Interdomain Traffic Engineering in a Loc/Id Separation Context

INM'08
October 19, 2008

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

D. Saucez, B. Donnet, L. Iannone, O. Bonaventure

Université catholique de Louvain

Path Selection Problem?

IPv4 vs IPv6 Dual Stack (DS)

- Dual stack hosts/routers will exist for many years
  - IPv4 and IPv6 performance (e.g., reliability) are not equivalent [1]
- How to select the best stack ?
  - always prefer IPv6? RFC 3484 static selection?

=> determine the best path among several:

$$\{<s_{IPv4},d_{IPv4}>, <s_{IPv6},d_{IPv6}>, <s_{IPv4},d_{IPv6}>, <s_{IPv6},d_{IPv4}>,\}$$

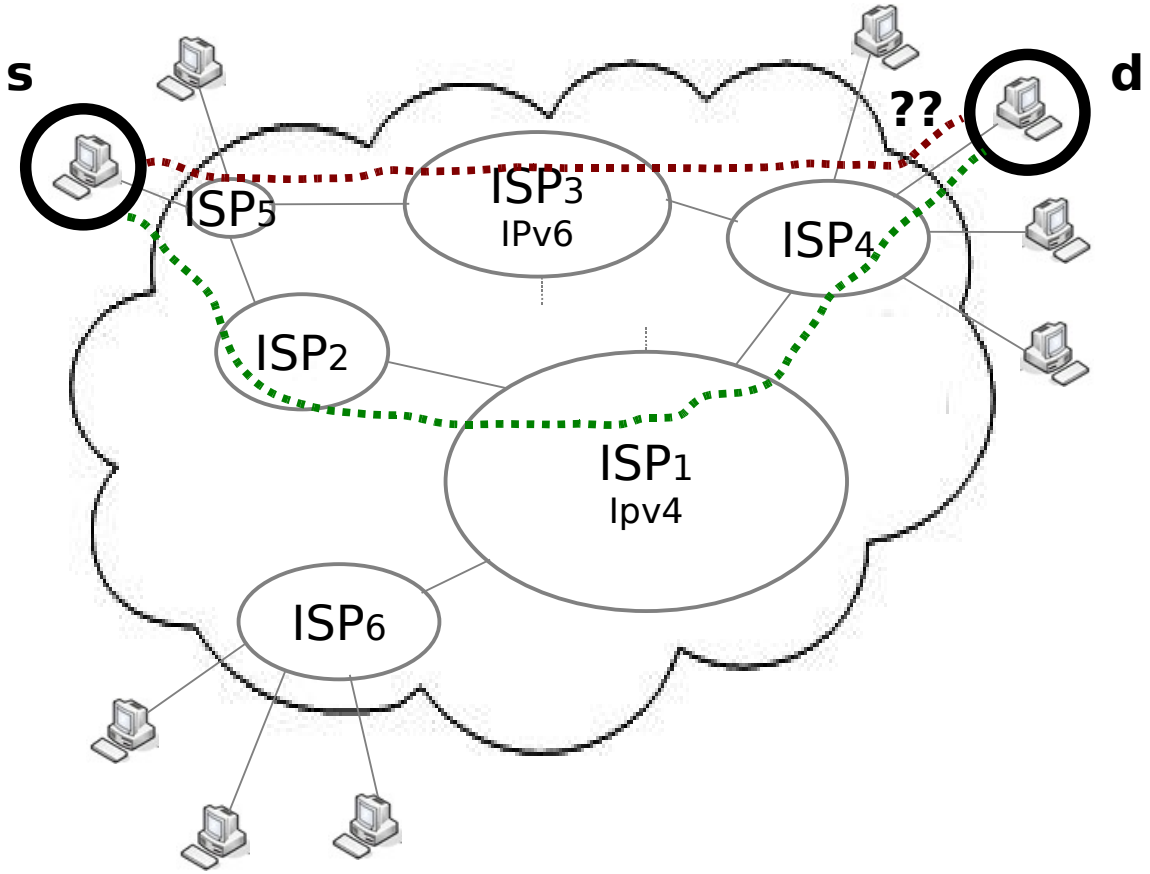

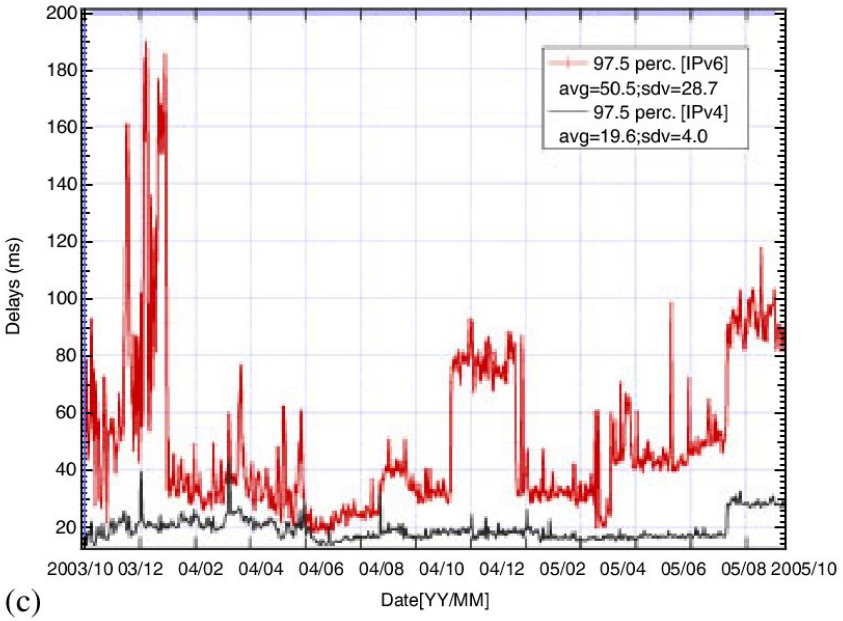

[1] X. Zhou et al., IPv6 delay and loss performance evolution, IJCS 2008

Multi-Homing (MH)

- Multi-homing implies choice among multiple feasible paths with much varying properties [2]
  - AS-based MH: how to select the best path (ISP-based objectives)
  - Host-based MH: how to select the best path (customer-based objectives)

=> determine the best path among several:

$$\{<s_1,d_1>, \ldots ,<s_1,d_n>, <s_2,d_1>, \ldots , <s_m,d_n>\}$$

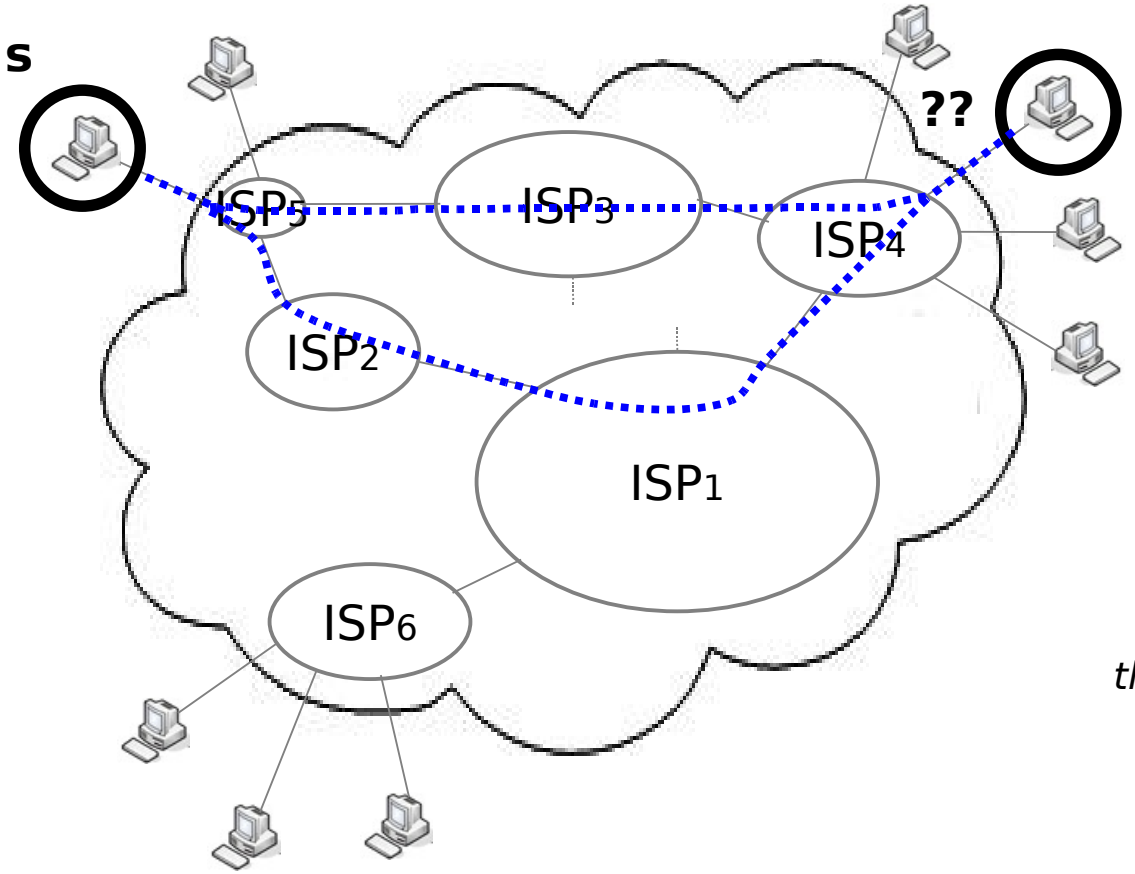

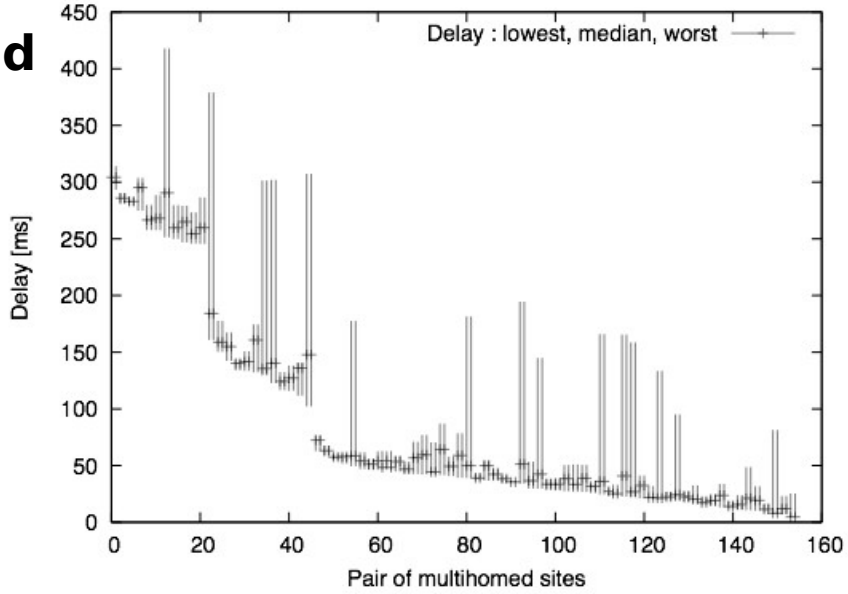

[2] B. Quoitin et al., Evaluating the Benefits of the Locator/Identifier Separation, MobiArch 2007

Server replicas

- How to select the best replicas
  - within set $\{d_a,d_b,d_c,d_d\}$
  - per source: s_1, s_2, s_3

=> determine the best replicaS among several:

$$\{<s_i,d_a>, <s_i,d_b>, <s_i,d_c>, <s_i,d_d>\} \forall i$$

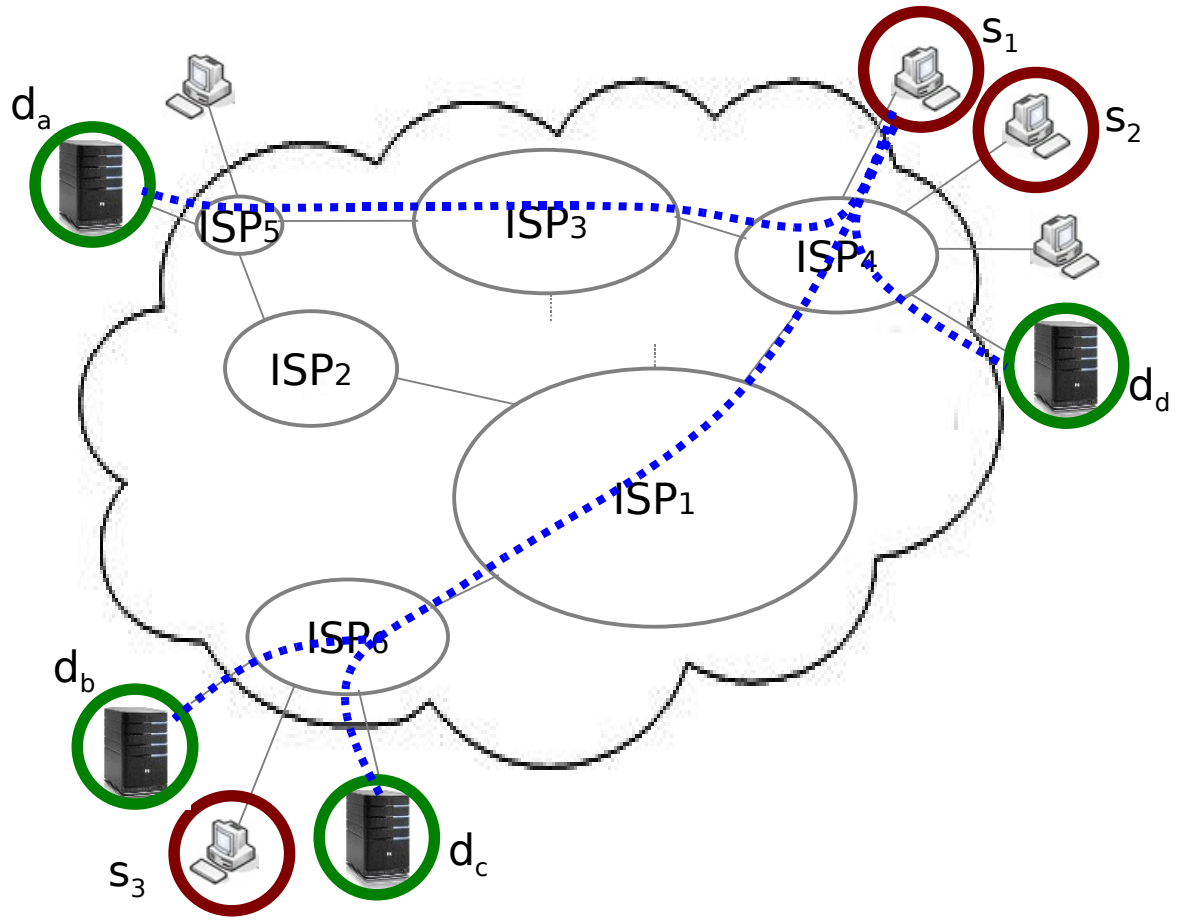

Best Peer Selection in P2P

- How to select the best peers set from the swarm
  - Example: selected peer set $\{p_e,p_f,p_h\}$ extracted from possible set $\{p_a,p_b,p_c,p_d,p_e,p_f,p_g,p_h\}$
  - per source: s_1

=> determine the best peerS among several: $\{<s,p_a>, \ldots, <s,p_h>\}$

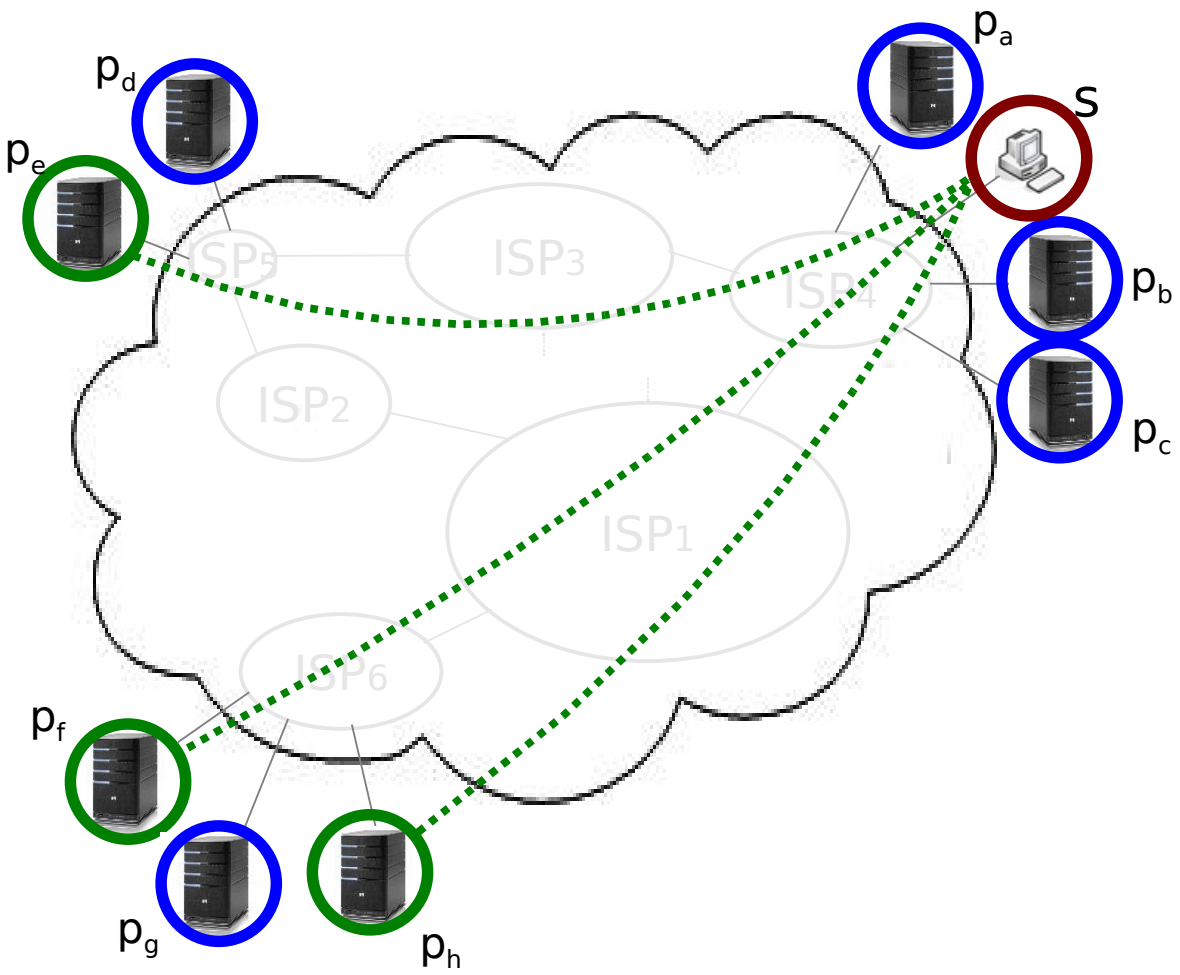

Selected peer

Possible peer

Problems are similar...

- IPv4 - IPv6 DS $\in$ {<s_{IPv4},d_{IPv4}>, <s_{IPv6},d_{IPv6}>,<s_{IPv4},d_{IPv6}>, <s_{IPv6},d_{IPv4}>}
- MH $\in$ {<s_1,d_1>, … ,<s_1,d_n>, <s_2,d_1>, … , <s_m,d_n>}
- Server replication $\subseteq$ {<s,d_a>, <s,d_b>, <s,d_c>, <s,d_d>}
- P2P Apps $\subseteq$ {<s,p_a>, …, <s,p_h>}

=> General problem $\subseteq$ {<s_1,d_1>, … ,<s_1,d_n>, <s_2,d_1>, … , <s_m,d_n>}

for any s,d representation

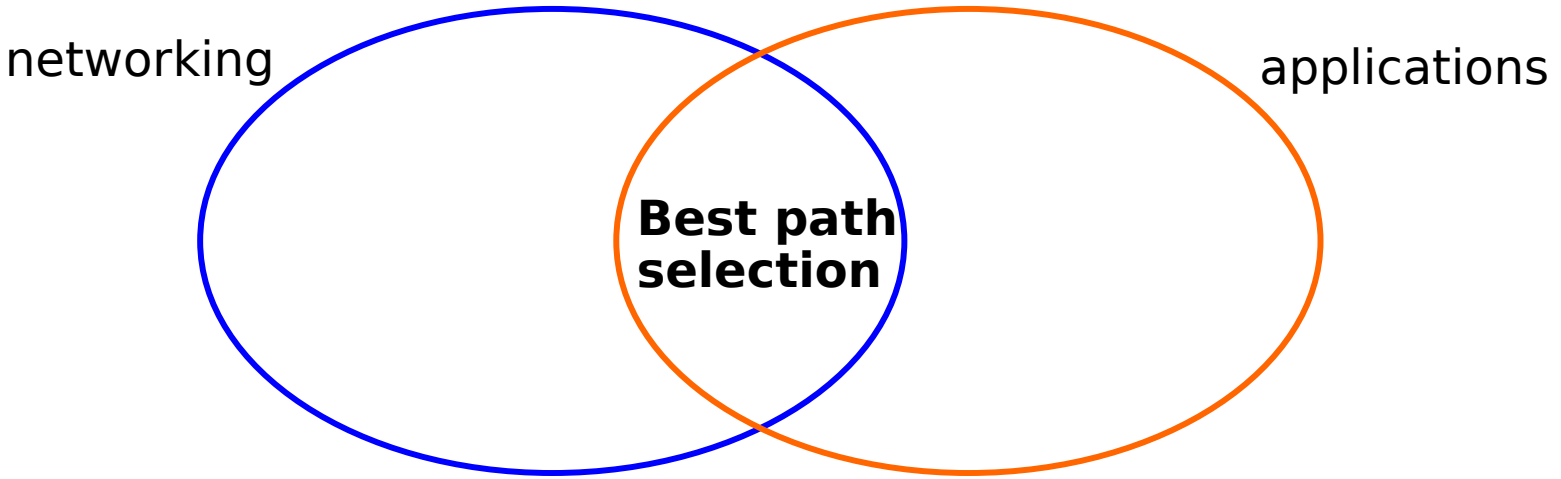

ALL share a common problem: how to efficiently make best path selection ?

… but knowledge is different

- Making the choice at the transport/application level is not always the best:
  - limited view of the topology
  - duplication of measurements
- Making the choice at the physical/network level is not always the best:
  - not aware of real application needs
  - expectations are different

=> Collaboration!

IDIPS: ISP-Driven Informed path Selection

We need a service:

- able to rank paths independently of the application
- that can inform applications about the ranks
- related works: Oracle [5], P4P [6]

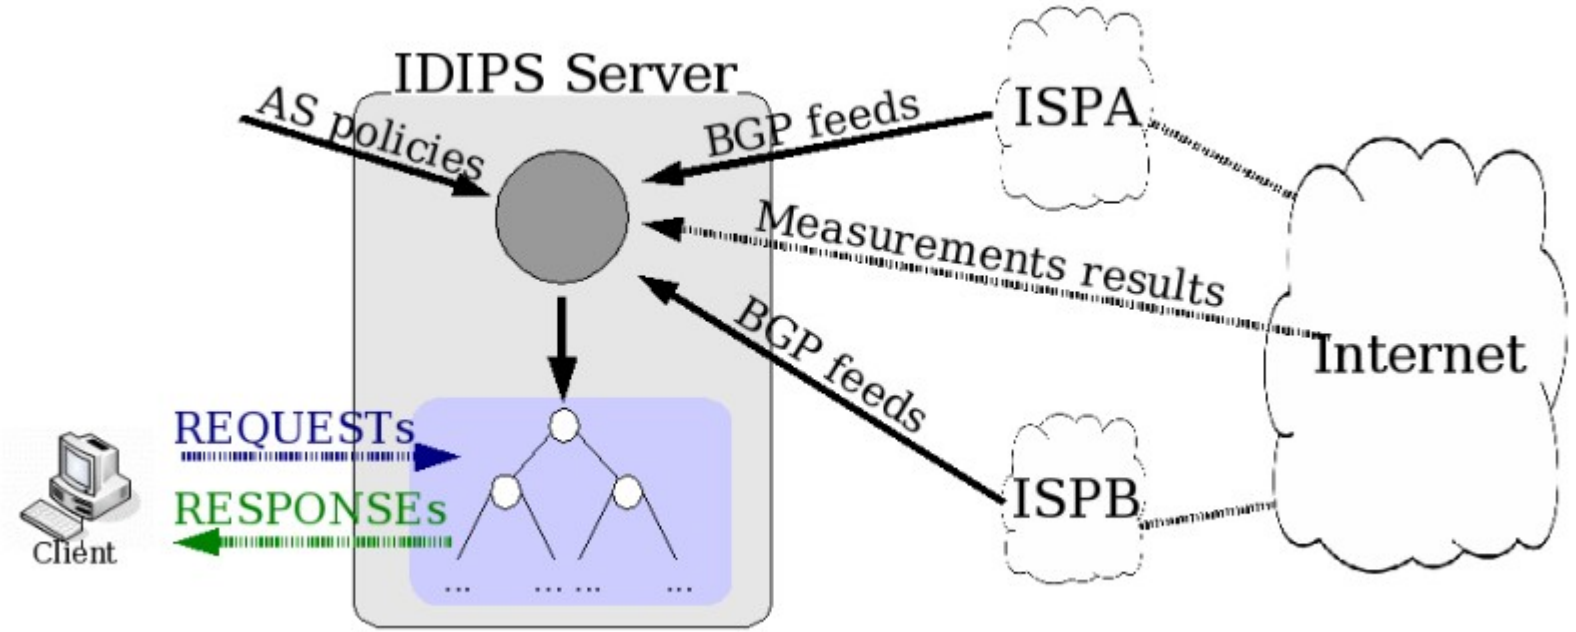

How IDIPS works?

0. The scenario

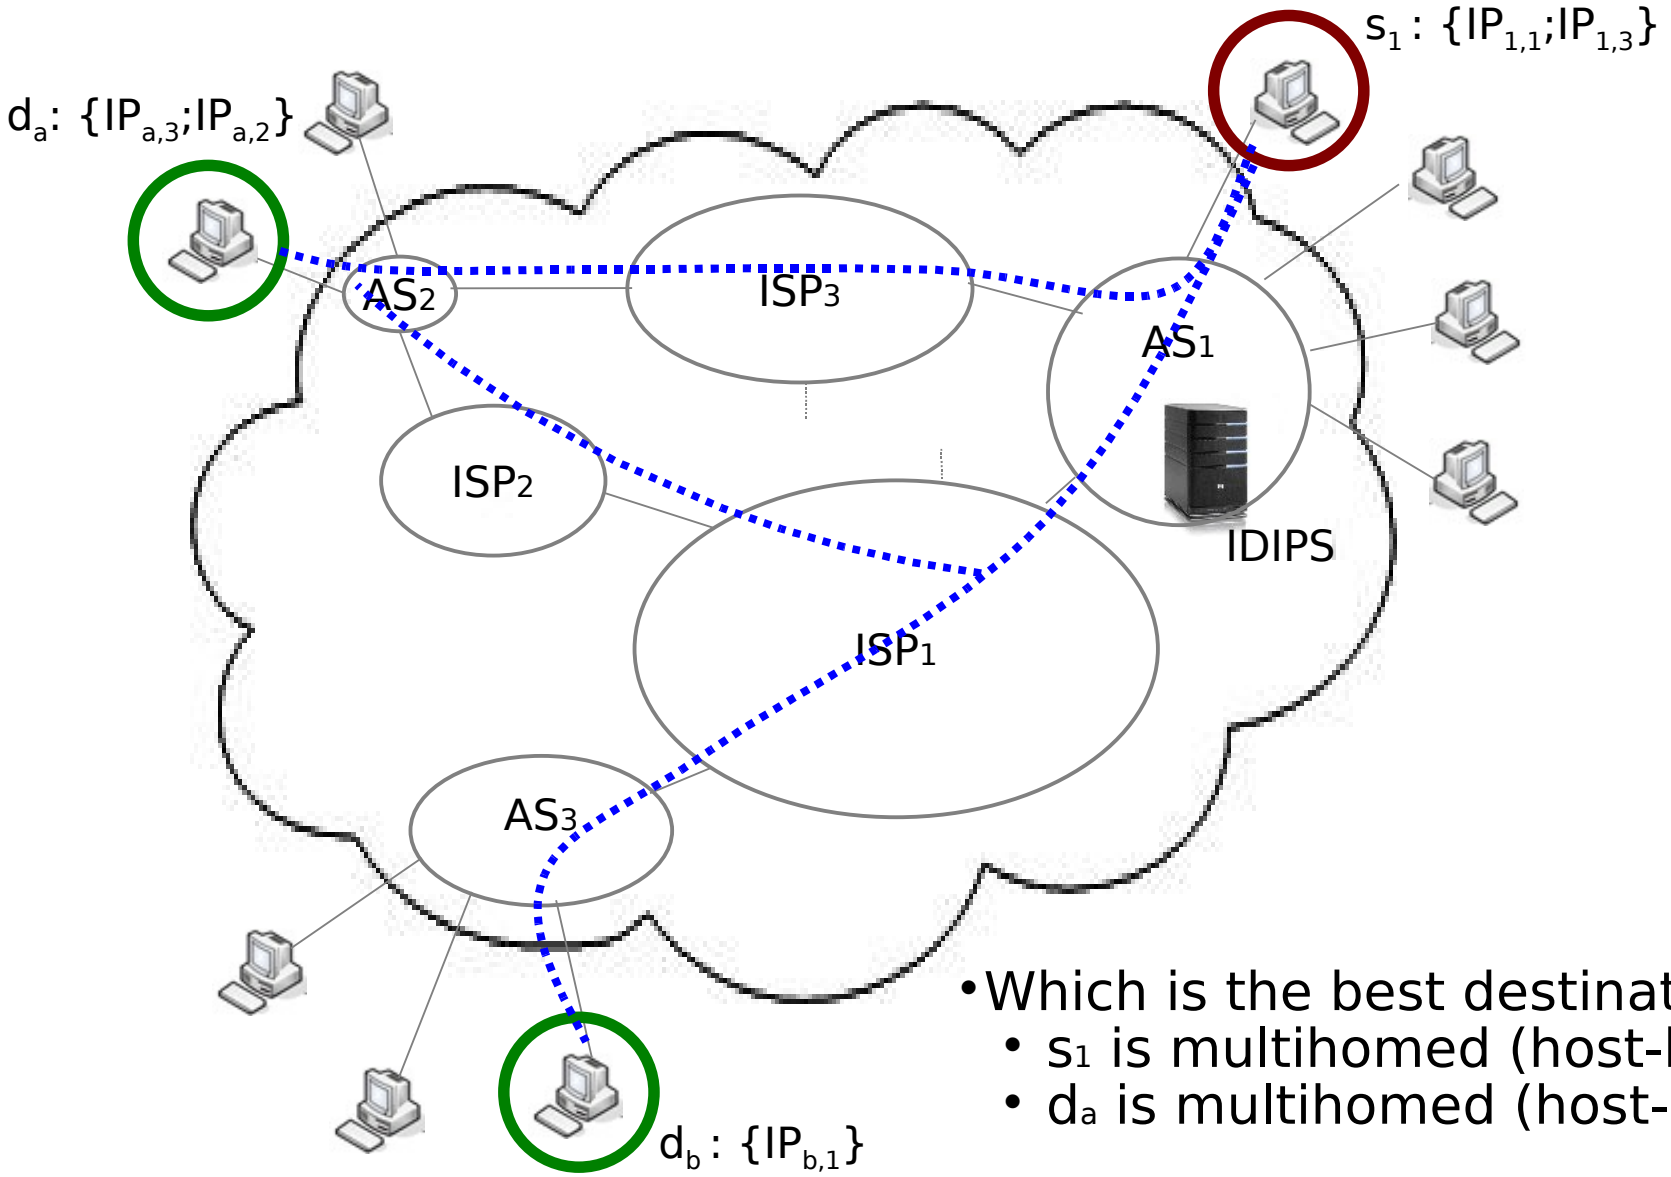

- Which is the best destination: d_a or d_b?
  - s_1 is multihomed (host-based)
  - d_a is multihomed (host-based)

How IDIPS works?

1. The Request

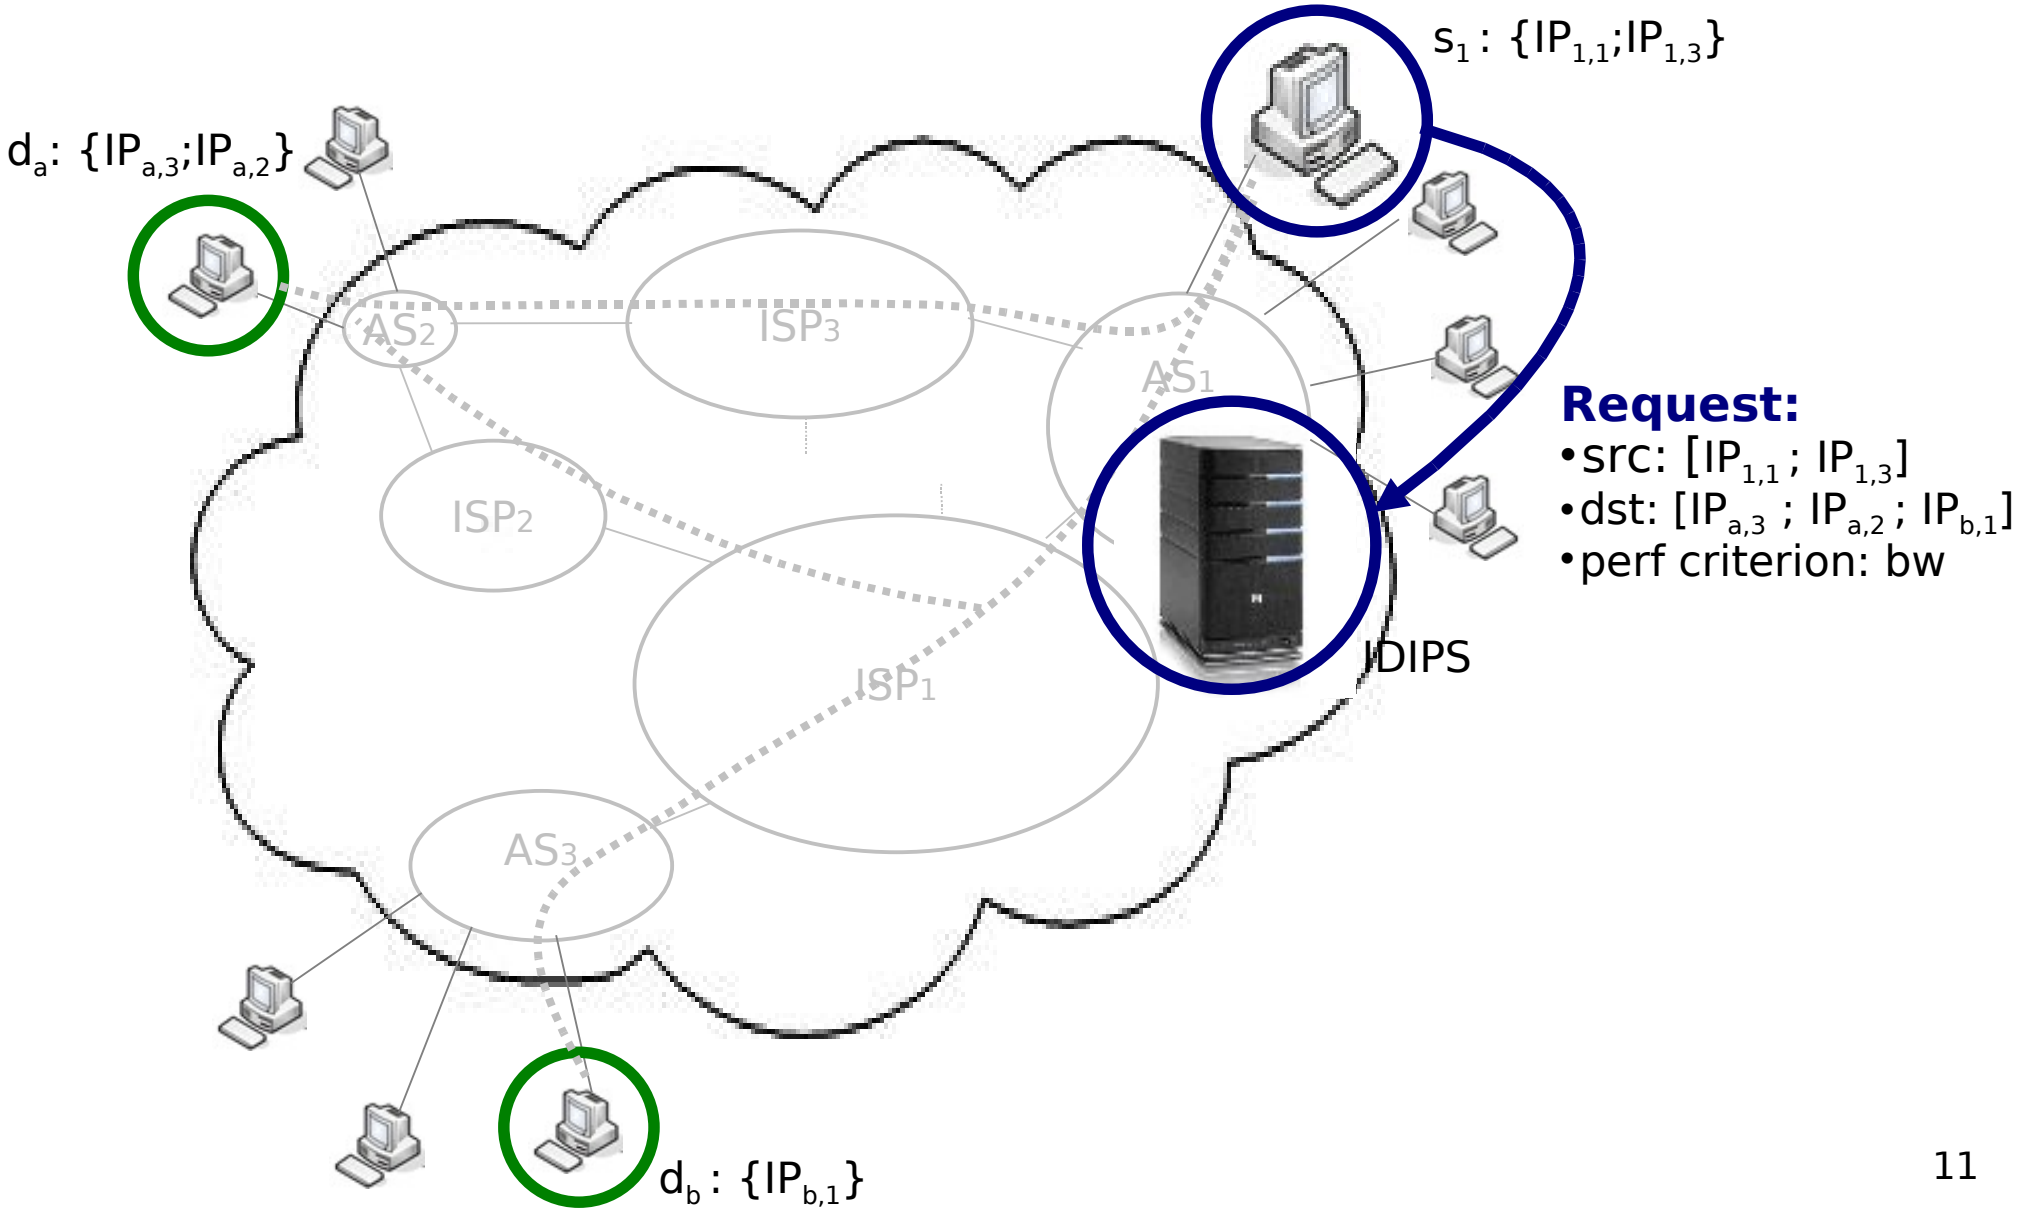

Request:

- src: [$IP_{1,1}$; $IP_{1,3}$]
- dst: [$IP_{a,3}$; $IP_{a,2}$; $IP_{b,1}$]
- perf criterion: bw

How IDIPS works?

1-2. The Paths Ranking

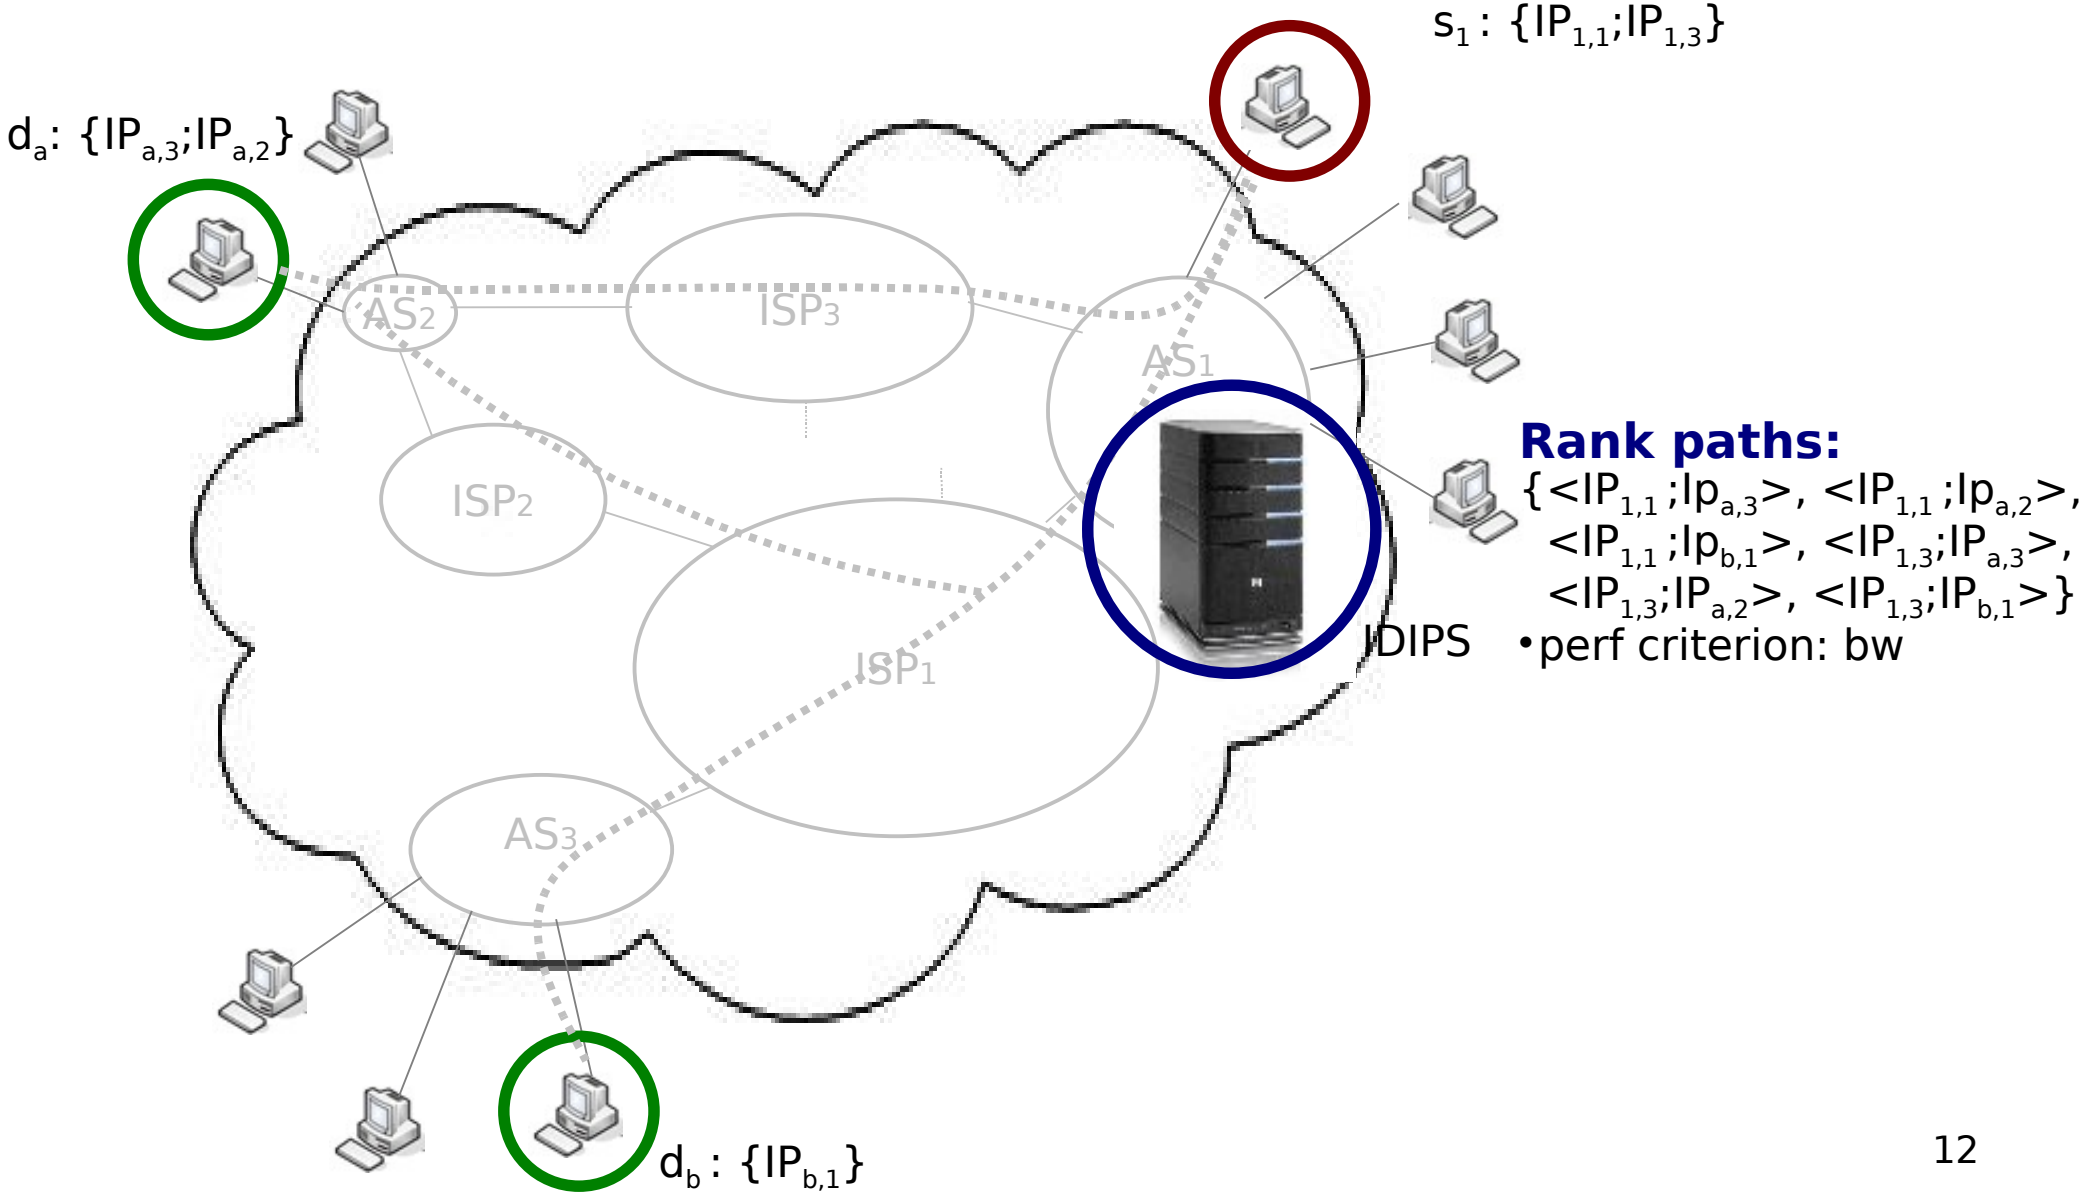

Rank paths:

{<$IP_{1,1}$;$lp_{a,3}$>, <$IP_{1,1}$;$lp_{a,2}$>,
<$IP_{1,1}$;$lp_{b,1}$>, <$IP_{1,3}$;$IP_{a,3}$>,
<$IP_{1,3}$;$IP_{a,2}$>, <$IP_{1,3}$;$IP_{b,1}$>}

- perf criterion: bw

How IDIPS works?

2. The Reply

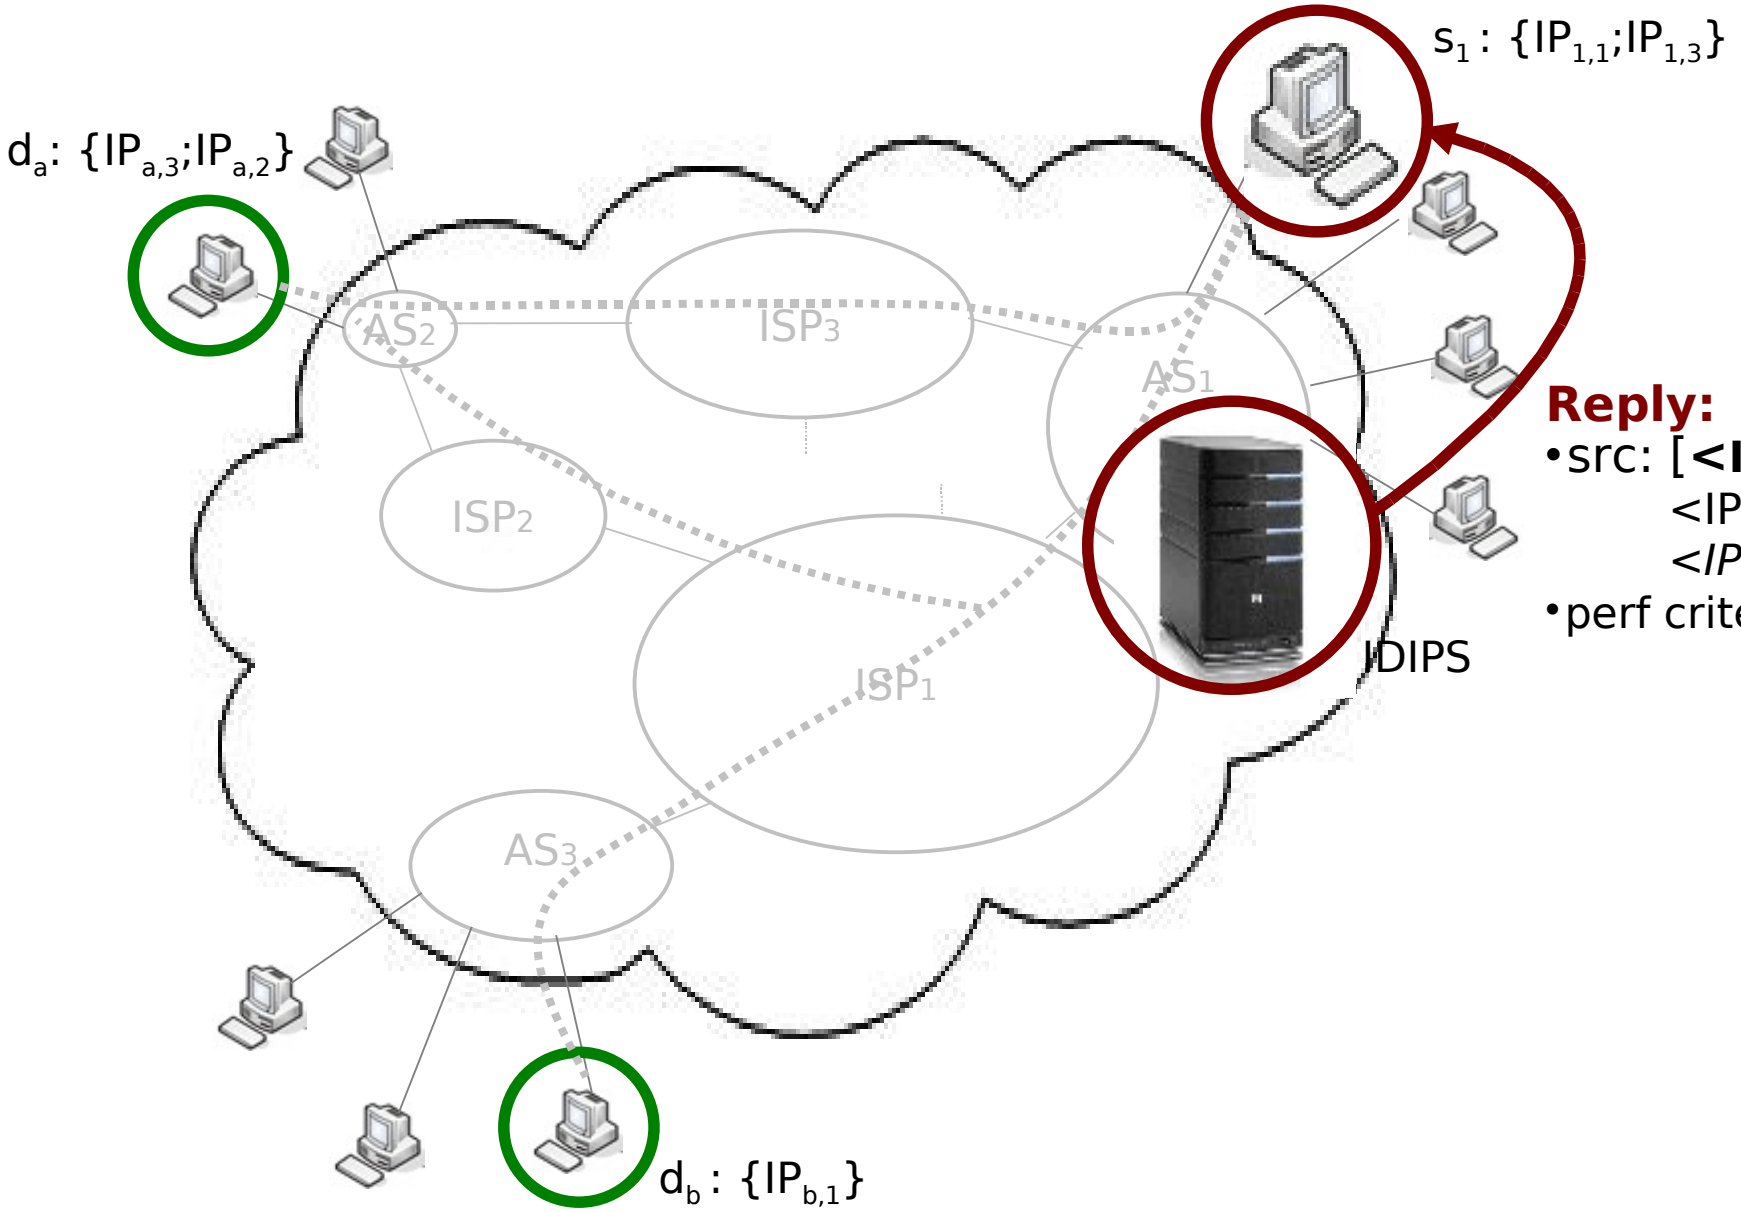

Reply:

- src: [**<$IP_{1,1}$; $IP_{b,1}$>:1**,
  <$IP_{1,3}$; $IP_{a,3}$>:2,
  <$IP_{1,3}$; $IP_{b,1}$>:2]
- perf criterion: bw

How IDIPS works?

3. The Choice

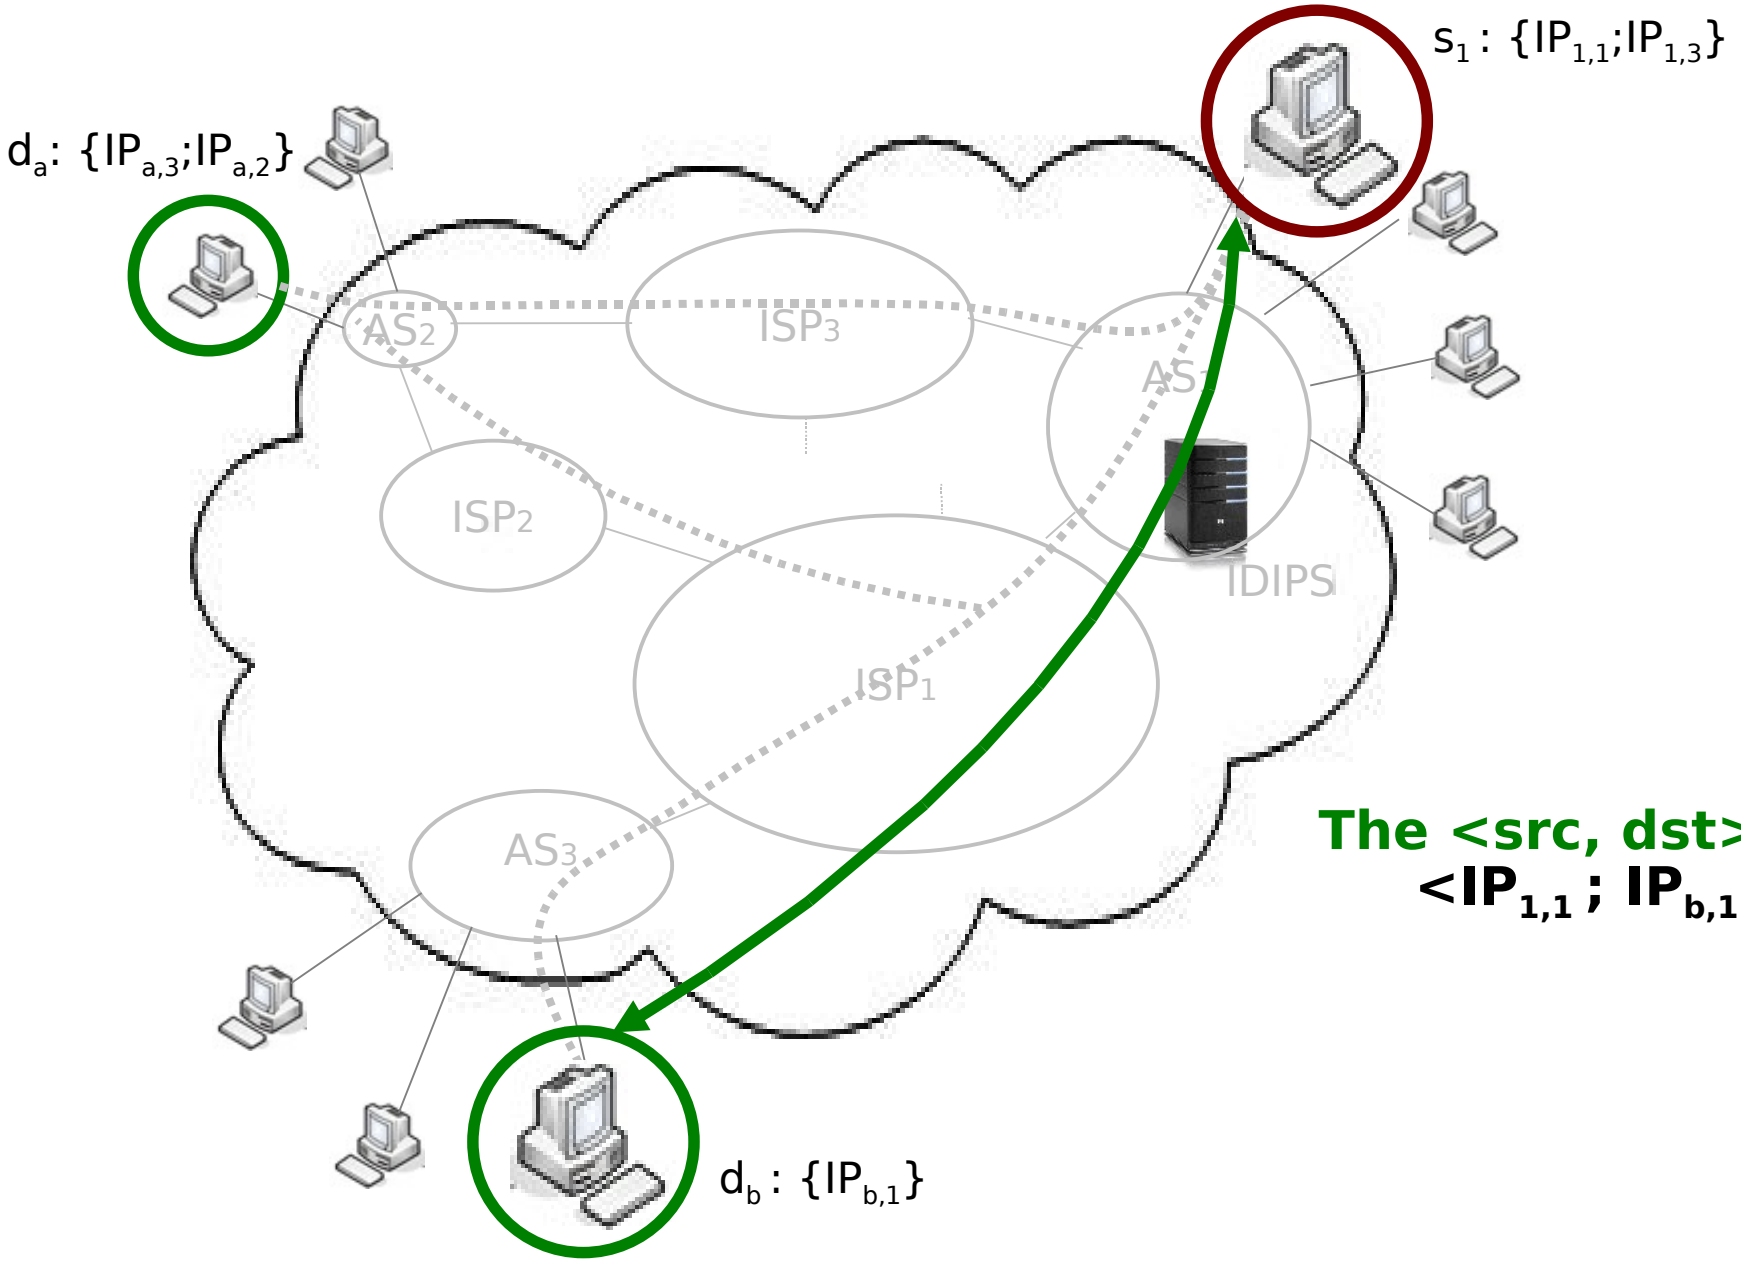

The <src, dst> pair:
<$IP_{1,1}$; $IP_{b,1}$>

How IDIPS works?
4. Summary

- You need to rank paths? Ask IDIPS!
  1. Send your addresses, the addresses of your possible destinations and your performance criterion
  2. IDIPS computes the possible paths and ranks them
  3. IDIPS replies with an ordered list of paths
  4. Take the first proposed path, it should be the best

Case Study

LISP (Locator/Identifier Separation Protocol)

- Separate IP space in two different spaces
  - EndPoint Identifiers (EIDs) are used to identify end-hosts.
    - Not globally routable
    - Hosts in a site are expected to use EIDs in the same prefix
  - Routing Locators (RLOCs) are used to identify EID locations
    - Globally routable
    - Attached to DFZ Border Routers

Case Study

LISP (Locator/Identifier Separation Protocol)

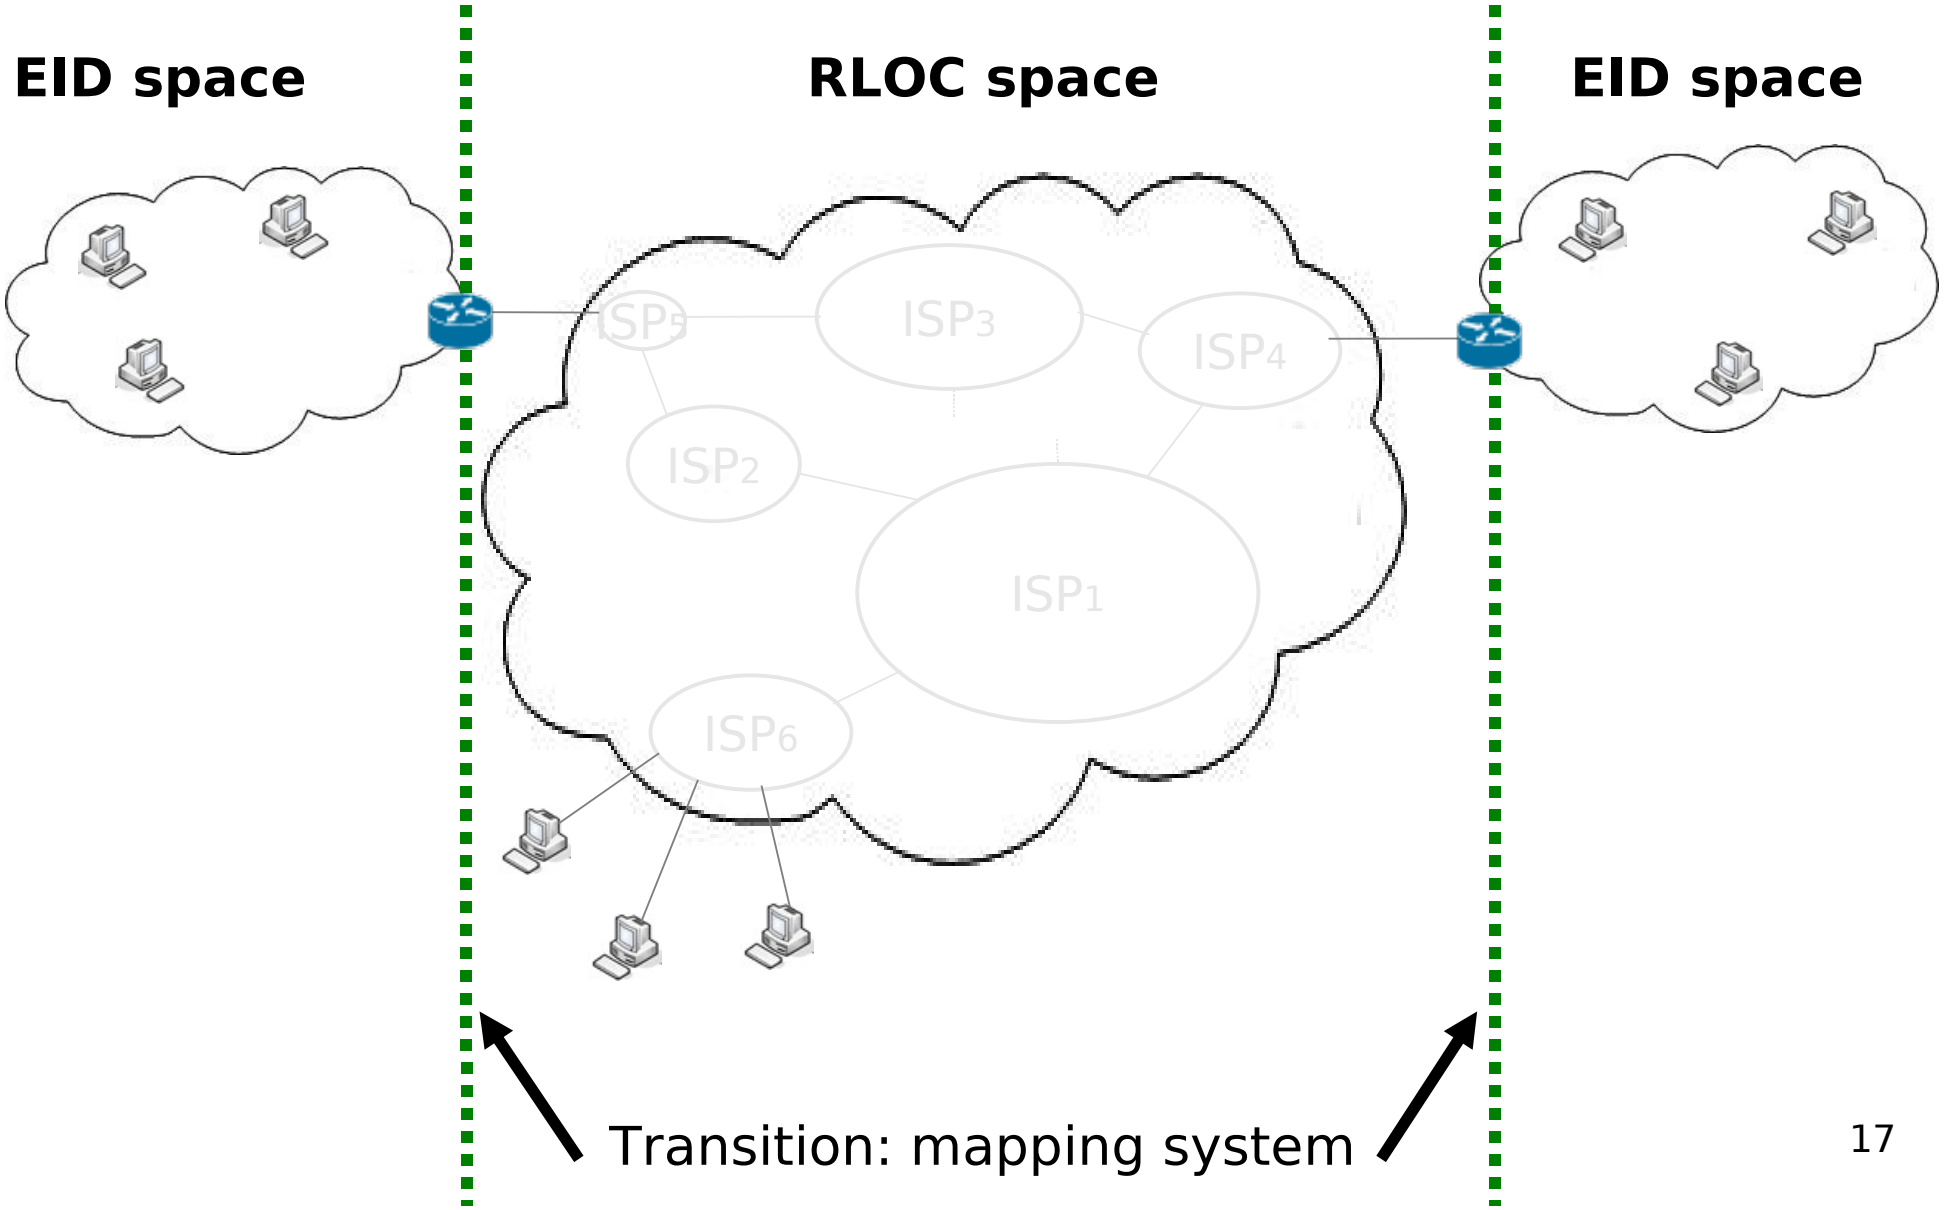

Case Study

LISP (Locator/Identifier Separation Protocol)

- EIDs and RLOCs are in disjoined spaces
- Each EID is associated to *n* RLOCs with priorities
- The RLOC with the lowest priority value is selected
- RLOCs can be changed dynamically during any flow (the EID is stable, not the RLOC)

Case Study
LISP (Locator/Identifier Separation Protocol)

- Tune RLOC priorities to make TE
- Set the lowest priority value to the RLOC that must be used
- Ask IDIPS to rank EID's RLOCs *(criteria can be different for each EID)*
- Let the mapping system translate IDIPS ranks into RLOC priorities
- Use the mapping system to distribute priorities

Conclusion

- Path selection is a general problem
- We need collaboration between the layers
  - IDIPS: a service able to rank paths based on its network knowledge
  - You need to rank a path? Ask IDIPS!
- TE with LISP
  - Tune RLOC priorities
  - Ask IDIPS to compute RLOC priorities

Further works

- How to reduce measurements?
- How to efficiently predict the future performances (machine learning)?
- How to avoid oscillations?

References

[1] Zhou et al., Ipv6 delay and loss performance evolution, IJCS 2008

[2] Quoitin et al., Evaluating the Benefits of the Locator/Identifier Separation, MobiArch 2007

[3] A. Akella, S. A., and R. Sitaraman, A measurement-based analysis of multihoming, in Proc. ACM SIGCOMM, 2003

[4] Saucez et al., draft-saucez-idips-00.txt, IETF draft, 2008

[5] Aggarwal et al., Can ISPs and P2P systems co-operate for improved performance?, ACM SIGCOMM Computer Communications Review (CCR), 37(3):29–40, July 2007

[6] Xie et al., P4P: Provider Portal for Applications, in Proc. ACM SIGCOMM, 2008

[7] R. Gao et al., Avoiding Oscillations due to Intelligent Route Control Systems, IEEE INFOCOM, 2006

[8] de Launois et al., Leveraging network performance with IPv6 multihoming and multiple provider-dependent aggregatable prefixes, Computer Networks, 2006

Questions? Remarks?

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

Backup Slides

Multihoming

- How to reduce the costs?
- How to finely control the costs (per customer? per flow? per ToS?)
- How to improve QoS experience without end-to-end reservation protocol?
- How to globally improve performances?
- How ISPs can control the Shim6 path selection algorithm?

=> determine the best paths among several