

Infrastructure to Application Exposure

- USE CASE: CDN -

Jan Seedorf

IETF 83, Paris

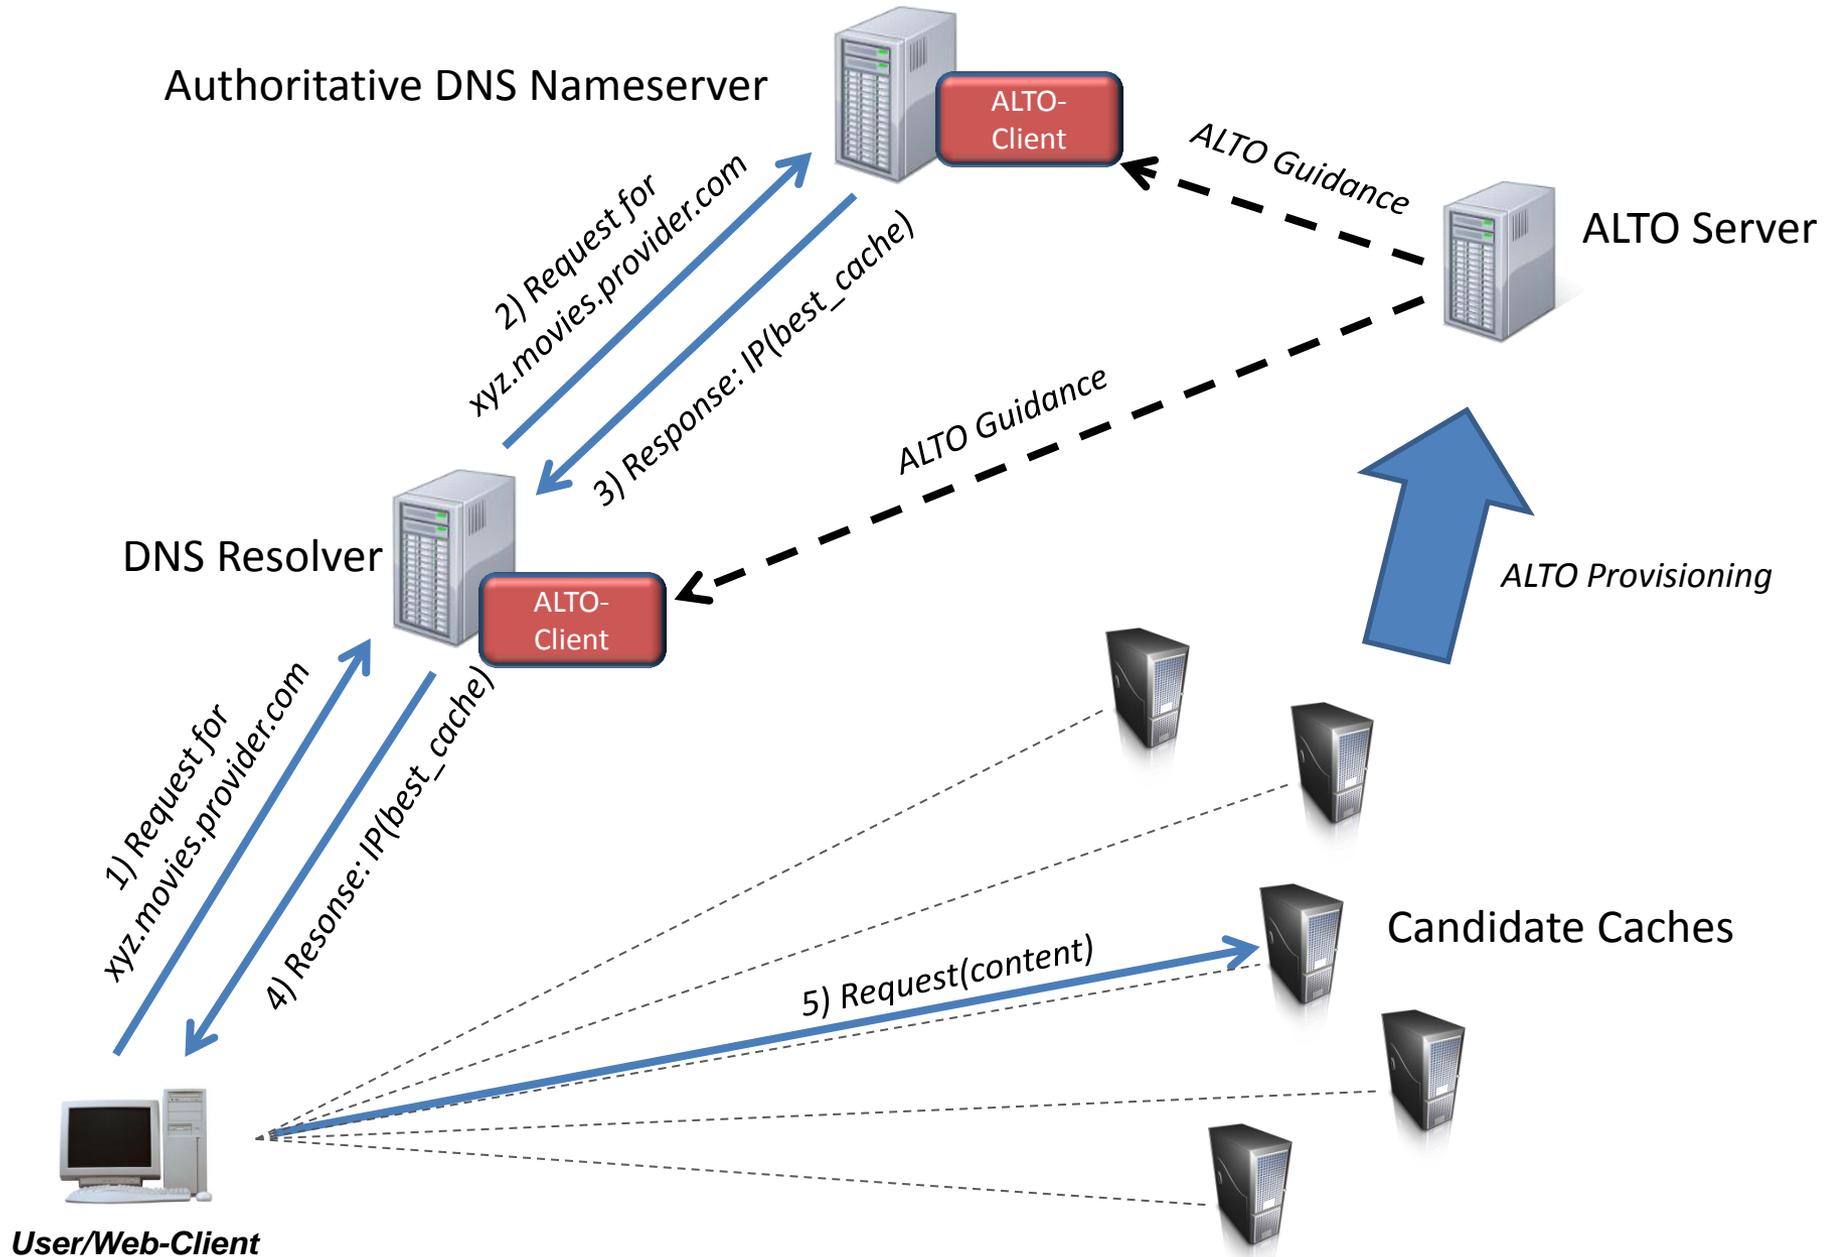
i2aex BoF

Monday, March 26 2012 1300-1500

Introduction

- ALTO (Application Layer Traffic Optimization)
 - Goal: “Better-than-random resource provider selection”
 - The same resource is available at multiple locations in the network
 - ALTO service provides network layer topology information to distributed applications, so that these applications can improve their resource provider selection
 - Use Cases
 - Initial focus on P2P applications
 - Recently, CDN investigated as key use case for ALTO
 - ALTO for CDNs
 - ALTO Client in CDN Request Router
 - Redirecting request to “best” CDN service node (surrogate)
 - E.g using http-redirect or DNS-based
 - ALTO can provide network layer topology information to the cache selection process in the request router
- (draft-penno-alto-cdn, draft-jenkins-alto-cdn-use-cases, draft-ietf-alto-deployments)

ALTO-Guidance within CDN Request Routing (DNS Example)



Current ALTO

Client can request only complete ALTO maps (or use ECS)

- May be ok if these maps are not very large (i.e. suitable for P2P apps)
- Alternatively, ALTO Endpoint Cost Service (ECS) enables to query costs between individual endpoints

Client queries Server

- Server has no way to inform a client about an important change
- Ok for P2P if information is slightly outdated, as ALTO provides only “better than random initial peer selection”

Useful Enhancements from CDN Use Case Perspective

Incremental updates for ALTO network and cost maps

- ALTO-maps for CDNs are likely to be more fine-granular (=larger) than for P2P apps
- Incremental updates are very useful to avoid huge amount of traffic when ALTO maps grow large

Server can notify client about changes

- i.e. publish/ subscribe model
- More useful in controlled environments like CDNs: CDN request router can be sure that it always has the most up-to-date information

Current ALTO

ALTO delivers network topology information

- Can deliver abstract “cost” between source-location and each cache according to a certain “cost type”

Useful Enhancements from CDN Use Case Perspective

Delivering new types of CDN-relevant information

- load on a caching server
 - Enables to take this information into account in CDN request routing, e.g. for load balancing or guaranteeing QoS
- content availability
 - i.e. what content is stored on what cache
- storage capacity
 - Enables advanced content placement strategies
- Not only “network topology” information, but also information about caches

Potential new Use Case: ALTO for CDN Interconnection

- Recently, ALTO has been suggested within the CDNI WG
 - to facilitate the selection of a downstream CDN (dCDN)
 - Advertisement of “footprint” and “capabilities” towards an upstream CDN
- Use of ALTO not agreed within CDNI WG
 - but ALTO mentioned in charter and currently discussed within WG
- If ALTO were used for downstream CDN Selection within CDNI, similar enhancements (as for CDN use case) would be very useful / needed
 - Incremental Updates / Server-Initiated Updates, to convey when footprint or the capabilities of a dCDN change abruptly / partially
 - More information than just network topology, i.e. other capabilities of a dCDN such as content availability or capabilities on caches / links
- Drafts:
 - **draft-stephan-cdni-alto-session-ext**
 - **draft-seedorf-i2aex-alto-cdni-perspective**
 - **draft-seedorf-cdni-request-routing-alto**

Summary

Implications of ALTO-CDN Use Case and corresponding Requirements for new Protocol Solutions

- ❑ ALTO maps can grow large
 - Necessary to have incremental updates if ALTO maps (and not ECS) are used
- ❑ Minor changes in ALTO maps may be important
 - Necessary to have a mechanism for the server to inform the client when new information is available
- ❑ Not only network layer topology, but also information about caches/surrogates is useful
 - Necessary to have a way to convey information about caches (status, load, capacity, ...)

Acknowledgements

- Many thanks for valuable input to this presentation goes to
 - Ben Niven-Jenkins
 - Stefano Previdi
 - Tina Tsou
 - Emile Stephan

- Acknowledgement: Jan Seedorf is partially supported by the COAST project (COntent Aware Searching, retrieval and sTreaming, <http://www.coast-fp7.eu>), a research project supported by the European Commission under its 7th Framework Program (contract no. 248036). The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of the COAST project or the European Commission.