Research issues raised by edge services

Sally Floyd

WCW 2002

August 15, 2002

Three topics:

- Data integrity and privacy.
- Balancing competing interests.
- Internet research needs better models.

Data integrity and privacy:

- Edge services involve modifying data between the content provider and the client, and raise new architectural issues.

- Data integrity:
How can data integrity be provided?
What are the requirements for data integrity?

- Privacy:
How can edge services be compatible with privacy, when that is desired?

Data integrity and privacy:

- Each edge service might have its own data integrity and privacy issues.

- Client-centric services:
  - Virus scanning.
  - Language translation.
  - Transcoding for limited-bandwidth clients.
  - Content filtering of requests.

- Server-centric services.
  - Ads.
  - Regional information.
  - Personalized web pages.

IAB Architectural and Policy Considerations for OPES:

Notification: The overall OPES framework needs to assist content providers in detecting and responding to client-centric actions by OPES intermediaries that are deemed inappropriate by the content provider.

Notification: The overall OPES framework should assist end users in detecting the behavior of OPES intermediaries, potentially allowing them to identify imperfect or compromised intermediaries.

IAB Architectural and Policy Considerations for OPES:

Non-blocking: If there exists a ”non-OPES” version of content available from the content provider, the OPES architecture must not prevent users from retrieving this ”non-OPES” version from the content provider.

Privacy: The overall OPES framework must provide for mechanisms for end users to determine the privacy policies of OPES intermediaries.

Balancing competing interests:

- Frameworks for edge services have to consider competing interests:
  - Clients.
  - Content providers.
  - Internet service providers.
  - Caching/content distribution providers.
  - Vendors for edge services.

- How and when should these competing interests be taken into account by the architecture?

– Clark, D., Wroslawski, J., Sollins, K., and Braden, R., ”Tussle in Cyberspace: Defining Tomorrow’s Internet”, SIGCOMM 2002.

Internet research needs better nodels:

- Each research question requires its own models.
  - Traffic models:

    Web traffic? Peer-to-peer traffic? TCP or UDP? Multicast?

    Flash crowds? Denial of Service attacks?
  - Topology models (with link bandwidths and propagation delays):

    Router-level topologies? AS-level topologies?

    Topologies of web caching infrastructures?
  - Models of congestion (packet loss rates, etc.)?
  - ...

- How do we understand which aspects of a model are critical for a particular research question, and which are not?

- How do we use measurement to improve our models?

URL: http://www.icir.org/models/bettermodels.html