

P2P Application Classes and the IETF

What do we have? : What are we missing?

David Bryan
P2PSIP WG co-chair
Polycom
July 28, 2011

Introduction

- Cover major classes of P2P applications
 - History of them, big problems, how they touch the IETF
- What I mean here by “P2P”
 - P2P isn’t new. Using the modern sense of the word
 - P2P between clients, includes hybrid architectures
 - Excluding early things like SMTP, BGP
- This is a very simplistic and incomplete view
 - Many other P2P things out there, touching on major ones
- High level here – major common tasks of P2P apps

Outline

- Cover major classes of applications
 - Discussion of transport considerations of each
- What are the major “tasks” P2P applications perform, and which application classes use each?
 - Where do we have an existing IETF protocol for each task?
 - Where is work going on that ***could*** address the task?

File Sharing

- This was where the “new” P2P got started
 - Napster, Kazaa, Gnutella, etc.
- Have been a variety of approaches
 - (semi)-central server, direct data exchange
 - Napster (index server), BitTorrent (tracker)
 - Fully decentralized, flood based
 - Gnutella
 - Structured, DHT (Distributed Hash Table) based
 - Not widely used in “commercial” file sharing
- Originally music, most volume now video. Many uses
- Major transport considerations: large volumes of data between endpoints, often “odd” techniques to game transport

Video Streaming/P2PTV

- Some fuzzy lines with file sharing
 - Some P2P VoD works like file sharing (not real time)
 - True streaming share live, real-time video
 - Many services offer both
- PPLive, Octoshape, Joost (early)
- Major transport considerations: fairly tight real-time constraints, often many simultaneous participants

Communications

- Skype
 - Directly connects end users for voice, video
 - Massive scale (170m users, 207b minutes)
- Cisco Intercompany Media Engine
 - Uses a structured P2P approach to connect enterprises to complete calls
 - Based on P2PSIP RELOAD and ViPR work
- Major transport considerations: need for bi-directional media, very tight real-time constraints, relay considerations

Other Applications

- P2P CDNs
 - Either P2P is used in the back, or hybrid P2P and CDN (Pando)
- P2P “money”
 - BitCoin is a decentralized currency
 - Interesting aspects are primarily crypto/computational (proof-of-work), not networking
 - Clients do connect with a P2P structure to authenticate coins and exchange information
- P2P DNS

IETF Efforts

- What are some common P2P functions (“tasks”), and where/are these addressed in the IETF?
 - Broader uses for the application-specific P2P groups (PPSP, P2PSIP) products
 - Applicability of the generic P2P groups (ALTO, DECADE, LEDBAT)
 - Non-P2P protocols that are being reused
 - Future work that might be applicable (if WGs go that way)
- How do the tasks and IETF P2P efforts (mostly now in TSV) touch these application classes?

Task: Forming/Maintaining an Overlay

- Overlay structures actually differ quite a bit
 - DHTs
 - Broadcast Tree
 - Unstructured meshes, etc.
- Common to all apps is a need to learn about and exchange information about other peers
- Today's situation
 - Applications use many different structures
 - Most use proprietary on-the-wire encoding

Forming/Maintaining an Overlay

- RELOAD in P2PSIP
 - Designed generically, targeting DHTs
 - Modularity of DHTs considered, pluggable
 - Standardized mechanism for exchanging messages
 - Attempted to support unstructured overlays in design
 - To date, has not been tried
- PPSP Peer Protocol proposals
 - WG looking at connecting peers for streaming
 - May use RELOAD, may build a new (lighter?) protocol
 - May also only cover information exchanged (not address organization of the peers themselves)
 - WG still debating direction

Task: Search (Peer and Data Location)

- Again, cross-application need to locate peers
- Method can have large impact on network (flood)
- Locating using DHT (structured search)
 - RELOAD supports DHT hash-based search
 - No explicit support for other forms of search
- Central index server / learn from peers
 - PPSP is considering a tracker/index server protocol to locate peers
- Unstructured search
 - PPSP also considering allowing peers to learn from other peers in peer protocol
- Again, PPSP work is early – unclear how modular or what will actually be work items

Task: Exchanging Data

- Most significant implications for transport
- Many apps applications use naïve (or harmful) mechanisms today (i.e. many simultaneous HTTP streams)
- Efforts in IETF to address this:
 - LEDBAT: provide efficient data exchange without adding delay
 - DECADE: provide generic in-network storage to minimize last-mile traffic
 - ALTO: Locate optimal peers/hosts to obtain data from candidates
- Need to identify which mechanism/protocols are appropriate for applications to transfer data
 - DECADE currently plans to consider existing mechanisms
 - Can that effort be used generically?
 - Should the IETF provide guidance based on transfer patterns?

Task: Identity, Reputation and Security

- Strong need to identify peers, verify identity, decide who to trust, and secure connections
- Some specific P2P work in IETF
 - RELOAD (P2PSIP): provide a centrally-issued certificate identity system for overlays
 - DECADE: mechanism to centrally issue tokens used to authorize access to data in other locations
 - ViPR group is working on a very specific identity problem (P2P proof of phone number ownership)
- To date, no IETF P2P-work on reputation
 - Is there work in security area? New work to be done?
- Many non-P2P efforts of IETF applicable
 - TLS etc.

Task: NAT Traversal

- Today, P2P NAT traversal is generally either tunneling or uses IETF suite
- Work on this issue is being done in IETF, but not in P2P specific group
 - Problem is exacerbated by, but not uniquely P2P

What is the IETF missing?

- Some “pieces” that we don’t currently have solutions for
 - Places where work is underway that might be used for this (depending on what WG does)
- Clearly agreed upon mechanism to exchange data between peers
 - DECADE planning to identify mechanism for their use
- Protocol for non-DHT peer structures
 - Current work on PPSP Peer Protocol might address this
- Protocol for central tracker/index server
 - PPSP Tracker Protocol looks likely to address this issue
- One-to-many transport issues (multicast systems for streaming video)
- More distributed identity solutions, reputation systems
 - Some of this is still not well understood (research issues)

Conclusion

- IETF is addressing many critical issues to P2P
- Transport related issues are very important in design of P2P protocols
- Still lots of work before one could make a complete P2P solution using IETF pieces

Conclusion

- P2P is a set of techniques/implementation approach, not an application per-se
- Some work at IETF on generic P2P techniques, some on specific applications that happen to be P2P
- We are assembling a suite of protocols that can be used in P2P applications for more rapid development, better interoperability