

Edge Caches and Localization

Nicholas Weaver
International Computer Science
Institute
(Apologies for my absence)

Bulk data P2P *shifts* costs

- By default, bulk-data P2P *shifts* the deliver costs from the content provider to the retail ISPs
 - All bits *must* come from someplace: in bulk-data P2P, it comes from other customers rather than the content provider
- Creates a significant incentive for commercial bulk-data P2P
 - E.G. CNN now uses P2P extensions to flash for live streaming videos: Reportedly saves 30% in bandwidth costs for the content provider.

Localization alone may be insufficient

- Localization removes **transit** costs for bulk-data P2P, but **does not affect** last-mile uplink costs
 - For many networks with shared uplinks (e.g. cable modems, wide-area wireless), last-mile uplink may be substantially more expensive than transit when under congestion
 - Even with localization, bulk-data P2P may substantially increase **total** costs for content delivery
- Edge caches reduce last-mile uplink costs as well as transit costs
 - An edge-cache is simply a P2P node located in the ISPs network in the **least-cost** position for the ISP to deliver content
 - Edge-caches with P2P minimizes **total** network costs as well as content-provider costs
- Unlike web caches, edge-based CDNs, or proxies, P2P edge-caches are fail safe, transparent to the user, and can be incrementally deployed

Edge Caches need Localization

- Localization services should be aware of any edge caches
 - Drive user traffic to the edge-caches
- Localization services should provide content discovery for edge caches
 - Localization should know what users are requesting, and edge-caches need to know this information to determine what to cache
- Edge caches need localization services
 - An edge-cache should only allow free-riding peers which are in the ISP's own network: localization provides this mechanism

And Don't Worry (much) About Privacy in Localization Services...

- Edge caches need to know and will discover who is requesting what files
 - Otherwise, this can't work
- **Direct** bulk-data P2P can never be privacy preserving!
 - Peers can always know a subset of other peers sharing the same file
 - Thus anyone who **really** wants to know can create a bunch of sibyls
- As a result, privacy preserving for localization is probably overrated
 - If there is a tradeoff between privacy and utility, favor utility

Interaction with Local Trackers

- Local trackers are just one possible protocol-specific localization service
 - All the requirements for a generic localization service still apply
- It would be sensible to integrate a local tracker and a cache into the same system
 - Local trackers requires little state and little computation
 - Edge caches require a lot of state but little computation