

Enabling ISP-CDN Collaboration:

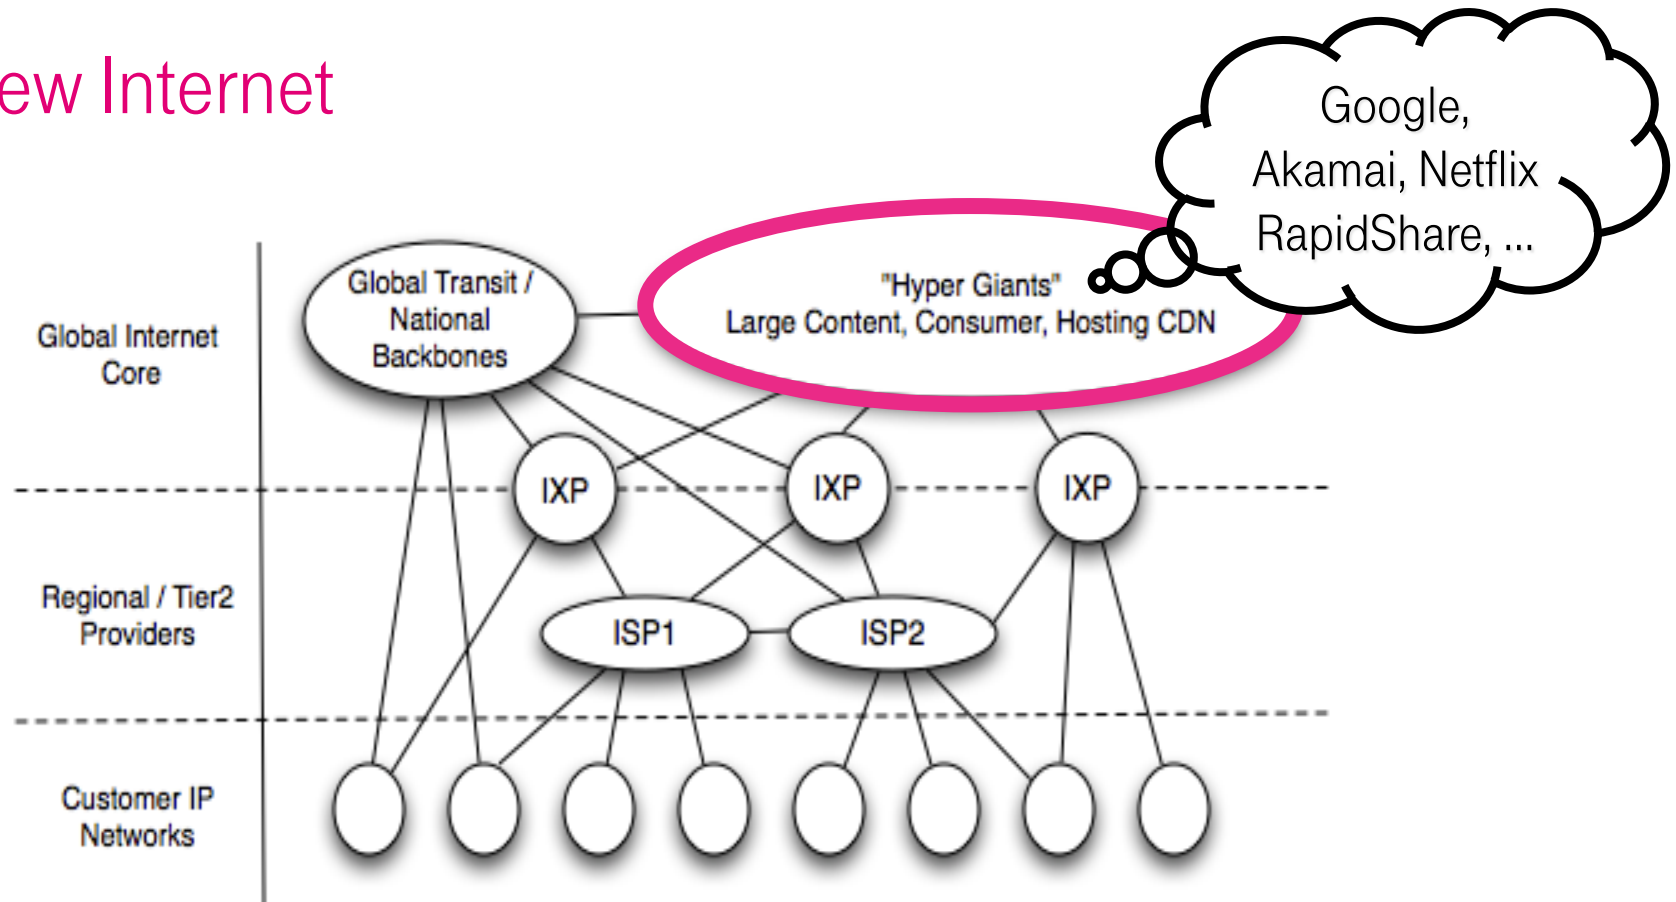
Turning Challenges into Opportunities

Georgios Smaragdakis (T-Labs/TU Berlin)

Joint work with Benjamin Frank, Ingmar Poesse, and Anja Feldmann (TU Berlin)
and Bruce Maggs (Akamai/Duke University)



The New Internet



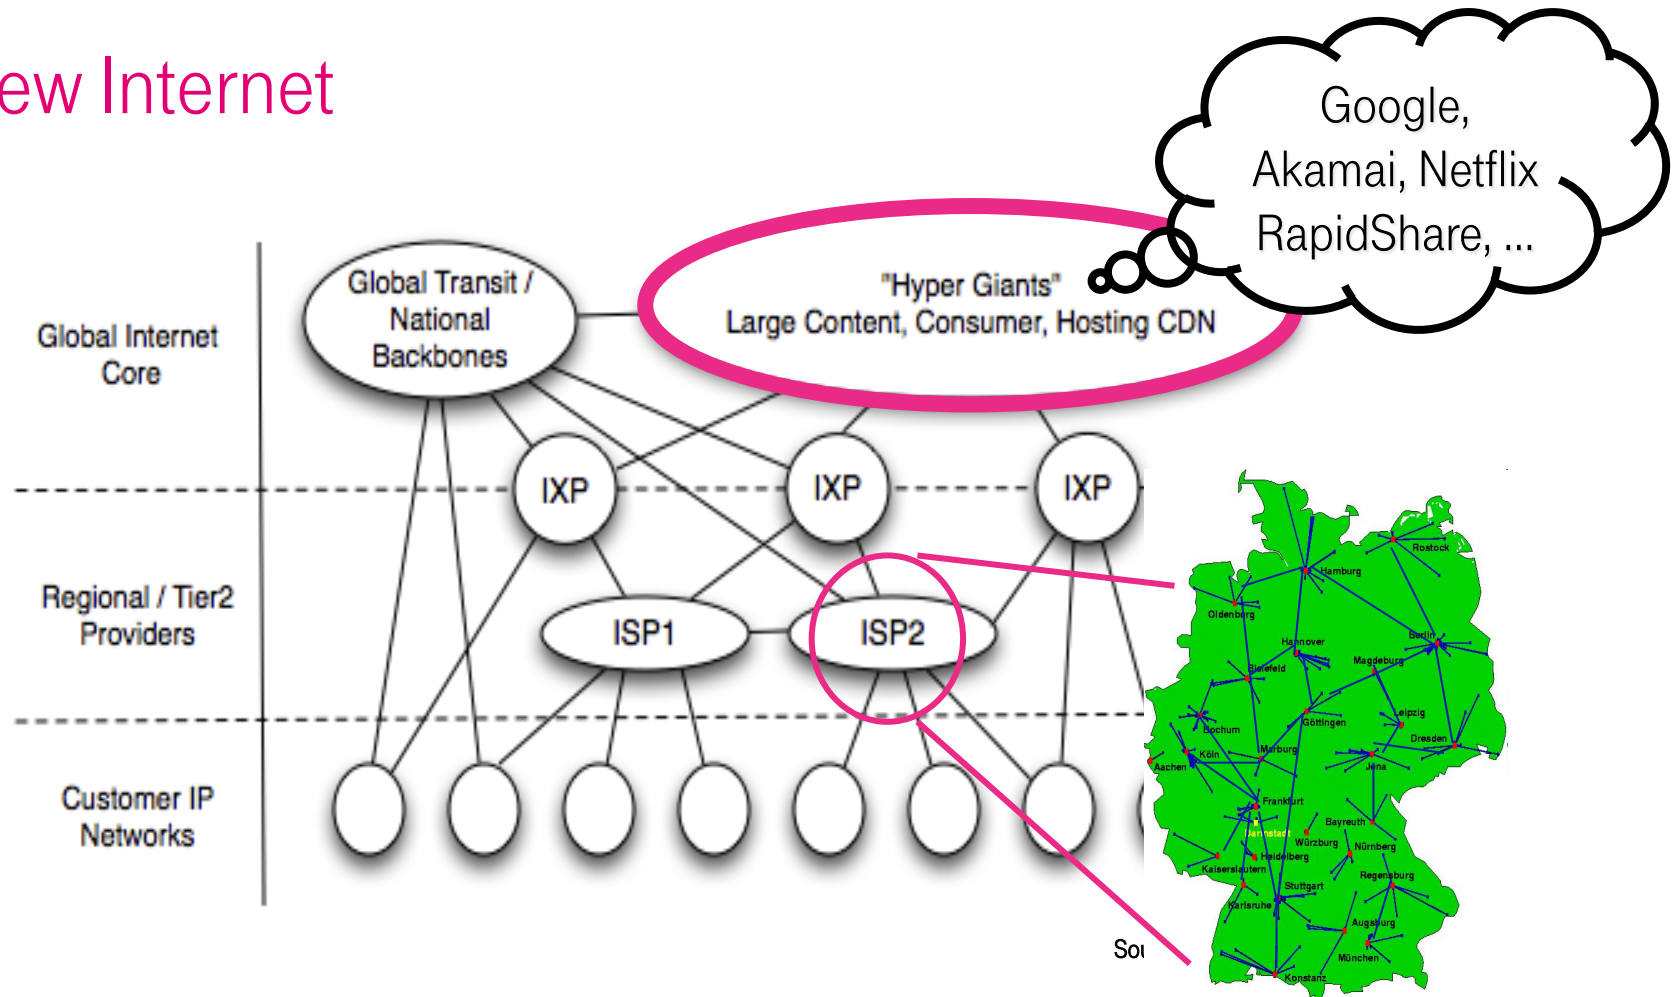
Source: Arbor Networks 2010

→ New core of interconnected content and consumer networks
[1]

[1] "Internet Interdomain Traffic", Labovicz, Lelak-Johnson, McPherson, Oberheide, Jahanian, SIGCOMM 2010



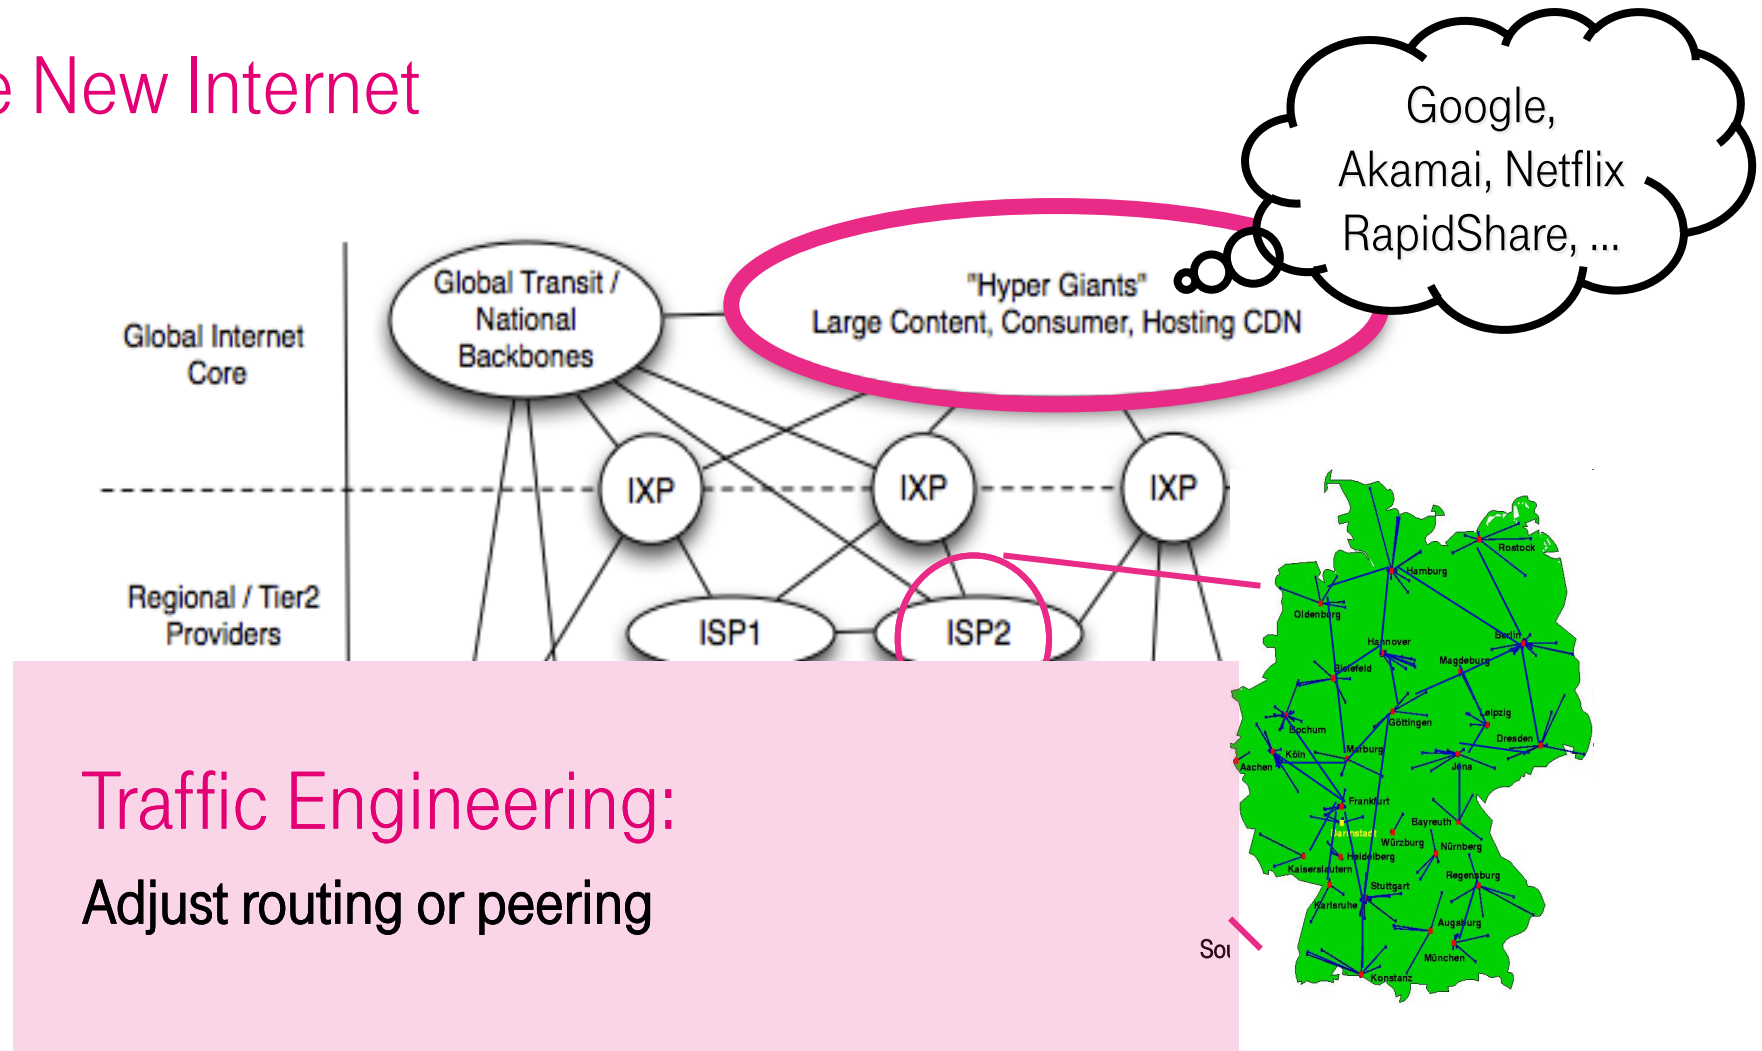
The New Internet



- New core of interconnected content and consumer networks
- Over-provisioning?



The New Internet



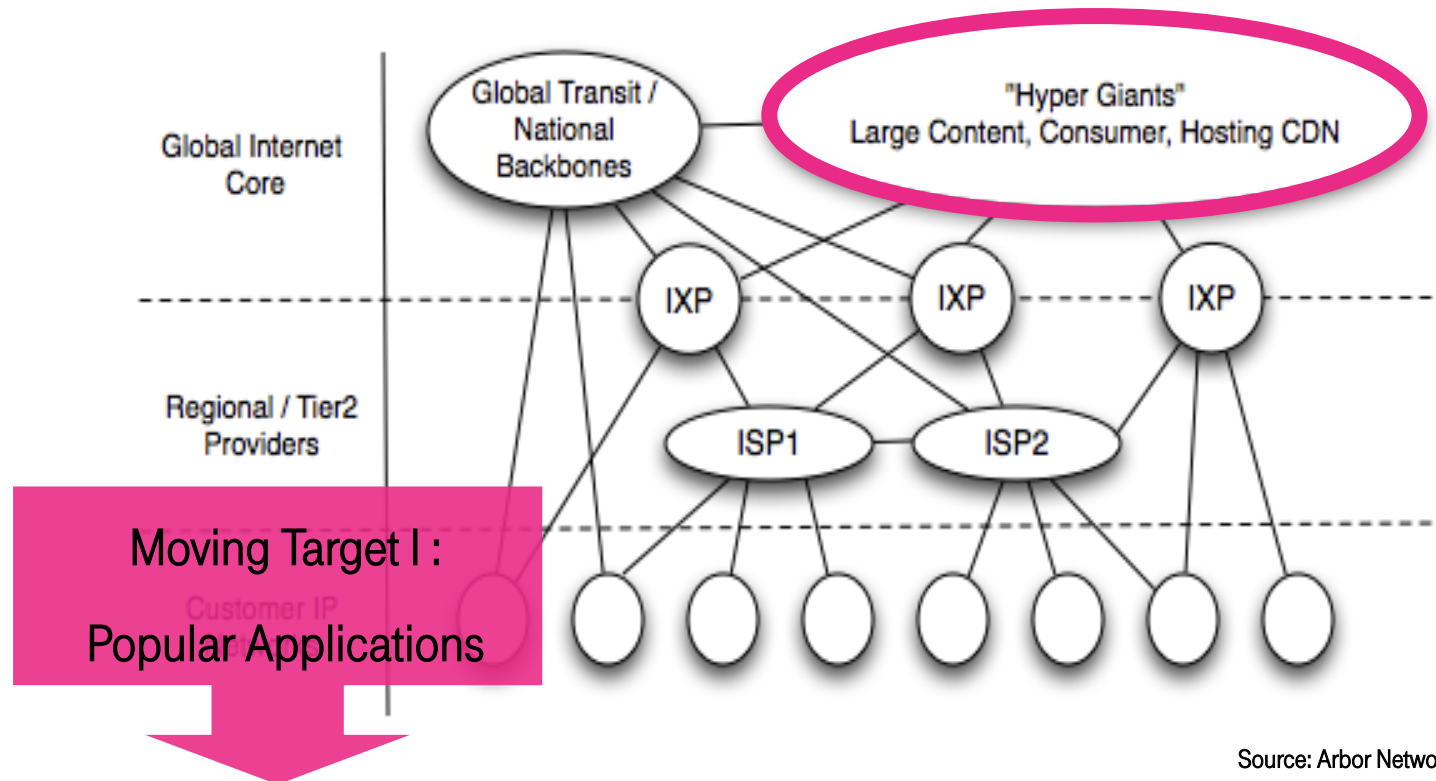
Traffic Engineering:
Adjust routing or peering

- Offline Process, or at best, time-scale of several hours
- Few changes take place to avoid oscillations ^[1]

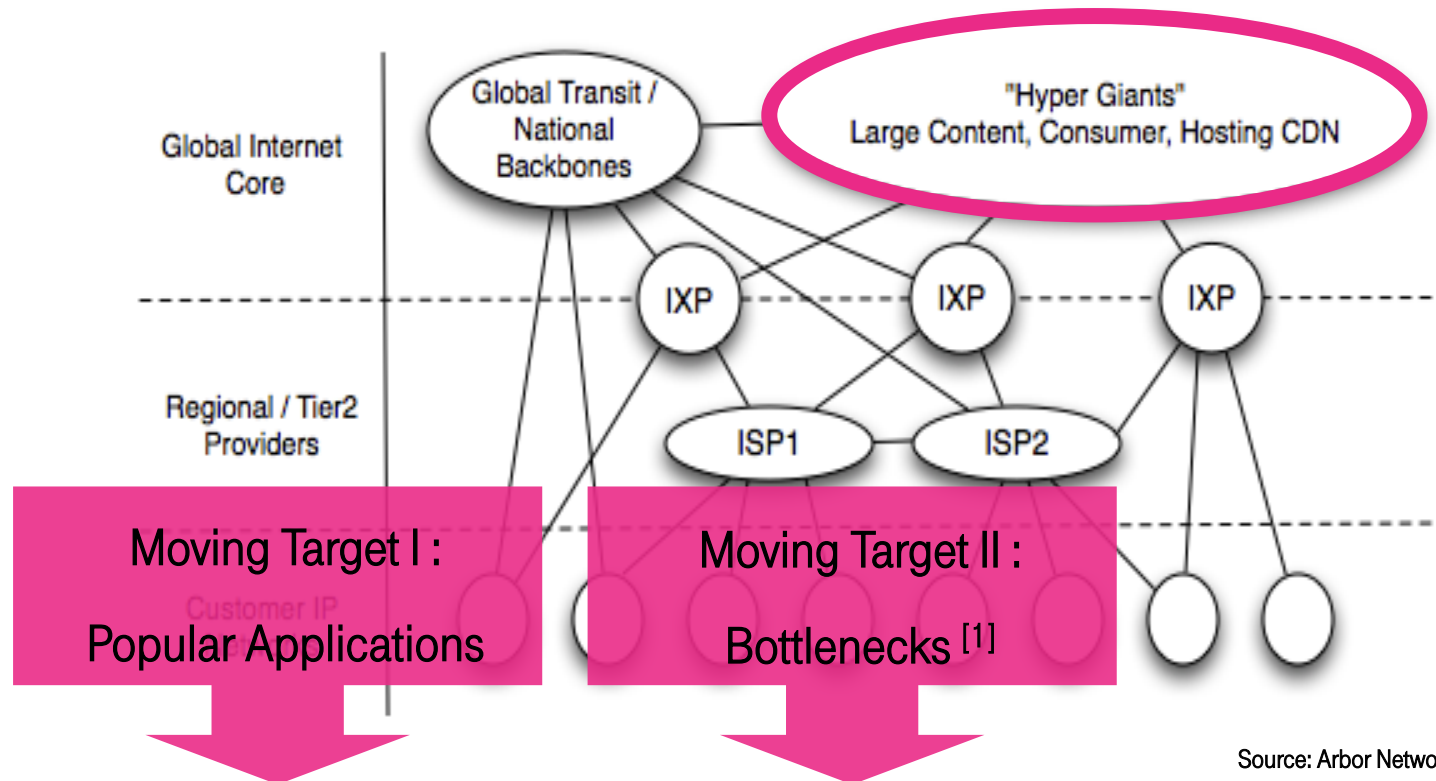


[1] "Internet Traffic Engineering by Optimizing OSPF Weights", Fortz, Thorup, INFOCOM 2000

New Challenges for ISPs



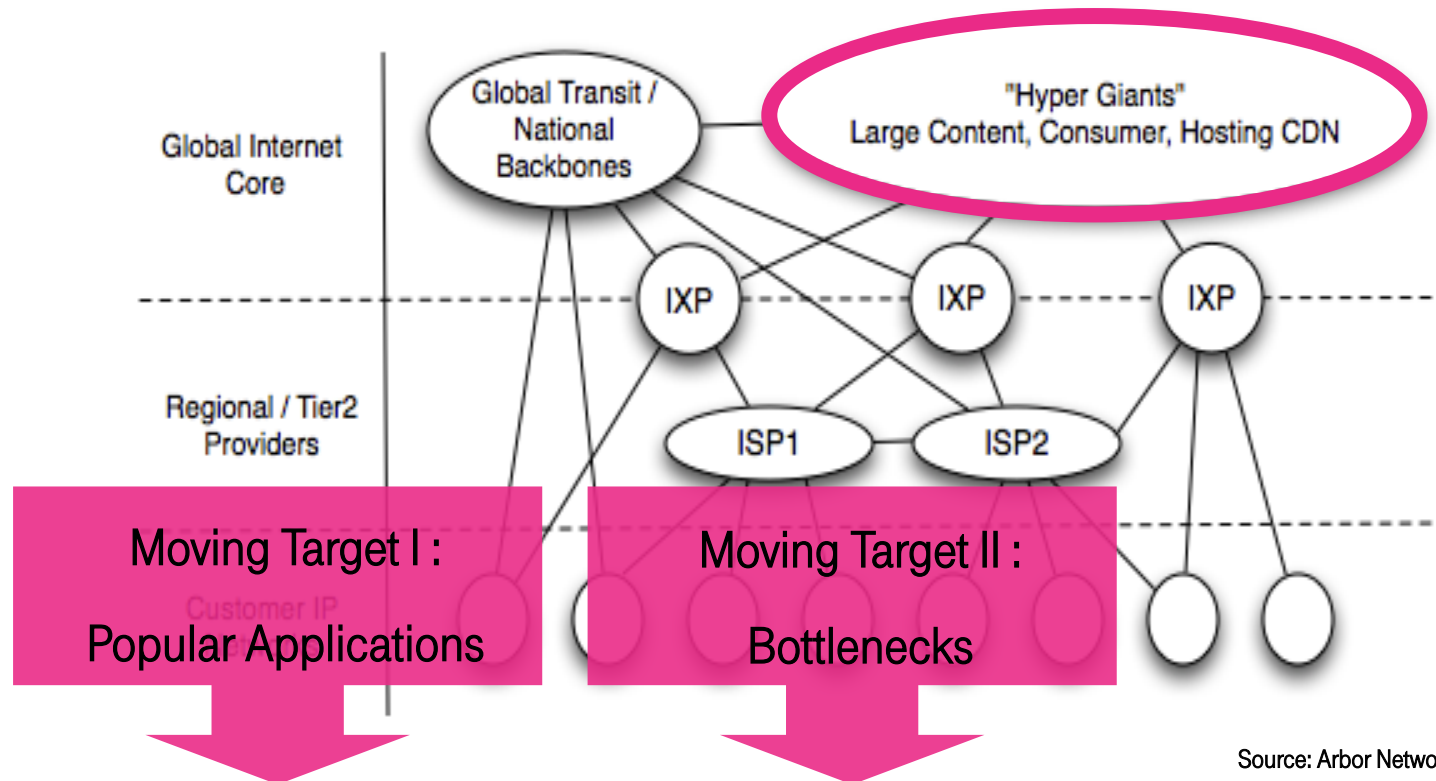
New Challenges for ISPs



[1] "Improving Performance on the Internet", Leighton, CACM 2009




New Challenges for ISPs



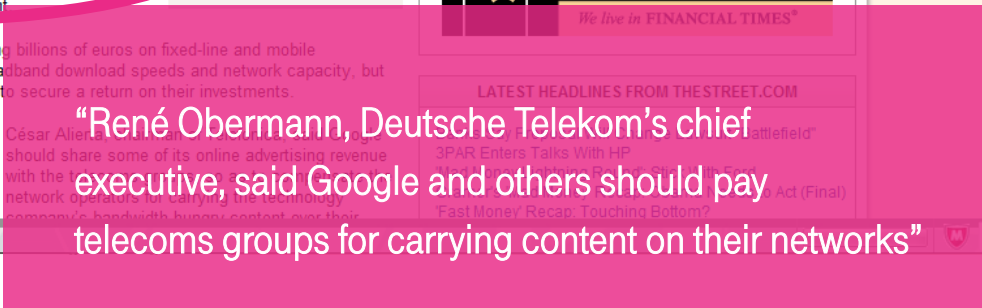
- Traditional Traffic Engineering too slow to react
- ISPs lose control of their network



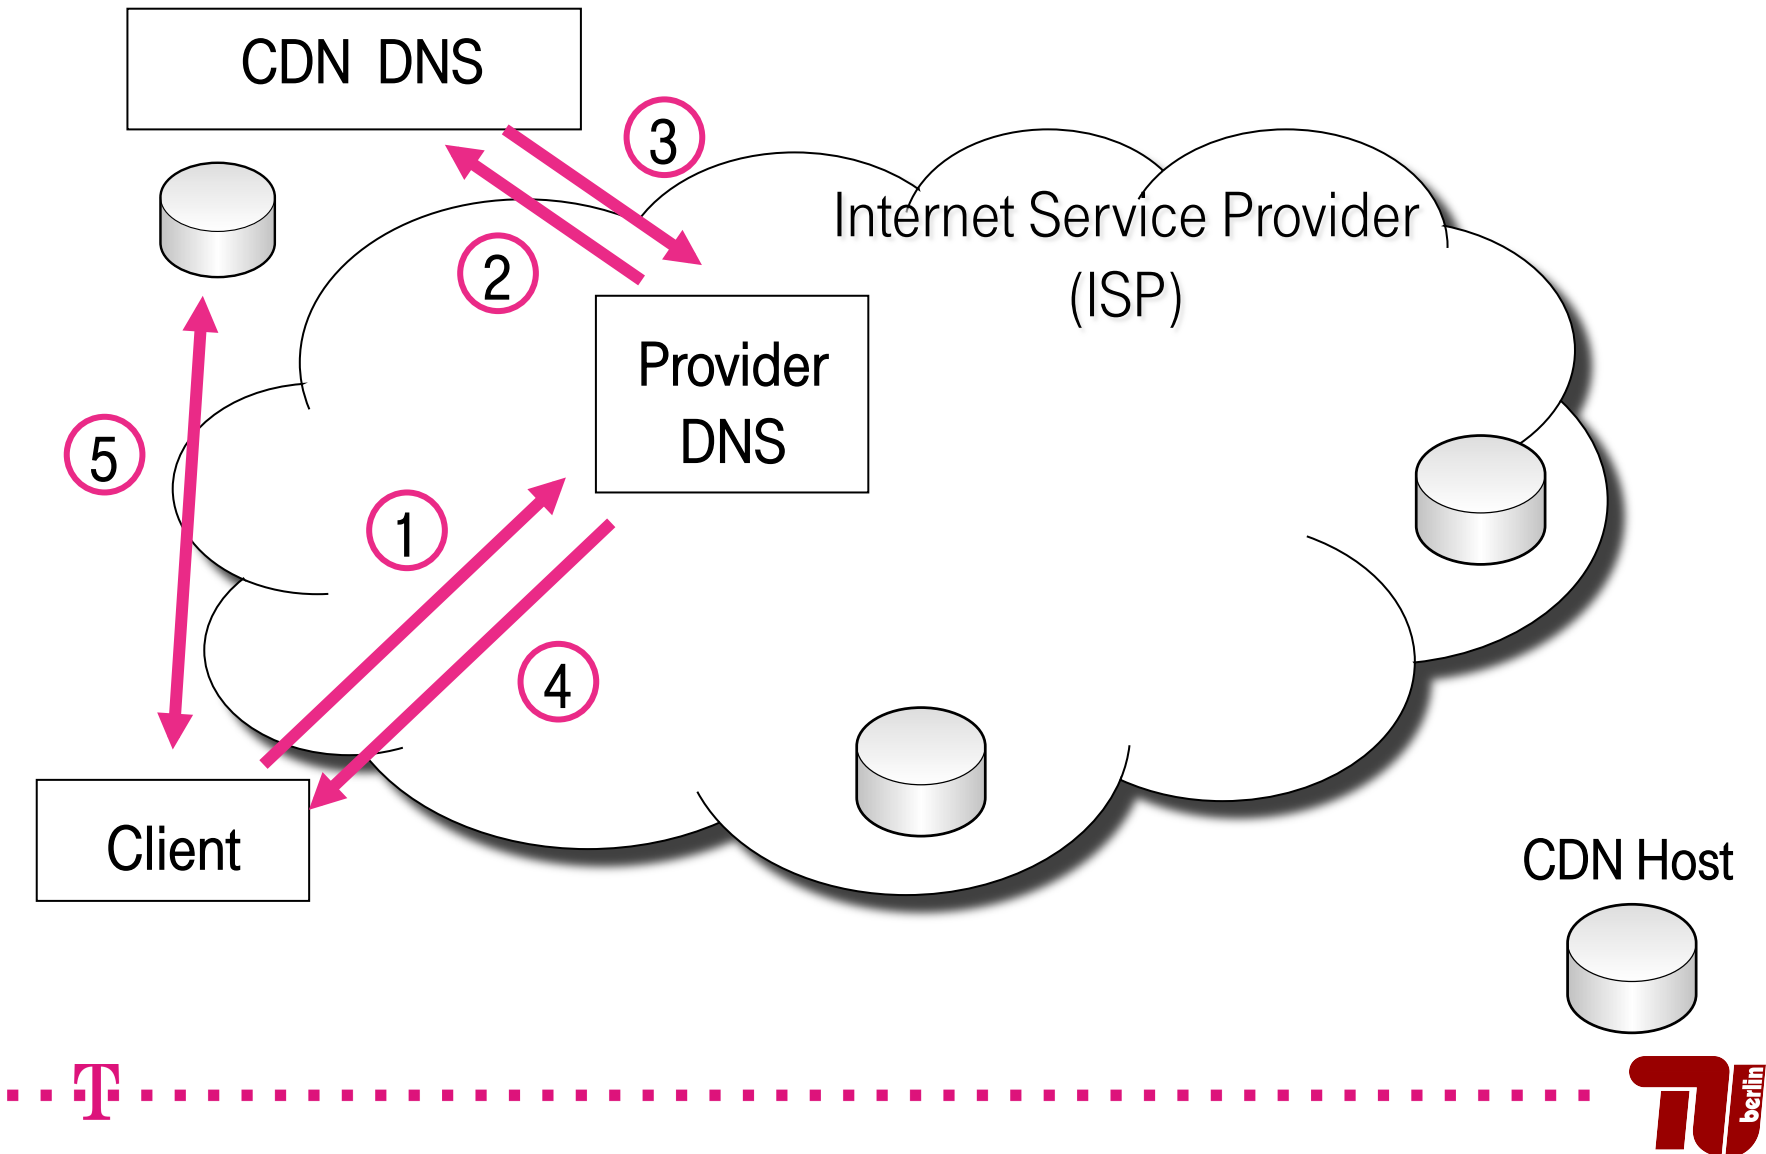
T



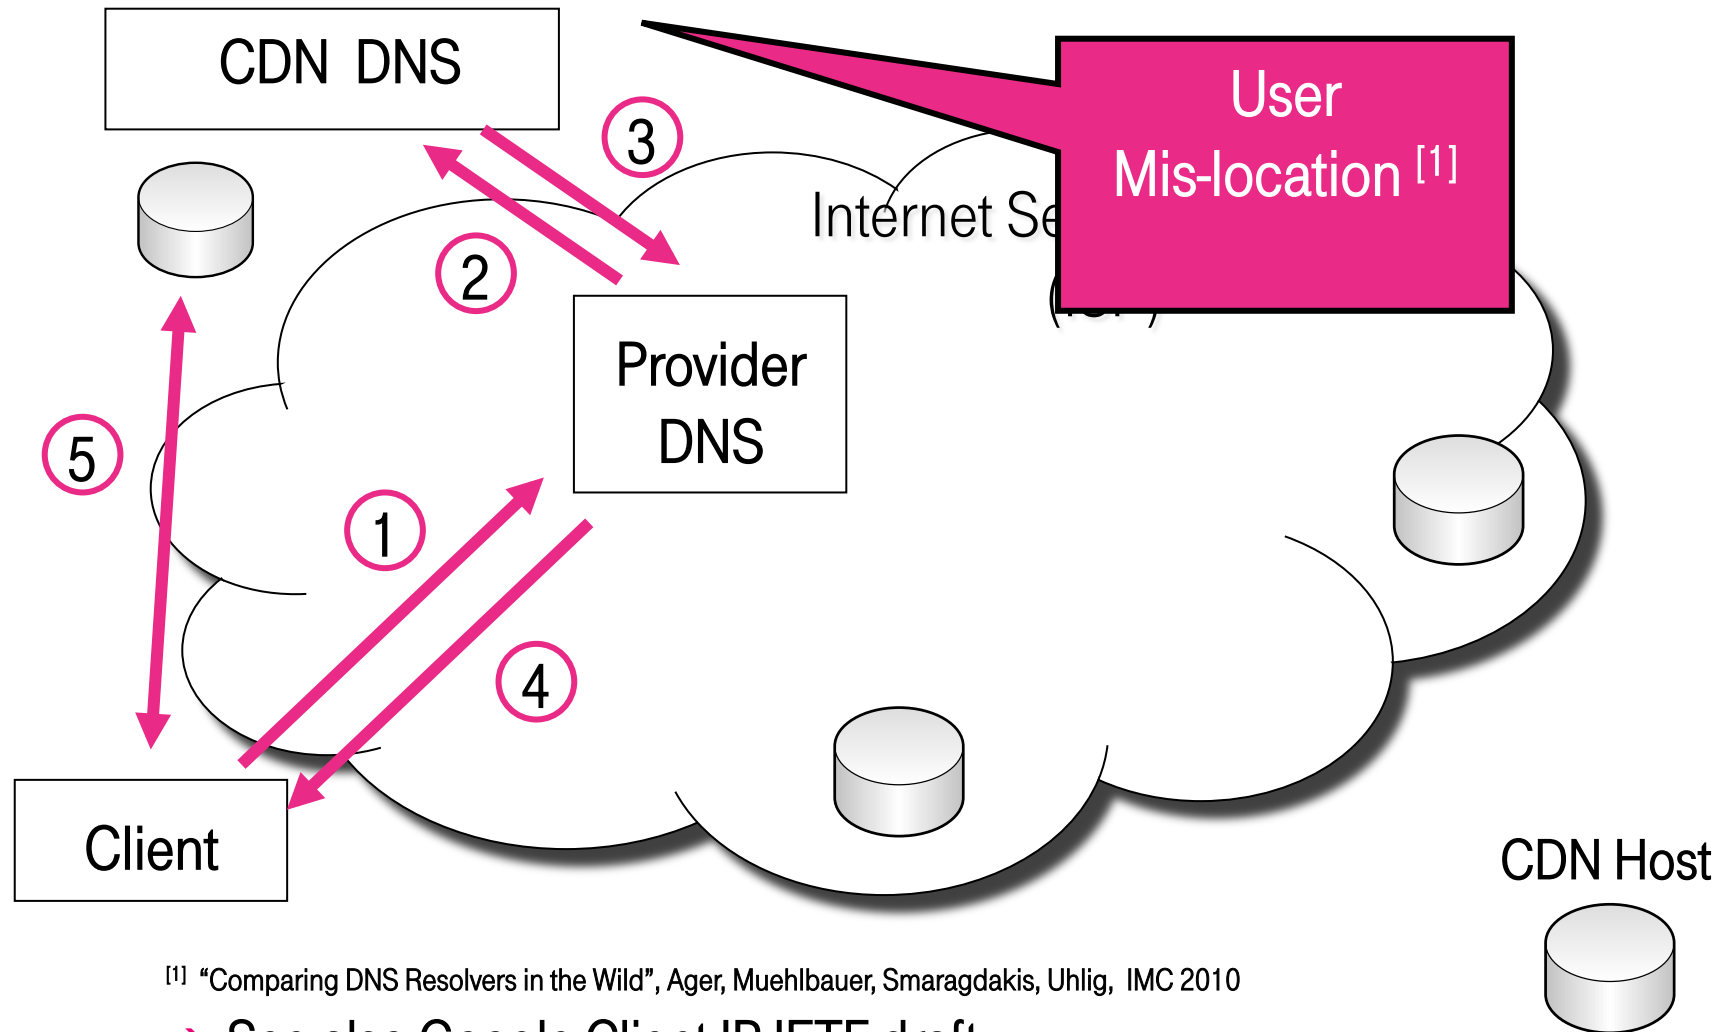
→ ISPs lose



Content Distribution Prime



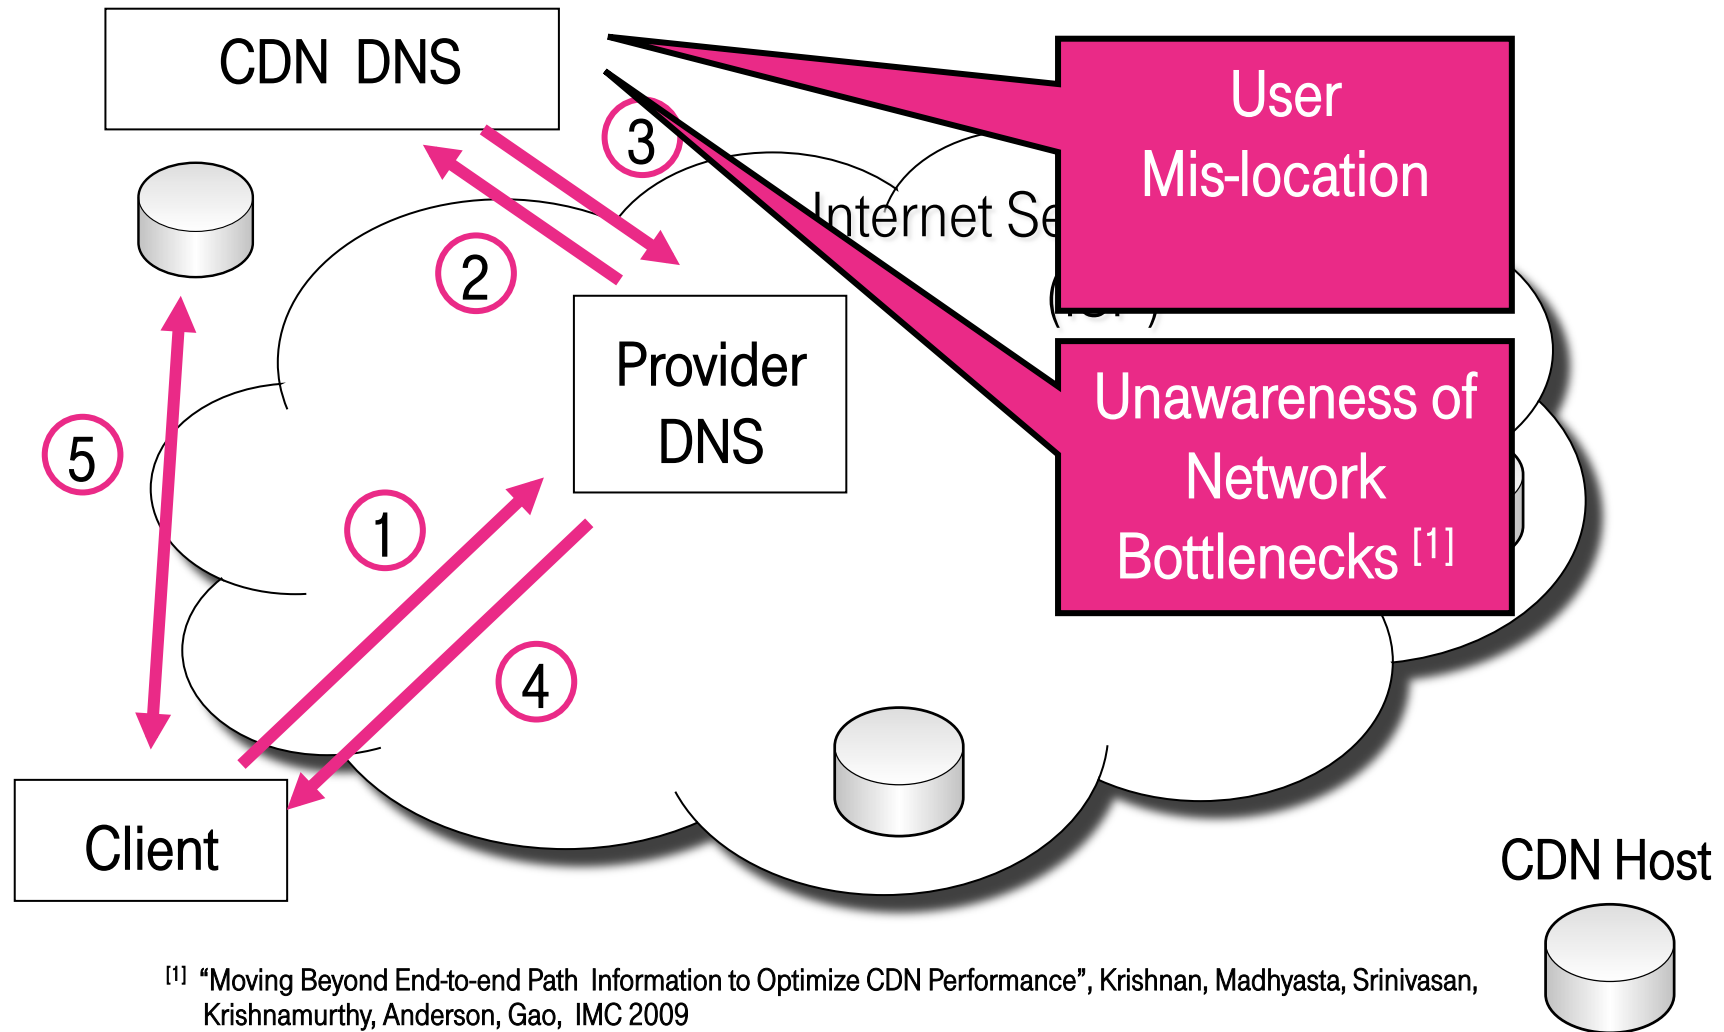
Content Distribution Challenges



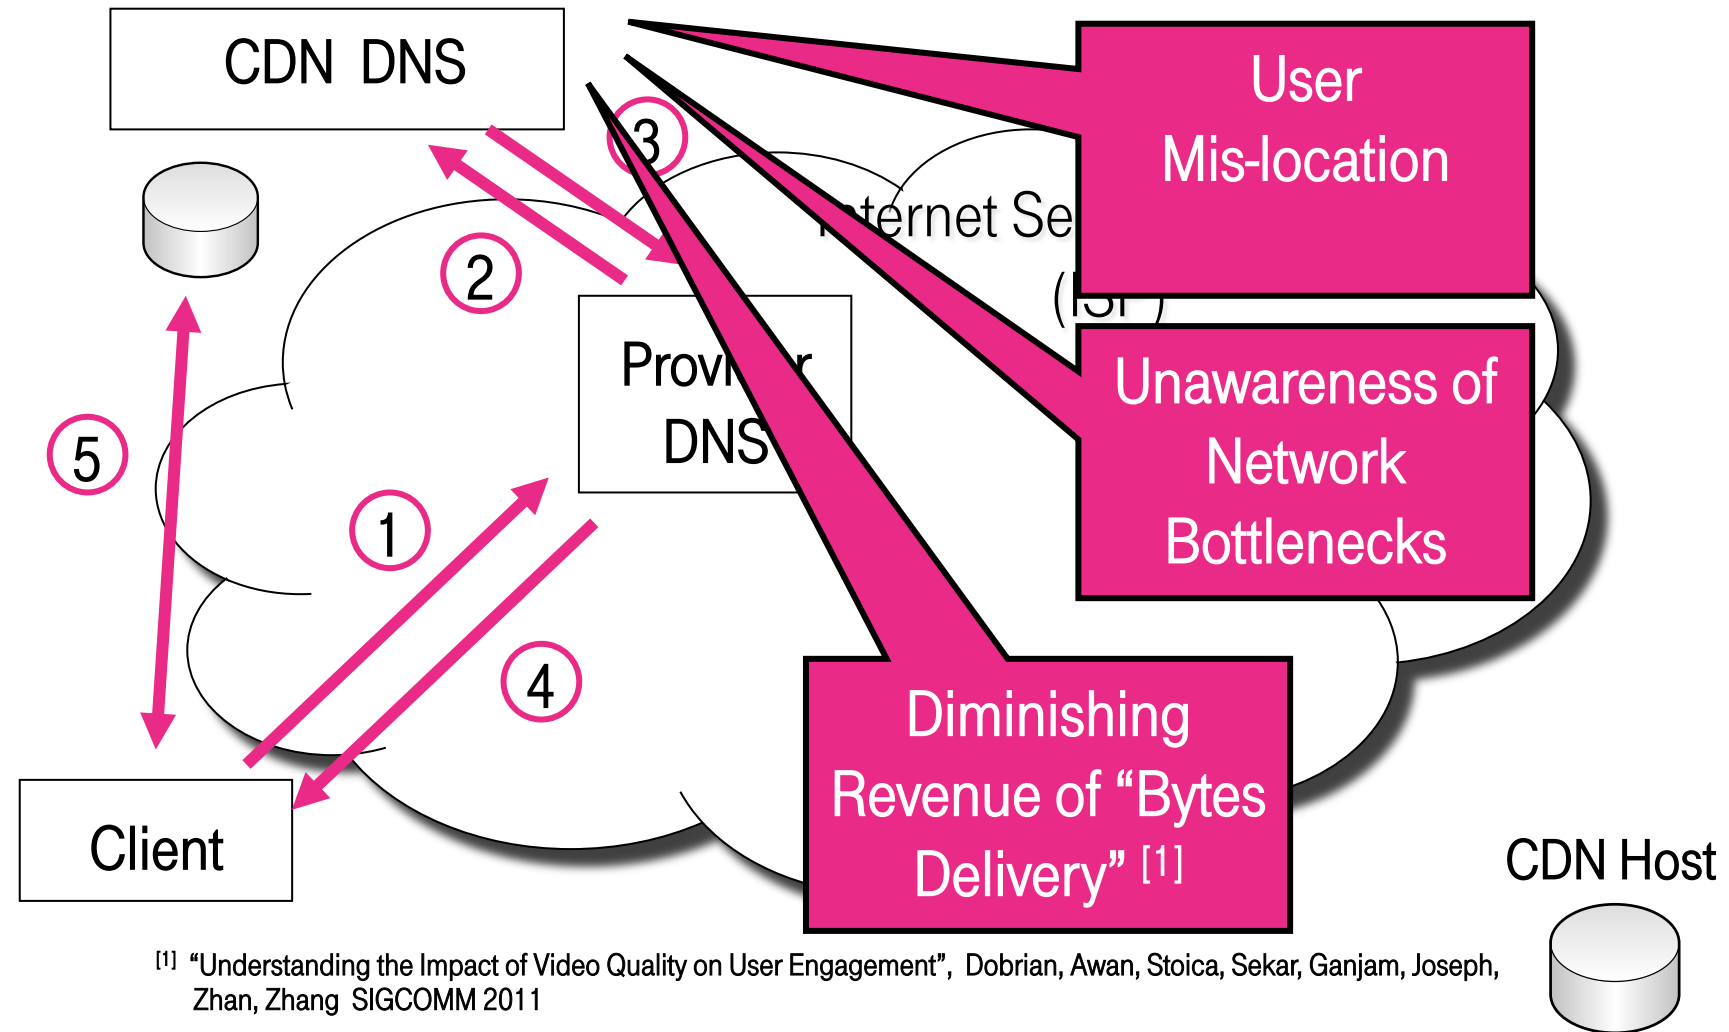
[1] "Comparing DNS Resolvers in the Wild", Ager, Muehlbauer, Smaragdakis, Uhlig, IMC 2010

→ See also Google Client IP IETF draft

Content Distribution Challenges




Content Distribution Challenges



→ See also Hybrid-CDNs, e.g. Akamai NetSession

Grand Challenge: Content-aware Traffic Engineering

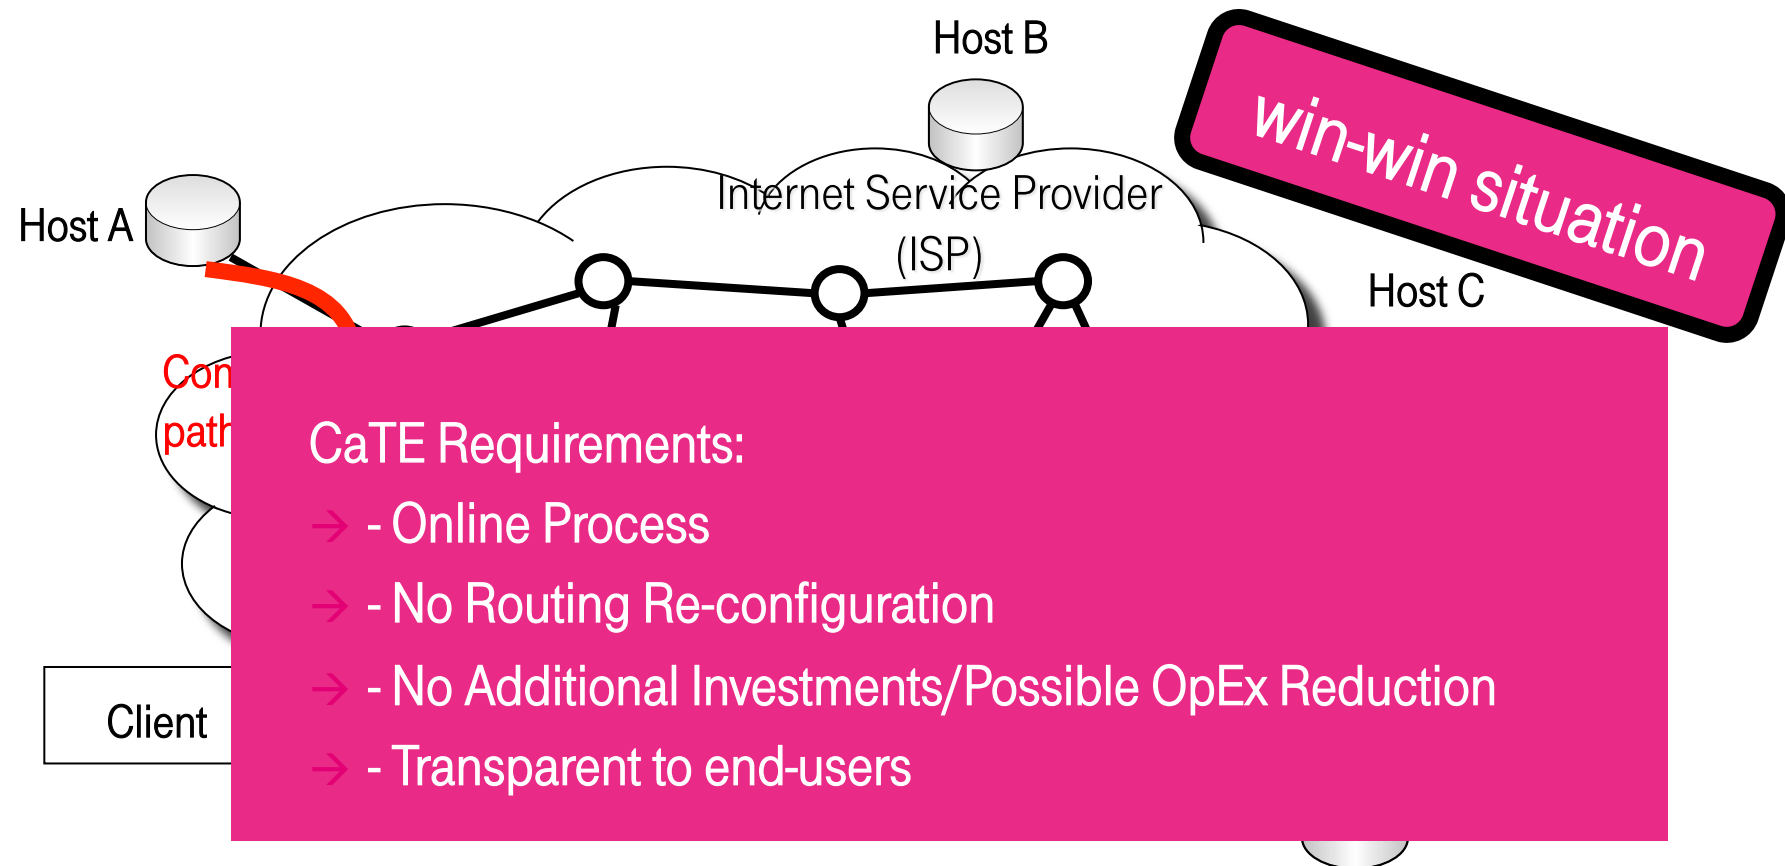
Dynamically adapts traffic demand by taking advantage of server and path diversity, and ISP information!



[1] "Content-aware Traffic Engineering", Frank, Poese, Smaragdakis, Uhlig, Feldmann SIGMETRICS 2012 (extended abstract)

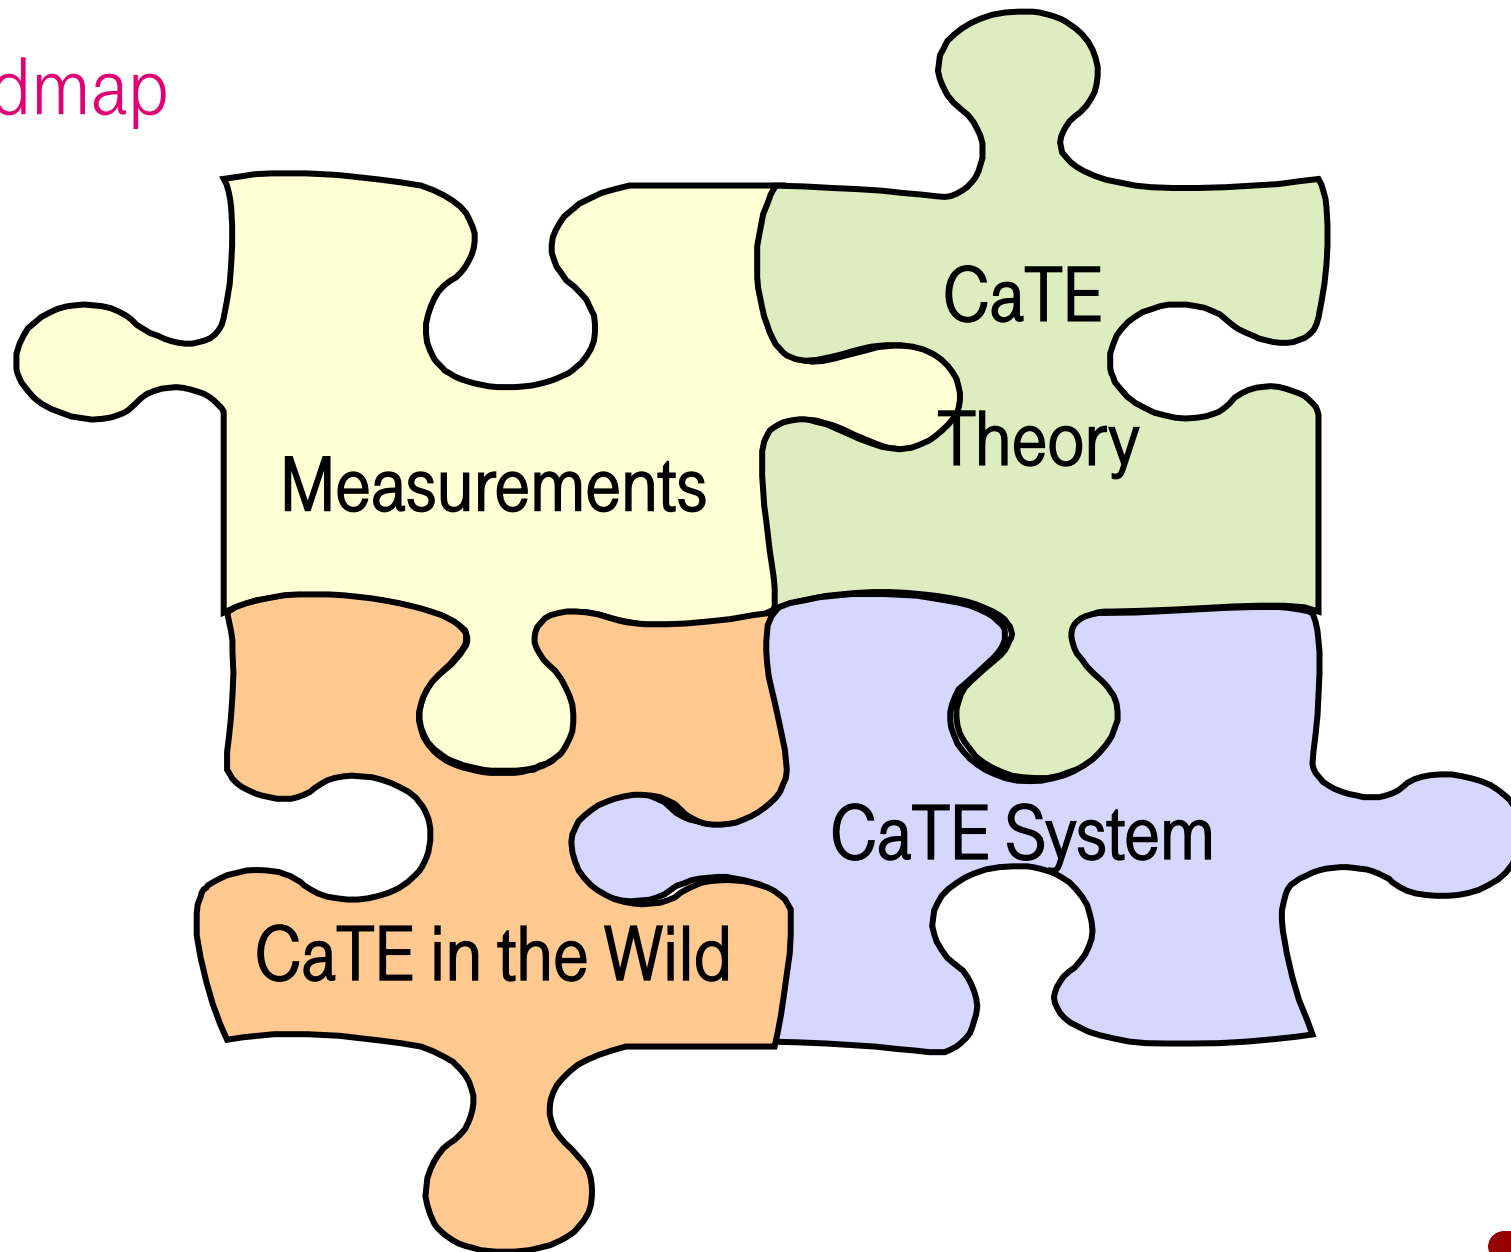
Grand Challenge: Content-aware Traffic Engineering


Dynamically adapts traffic demand by taking advantage of server and path diversity, and ISP information!



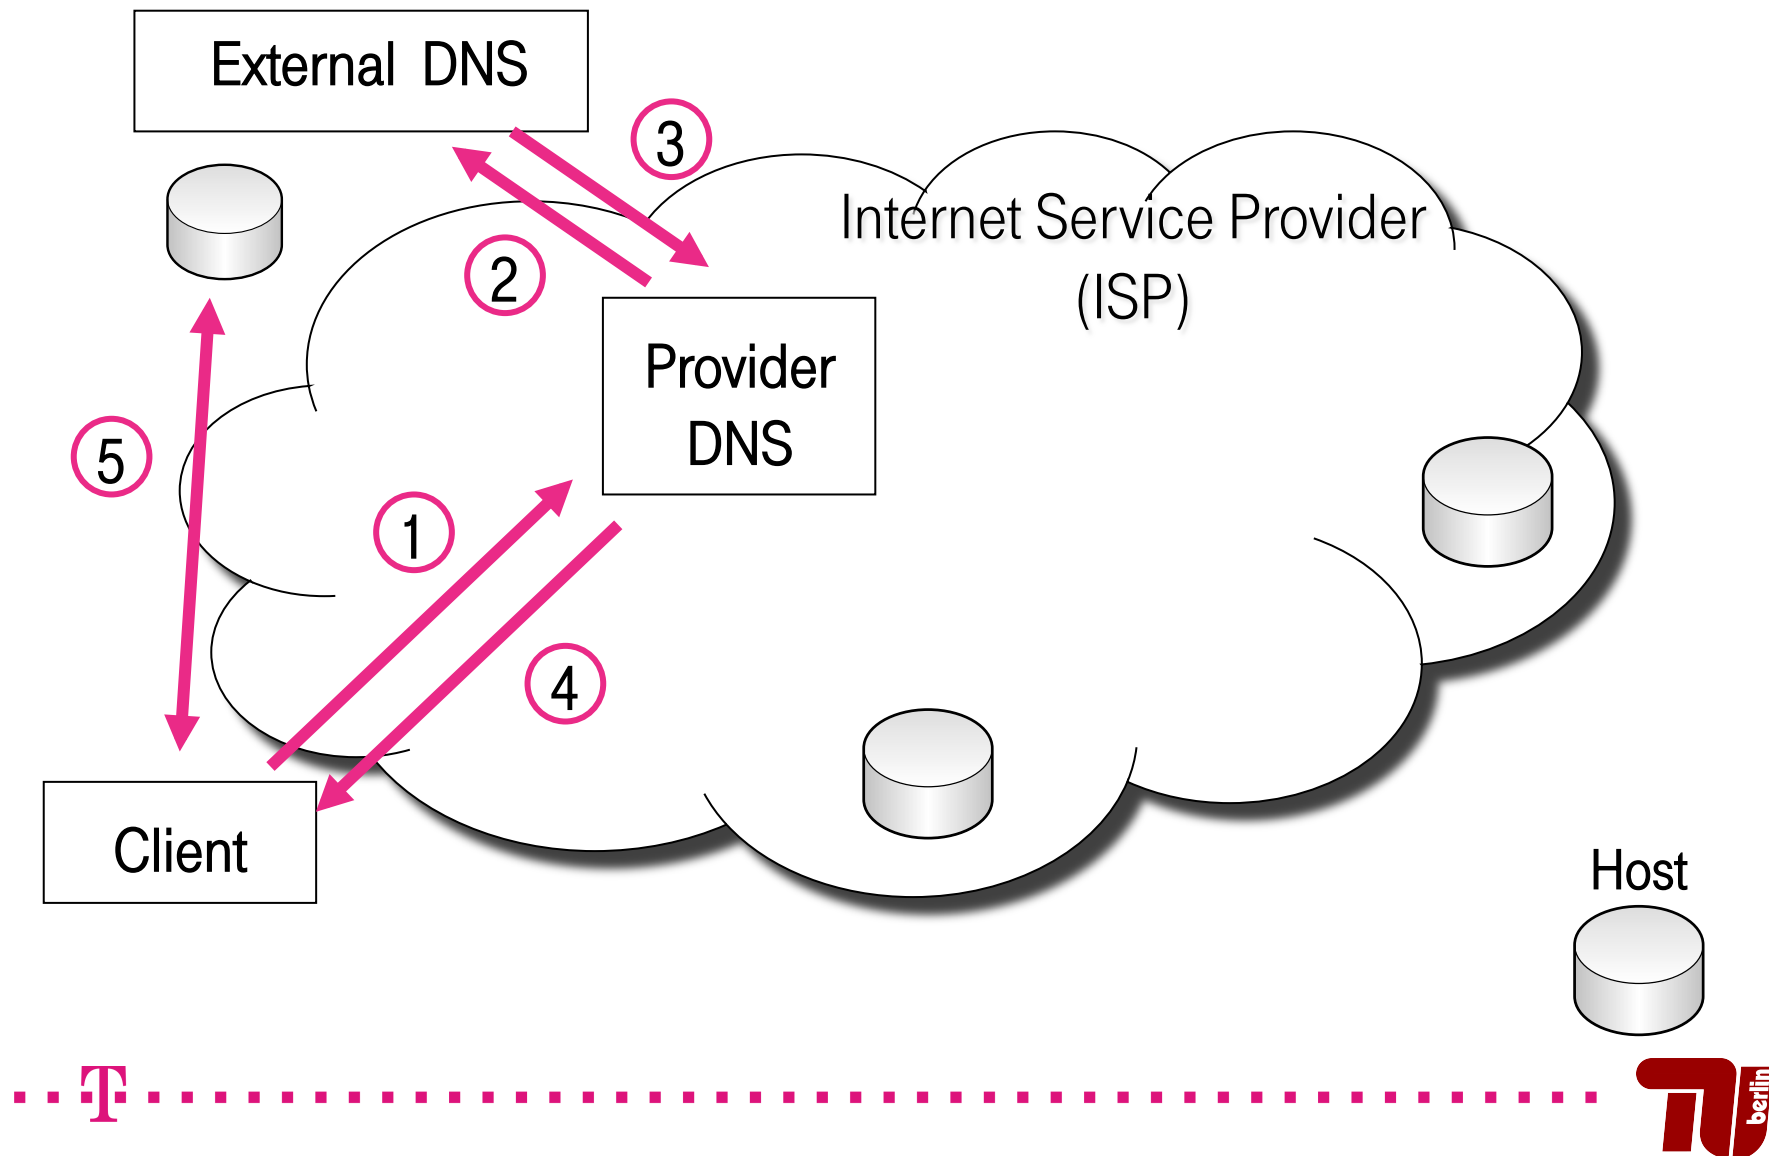
[1] "Content-aware Traffic Engineering", Frank, Poesse, Smaragdakis, Uhlig, Feldmann, Maggs, CCR Juny 2012, SIGMETRICS 2012.

Roadmap

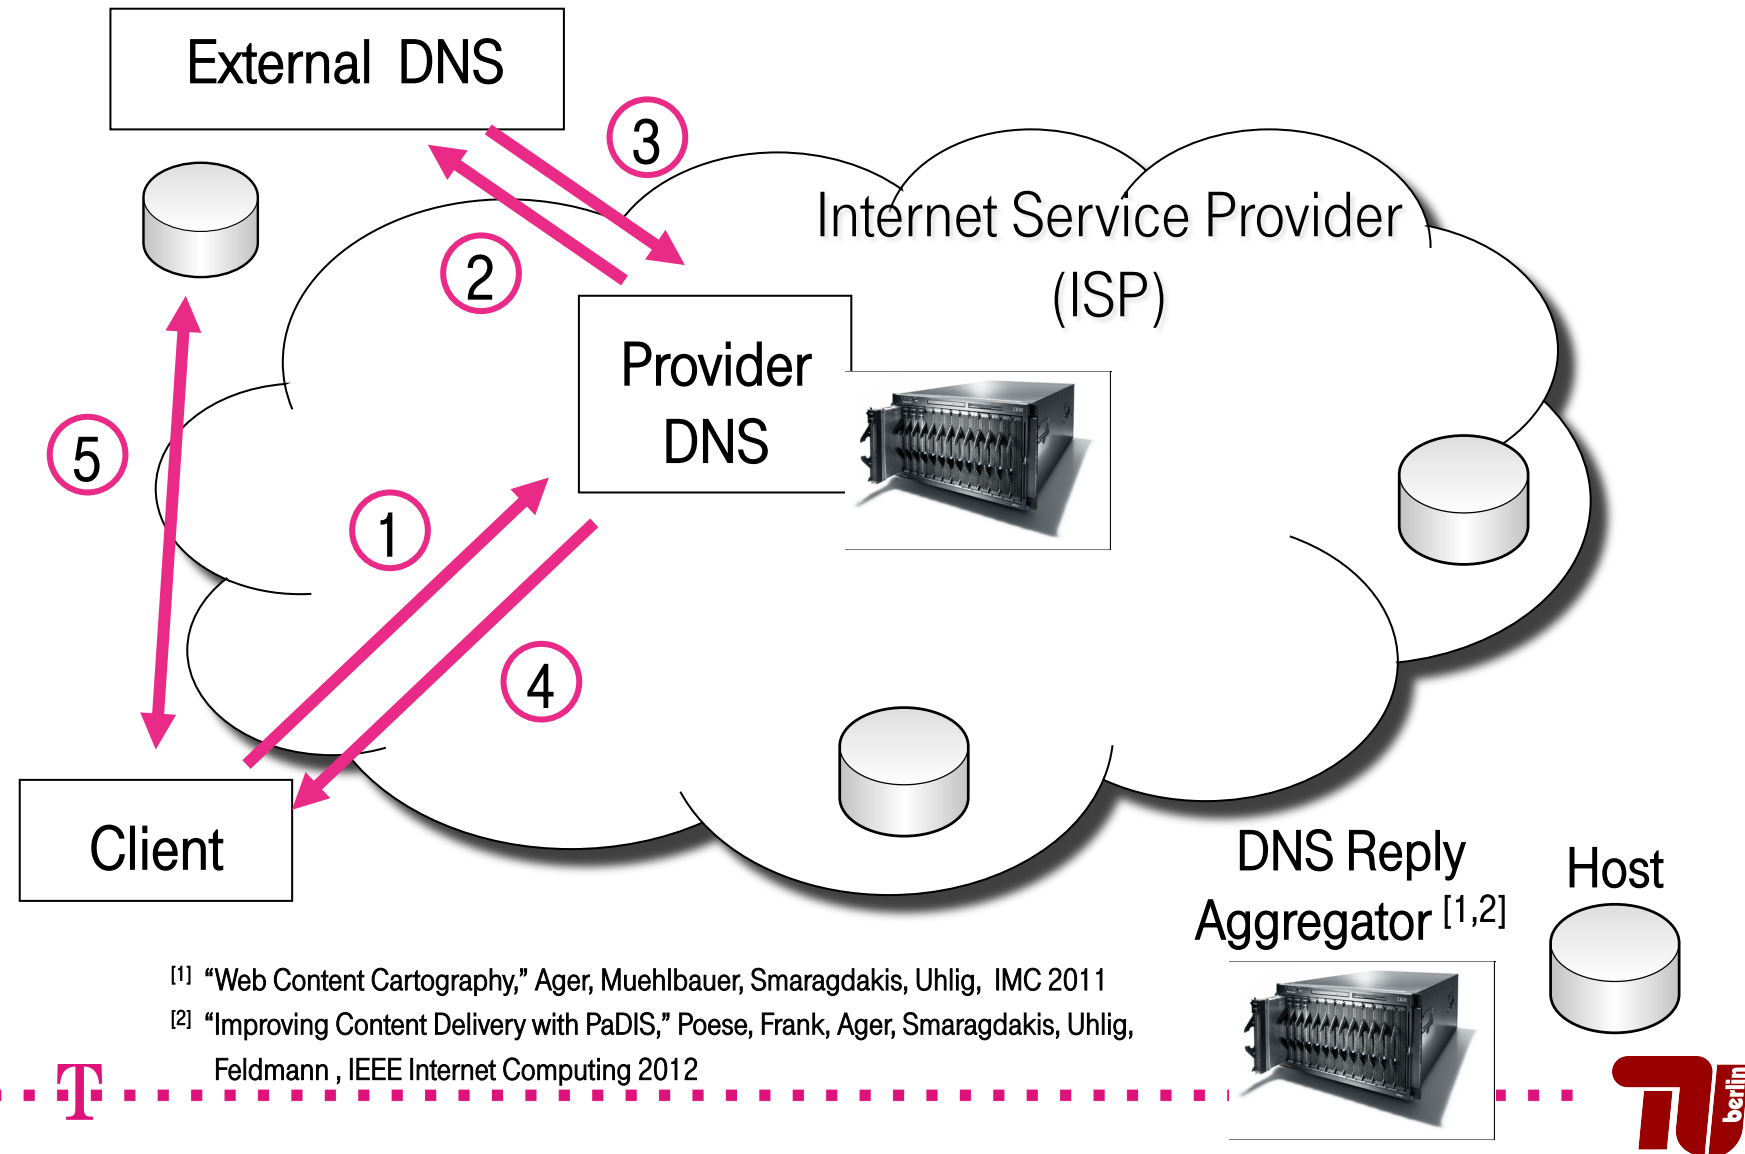




Monitoring CDN Server Diversity



Monitoring CDN Server Diversity



Reply Anatomy

→ Requesting a photo from Facebook

More than 60% of websites (>30% of traffic) redirect to at least 5 non-original servers [2]

More than 60% of the traffic is HTTP! [1]

```
$ dig photos-h.ak.fbcdn.net
```

```
; <<>> DiG 9.7.0-P1 <<>> photos-h.ak.fbcdn.net
```

```
:: QUESTION SECTION:
```

```
;photos-h.ak.fbcdn.net.      IN      A
```

2nd Level Domain → Application **facebook**

```
:: ANSWER SECTION:
```

```
photos-h.ak.fbcdn.net. 6099      IN      CNAME  photos-d.ak.facebook.com.edgesuite.net.
```

```
photos-d.ak.facebook.com.edgesuite.net. 20492 IN CNAME a998.mm1.akamai.net.
```

```
a998.mm1.akamai.net. 7      IN      A      62.41.85.74
```

```
a998.mm1.akamai.net. 7      IN      A      62.41.85.90
```

```
...
```



Redirection → Content Delivery Network




[1] "On Dominant Characteristics of Internet Traffic", Maier, Feldmann, Paxson IMC 2009

[2] "Understanding Web Complexity," Butkiewicz, Madhyastha, Sekar, IMC 2011




CDNized Traffic Dynamics



→ Popularity of content providers seems to have a diurnal pattern




Diversity of Paths



→ More than 40% of the HTTP traffic can be download from at least 3 different network locations



Network Diversity of CDN Servers




- Significant Network Diversity of servers over time for top content providers
- During peak hours more traffic is delivered and a more diverse set of servers is used by content providers

→ Typically only one location is returned to the end user (low TTL) by the CDN



Network Diversity of CDNs



- Significant Network Diversity of servers over time for top content providers
- During peak hours more traffic is delivered and a more diverse set of servers is used by content providers
- Typically only one location is returned to the end user (low TTL) by the CDN

Turning Challenges into Opportunities

Bias the host selection by exploring diversity!



Turning Challenges into Opportunities

Bias the host selection by exploring diversity!

Utilize the DNS infrastructure!

- Transparent to users and applications
- Online Process (per request or up to TTL)
- No Routing Re-configuration



Turning Challenges into Opportunities

Bias the host selection by exploring diversity!

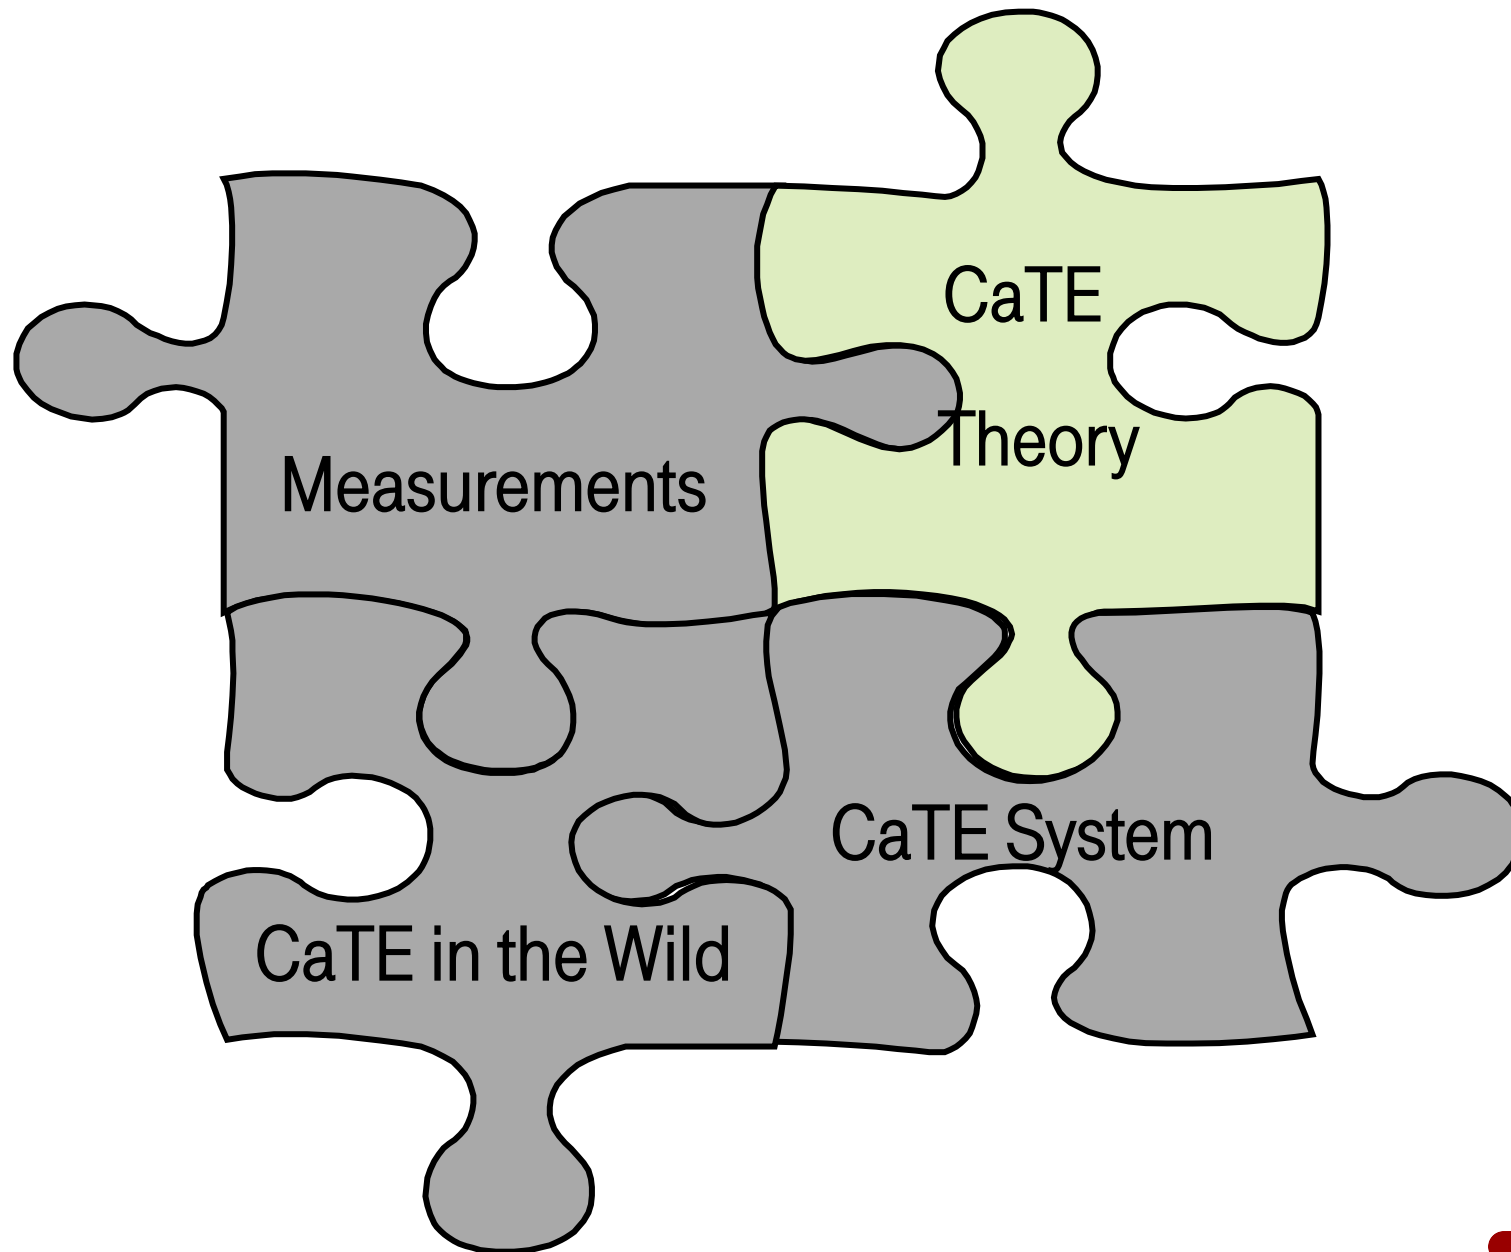
Utilize the DNS infrastructure!

- Transparent to Users and Applications
- Online Process
- No Routing Re-configuration


Utilize Strategic Advantage!

- Utilize User Location and Network Information
- Enable Revenue Sharing Negotiations
- Reduce Delivery and Network Expansion Costs





Traffic Engineering Model




$$\begin{array}{c} \text{Link Vector} \\ \mathbf{Y} \\ \left[\begin{array}{c} y_1 \\ y_2 \\ \dots \\ y_n \end{array} \right] \end{array} = \begin{array}{c} \text{Routing Matrix} \\ \mathbf{A} \\ \left[\begin{array}{c} 0 \ 0 \ \dots \ 1 \\ 1 \ 0 \ \dots \ 0 \\ \dots \\ 1 \ 0 \ \dots \ 0 \end{array} \right] \end{array} \cdot \begin{array}{c} \text{Demand Vector} \\ \mathbf{X} \\ \left[\begin{array}{c} x_1 \\ x_2 \\ \dots \\ x_n \end{array} \right] \end{array}$$

Change A (routing) such that a traffic engineering goal is achieved



Content-aware Traffic Engineering Model



$$Y = A \cdot X$$


$$\begin{bmatrix} y_1 \\ y_2 \\ \dots \\ y_n \end{bmatrix} = \begin{bmatrix} 0 & 0 & \dots & 1 \\ 1 & 0 & \dots & 0 \\ \dots & & & \\ 1 & 0 & \dots & 0 \end{bmatrix} \cdot \begin{bmatrix} x_1 \\ x_2 \\ \dots \\ x_n \end{bmatrix}$$

Given the routing and content demand, change the flows such that a traffic engineering goal is achieved

... T ...



Content-aware Traffic Engineering Model



$$Y = A \cdot X$$


$$Y_d + Y_s = A \cdot (X_d + X_s)$$

Set of CDNs to consider


Background Traffic

Content-aware Traffic Engineering Model

$$Y_d + Y_s = A \cdot (X_d + X_s)$$



+




Volume of Content Providers


Content Providers Server
Diversity




Network Load Balancing



Network Load Balancing



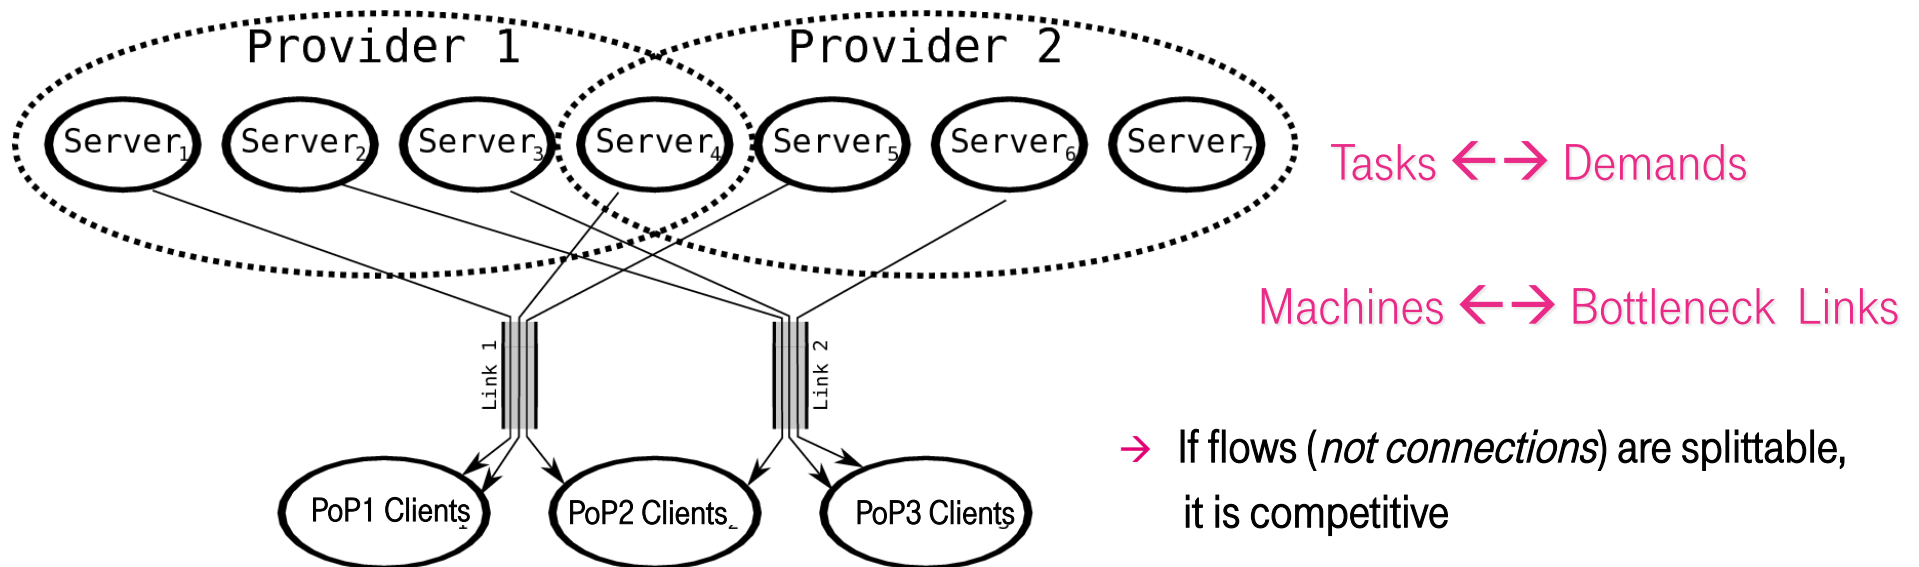
Network Load Balancing



CaTE: The Online Problem

Assign user requests to servers such that the flows from servers to users are well balanced

Reduction to the **Restricted Machine Load Balancing Problem** ^[1]



^[1] "The Competitiveness of Online Assignments", Azar, Naor, Rom, J. Algorithms 1995

CaTE: The Offline Problem

Given aggregated network statistics and content provider mix,
estimate the potential gain using CaTE

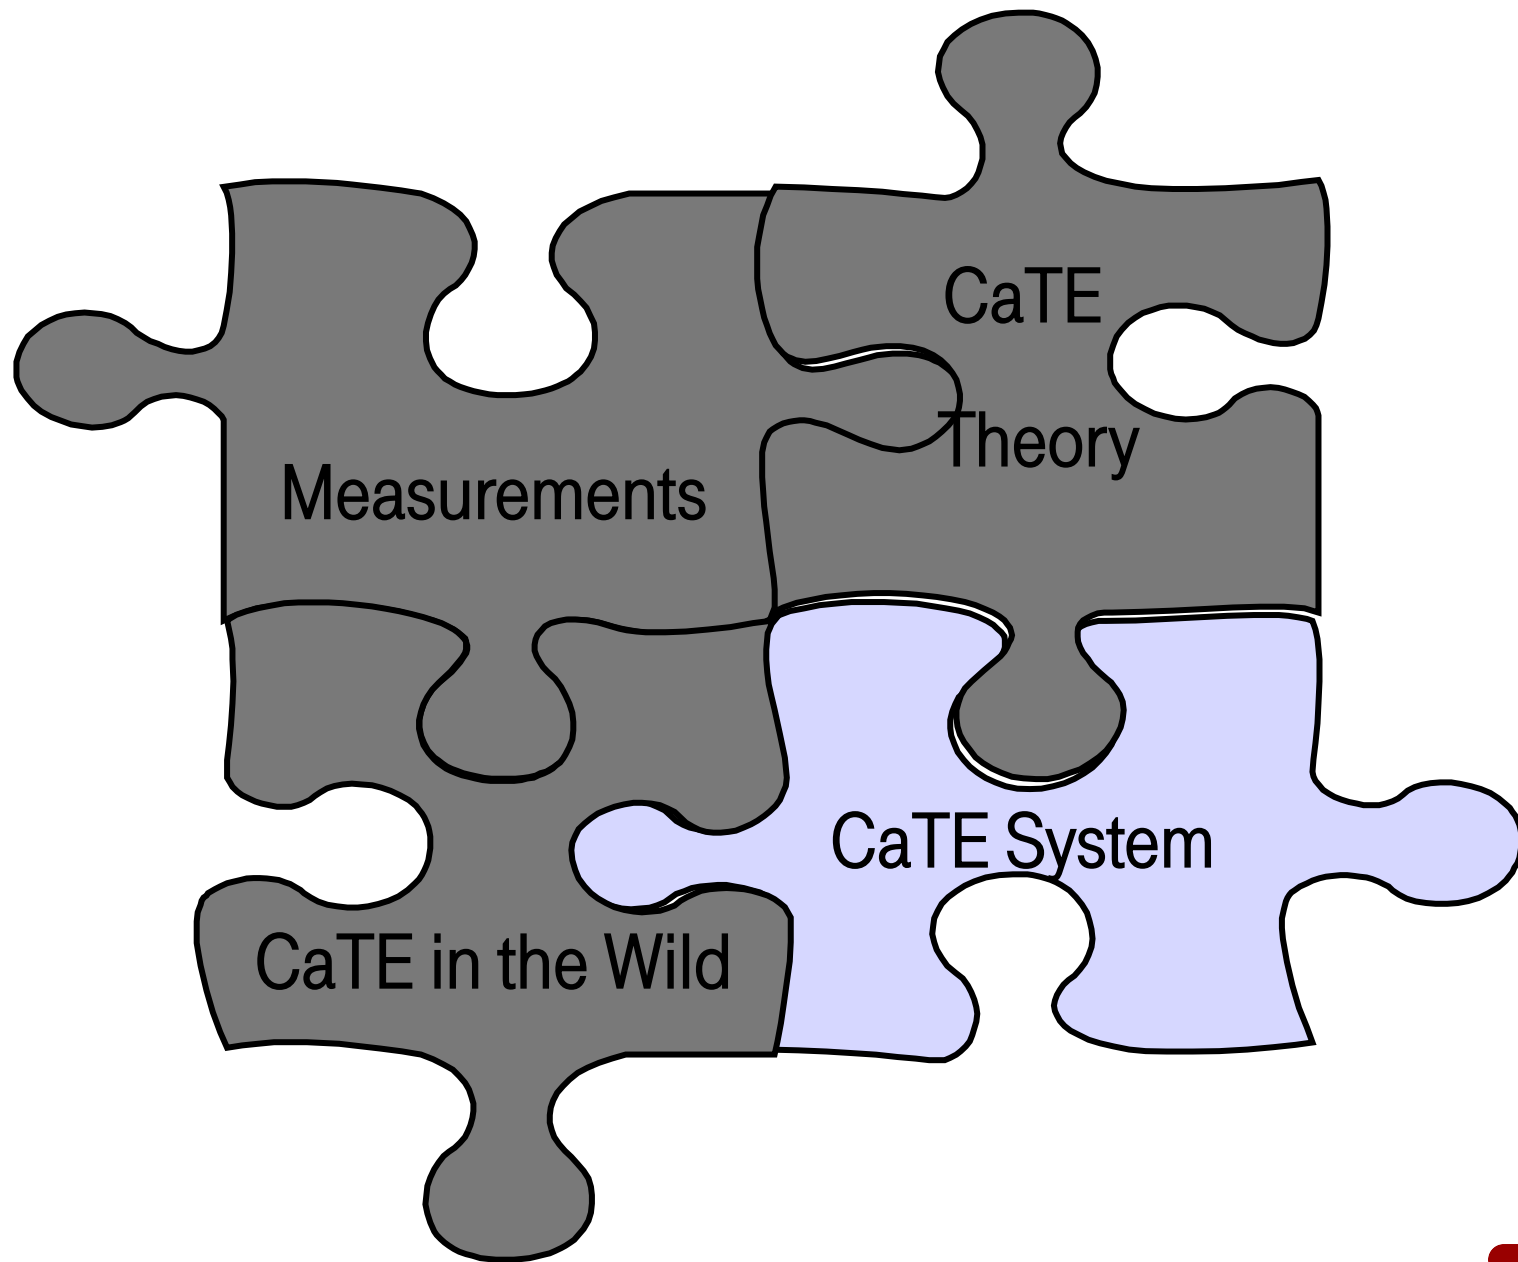
Linear Programming (LP):

- If flows are splittable (fractional LP), in polynomial time – slow for large networks and traces that span multiple weeks
- If flows are not splittable then it is NP-hard and a 2-approximation polynomial time rounding algorithm exists

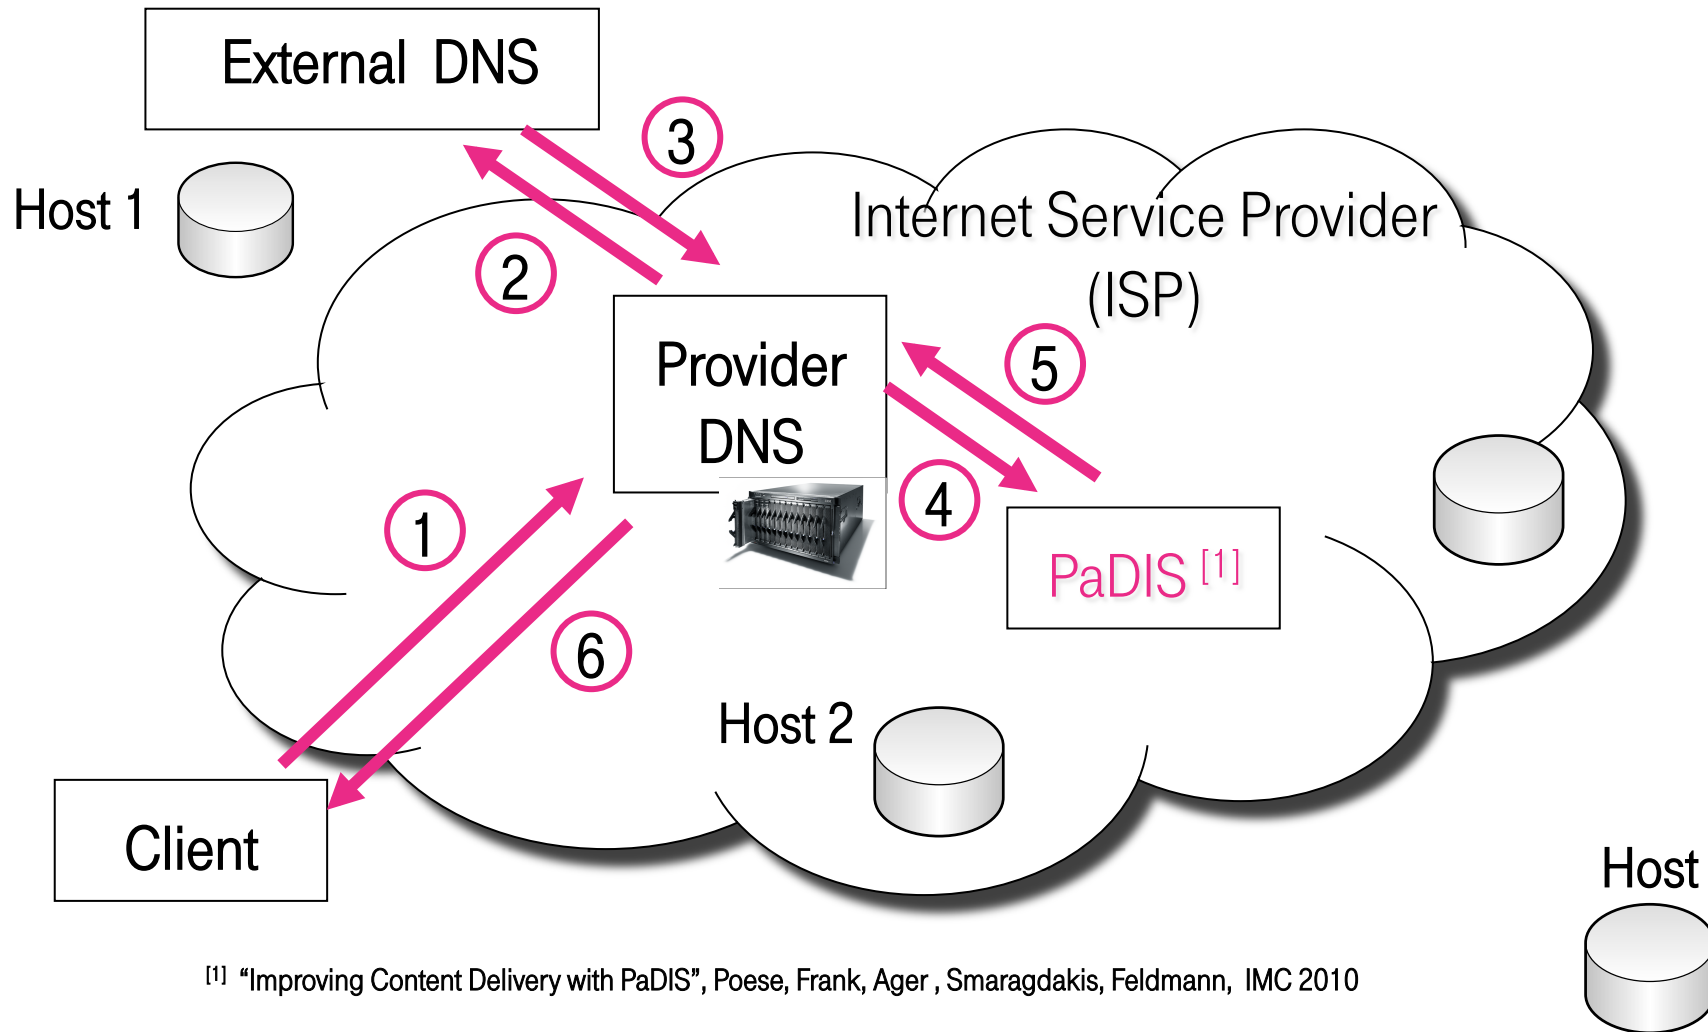
Greedy Heuristic:

- Speed up convergence: Sort flows in decreasing order
- Typically the largest flows are those of the largest providers
- Re-balance until convergence
- Very fast convergence, very good approximation





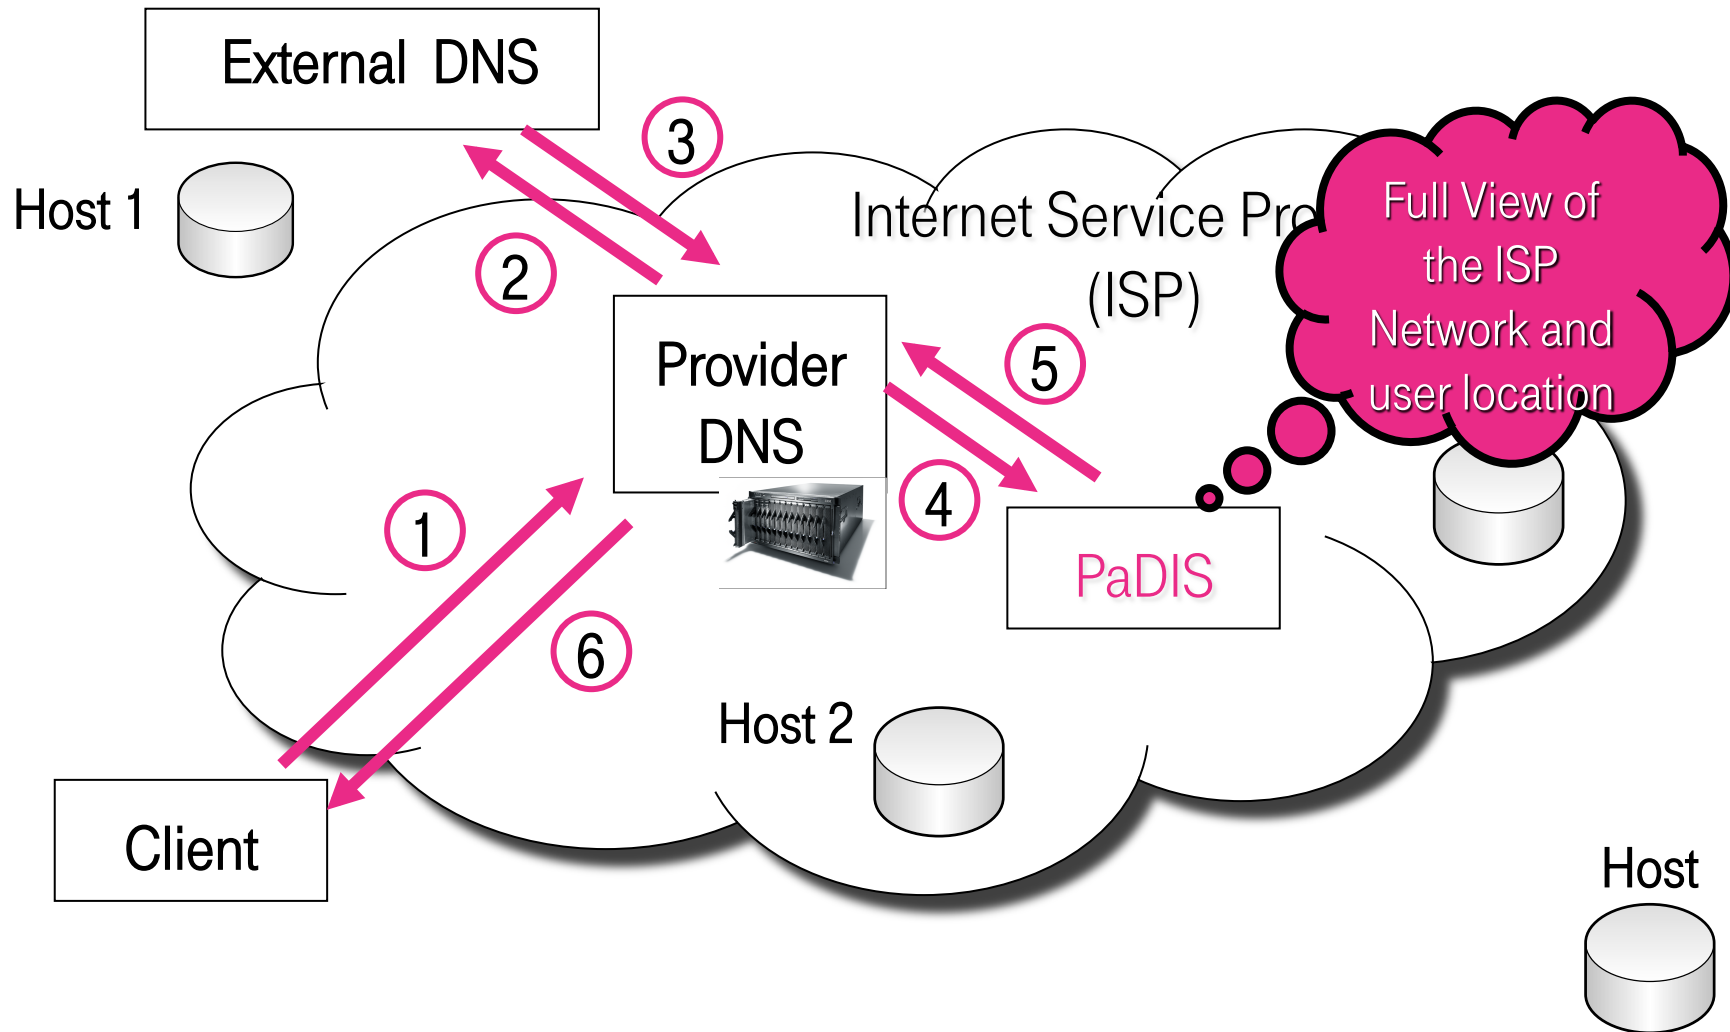
PaDIS: Provider-aided Distance Information System



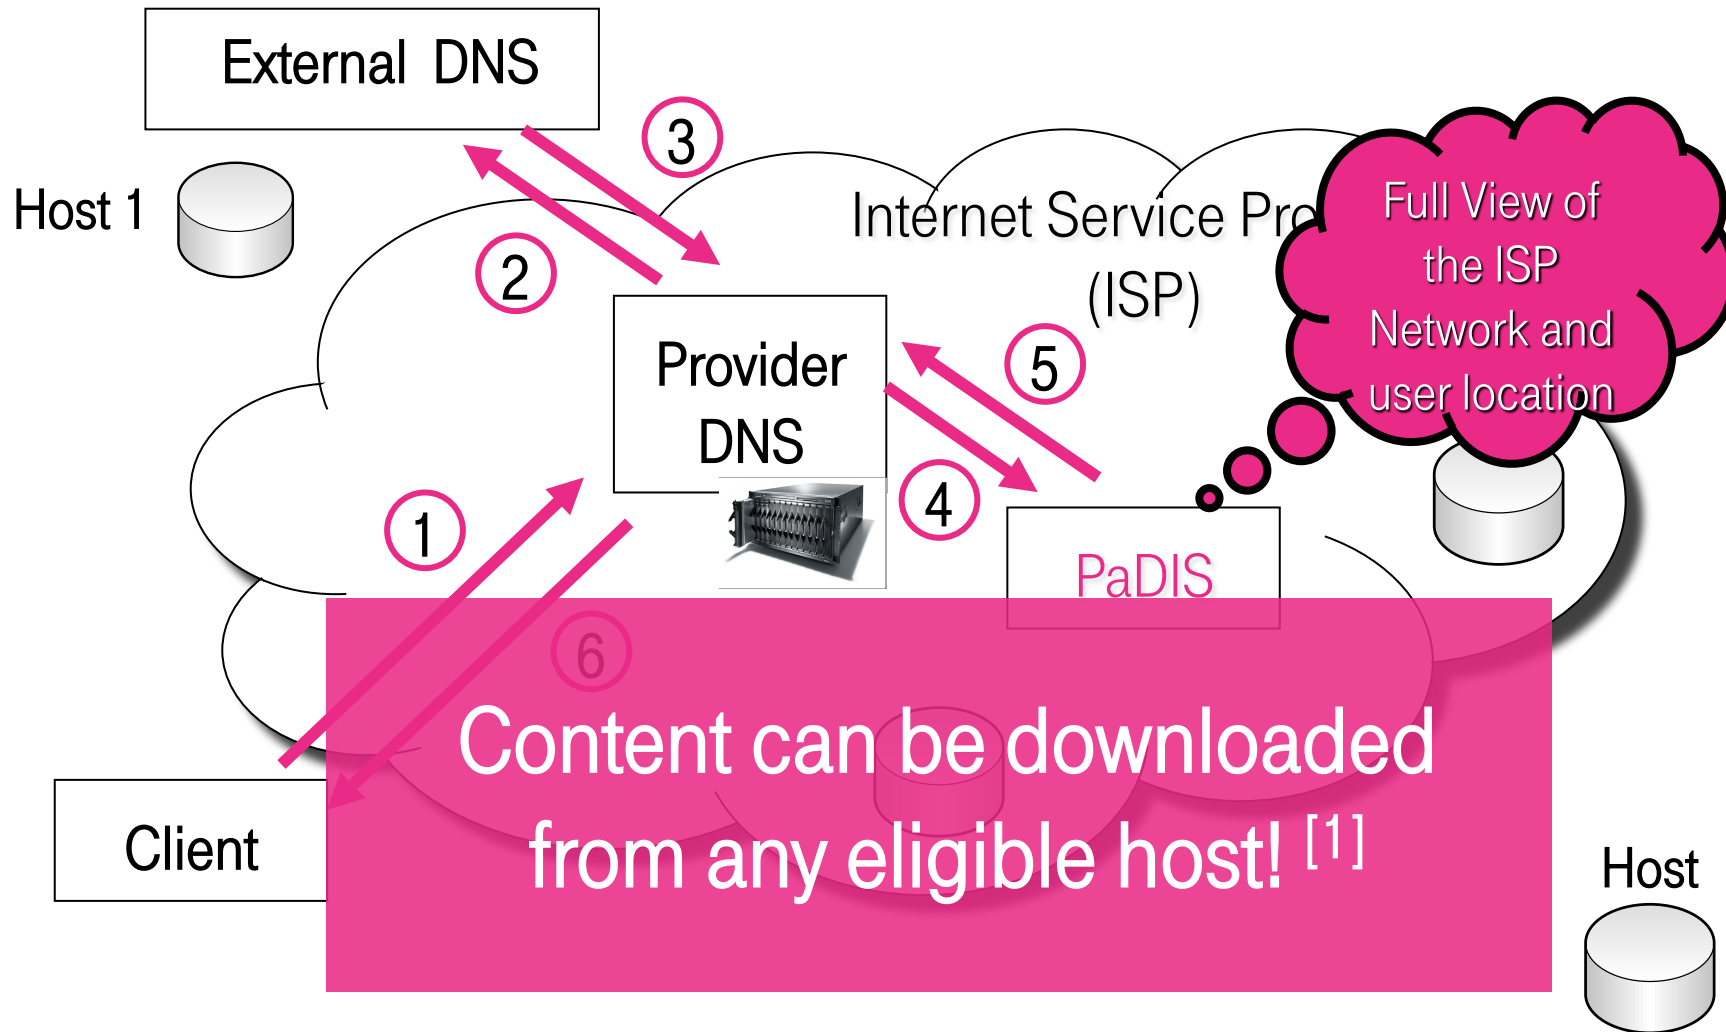
[1] "Improving Content Delivery with PaDIS", Poese, Frank, Ager, Smaragdakis, Feldmann, IMC 2010



PaDIS: Provider-aided Distance Information System

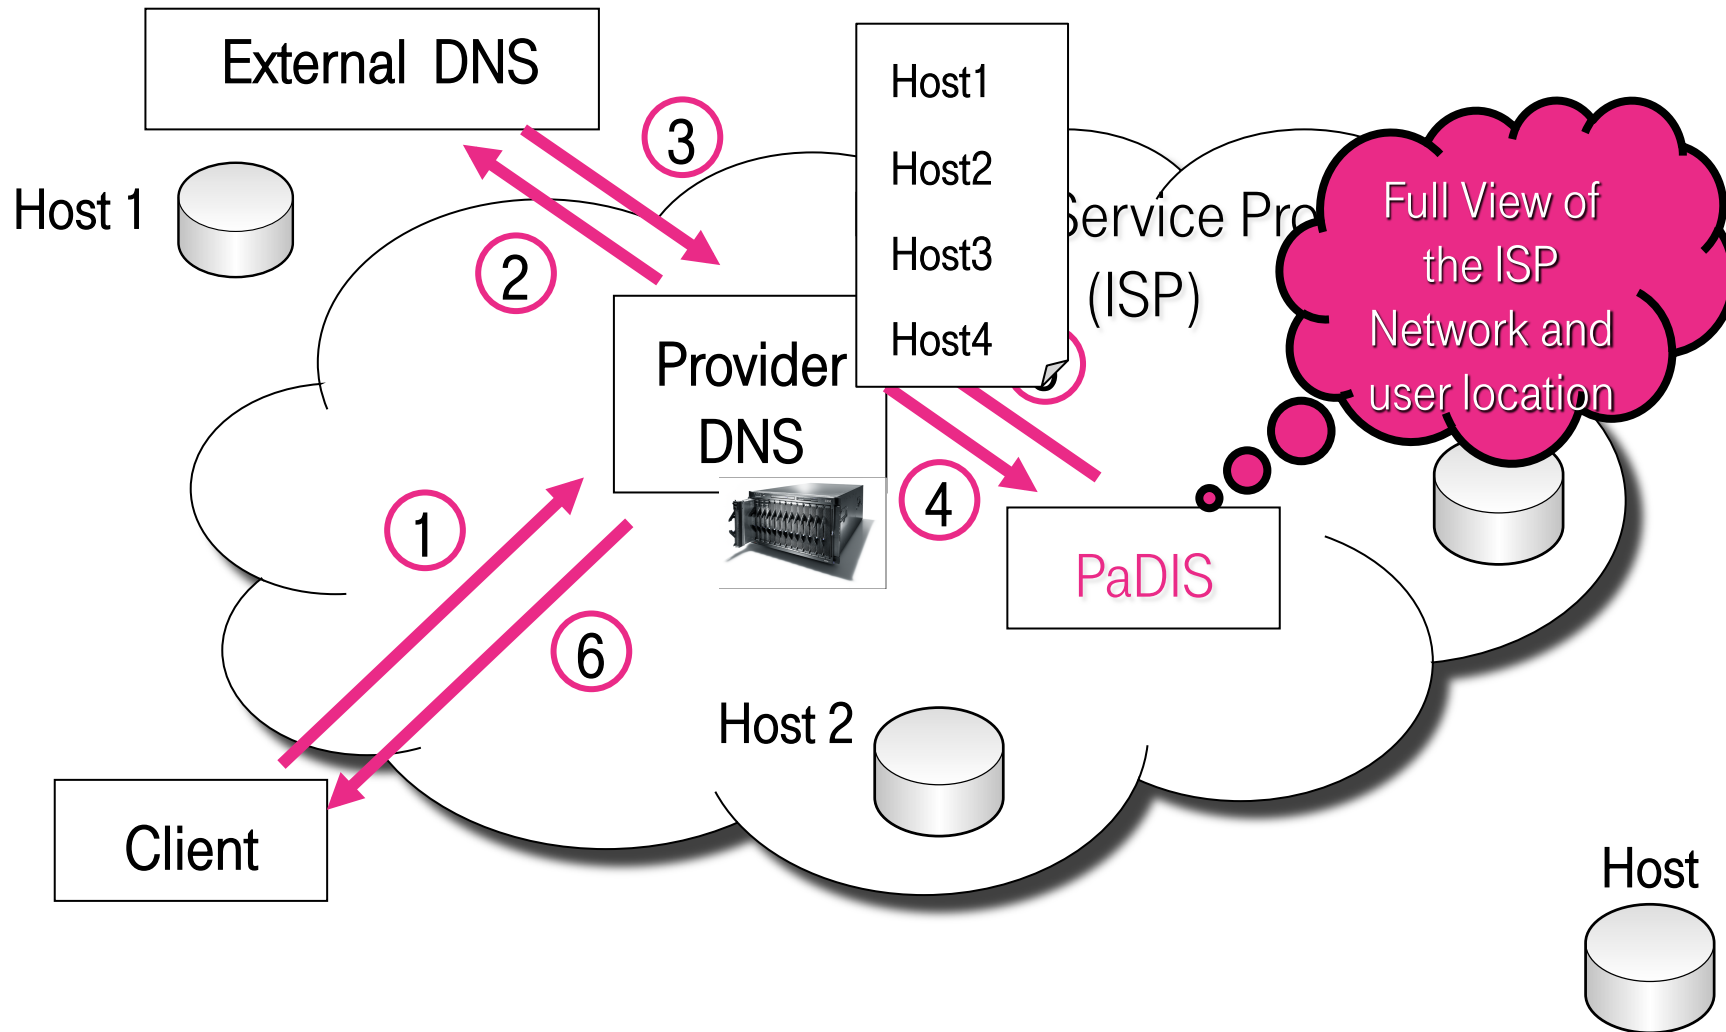


PaDIS: Provider-aided Distance Information System

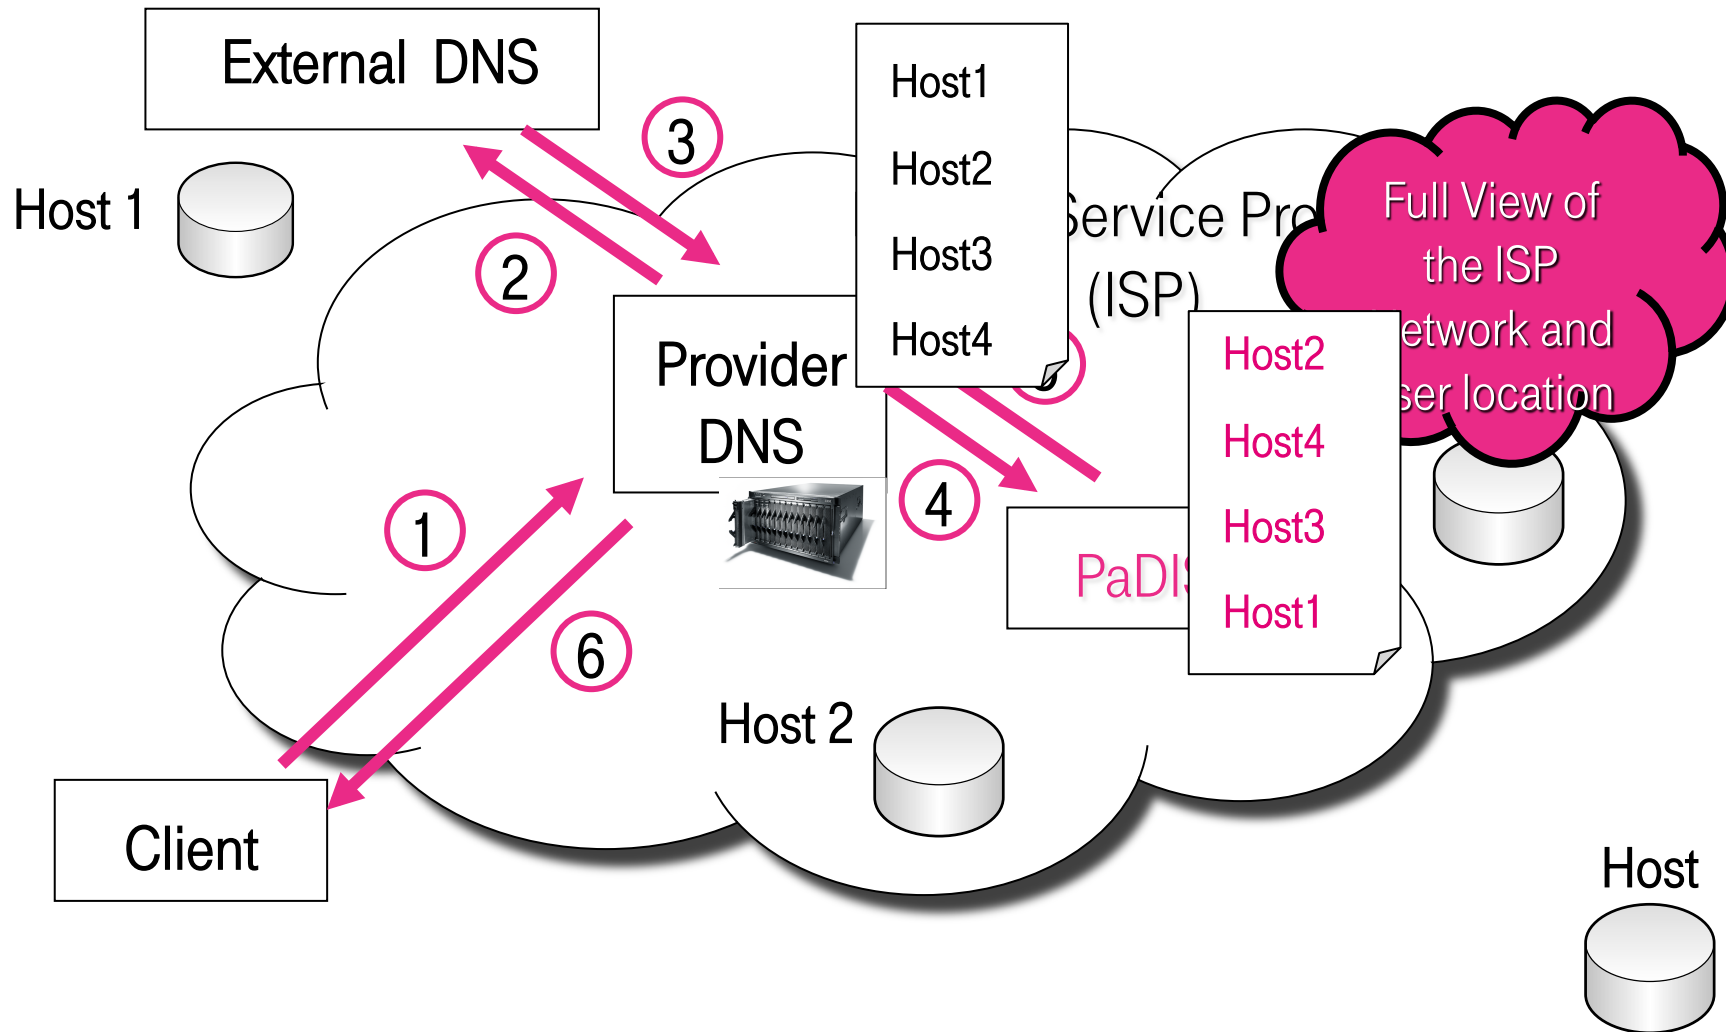


^[1] "Content Delivery Networks: Protection or Threat?", Triukose, Al-Qudah, Rabinovich, ESORICS 2009

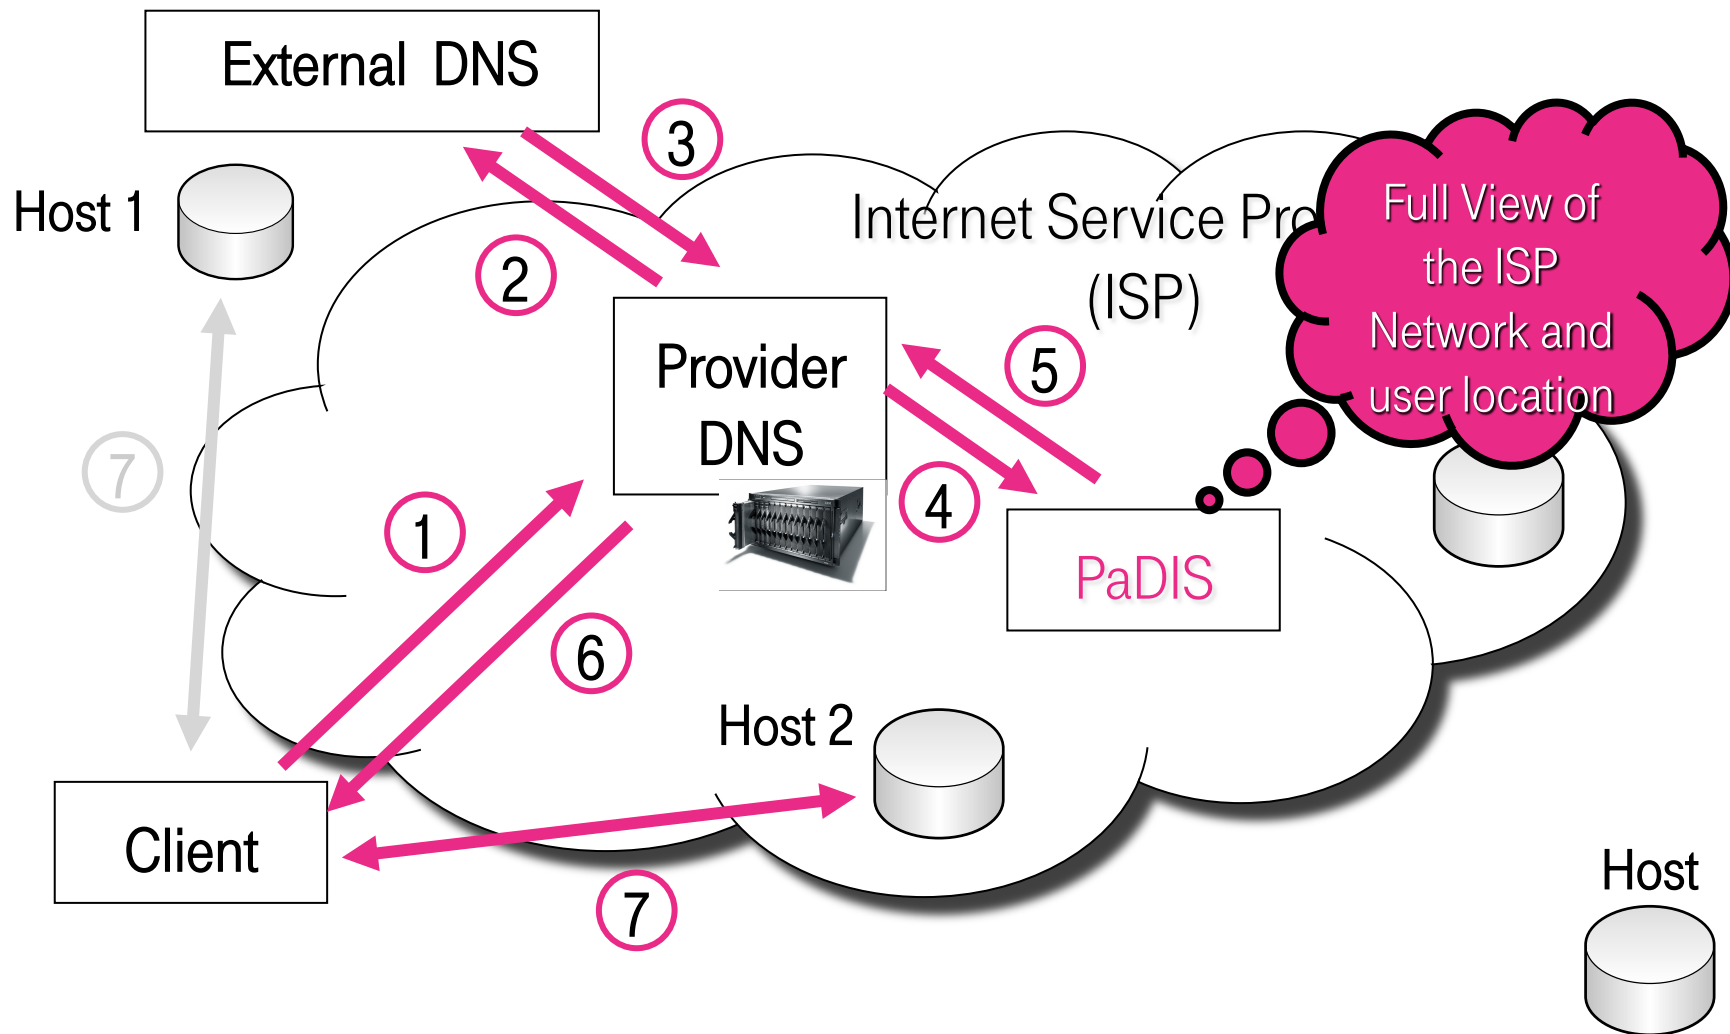
PaDIS: Provider-aided Distance Information System



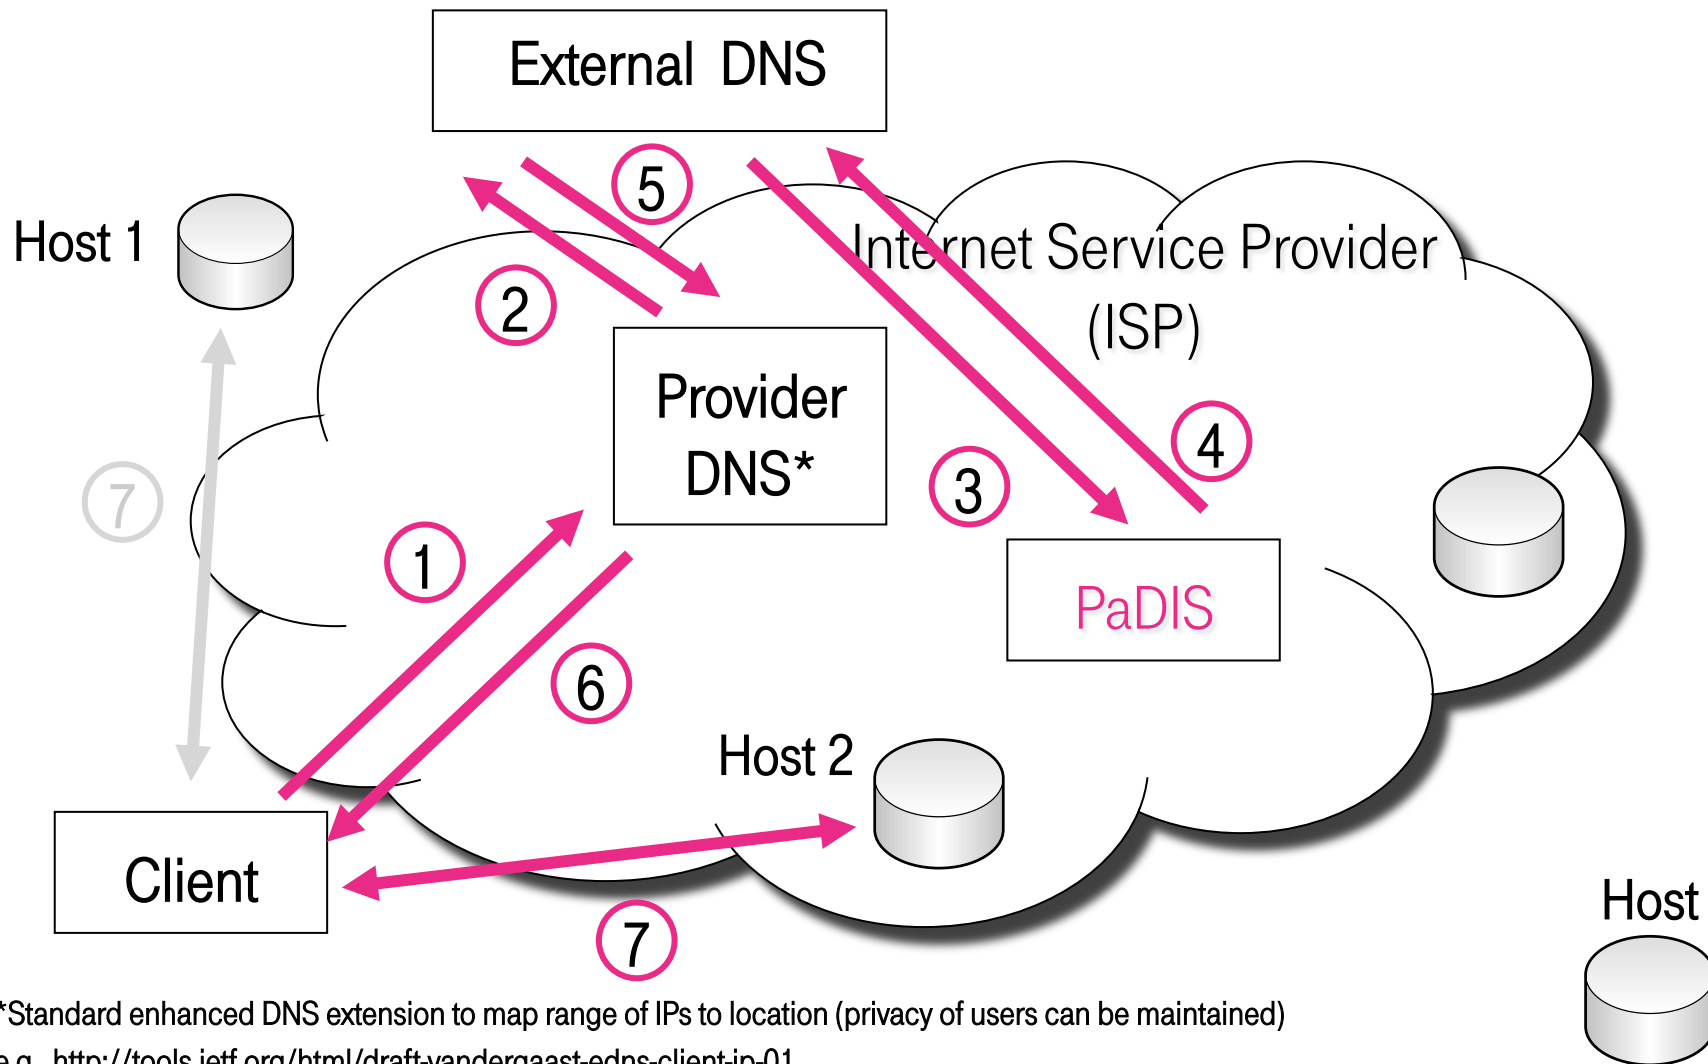
PaDIS: Provider-aided Distance Information System



PaDIS: Provider-aided Distance Information System

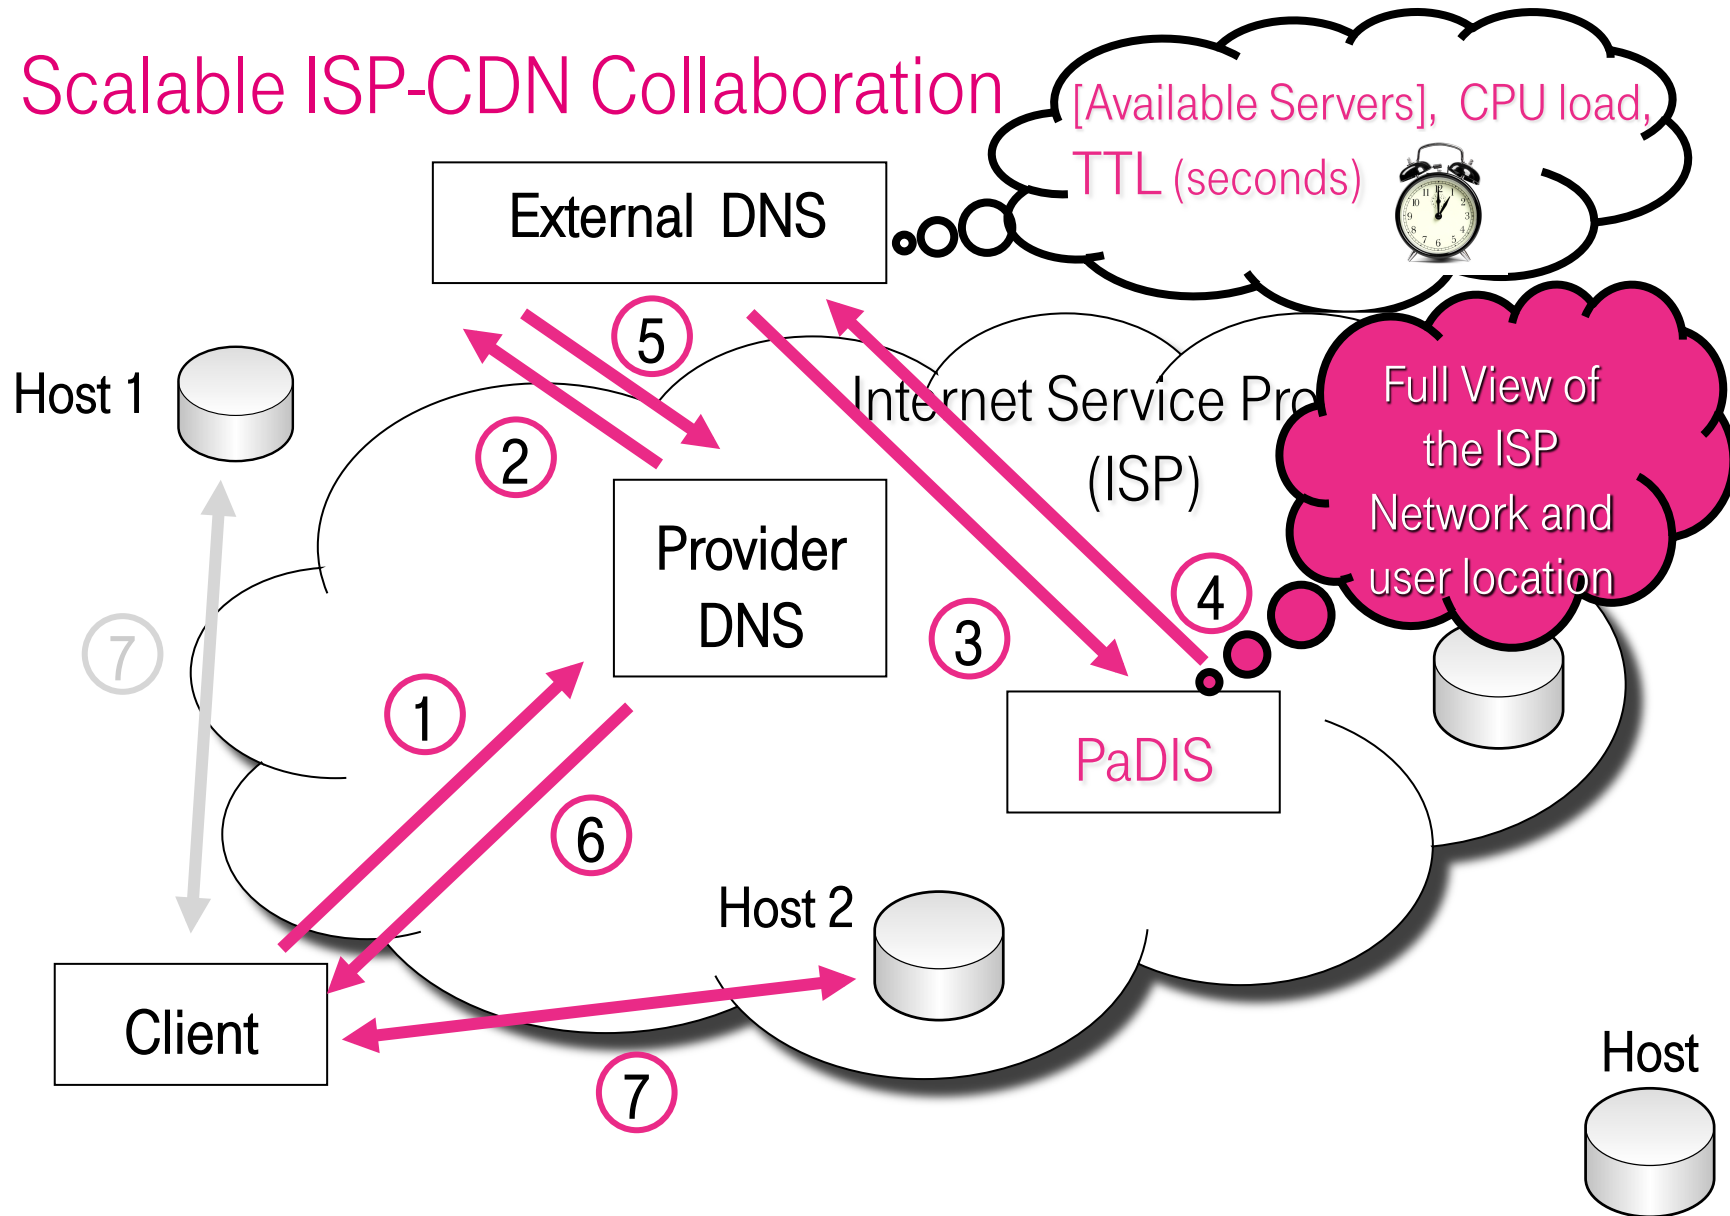


ISP-CDN Collaboration

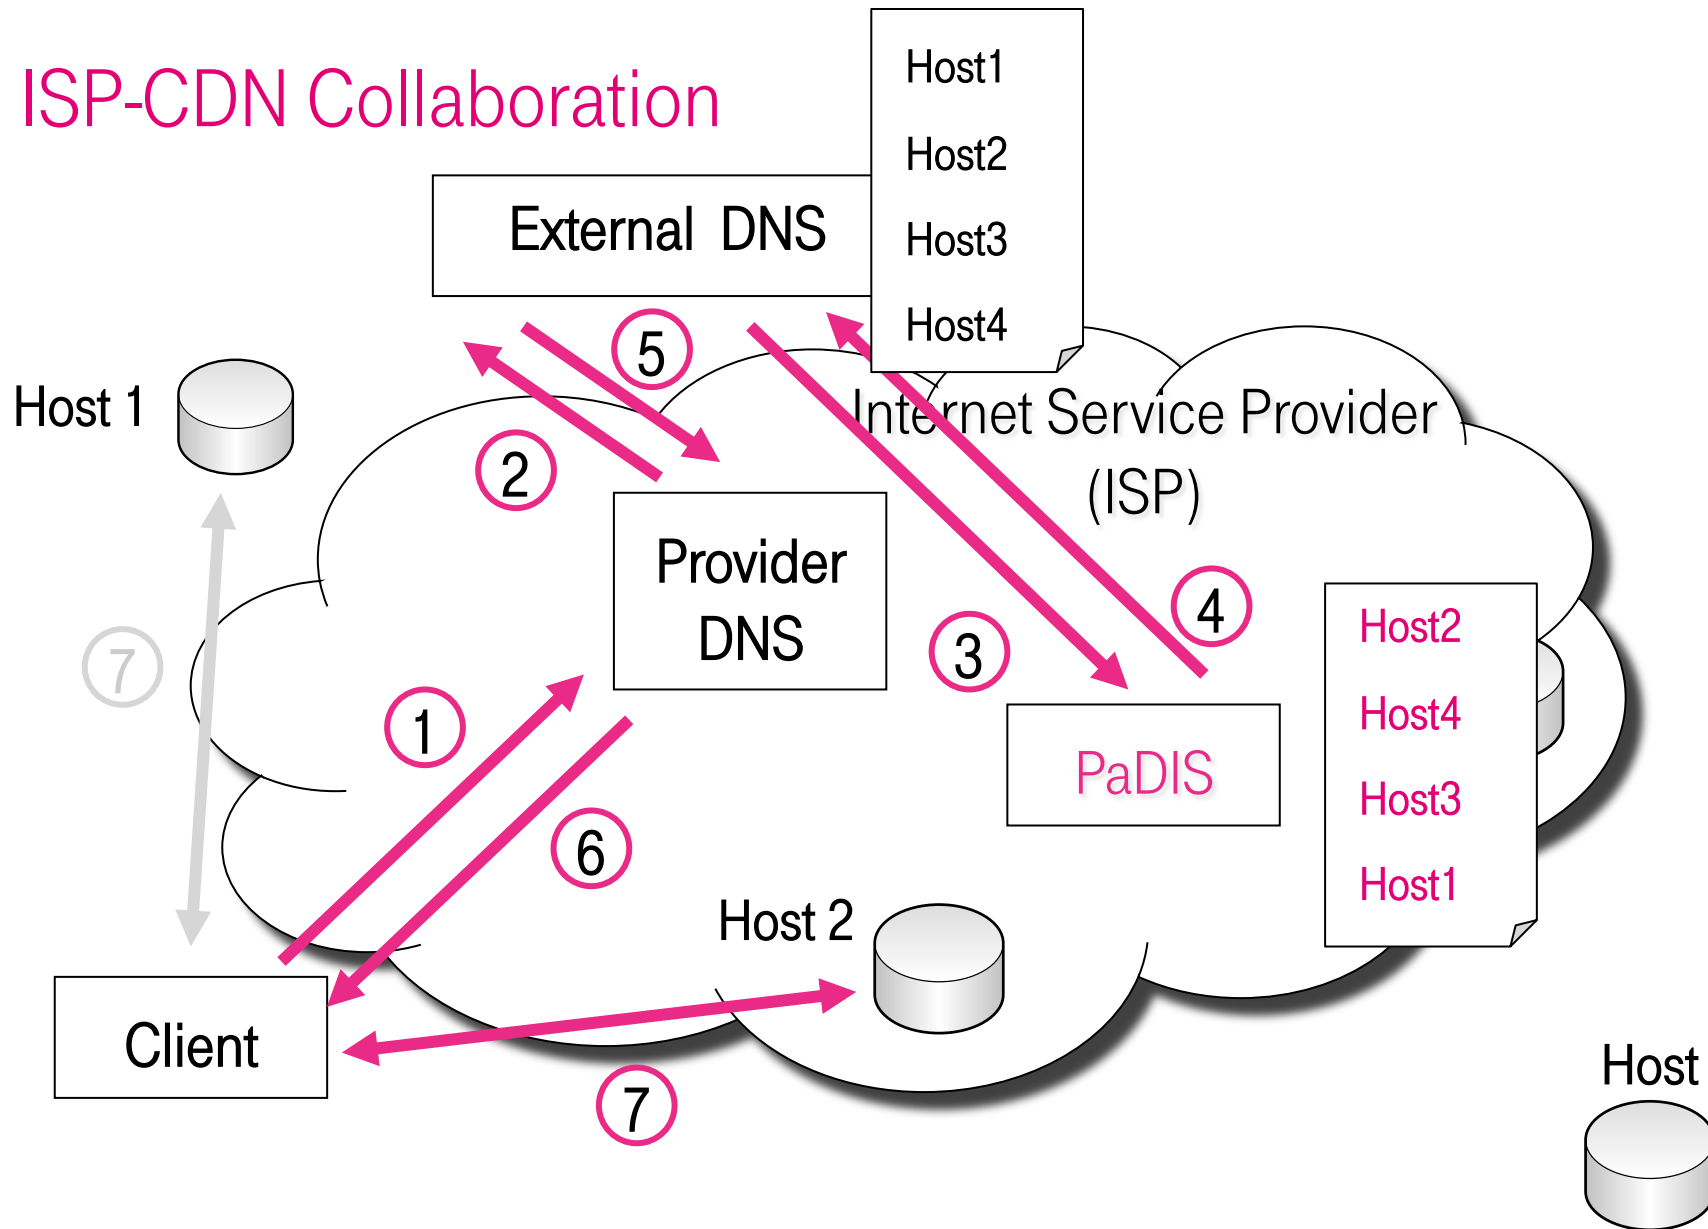


*Standard enhanced DNS extension to map range of IPs to location (privacy of users can be maintained)
e.g. <http://tools.ietf.org/html/draft-vandergaast-edns-client-ip-01>

Scalable ISP-CDN Collaboration



ISP-CDN Collaboration



ISP-CDN Collaboration

The diagram illustrates the collaboration between an Internet Service Provider (ISP) and a Content Delivery Network (CDN) to optimize content delivery. The components and their interactions are as follows:

- Client:** The user's device that initiates the request.
- External DNS:** The public DNS service that provides the initial IP address for the content.
- Provider DNS:** The ISP's internal DNS service that receives requests from the Client and External DNS.
- PaDIS:** The ISP's internal DNS service that receives requests from the Provider DNS and returns the IP address of the nearest Host.
- Hosts:** The servers within the ISP's network that serve the content. They are categorized into Host 1, Host 2, Host 3, and Host 4.
- ISP Cloud:** The network infrastructure of the Internet Service Provider.

The numbered steps (1-7) indicate the flow of the request and response:

- Client sends a request to the Provider DNS.
- Provider DNS sends a request to the External DNS.
- External DNS returns the IP address to the Provider DNS.
- Provider DNS sends a request to the PaDIS.
- PaDIS returns the IP address of the nearest Host to the Provider DNS.
- Provider DNS returns the IP address of the nearest Host to the Client.
- Client sends a request to the nearest Host.

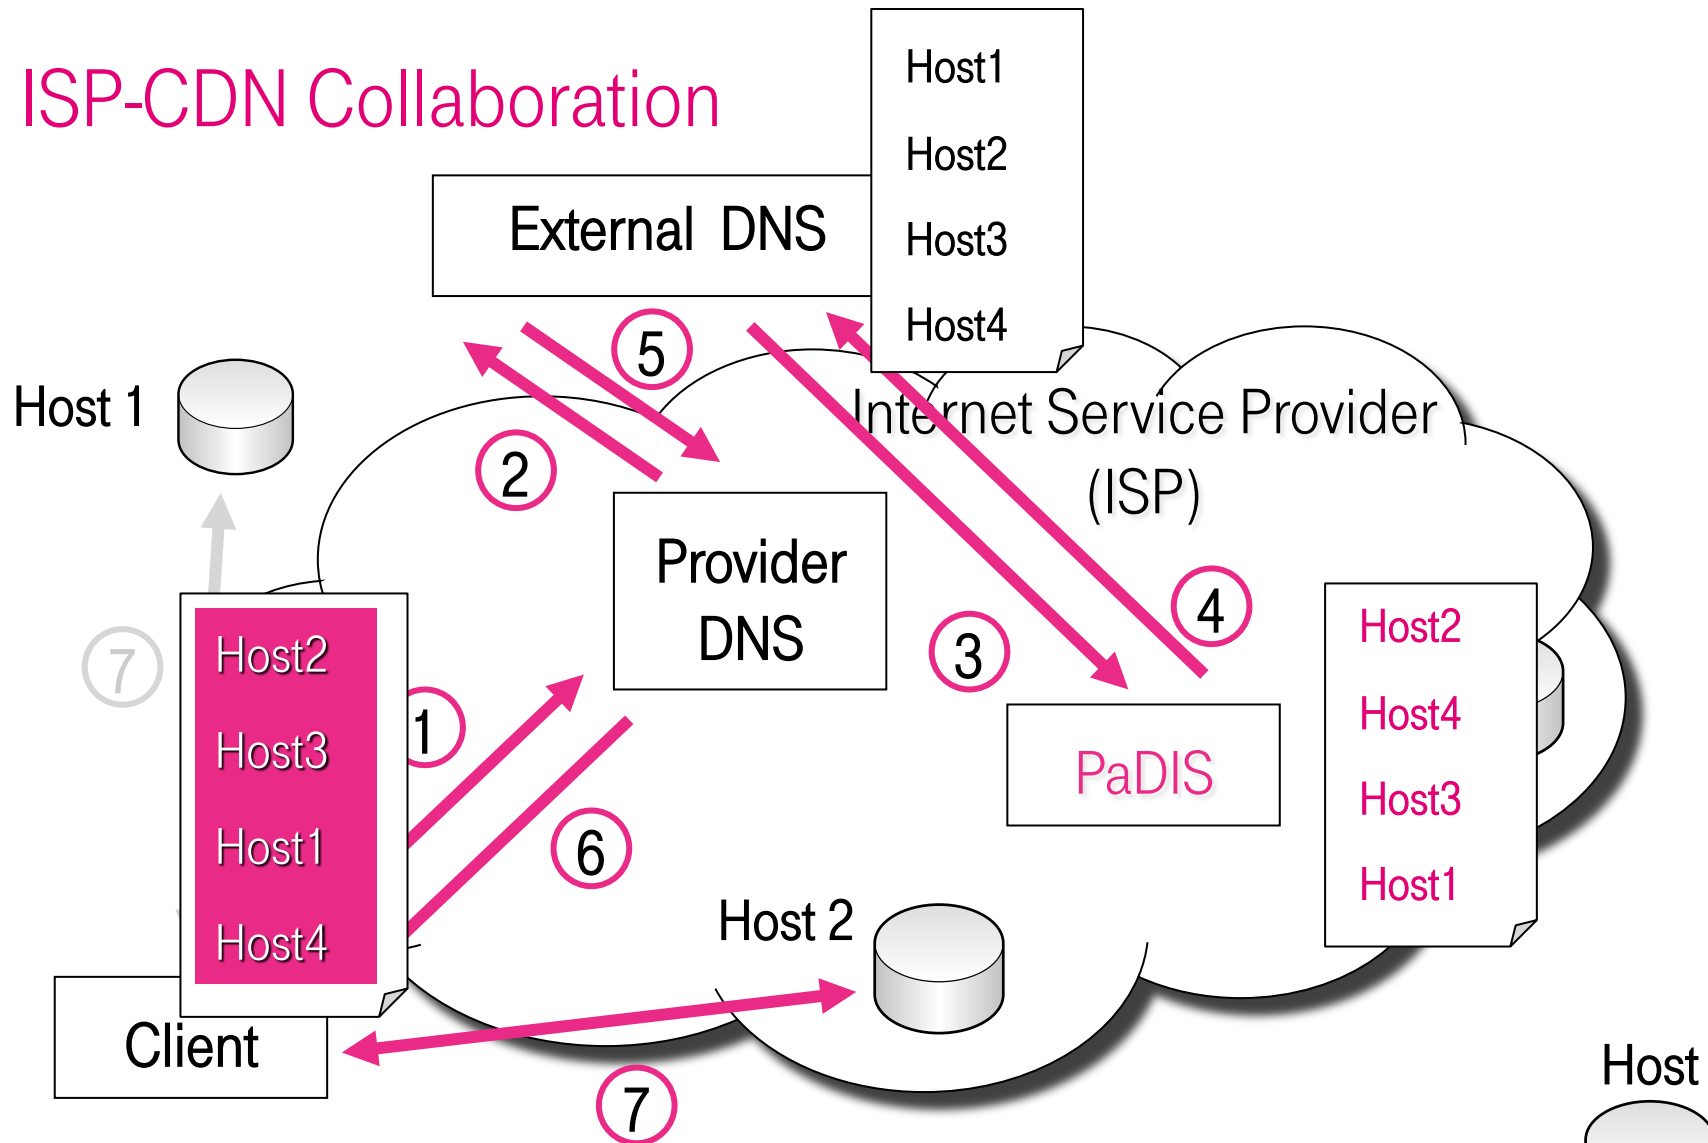
Additional details include:

- Host 1 is associated with Host 2, Host 3, and Host 4.
- Host 2 is associated with Host 1, Host 3, and Host 4.
- Host 3 is associated with Host 1, Host 2, and Host 4.
- Host 4 is associated with Host 1, Host 2, and Host 3.

• CDN and ISP agree on the stable matching algorithm

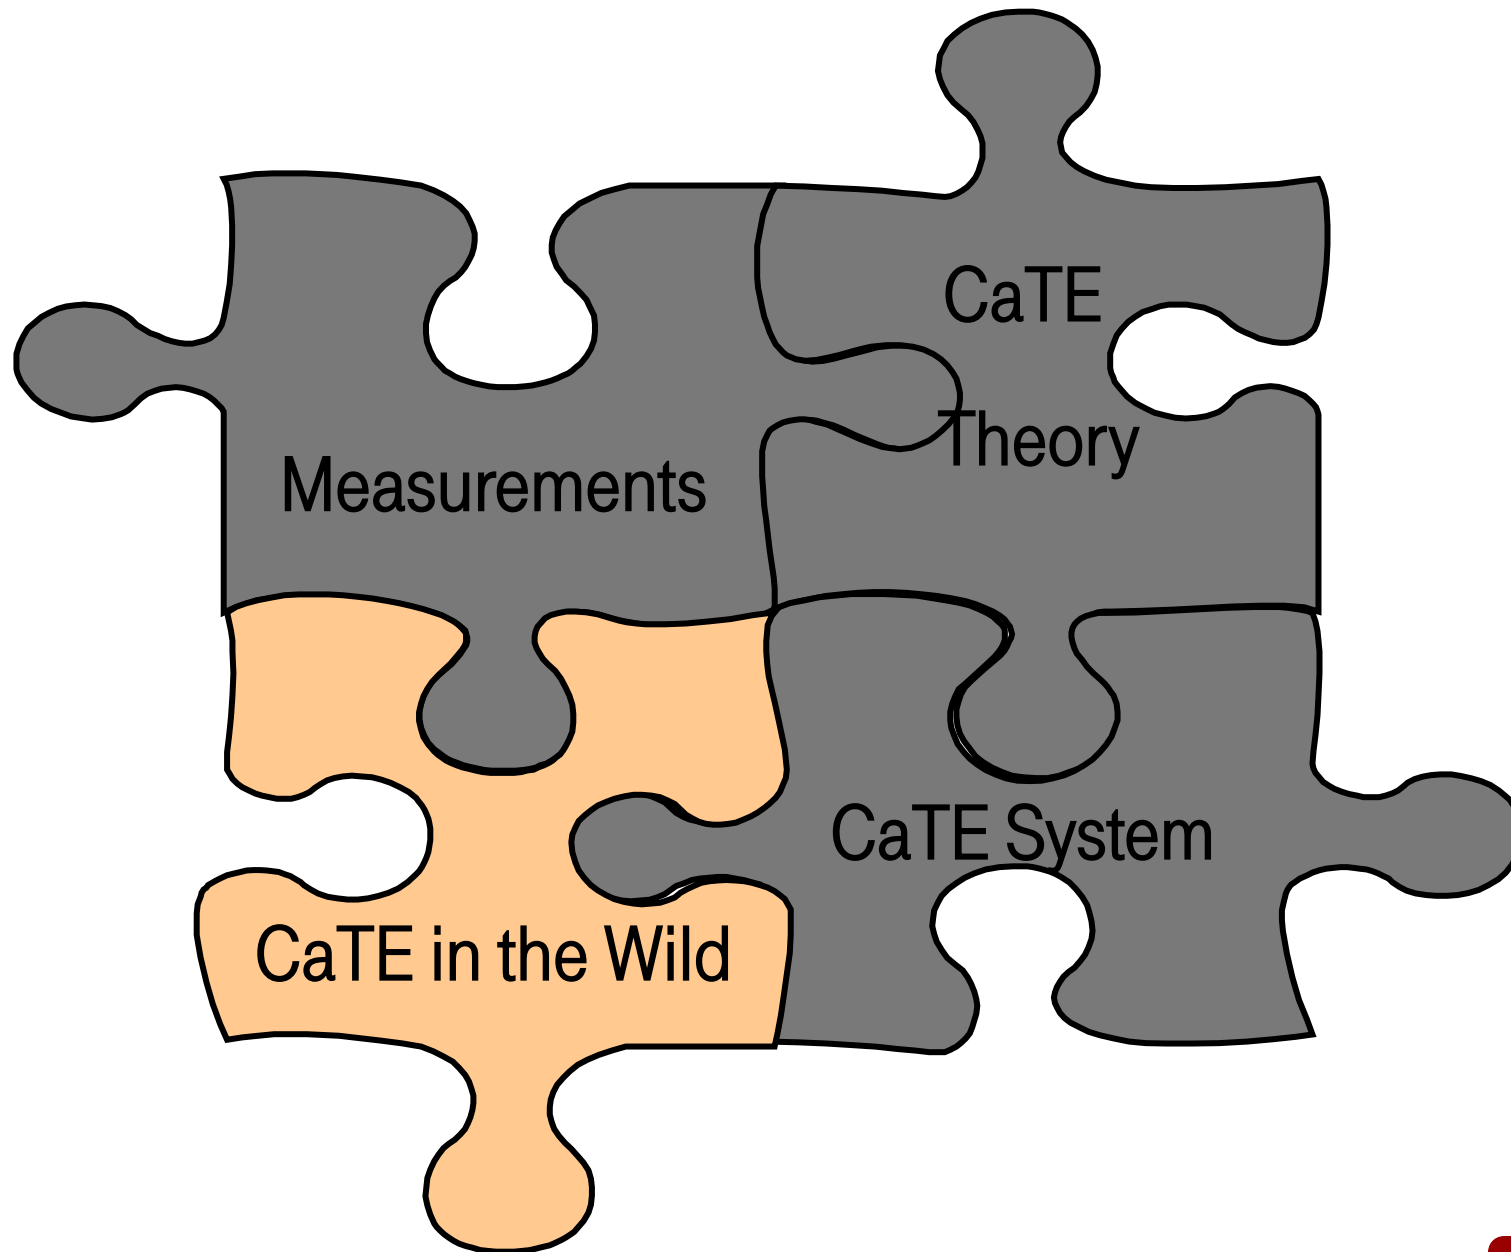
- Final server decision is made by the CDN (step 5)
- No packet priority (eg Google-Verizon) that raises net neutrality flags..

ISP-CDN Collaboration



- CDN and ISP agree on the stable matching algorithm
- Final server decision is made by the CDN (step 5)
- No packet priority (eg Google-Verizon) that raises net neutrality flags..





Optimization Functions


- CaTE supports a number of optimization functions per pair of ISP-CDN.
- The optimization function can change on-the-fly

In our evaluation we focus on:

- Top-10 CDNs and
- A number of *optimization functions*:
 1. “**utilization**”: CaTE selects the path that has the minimal maximum link utilization.
 2. “**delay**”: CaTE chooses the path that yields the minimal overall delay.
 3. “**path**”: CaTE picks the shortest path among all possibilities.




Reduction in Link Utilization



→ Reduction up to 40% in the most congested link, during the peak time!




Reduction of Network Traffic



→ Up to 15% reduction in the total traffic
(Petabytes of traffic)




Reduction in End-to-end Delay



→ Up to 60 msec delay reduction for more than 25% of the traffic



An Opportunity for Better Traffic Engineering



- Moving traffic from congested link to less congested
- Improvement of HTTP traffic locality from 25% → 50%




New Challenges

Web-based applications and services:

→ Significant part of today's Internet traffic




→ Increasing Complexity



→ Volatile demand

→ Over-provisioning comes at a high cost

→ Deployment is not flexible



source: Google

New Challenges

Web-based applications and services:

→ Significant part of Internet traffic




→ Increasing Cost

→ Volatile demand

→ Over-provisioning comes at a high cost

→ Deployment is not flexible

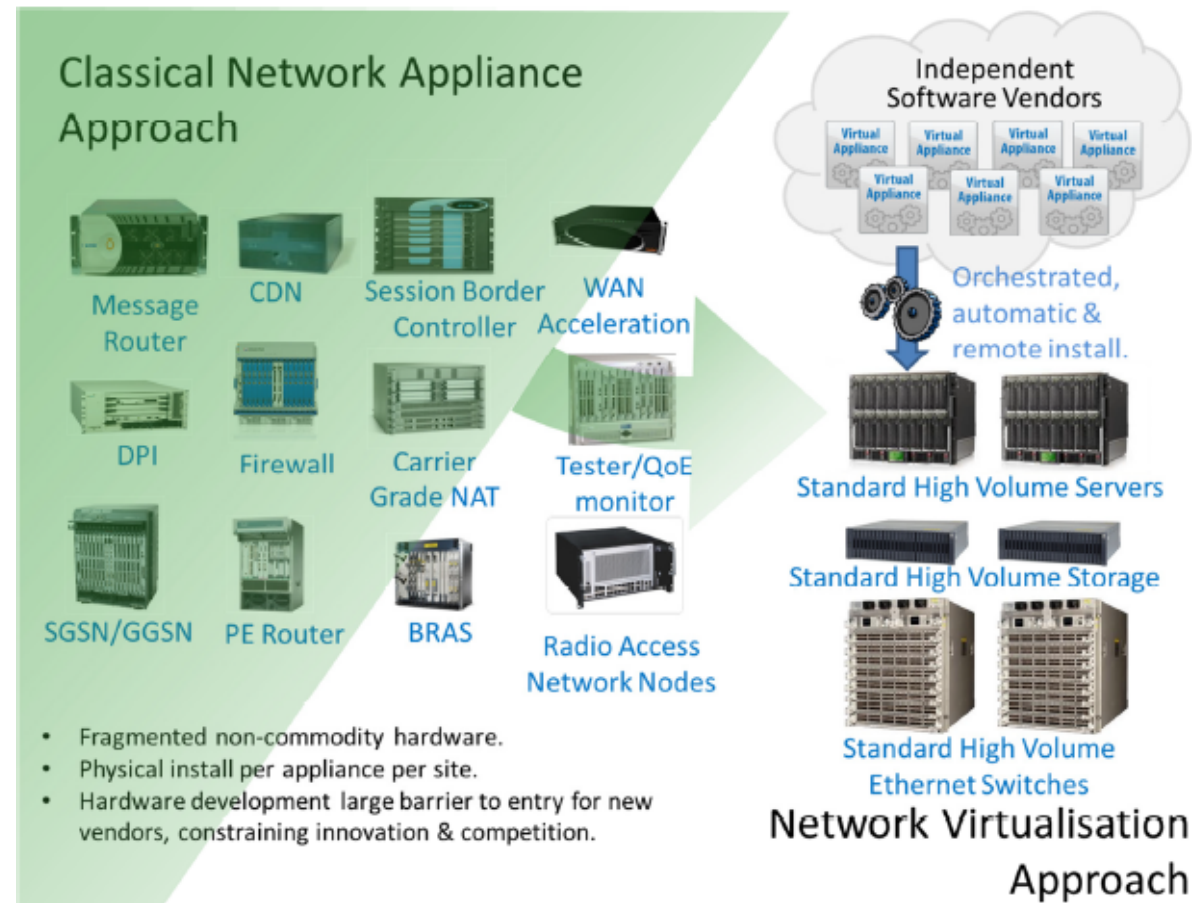
On-demand Service Deployment



source: Google



New Opportunities: Network Functions Virtualization



Supported by AT&T, BT, CentruryLink, China Mobile, Colt, Deutsche Telekom, KDDI, NTT, Orange, Telekom Italia, Telefonica, Telstra, Verizon, ..

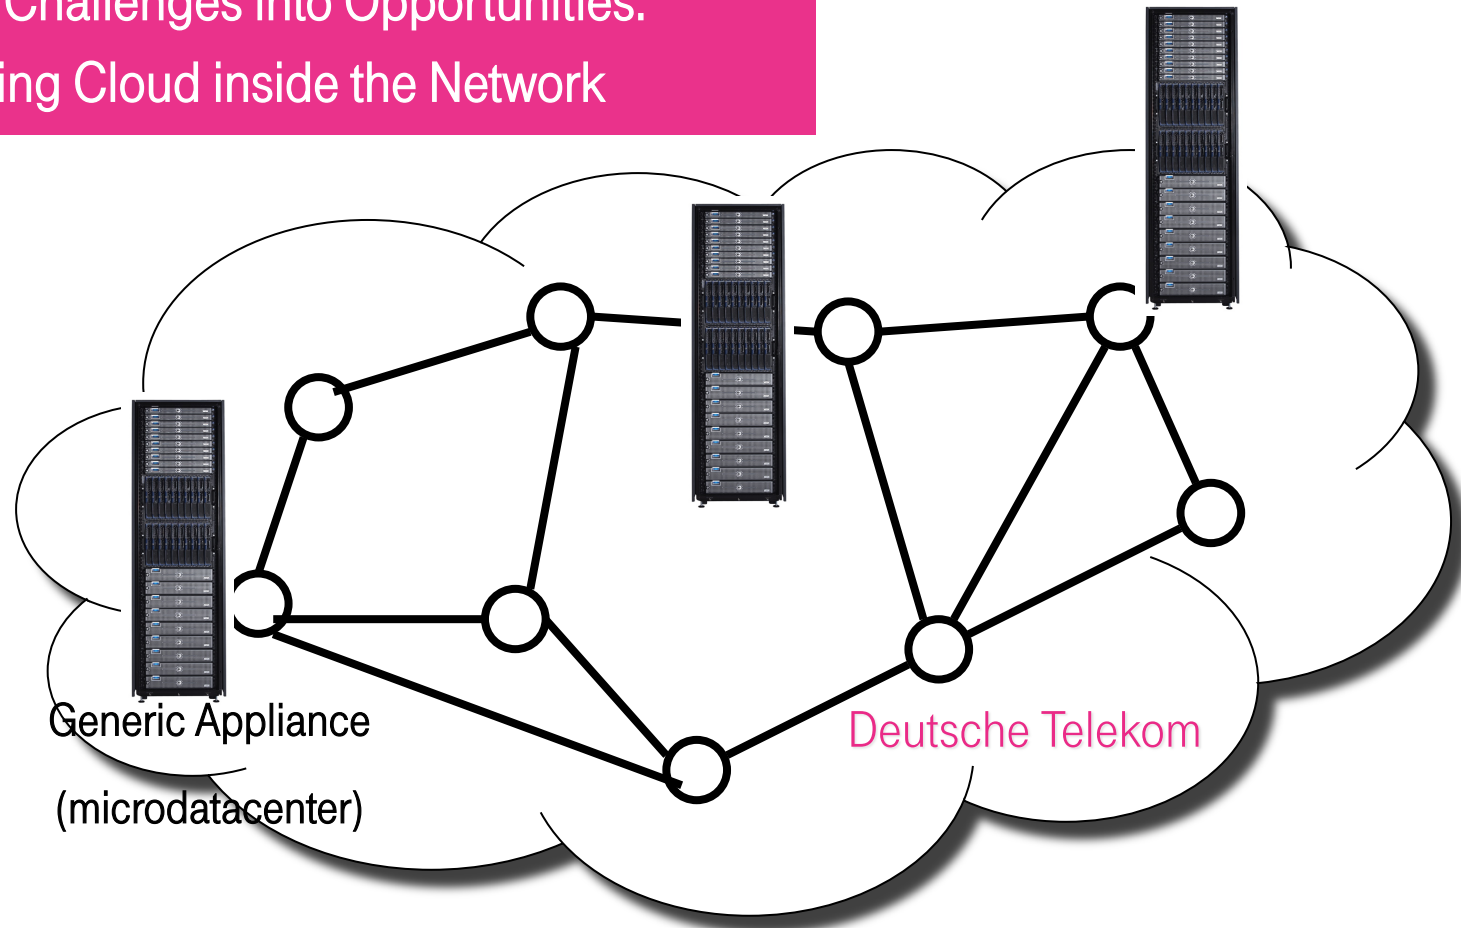


[1] "White Paper: Network Functions Virtualisation", SDN and OpenFlow World Congress, Oct 2012



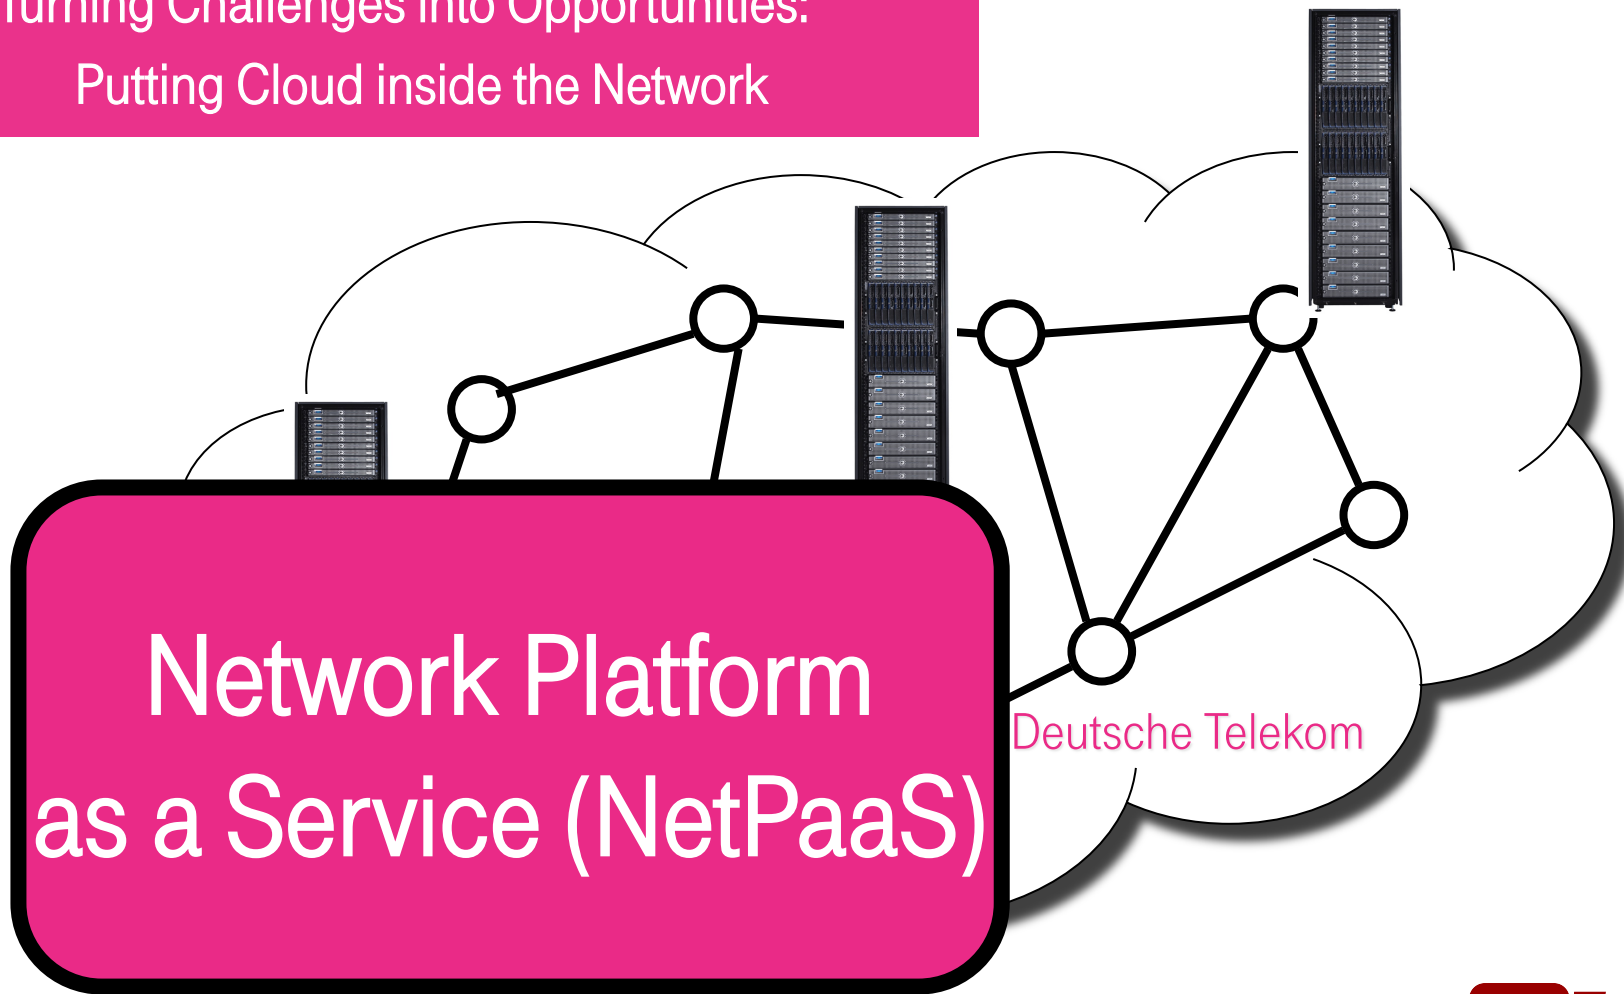
Vision: On-demand Service Deployment inside the Network

Turning Challenges into Opportunities:
Putting Cloud inside the Network

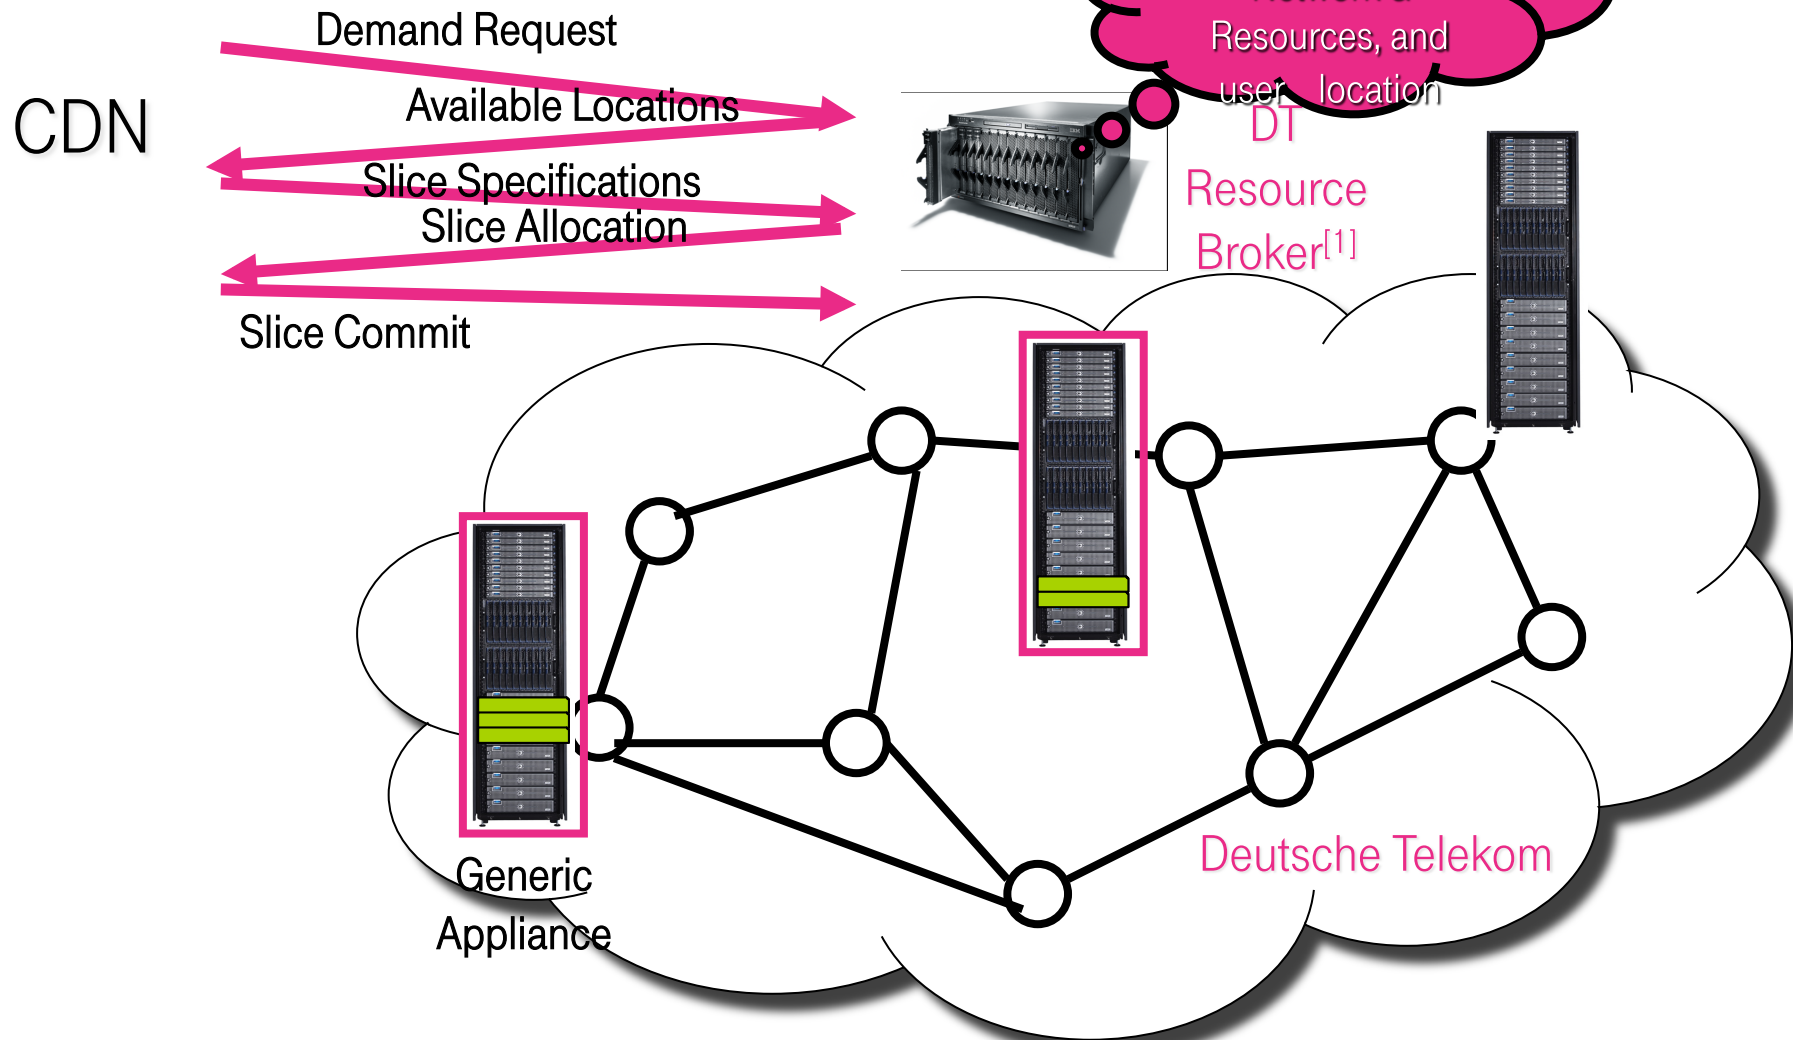


Vision: On-demand Service Deployment inside the Network

Turning Challenges into Opportunities:
Putting Cloud inside the Network



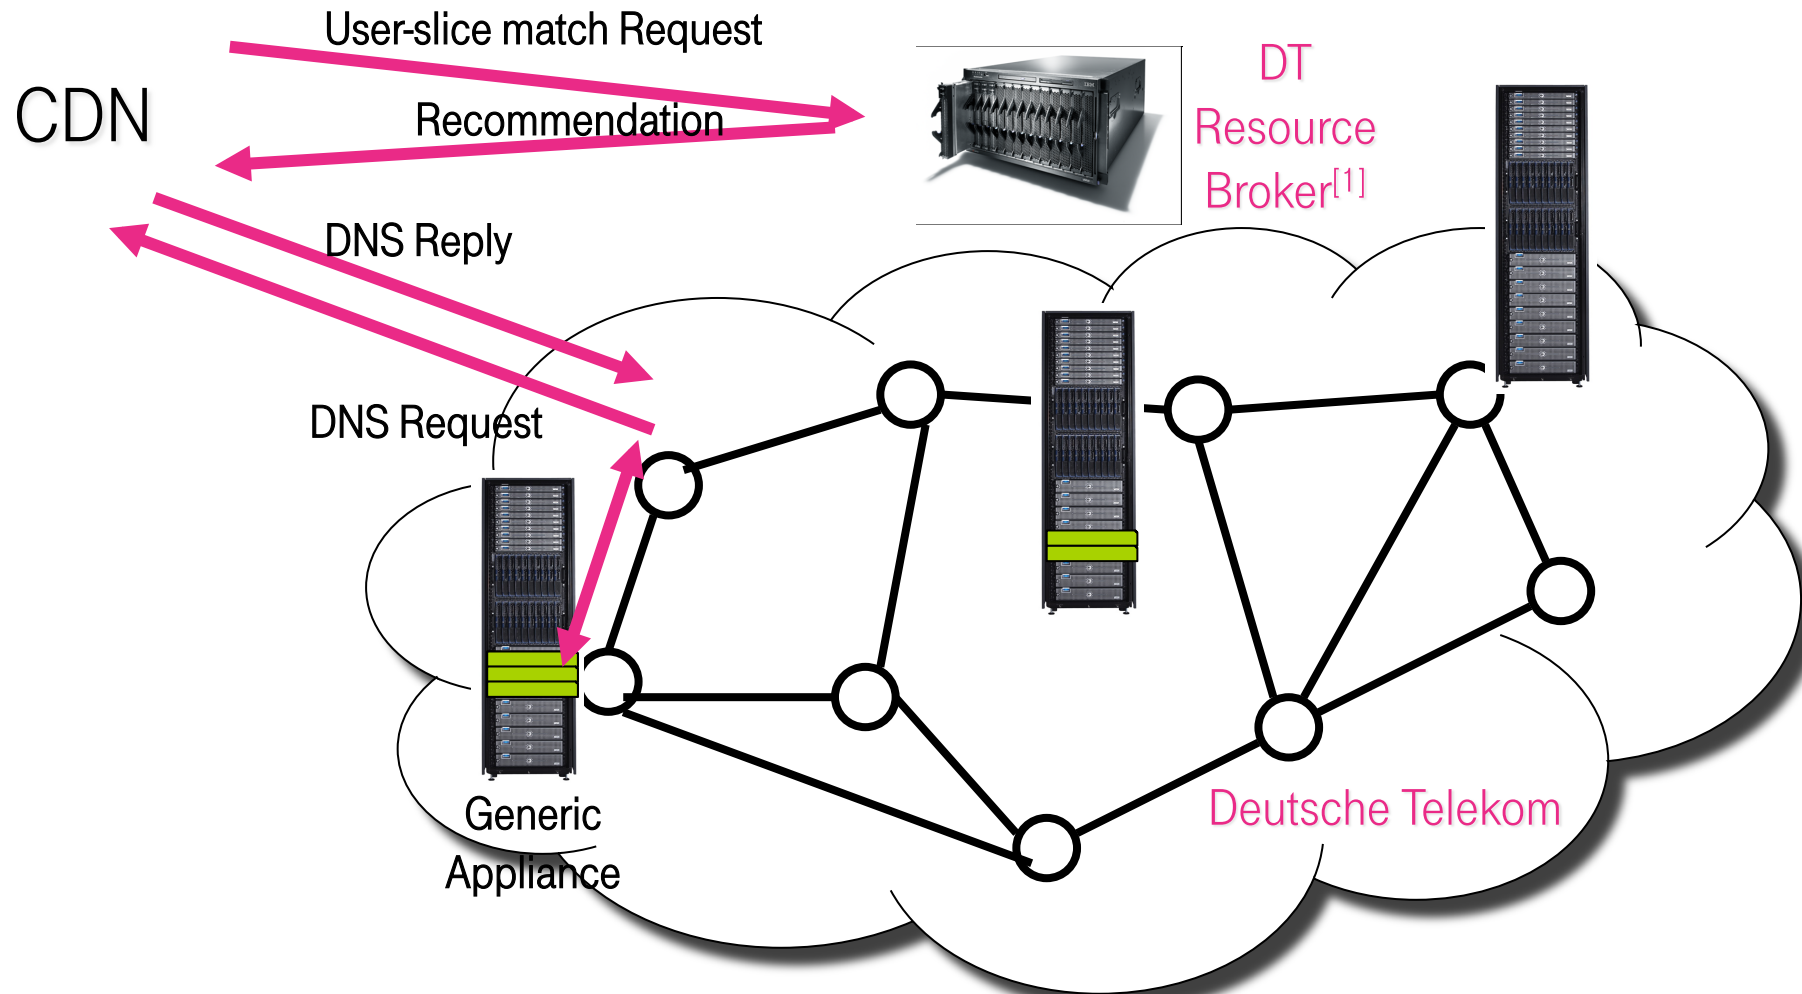
Operation: Slice Allocation



^[1] "Pushing CDN-ISP Collaboration to the Limit", Frank, Poesse, Yin, Smaragdakis, Feldmann, Maggs, Rake, Uhlig, and Weber, ACM SIGCOMM CCR, July 2013




Operation: User Assignment



[1] "Pushing CDN-ISP Collaboration to the Limit", Frank, Poesse, Yin, Smaragdakis, Feldmann, Maggs, Rake, Uhlig, and Weber, ACM SIGCOMM CCR, July 2013

Evaluation



Utilizing up to 50 servers out of more than 900
available servers in 80 PoPs



Summary

- A large fraction of Internet Traffic is due to a small number of CDNs
- Opportunity for joint CDN deployment and operation by ISP and CDN by utilizing:
 - (1) server and path diversity
 - (2) knowledge about the network and user location
- An Opportunity for better Traffic Engineering
- Benefits for all involved parties: ISPs, CDNs, Content Providers and end-users.



Thank you!

Learn more about our research:

<http://www.smaragdakis.net/research/Collaboration>



Acknowledgments

This work would not have been possible without the help and contributions of:

Michael Duesser, Jannis Rake-Revelant,
Andreas Gladisch, Marten Schoenherr (T-Labs),
Steve Uhlig (Queen Mary, U. London),
Yin Lin (Duke University),
Richard Weber (Akamai Technologies).

This work was partially funded by EU Projects:

- CHANGE (<http://www.change-project.eu>)
- BigFoot (<http://bigfootproject.eu>)

