Network Neutrality and the Future Internet

G. Kesidis
CSE and EE Depts
The Pennsylvania State University
gik2@psu.edu

Work supported by NSF grants 0916179, 1116626:
NeTS: Inter-provider dynamics in neutral and non-neutral networks;
in collaboration with S. Sarkar, U. Penn.

NSF US/Mideast Workshop on Trustworthiness in
Emerging Distributed Systems and Networks
Istanbul, June 4-6, 2012

Definition and motivation

- A neutral network does not unilaterally discriminate on applications based on type or origin, and does not charge remote "content" providers (CPs) for access to end-users (*i.e.*, no side payments)

- Network neutrality includes issues of pricing and revenue among all networking participants; significant revenue that can fuel network architectural change.

Origins of the NN debate

- Massive copyright infringements led copyright holders to seek remuneration from access providers (ISPs) and certain content providers.
- While access congestion due to P2P file sharing led some ISPs to
  - adopt policies that are not application neutral (*e.g.*, throttling BitTorrent traffic) and
  - consider usage pricing as a congestion penalty.
  - Note that usage pricing also raised in the contexts of overage of a monthly quota and premium (not best-effort) access service.
- Arguments against non-neutral action often invoke freedom of expression and freedom of competition.
- Also declared fundamental rights to Internet access, a heavily regulated industry with enormous infrastructure costs.
- NN debate highlights competition between ISPs' "managed" services ("eyeball" ISPs) and CPs over public commodity Internet (*e.g.*, Skype, Netflix).

Application neutrality

- Again, application neutrality involves non-discrimination of applications based on type or origin.
- *End-user* selected service differentiation is neutral.
- Note that certain content providers themselves not immune from neutrality complaints, *e.g.*,
  - Web search is "name resolution" infrastructure for the web, and
  - the manner in which search results are presented may be subject to neutrality interpretation.

NN and traffic volume

- Neutrality prohibits discrimination based on content and origin, but not necessarily on other traffic attributes such as volume.
- Deployed traffic throttles based on aggregate traffic volume are neutral.
- In a neutral network, ISPs can gain more revenue from usage-priced access modalities, *e.g.*, quota overages by aggregate traffic volume.
- Note that download quotas are becoming prevalent in wireless access where consumers may be more accepting of quotas, possibly owing to the perceived convenience of the service.

A consequence of broadband-access quotas

- A consumer may engage in usage-priced access simply after a quota has been reached.

- Considering the potential for misuse, *e.g.*, malware activity precipitating quota overages, it is important that traffic is *authorized* by the human end-user in a usage-priced context.

- That is, "network-layer" authenticated "sessions" by access ISP.

From broadband-access quotas to end-to-end middleware

- There are detailed standards for middleware (*e.g.*, midcom) that can indicate whether a session is thus authorized throughout the Internet.
- Given deployment of such a secure end-to-end middleware system, consumers may be enticed to engage in authorized access services to improve service availability end-to-end.
- For example, quickly upon detection of the onset of a DDoS attack on a public server, it may *reactively* signal upstream firewalls to block all traffic except that which is thus authorized.
- This would be an improvement over reactive challenge-response.
- Such a deployed middleware system may
  - "ignite" deployment of end-to-end differentiated-services architecture where
  - users may engage in usage-priced *premium* access services for increased QoS for particular sessions.