ALTO and Content Delivery Networks

draft-penno-alto-cdn

Stefano Previdi, sprevidi@cisco.com
Richard Alimi, ralimi@google.com
Jan Medved, jmedved@juniper.net
Reinaldo Penno, rpenno@juniper.net
Richard Yang, yry@cs.yale.edu

Alto Server and the CDN

- Goals
  - Optimize traffic delivered from CDN
  - Improve global and local server load balancing
  - Facilitate topology & status information sharing between domains (e.g. CDN – ISP)
- Non Goals
  - Standardization of particular server selection criteria

Alto Server and the CDN

- Scope:
  - Integration of the ALTO Service with the main CDN request routing / content location selection techniques
    - HTTP Redirect
    - DNS Integration
  - Multi-Domain Use Cases

HTTP Request Routing

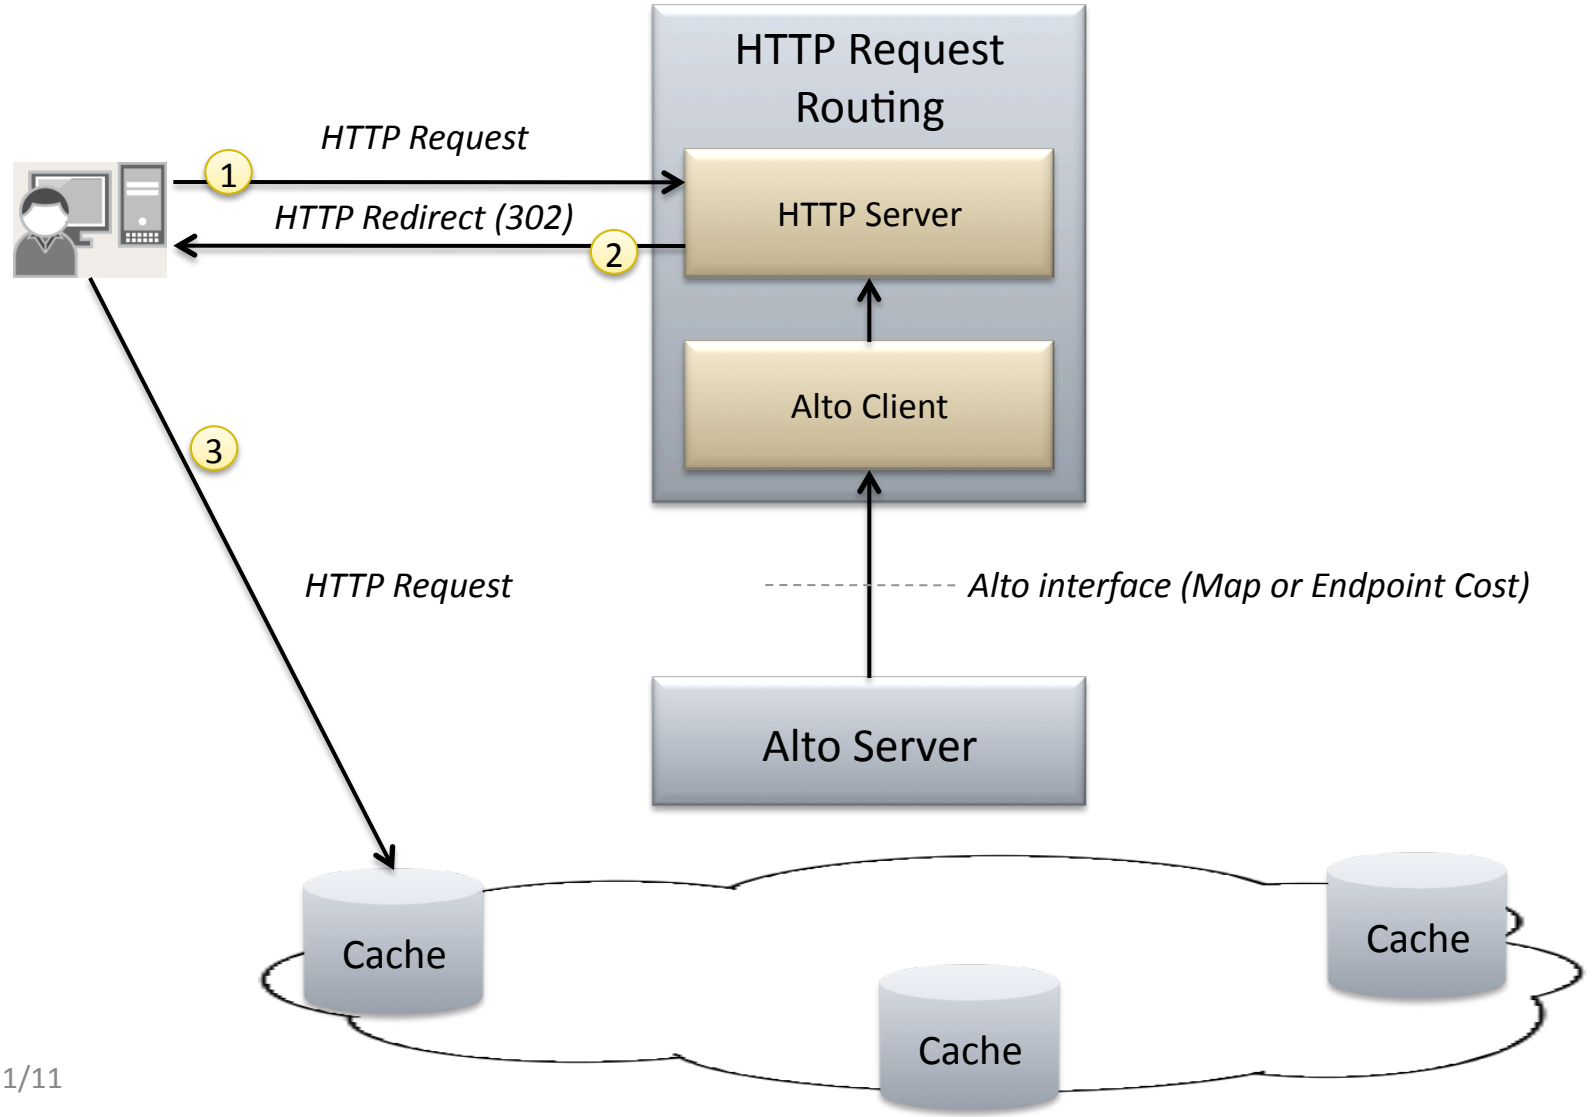

HTTP Request Routing
(Status Updates into Request Routing)

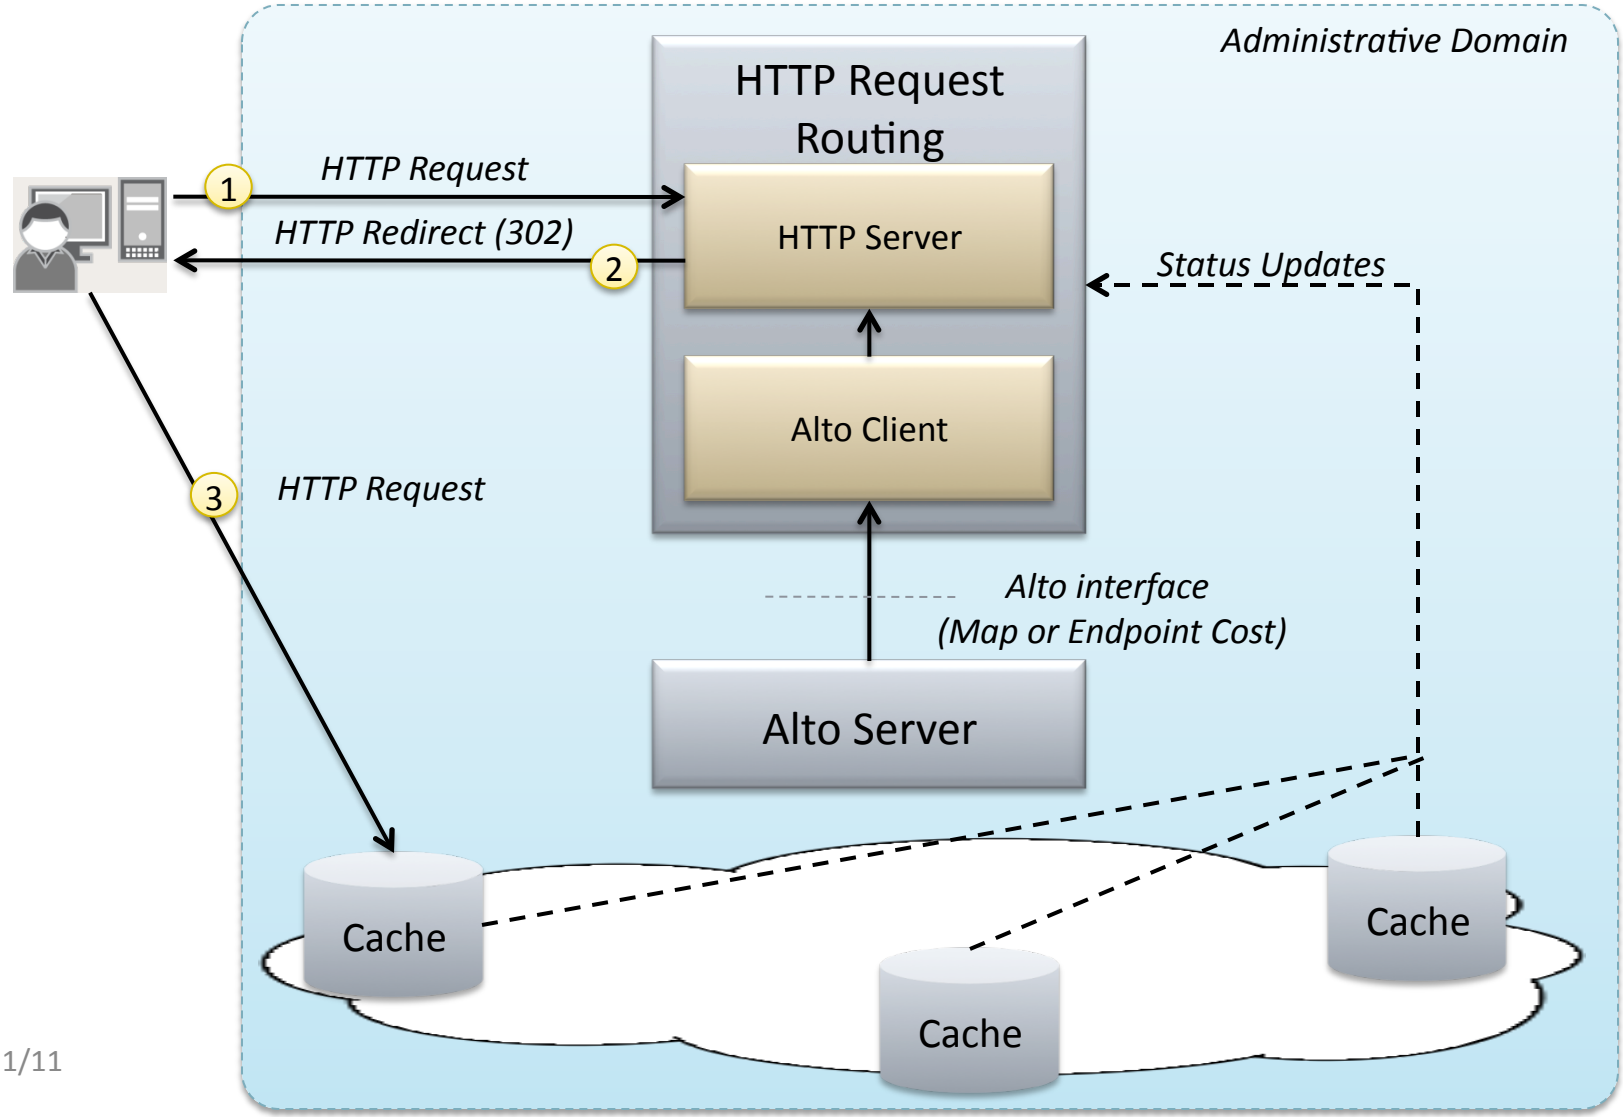

HTTP Request Routing
(Status Updates into ALTO Server)

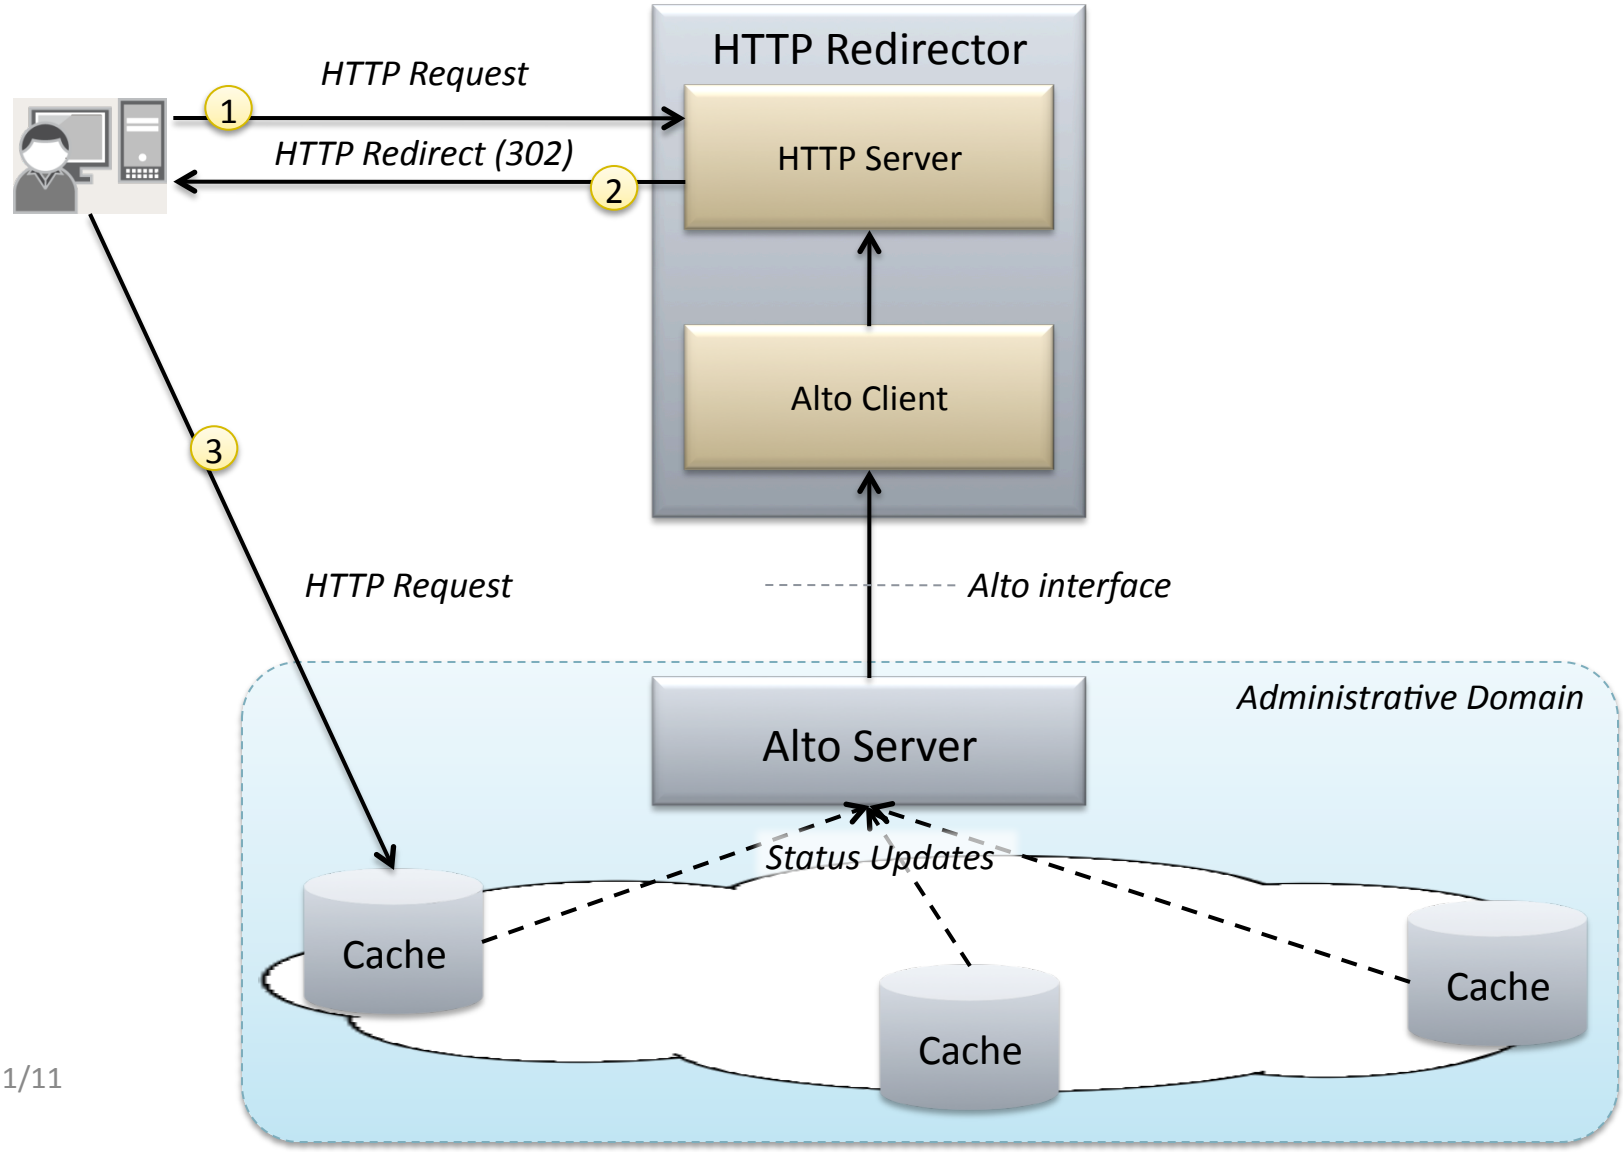

DNS Resolution

Single Administrative Domain

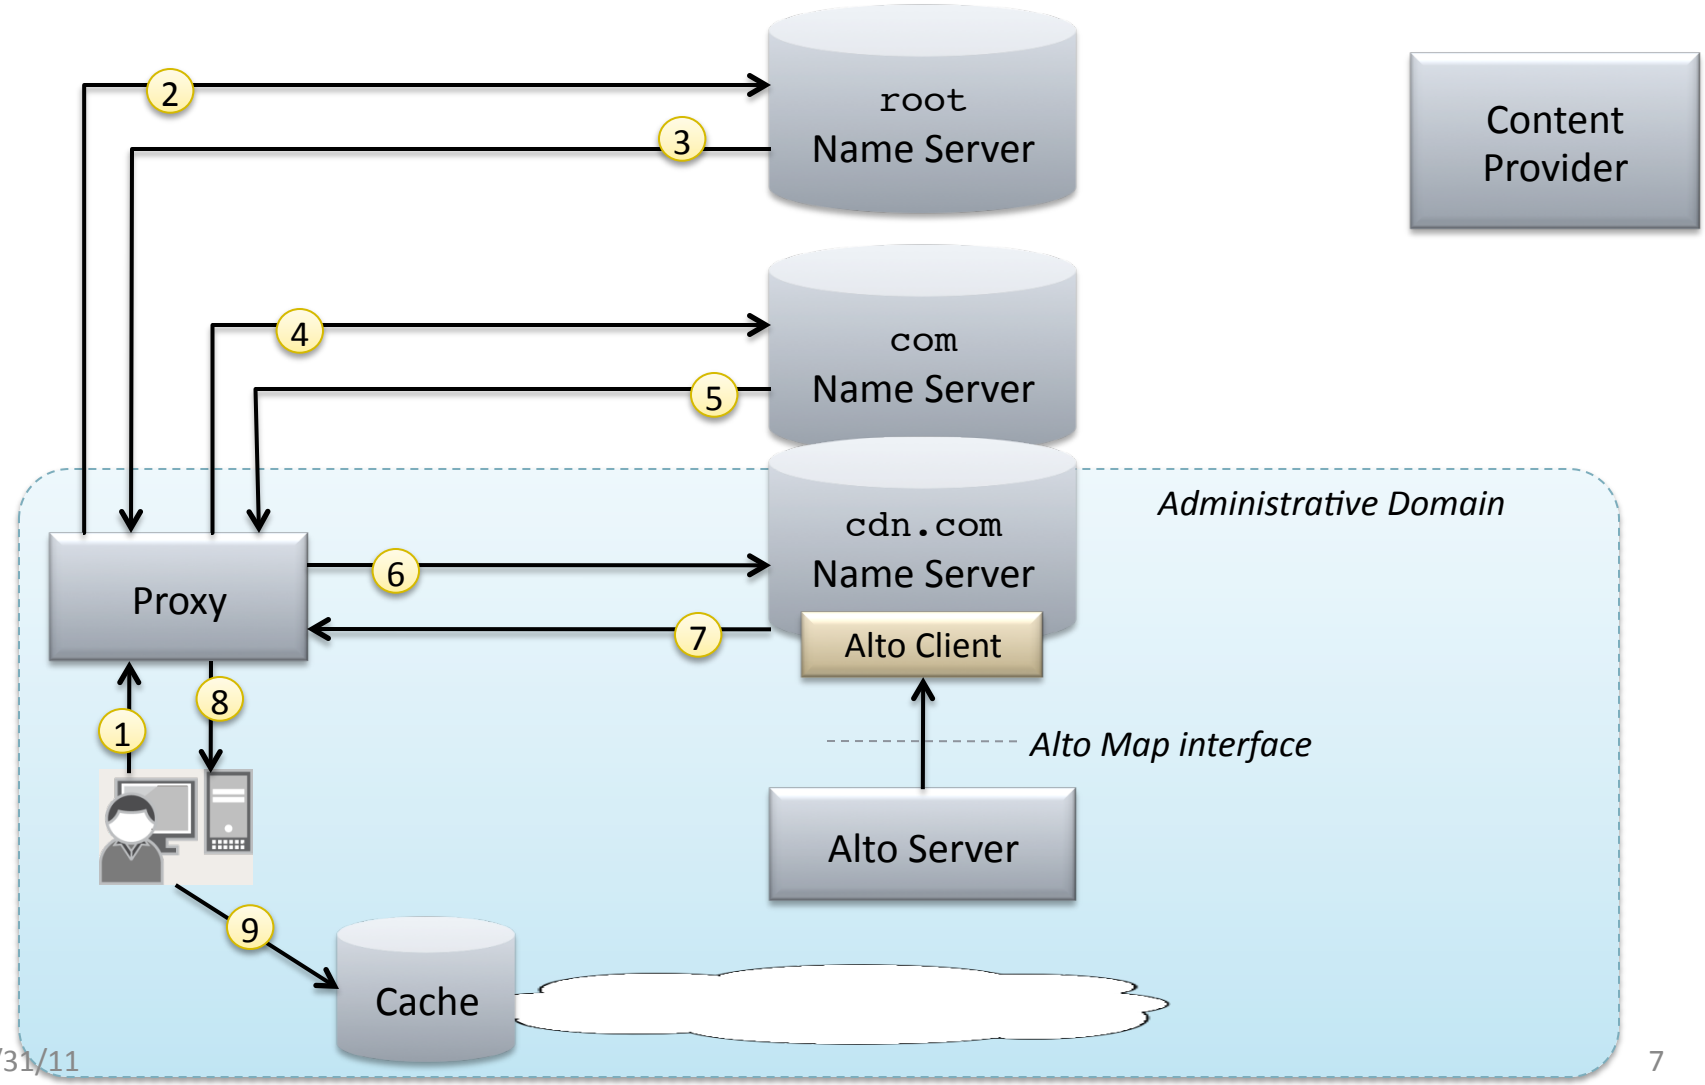

Request Routing: Map Service

- Host PIDs: host subnets
- CDN Node PIDs: IP addresses of available CDN Nodes
- If Request Router knows the subset of host and CDN PIDs that it will serve, it can download relevant Network and Cost Maps in advance:
  - Use the Filtered Map Service to only get data about relevant PIDs
  - Or ALTO Server should support multiple views
- HTTP Cache Control headers to determine how long the client keeps the Network and Cost Maps

Suggested Protocol Additions/Extensions

- PID Attributes (and Query):
  - Needed to disambiguate between PIDs containing endpoints of a specific class
  - PIDs classified as containing "CDN nodes", "Mobile Hosts", "Wireline Hosts", etc.
  - Allows ALTO Network Map to provide simple resource discovery/ location.

- Incremental Map Updates:
  - Beneficial is maps are large and/or change often
  - The map contained in the reply is the delta from the previous version
  - Note: pertinent for ongoing discussions regarding mobility.

Request Routing: Endpoint Cost Service

- ECS is a ranking service delivered on demand by ALTO server
  - ALTO client resides in CDN and request ECS based on endpoint addresses
  - ALTO server collects topology information, computes rankings
- Network Information sources
  - By default, ranking is based on routing distance
  - ALTO server may integrate different information sources: geo-location, network performance, resources utilization, policies, …
- Caching
  - ALTO server may pre-compute topology information and store different abstractions of network infrastructure/topology
  - ALTO client may cache ECS results for further use

Request Routing: Endpoint Cost Service

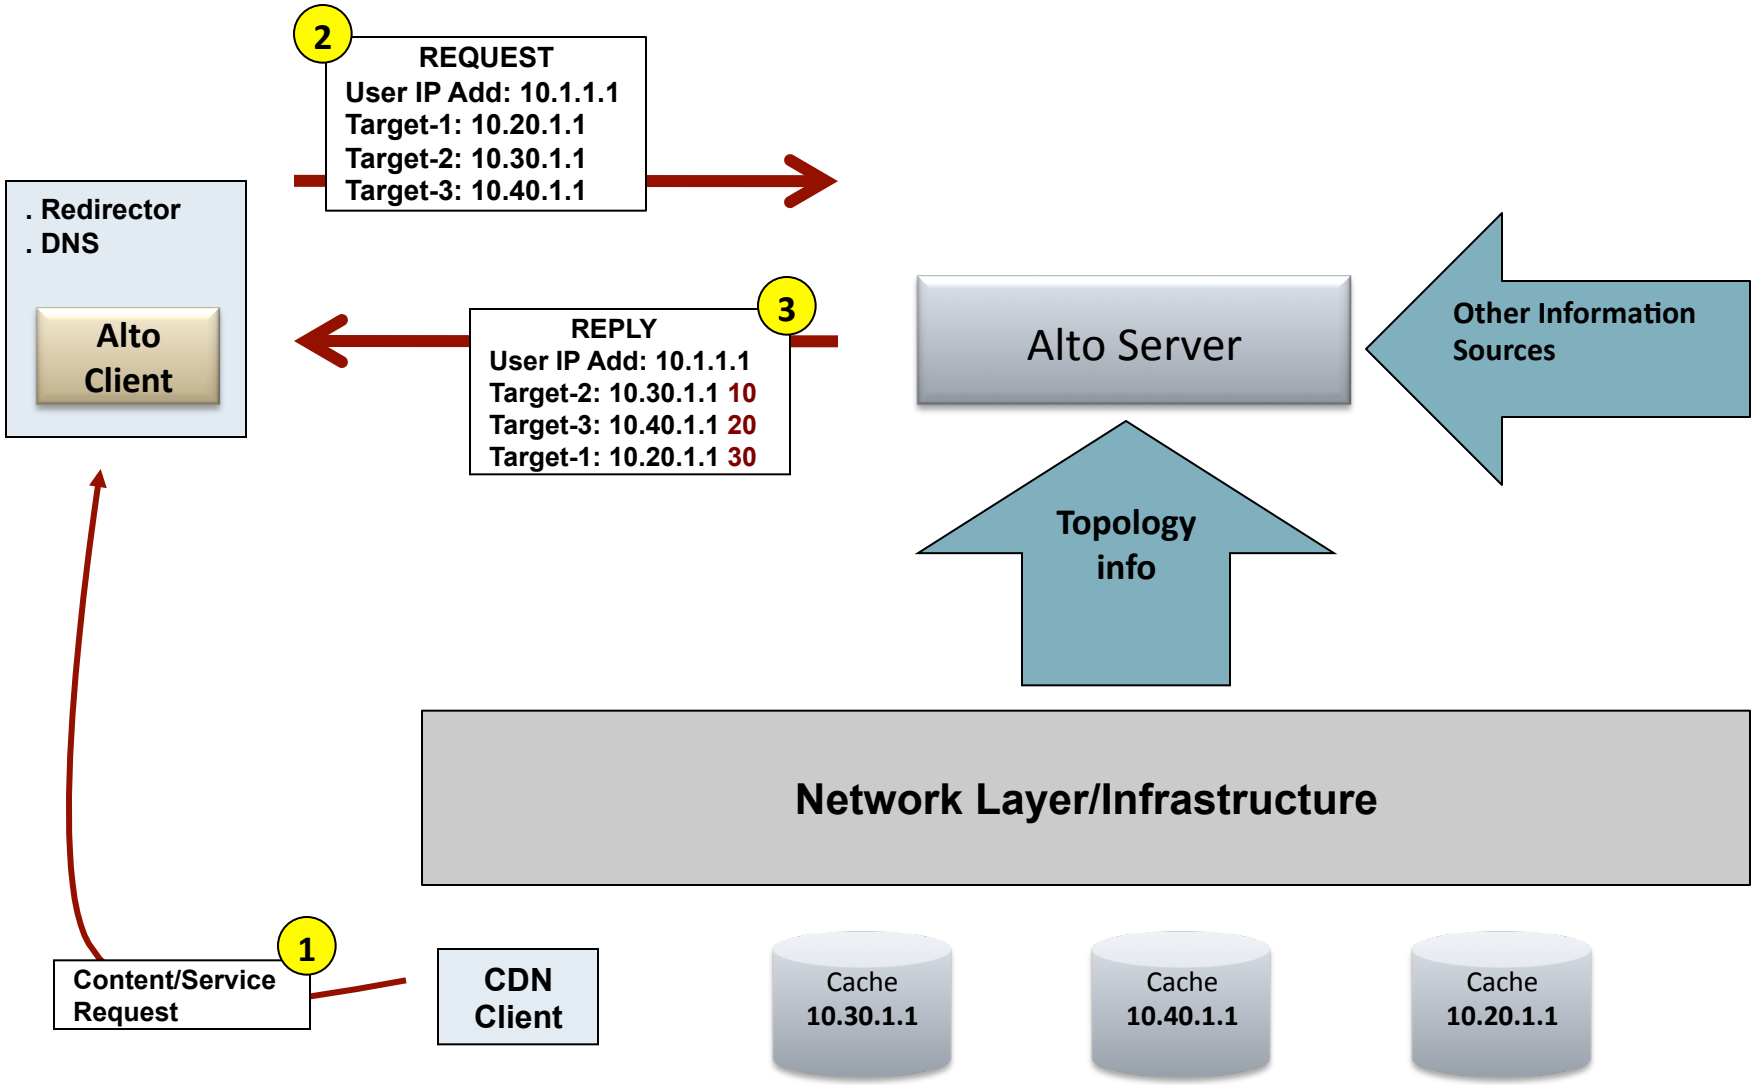

Request Routing: ECS and Maps

- Security & confidentiality ALTO requirements are addressed by both Maps and ECS
- Mechanisms through which an ALTO server collects network infrastructure topology information are common to Maps and ECS
- ECS and Maps only differs in the way service is delivered
- When ECS is used, the ALTO client (CDN) doesn‘t hold any topology information
- Intelligence of topology computation, abstraction, ranking is placed inside the ALTO server

Multiple Admin Domains

Cost Advertisement

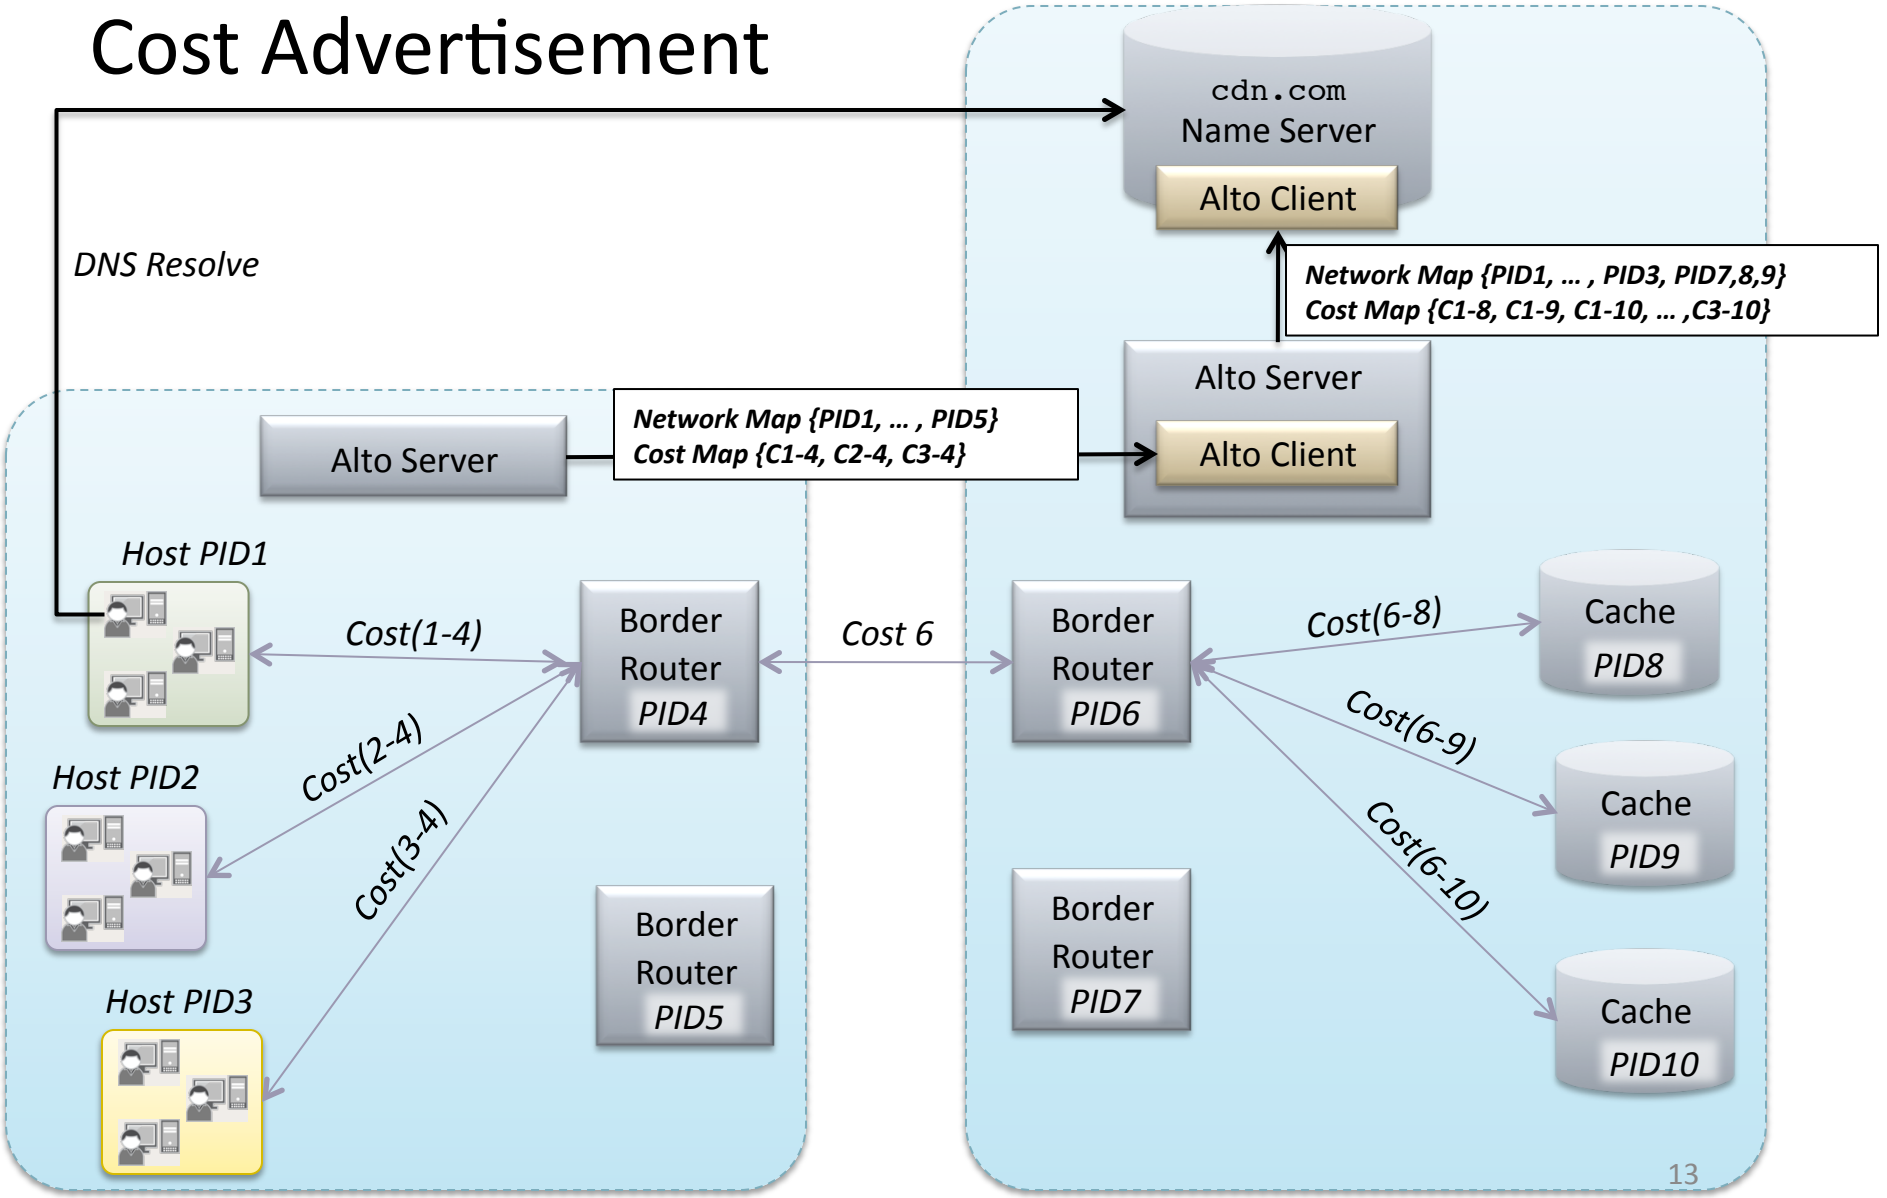

Additional Suggestions

- Federated deployment of ALTO Servers
  - Is anything beyond standard ALTO Protocol needed for distributing information amongst servers?
- Extensible Cost Maps
  - Does it make sense to store “opaque” information for paths within ALTO Maps, instead of just numeric/ordinal costs?
- Discovery
  - Discover CDN's ALTO Server instead of ISP's ALTO Server. Input to discovery discussion

Open Questions

- Maps with incremental updates
- Conceptually, maps are no different from topology DBs
  - Analogy between an OSPF/ISIS LSDB and Maps
- Does it make sense to apply same mechanisms ?
  - Incremental updates based on Maps objects
  - Routing protocols: node-ID, adjacency, cost
  - Maps: PID, neighboring PID, cost
- ALTO server may establish sessions/adjacencies between them and exchange information in a reliable incremental fashion
  - Minimal state is required
  - Scalable and efficient