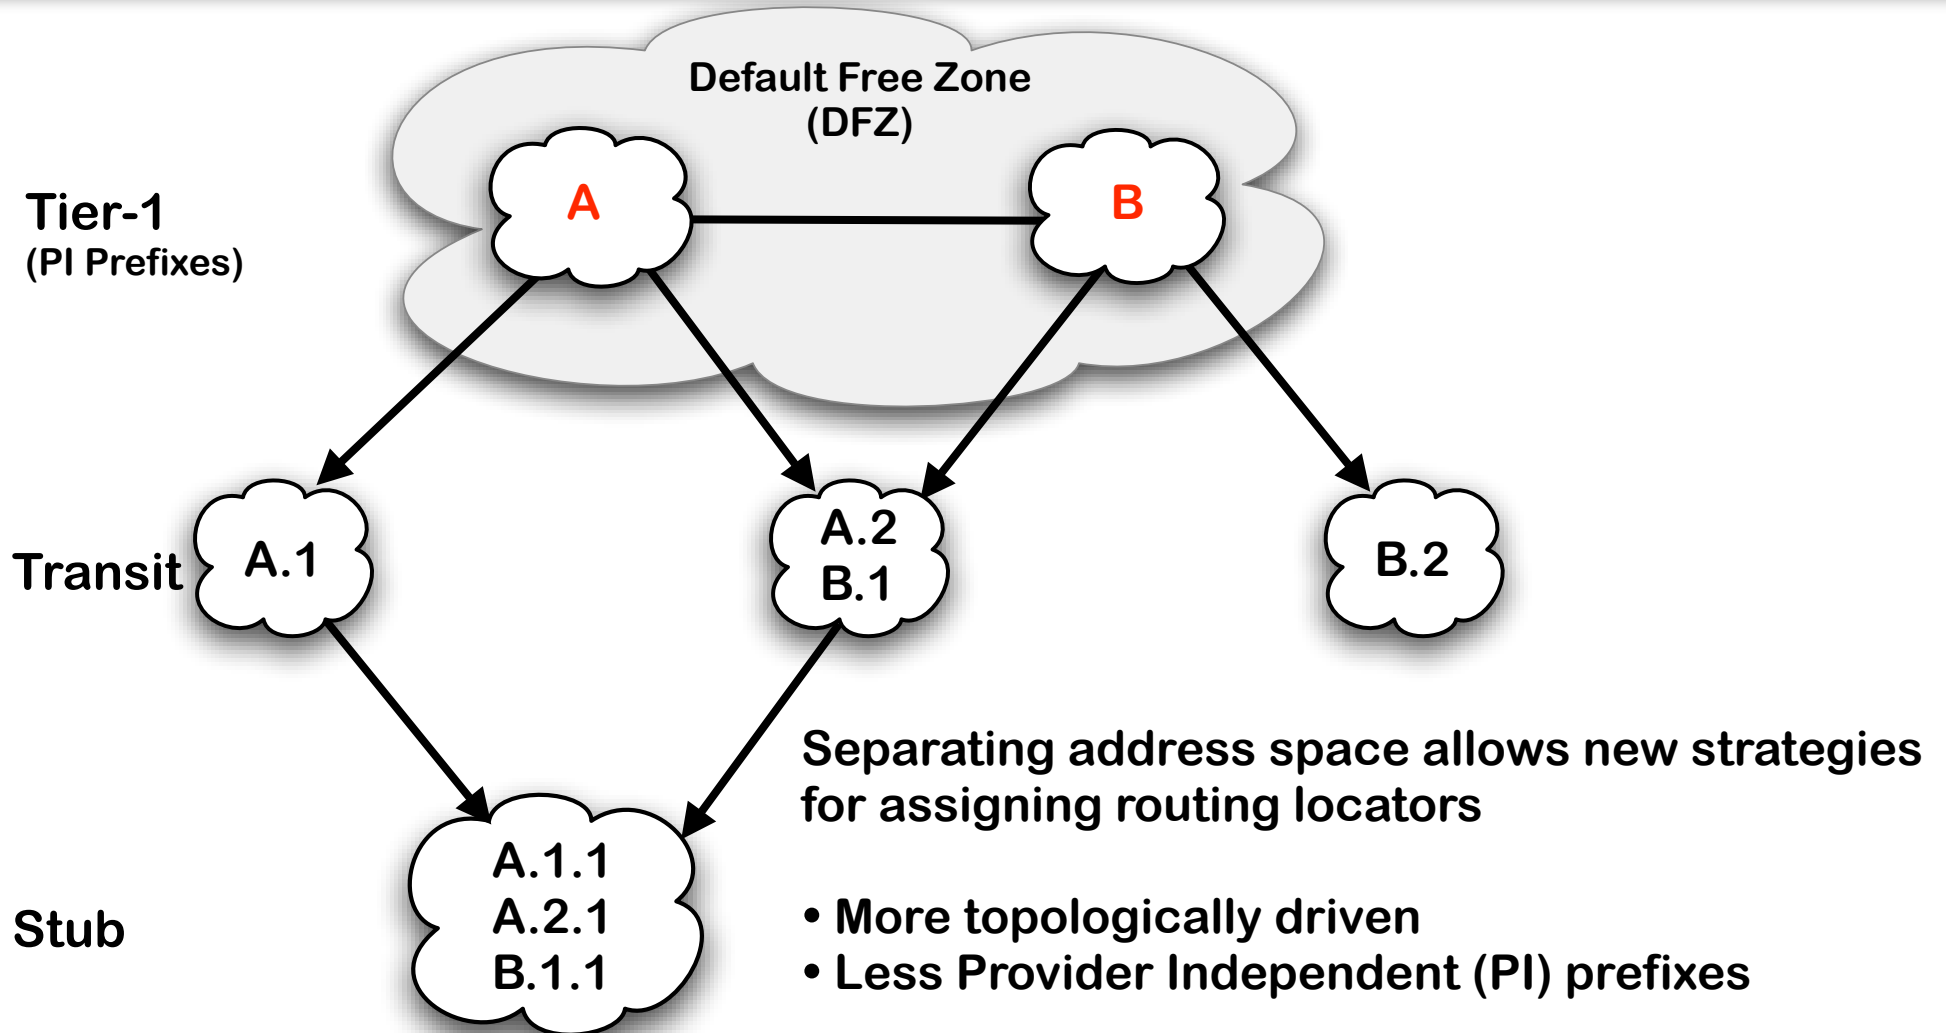
Roadmap

- Motivation
- Loc/ID Separation
 - How does it work
- **Benefits**
 - Improved QoS
 - **FIB Reduction**
- Costs

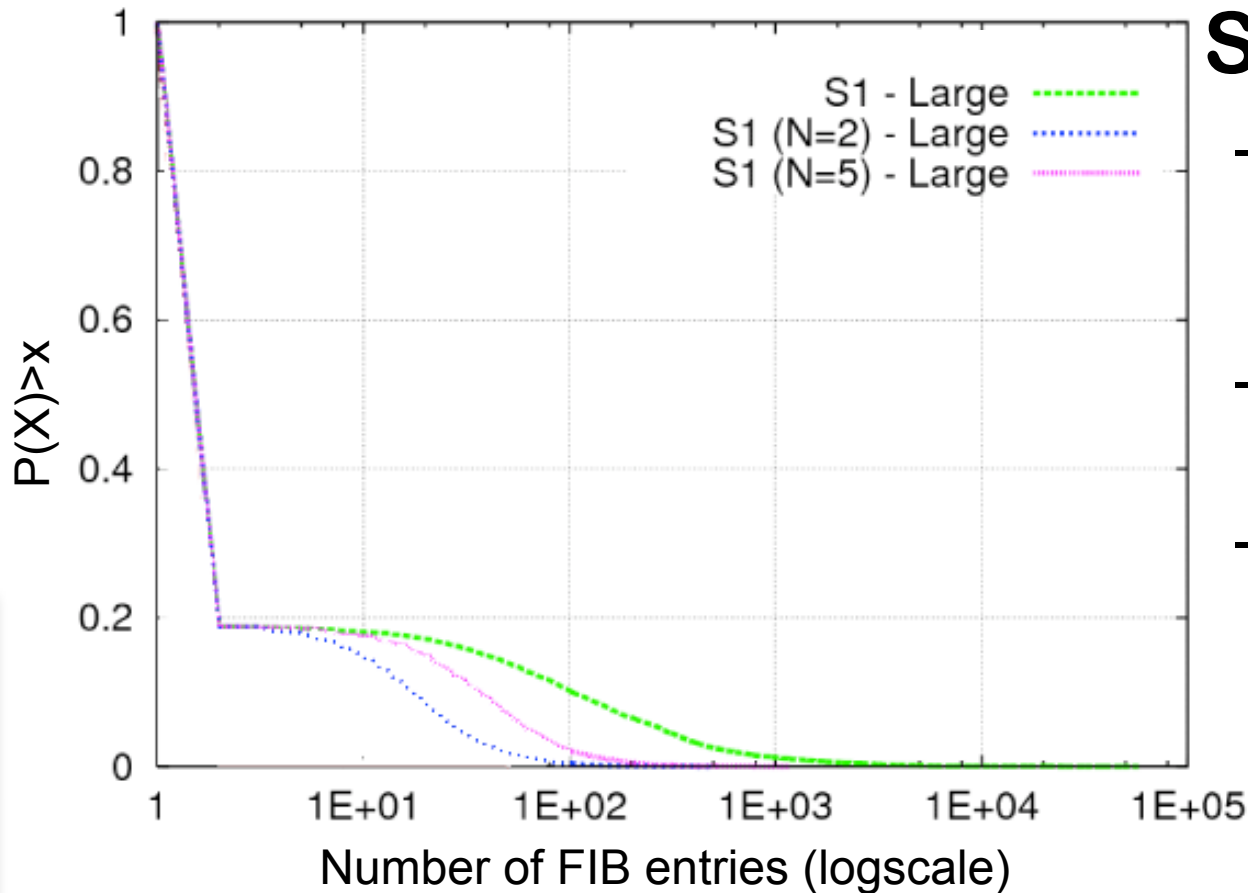
Advantage 1: shrinking the FIBs



Shrinking the FIB (I)




Shrinking the FIB (II)



Simulations

- Synthetic topologies generated with GHITTE
- Hierarchical with business relationship
- 14965 ASes

Shrinking the FIB (II)



Simulations

- Synthetic topologies generated with GHITTE
- Hierarchical with business relationship
- 14965 ASes


Cool, but....

**There Ain't No Such Thing As A
Free Lunch!**

Roadmap


- Motivation
- Loc/ID Separation
 - How does it work
- Benefits
 - Improved QoS
 - FIB Reduction
- **Costs**

Netflow data collection

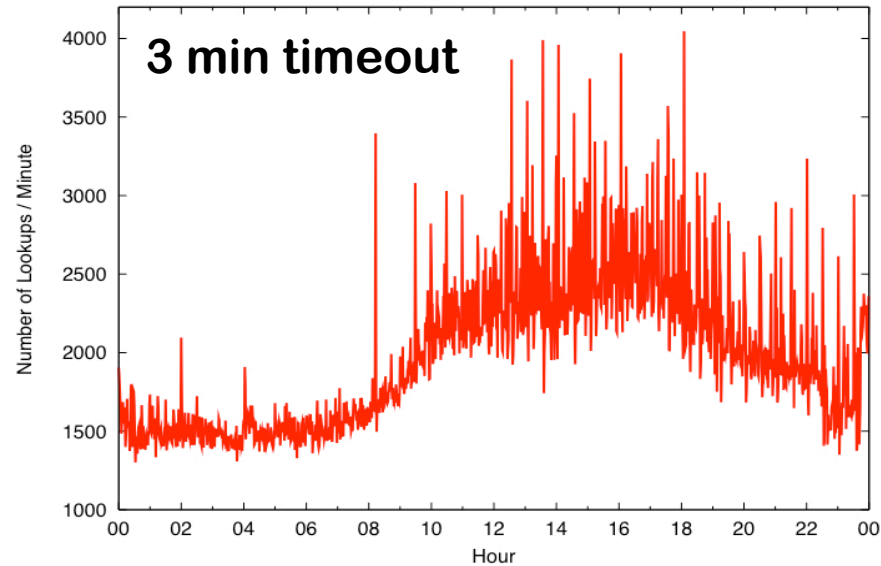


- Full Netflow (v7) on border router
- 1 Gigabit link to Belnet
- ~10000 users (/16 prefix block)
- Analysis: flow-tools + custom software
- /BGP Granularity of mappings
 - iPlane data set

Mappings' Cache Size



Lookups



Timeout	Period	1 RLOC	2 RLOCs	3 RLOCs
3 min.	Night	10.17 kbps	12.17 kbps	14.17 kbps
	Day	34.97 kbps	41.85 kbps	48.74 kbps
30 min.	Night	2.04 kbps	2.44 kbps	2.84 kbps
	Day	8.95 kbps	10.71 kbps	12.47 kbps
300 min.	Night	0.163 kbps	0.195 kbps	0.227 kbps
	Day	2.68 kbps	3.21 kbps	3.74 kbps

Conclusions

- **Loc/ID Separation approach allows:**
 - Exploiting Path Diversity
 - Reduce FIB size

- **Loc/ID Separation costs:**
 - Cache
 - Lookups & Tunneling overhead

Thanks

?? || /* */

<http://inl.info.ucl.ac.be>

<http://www.ist-agave.org>