

An Architecture for Content Routing Support in the Internet

Mark Gritter, David R. Cheriton
Stanford University

Relevance to mobile networking?

- What aspects of mobility are addressed by this architecture?
- Locating mobile host?
- Session continuity?
- Disconnection?

Content Routing

- Goal: Reduce time to access content
- Basic premise: replicated servers offer alternate routes to a piece of content




Figure 2: Content-Layer Routing

Existing approaches

- Proprietary schemes (Akamai, DD) use DNS
- Rely on centralized server for “cold” request
- Overall problem is difference in topology of name system vs. content storage system
- Result: non-uniform latency and availability

Content Routers

- Perform IP routing and name-based routing
- Clients: use INRP for name lookup
- Peer CRs: use NBRP for routing updates
- Firewall, gateway, and BGP-level routers are the targets for deployment

Internet Name Resolution Protocol

- Backwards compatible with DNS
- CR maps name to next-hop CR
- INRP requests burrow through network along the “best” path to the content
- Topology of name servers matches the content servers

Benefits of INRP

- Name requests gravitate to “nearest” data
- Tight coupling between name resolution and server availability
- Avoids false mappings (e.g. server failure)
- Trust in name lookups mirrors trust in packet delivery by network

Name-Based Routing Protocol

- Protocol for CRs to exchange routing data
- Route updates can be authenticated similar to Secure BGP
- IP packet routing policies must coincide with content routing policies (or WRAP)

Will it work?

- NBRP must be scalable (update overhead)
- CRs must be able to process requests with acceptable throughput
- Name lookups must have low latency

Aggregation

- To reduce traffic in communicating updates between CRs
- CR still needs one entry for each second-level domain in aggregates
- Only routing updates for the aggregate need to be sent between CRs

<i>site_threshold</i>	Affixes (1000s)	<i>aggregate_threshold</i>			
		3	5	10	20
2	1727	19.5 (6.7)	20.1 (5.6)	25.7 (4.4)	37.0 (3.4)
3	1692	14.9 (5.9)	16.1 (5.0)	20.6 (4.0)	30.1 (3.2)
10	1679	14.8 (5.9)	16.0 (5.0)	20.6 (4.0)	30.3 (3.2)

Table 1: Number of routes (and aggregates) in thousands for different site and aggregate threshold values.

Redirection

- A host outside of its zone tells the zone content router of its current location.
- Lookups on the hostname will receive a redirection response to the “care-of name”
- Re-issue lookup on “care-of name”
- Does this result in the same latency and stale data problems as with DNS?