CLive: Cloud-assisted P2P Live Streaming

Amir H. Payberah[1], Hanna Kavalionak[2], Vimalkumar Kumaresan[2]
Alberto Montresor[2], Seif Haridi[1]
{amir,seif}@sics.se[1], {name.family}@unitn.it[2]

The Problem

Existing P2P Streaming Solutions

Client-Server

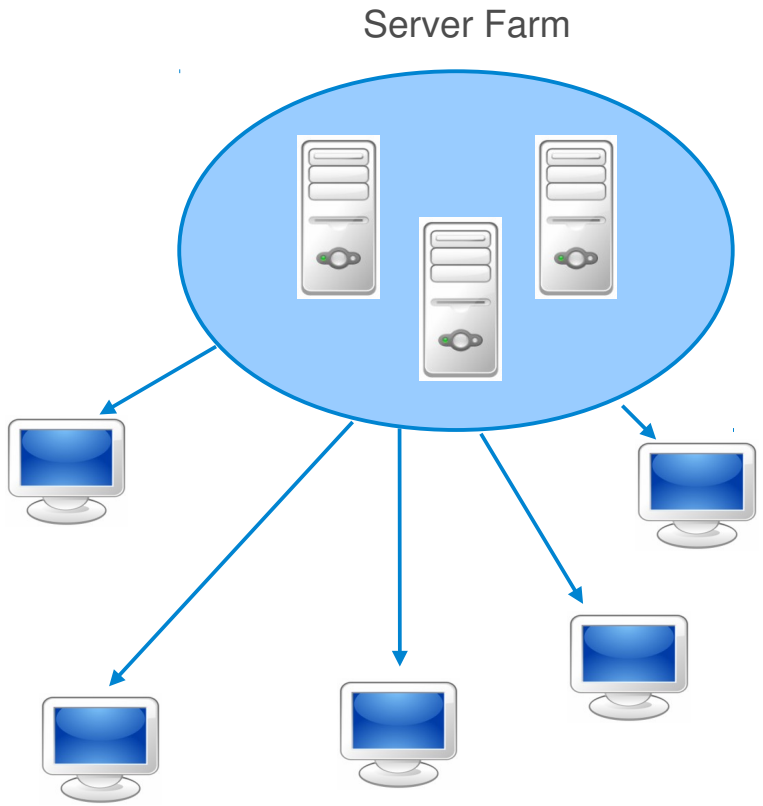

Peer-to-Peer

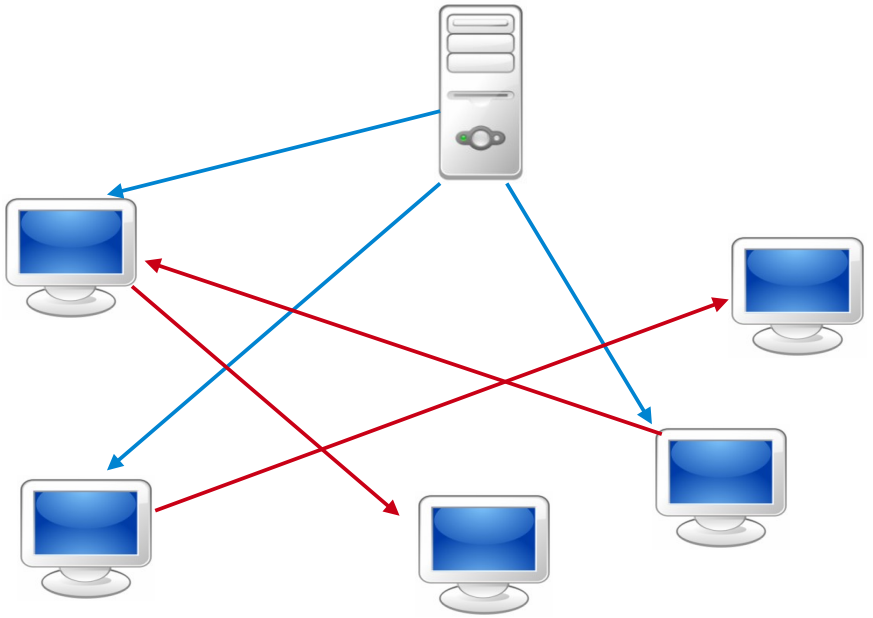

Problem?

- Bottlenecks in P2P video streaming systems: upload bandwidth

- A potential solution: P2P network is assisted by a cloud computing.

Problem?

- Bottlenecks in P2P video streaming systems: upload bandwidth

- A potential solution: P2P network is assisted by a cloud computing.

Problem?

- Bottlenecks in P2P video streaming systems: upload bandwidth

- A potential solution: P2P network is assisted by a cloud computing.

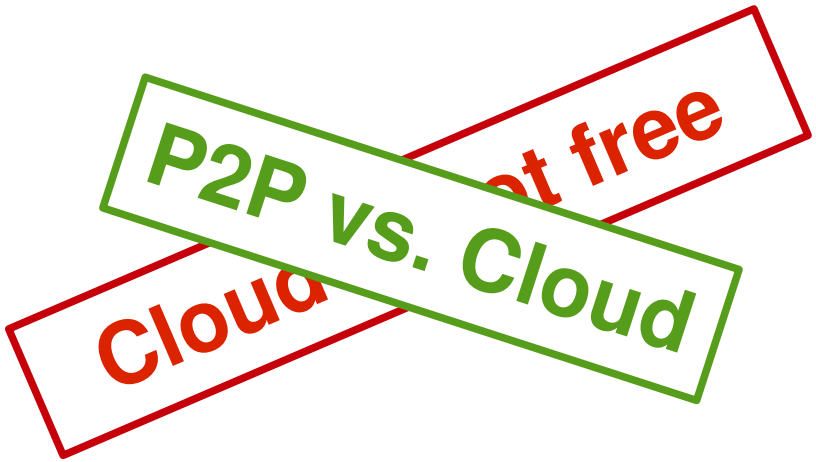

P2P vs. Cloud

- P2P pros and cons
  - P2P resources are cheap
  - Churn may compromise availability
- Cloud pros and cons
  - Superior availability
  - Cloud resources are not free

If You Cannot Beat Them, …? ;)

- The cloud as a support group for P2P.
- Reduce the number of (costly) cloud interactions as much as possible.

If You Cannot Beat Them, …? ;)

- The cloud as a support group for P2P.
- Reduce the number of (costly) cloud interactions as much as possible.

The problem to be solved becomes minimizing the economical cost, provided that a set of constraints on QoS is satisfied.

CLive

Baseline Model

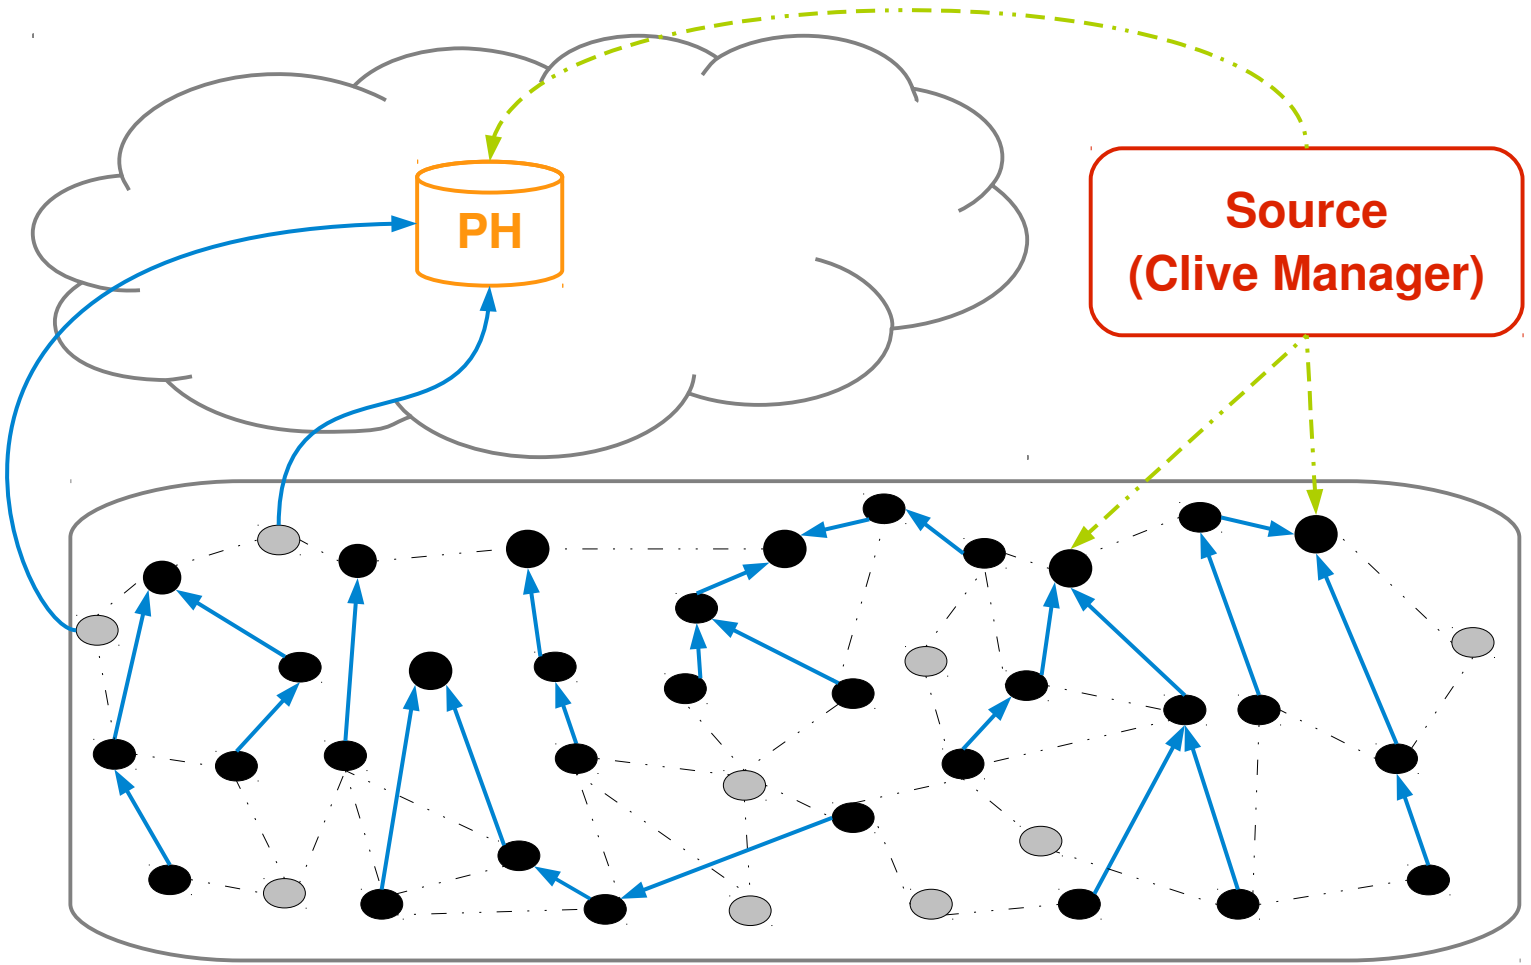

Enhanced Model

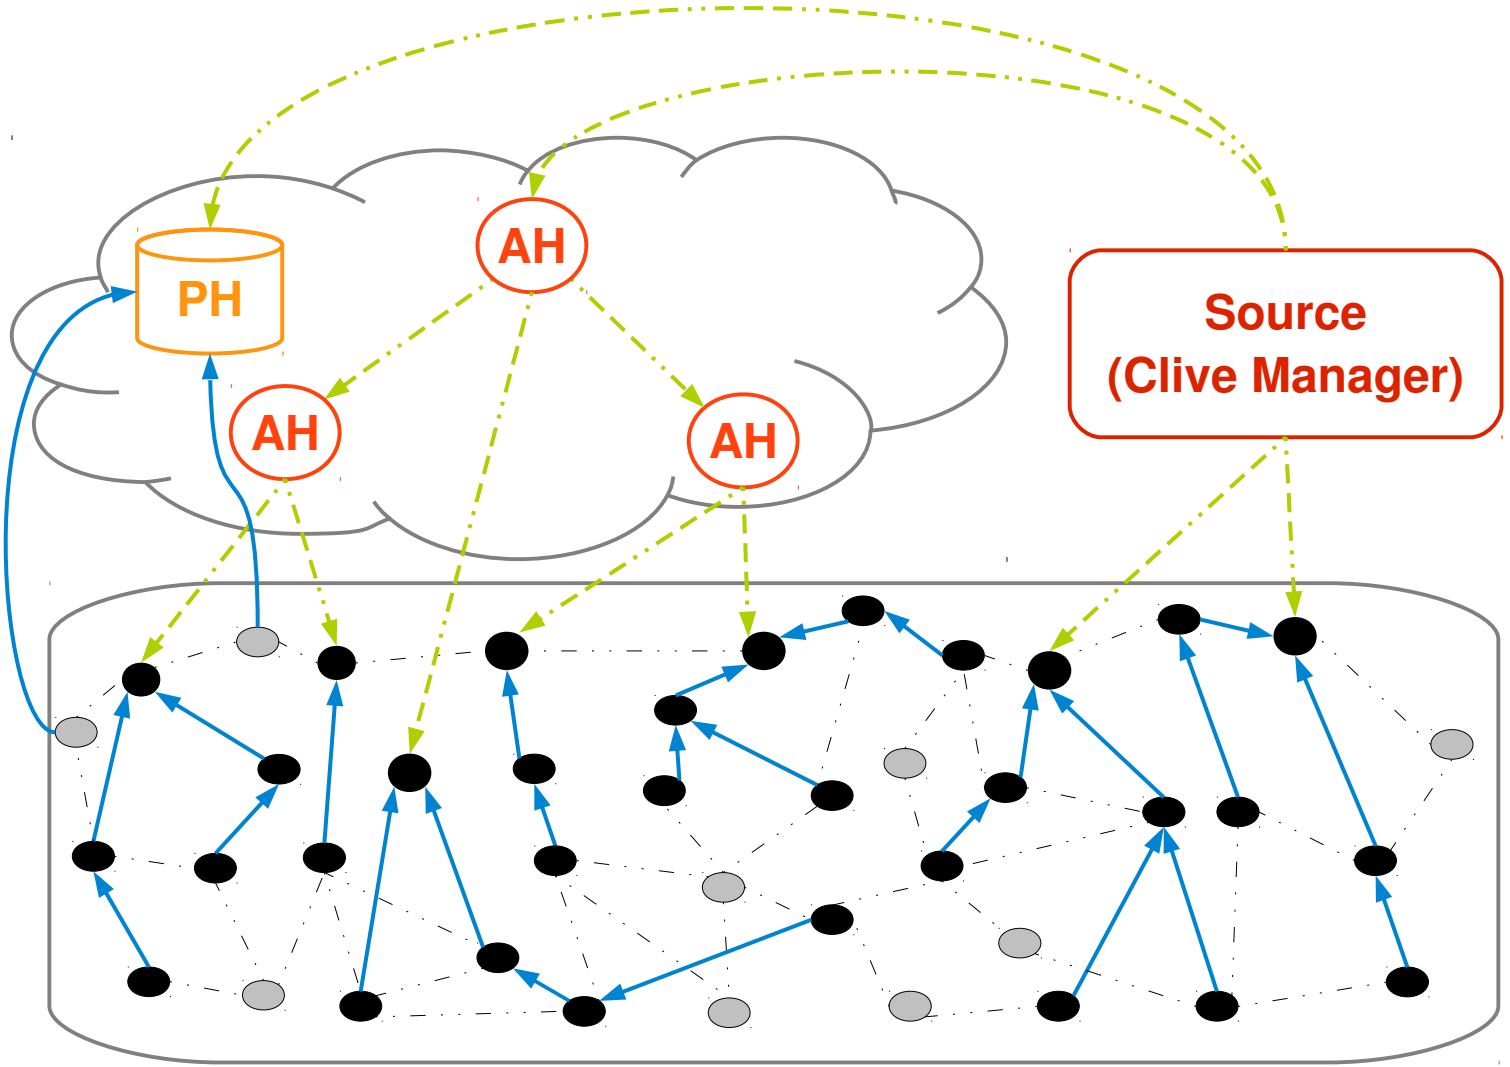

CLive Components

- P2P streaming overlay (swarm)
- Media source
- Passive Helpers
- Active Helpers
- Management component

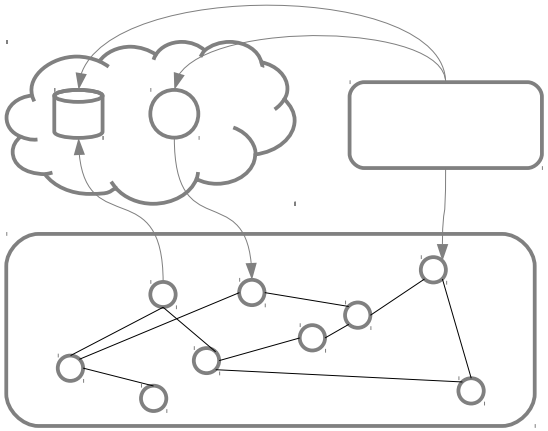

CLive Components: P2P Streaming Component

- We assume that nodes in the system use a mesh-pull model for data distribution.
- Nodes periodically send their buffer maps to their neighbours.
- The other nodes pull the required chunks from those nodes who own the chunks.

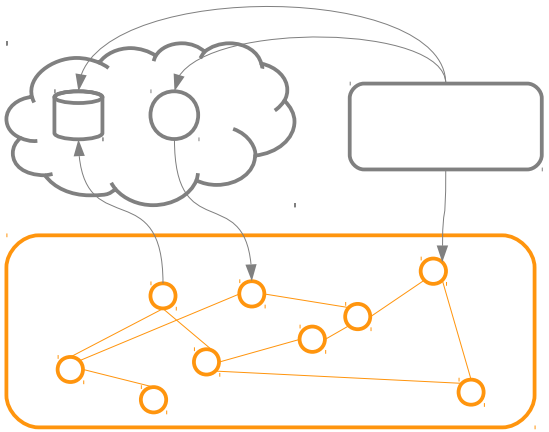

CLive Components: Media Source

- A media source is a node that generates data chunks and pushes them to the swarm.

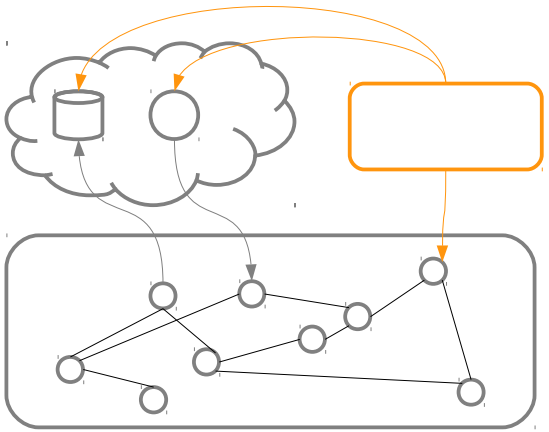

CLive Components: Passive Helper

- PH is a passive element that plays the role of a data storage, e.g., Amazon S3.
- The source pushes chunks to PH, as they are generated.
- The swarm nodes pull the missed chunks from it.
- We assume that a PH can serve as many requests as it receives.

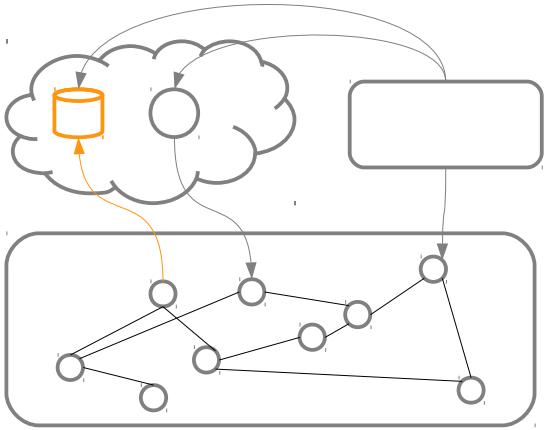

CLive Components: Active Helper

- AHs are active elements, e.g., Amazon EC2, that cooperate with other swarm nodes to accelerate the data dissemination.
- The source pushes chunks to AH, as they are generated.
- AHs forwards chunks to the swarm and other AH, as they are received.

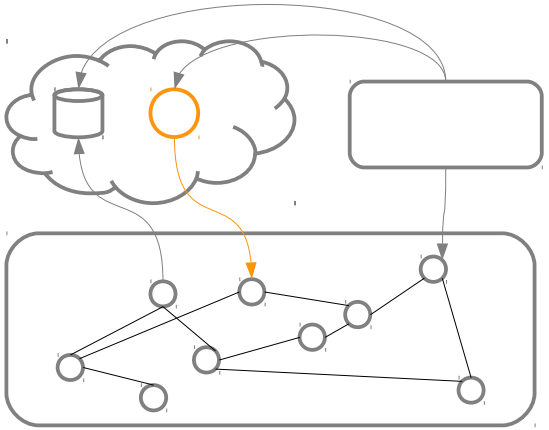

CLive Components: Management

- Participates in a gossip algorithm to estimate the available resources in the system.

- Adds/removes AHs to/from the system, based on the estimation.

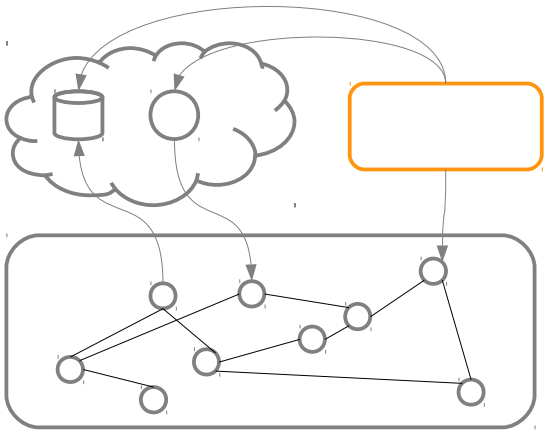

Two Main Questions in CLive?

- How to estimate the extra load in the overlay?
- How to relay the load to cloud with a minimum cost?

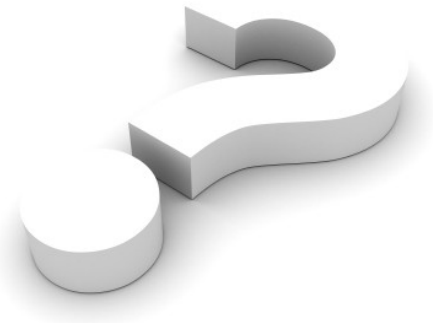

Two Main Questions in CLive?

- How to estimate the extra load in the overlay?
- How to relay the load to cloud with a minimum cost?

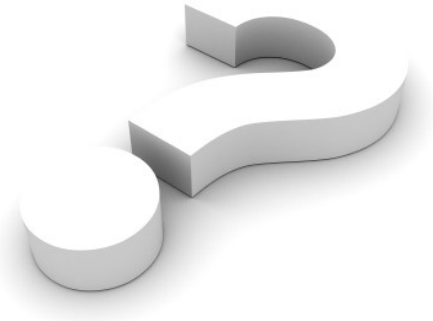

Main Idea

- Infected nodes: the number of nodes that can be served with the existing resources in the system.

- load = swarm size – infected nodes

The Swarm Size

- The swarm size estimation is easy: gossip-based aggregation

The Infected Nodes

- It is shown that each streamed chunk through mesh overlays follows a tree-based diffusion pattern [1].

- To compute the number of infected nodes, we model a diffusion tree for a chunk and estimate the number of nodes in that tree.

[1] B. Biskupski, M. Schiely, P. Felber, and R. Meier, “Tree-based analysis of mesh overlays for peer-to-peer streaming,” in Proc. of the 8th int. conf. on distr. app. and interoperable systems. Springer, 2008.

Estimate the Number of Nodes in a Tree

- Estimate the tree depth
- Estimate the fan-out distribution

Estimate the Tree Depth (1/2)

- Last Chance Window (LCW): A predefined number of data chunks that must be buffered ahead of the playback point.
- The missed chunks are downloaded directly from the PH, when the number of data chunks that are buffered ahead is less than LCW.
- However, if the LCW buffer is full, nodes pull the required chunks from the swarm.

Estimate the Tree Depth (2/2)

- T_{delay}: maximum acceptable latency.
- $T_{life} = T_{delay} - T_{lcw}$

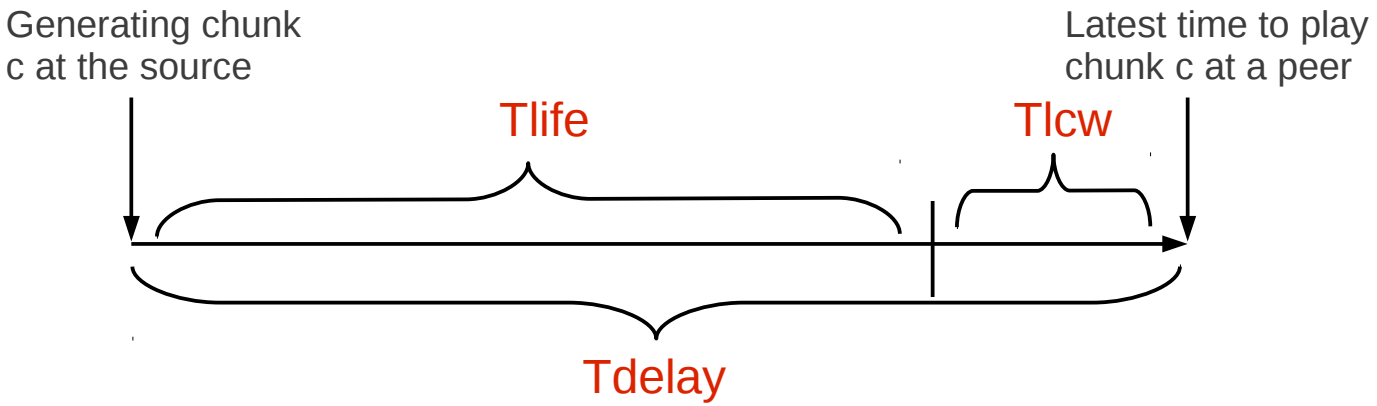

- If a node can not receive a chunk in T_{life}, it pulls that chunk from PH.
- The tree depth
  - $D_{max} = T_{life} / T_d$
  - T_d: average latency among the peers

Estimate the Fan-out Distribution

- The upload bandwidth distribution
  - Adam2
  - Gossip-based aggregation

Estimate the Number of Nodes in a Tree

- Estimate the tree depth
- Estimate the fan-out distribution

The Number of Infected Nodes

- N_{inf} = (the number of trees).n_{tree}
- The number of trees: source fanout + AHs fanout

Two Main Questions in CLive?

- How to estimate the extra load in the overlay?

- How to relay the load to cloud with a minimum cost?

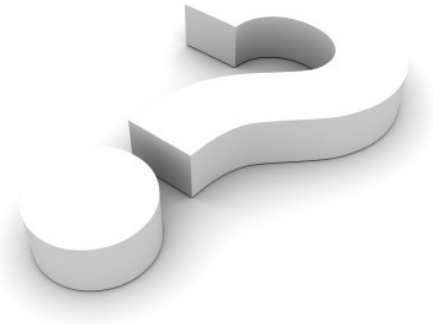

AH Cost

- AH cost in one round: $C_{ah} = C_{vm} + m.C_{chunk}$
  - C_{vm}: virtual machine cost
  - C_{chunk}: chuck transfer cost
  - m: number of chunks that one AH upload per round

PH Cost

- PH cost in one round: $C_{ph} = C_{storage} + r.(C_{chunk} + C_{req})$
  - $C_{storage}$: the storage cost
  - C_{chunk}: chuck transfer cost
  - C_{req}: chuck request
  - r: the number of retrieved chunks from PH in one round

∂

- We define ∂ as the number of peers that is economically reasonable to serve with PH utilization instead to run an additional AH for them.

- $\partial = C_{ah} / C_{ph}$

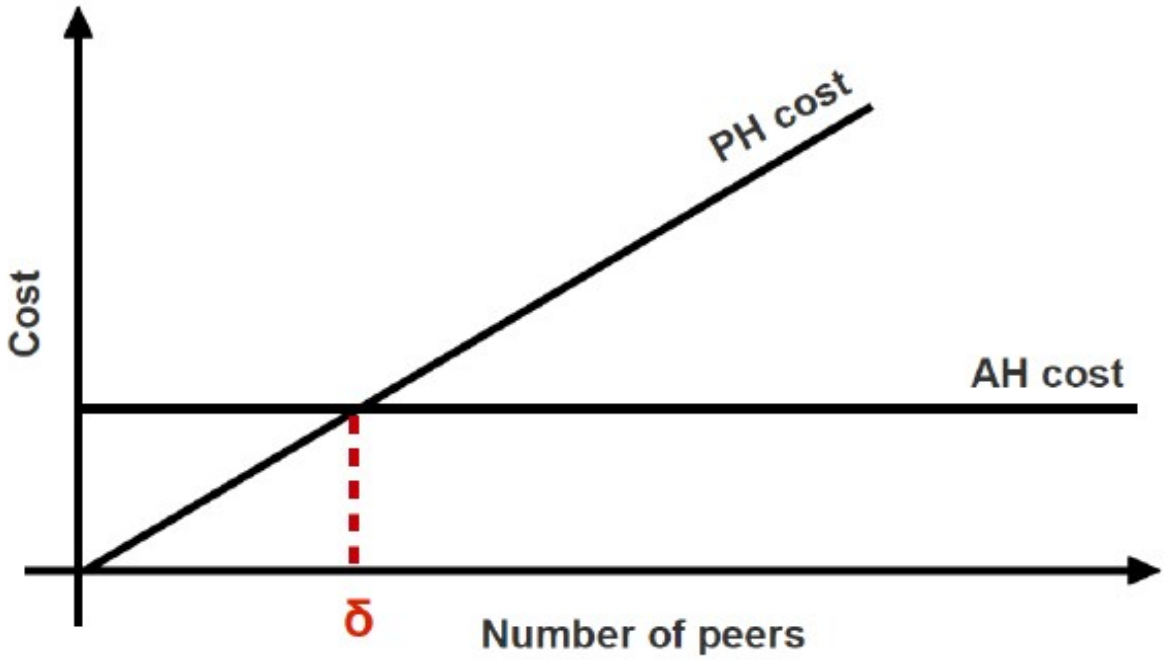

Manage the Cloud Resources

- If load > ∂: add AH
- If load < ∂ - H: remove AH
  - H: number of peers served by one AH.
- Otherwise don't change AHs.

Experiments

Different LCW size

Playback continuity

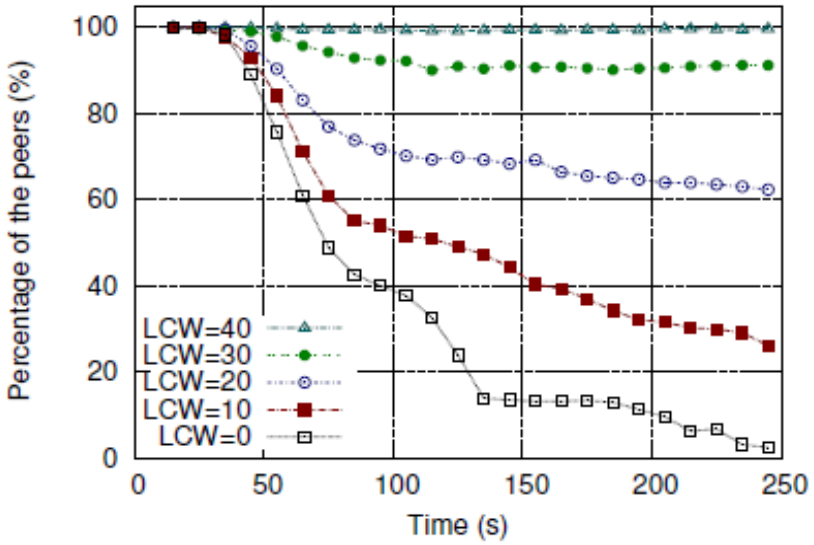

Playback latency

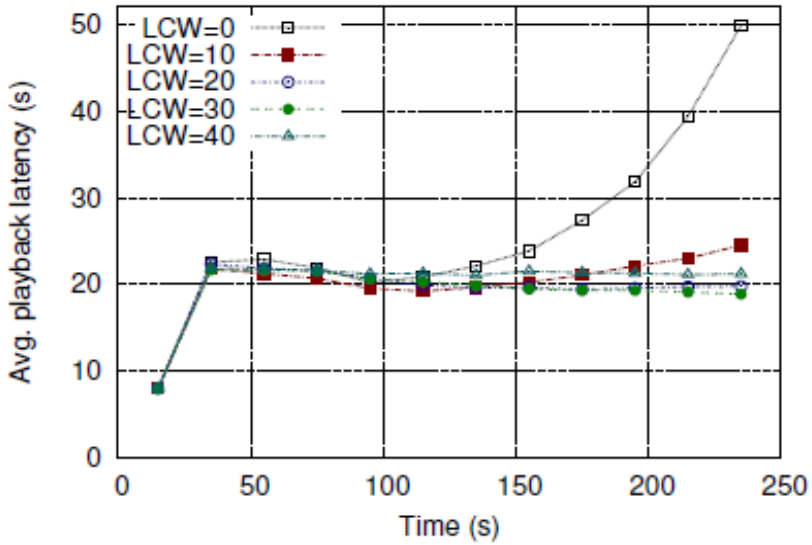

PH Load

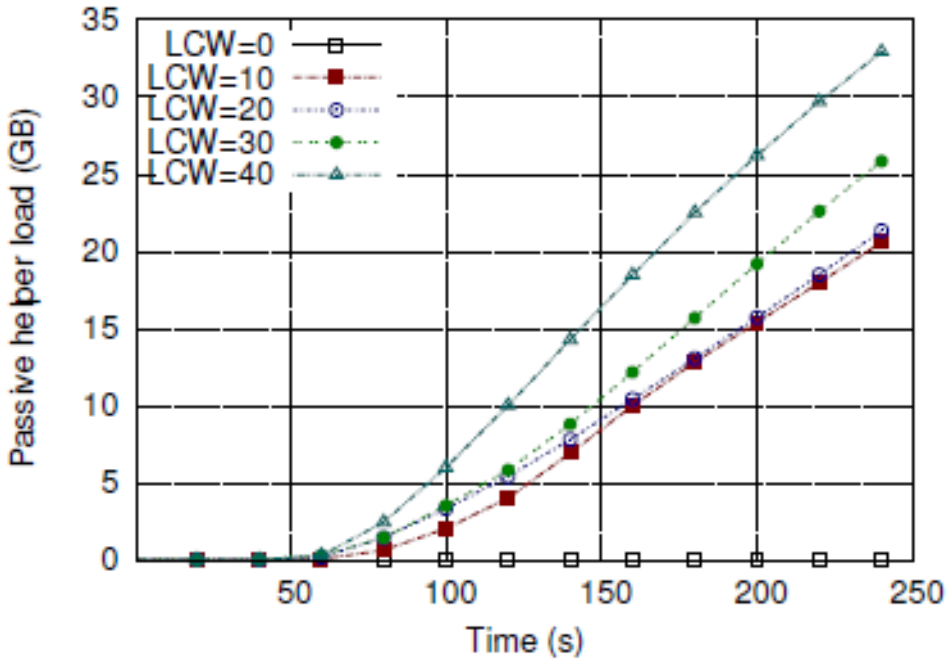
LCW=0
LCW=10
LCW=20
LCW=30
LCW=40
Passive helper load (GB)
Time (s)
0
5
10
15
20
25
30
35
50
100
150
200
250

PH Load in Different Churn Rates

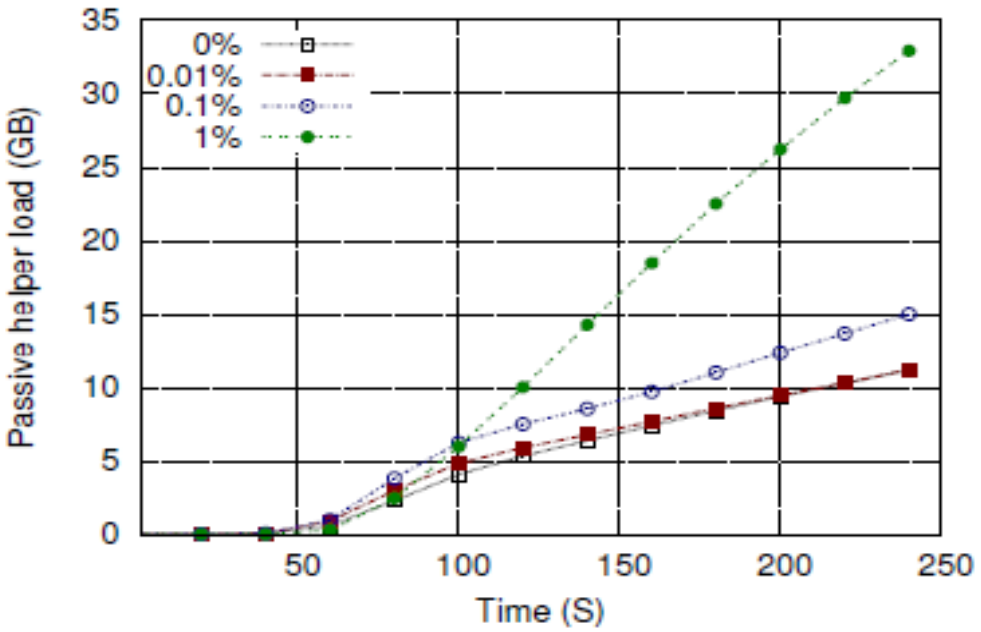

Slot Distribution Estimation

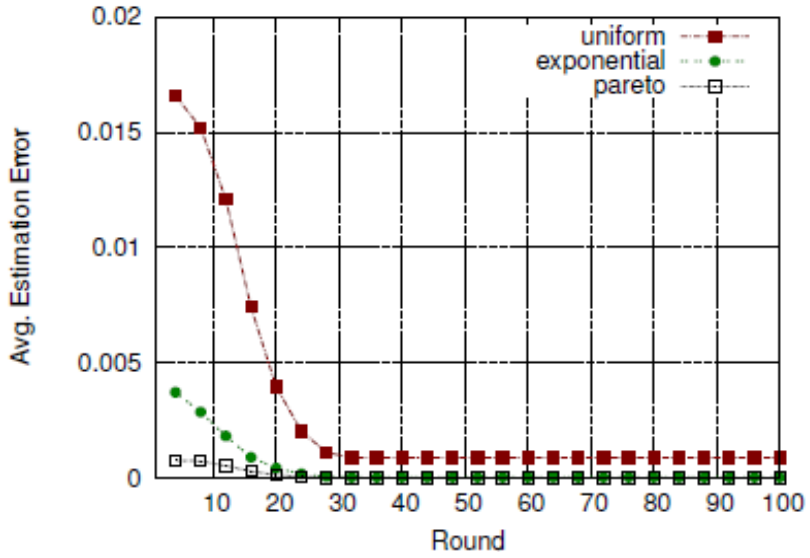

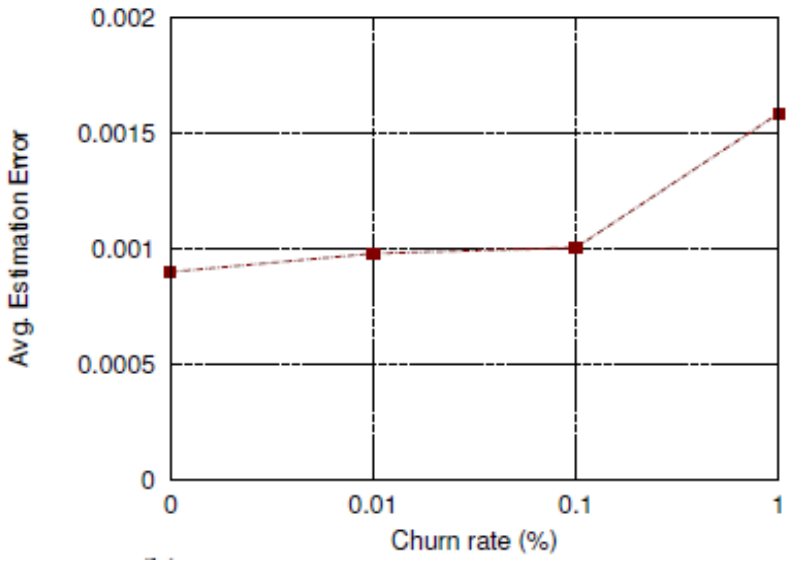

Cost

Number of AHs

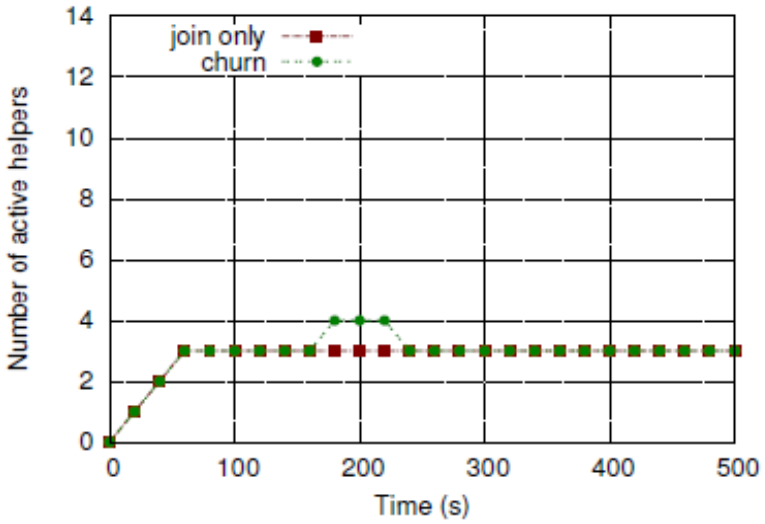

PH load

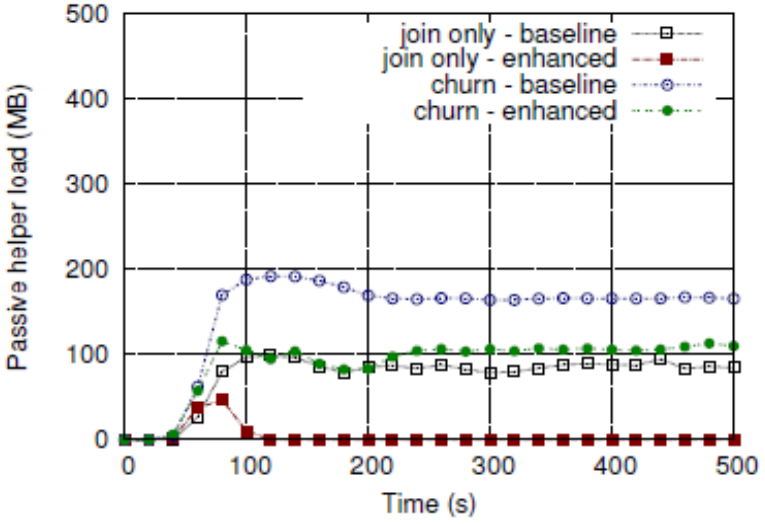

Total cost

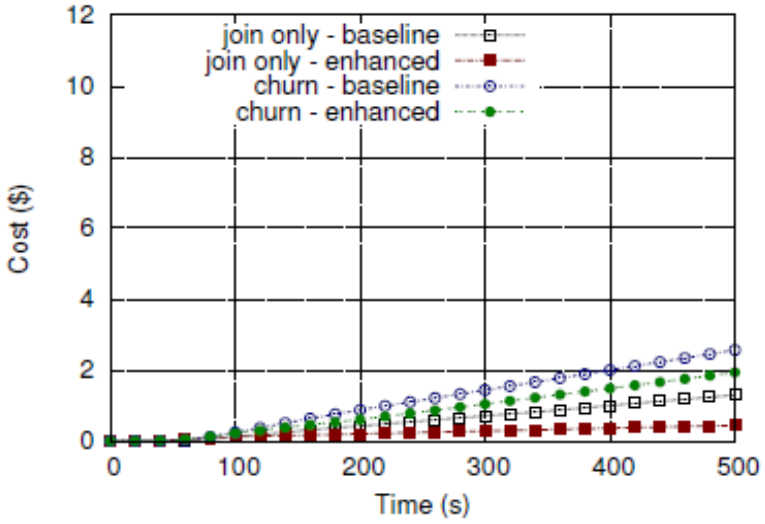

Cost – Real Trace Slot Distribution

Number of AHs

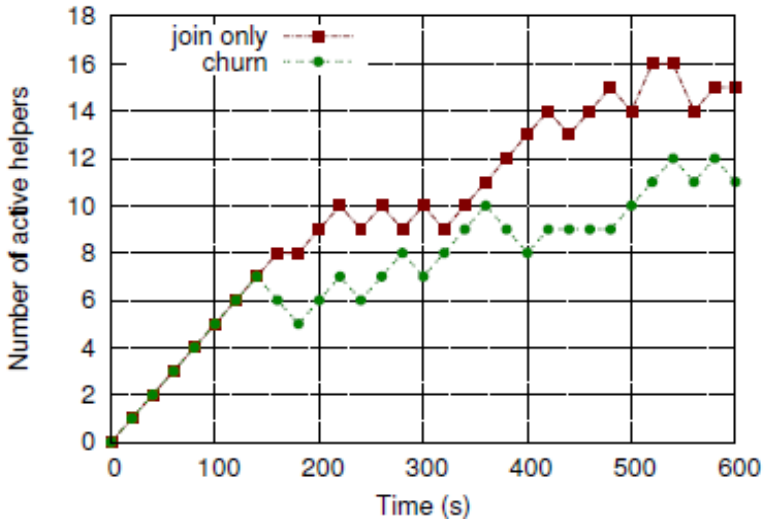

PH load

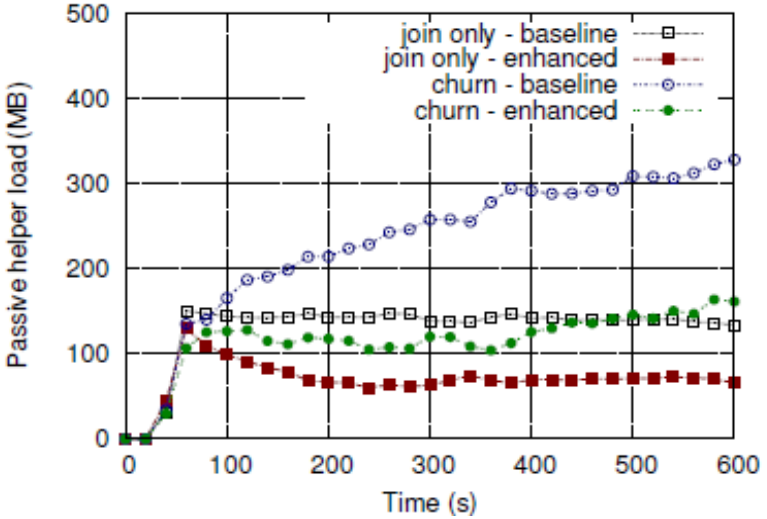

Total cost

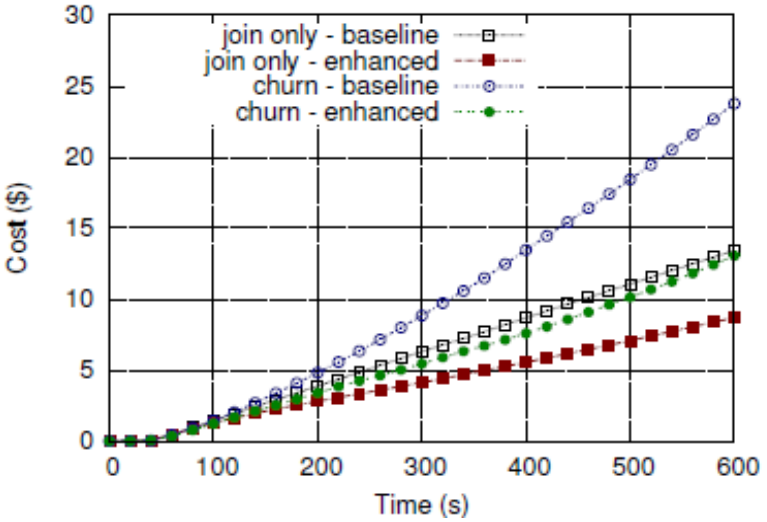

Summary


WHAT
IS OUR
ONE
DEMAND?
#OCCUPYWALLSTREET
SEPTEMBER 17TH.
BRING TENT.

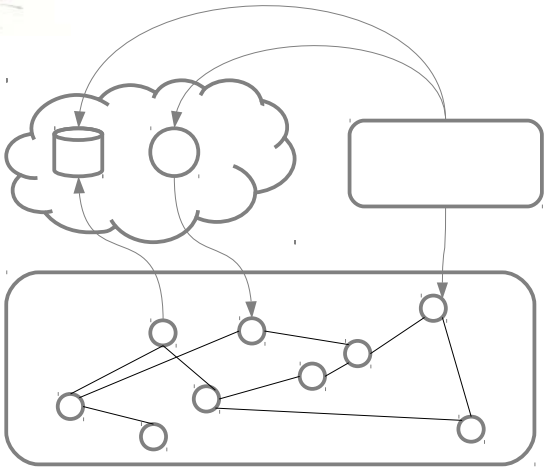

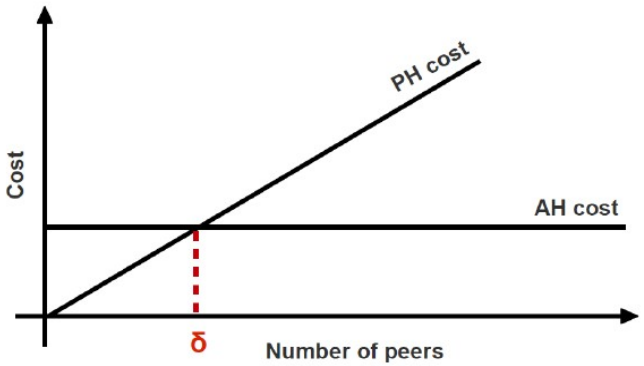
PH cost
AH cost
Cost
δ
Number of peers

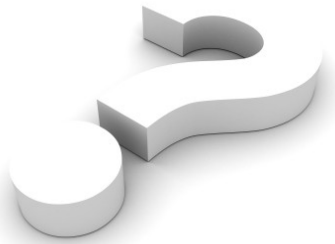

Any Questions?