



China Telecom ALTO/P4P Trial

draft-lee-alto-chinatelecom-trial-00


Kai Lee leekai@ctbri.com.cn

China Telecom



Contents

- Trial settings
- Trial results



Participants

□ Network: China Telecom

- The largest broadband service provider in China. Serving over 60 million fixed-line broadband subscribers
- More than 200 MAN networks managed by 31 provincial networks

□ P2P Application: Xunlei (also called Thunder)

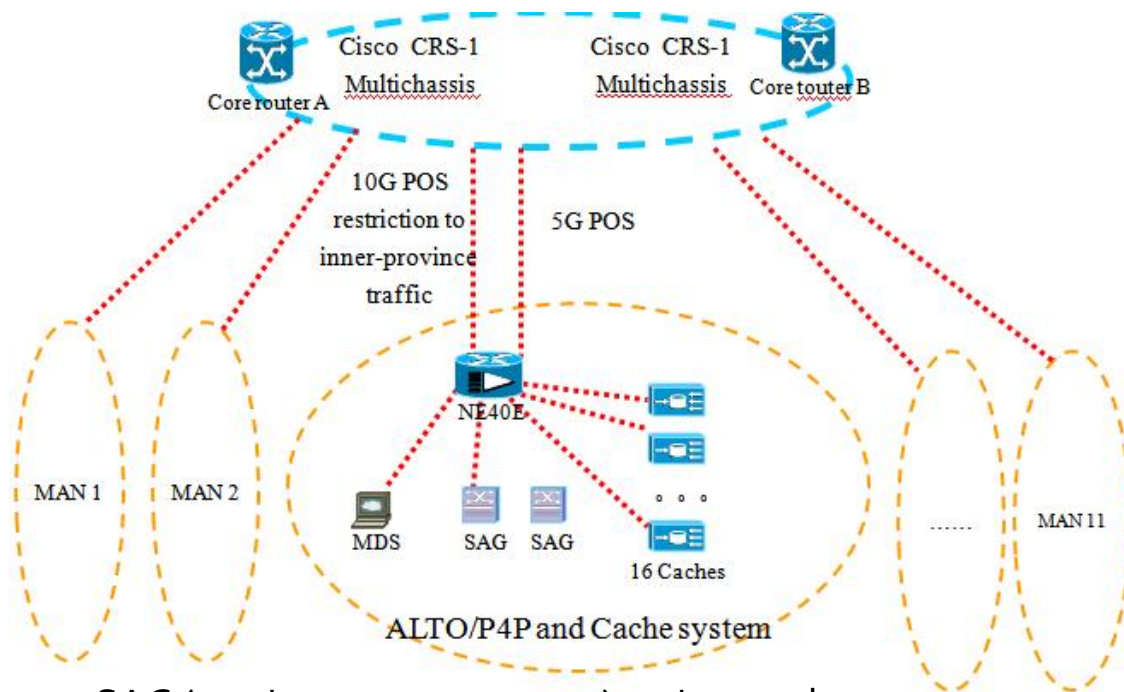
- Serving 20 million users each day
- Contributing ~20% inter-province traffic in the trial network



Trial Configuration: Network

- ❑ Trial conducted in one provincial network with 7 million broadband subscribers and 11 MAN networks
- ❑ Trial network deploys both ALTO/P₄P Server and Cache Servers


Trial Configuration: Network Components



- ❑ ALTO/P4P Server: SAG (service access gateway) acting as alto server
- ❑ Cache: 16 Caches with 1.8TB storage and 15 Gbps BW
- ❑ Management: MDS (Management and Deployment System) provides a browser-based interface network operators:
 - Update the information about network topology, geography location of IP address ,etc
 - Define traffic management policy
 - Monitor the SAG and p2p Cache running states
 - Analyze the statistics of xunlei P2P traffic

ALTO/P4P Network Configuration

- ❑ Report results of a simple setting
- ❑ PID Map: We define 11 PIDs
 - PID₁-PID₁₁ represent the 11 MANs of the trial network
 - PID₁₂ represents rest of the Internet
- ❑ Cost Map:
 - Bidirectional cost between any PIDs from PID₁ to PID₁₁ has the same value 1
 - Bidirectional cost between PID₁₂ and PID_i ($1 \leq i \leq 11$) has the same value 2






Trial Configuration: Xunlei



- ❑ Xunlei is a hybrid application utilizing both trackers and DHT
- ❑ About 85% of Xunlei traffic controlled by Xunlei trackers
- ❑ Trial modifies only Xunlei trackers

- ❑ Xunlei tracker peer selection algorithm before trial:
 - Utilizing two properties: ISP ID and UC (upload capability)
 - Peer selection priority:
 - Same ISP ID > different ISP ID
 - Higher UC > lower UC



Xunlei's Trial Algorithm: Utilizing ALTO/P4P Info

- ❑ With ALTO/P₄P information, the ISP ID of a peer in the trial province changes to "ISP ID_PIDx", where PIDx is the PID to which the peer belongs
 - Example: the original ISP ID of a peer in the trial province is "ct", and the peer's IP address belongs to PID₁, then the ISP ID of the peer is "ct_PID₁"
- ❑ Updated peer selection priority:
 - Same ISP ID > Different ISP ID > ISP ID with PIDx suffix
 - Higher UC > lower UC
- ❑ If a listing request is from the trial province, the maximum # of returned peers from xunlei tracker is set to 120, not the normal case of 500.



Data Collection Methods


- Measuring performance by both CT and Xunlei
 - CT measurement: using ISP's network OAM system and DPI system
 - OAM: Total (including Xunlei and others) inbound/outbound SNMP traffic flow statistic
 - DPI: Detect Xunlei traffic flow; monitor the uplink (only uplink direction) from the province to China Telecom's backbone.

 - Xunlei measurement:
 - Inter-province/intra-province traffic flow statistics [by sampled client logs]
 - Average user download speed



Configuration Differences Between This Trial and Comcast Trial (RFC5632)

	Comcast Trial	CT Trial
Contents	One content (22 MB) divided into 5 swarms for comparison	<i>All</i> Xunlei file/streaming contents in 4 months
ALTO/P4P conf.	Coarse-grain (22 PIDs) and fine-grain (1182 PIDs)	12 PIDs
P2P App/Peer Selection Alg.	Pando/Coarse-grain; Fine-grain; Generic	Xunlei's algorithm (slide 8)
Utilizing Cache	No	Yes (16 servers)
Monitoring	Pando client logs	Xunlei logs; CT's NMS system; CT's DPI detection



Trial Results: No Caching Servers

- ❑ Reduction of **Xunlei** inter-province outbound traffic
 - 42.77 Gbps (51% of outbound)

- ❑ Reduction of **total** inter-province outbound traffic
 - 31.58 Gbps

- ❑ Reduction of **total** inter-province inbound traffic
 - 12.86 Gbps

- ❑ Average user download speed
 - Slows down ~10%



Trial Results: with Caching Servers

- ❑ Adding caches (16 caches, 1.8TB each and 15 Gbps total BW)

- ❑ Reduction of **Xunlei** inter-province outbound traffic
 - 40 Gbps (~50%)

- ❑ Reduction of **total** inter-province outbound traffic
 - 39.18 Gbps


- ❑ Reduction of **total** inter-province inbound traffic
 - 28.3 Gbps

- ❑ Average user download speed
 - 5% increase with 4 Gbps cache BW, and 80% cache occupancy



Conclusions and Future Work

- ❑ ALTO/P4P information can substantially reduce Internet egress/ingress traffic
- ❑ Observed slight slowdown (~10%) in application speed without cache servers
- ❑ Network storage can recover lost speed:
 - 4 Gbps ISP cache BW => 40 Gbps egress BW reduction and 5% speedup
- ❑ Future work:
 - ❑ Evaluate more details on the slow down and determine possibility of improving speed
 - ❑ Broader deployment and more diverse types of applications
 - ❑ Collect ISP configuration and P2P application algorithm guidelines/best practices and contribute to ALTO.



Any comments or questions?