                        ALTO Deployment Considerations
                    draft-stiemerling-alto-deployments-06

Abstract

   Many Internet applications are used to access resources, such as
   pieces of information or server processes, which are available in
   several equivalent replicas on different hosts. This includes, but
   is not limited to, peer-to-peer file sharing applications. The goal
   of Application-Layer Traffic Optimization (ALTO) is to provide
   guidance to these applications, which have to select one or several
   hosts from a set of candidates, that are able to provide a desired
   resource. The protocol is under specification in the ALTO working
   group. This memo discusses deployment related issues of ALTO for
   peer-to-peer and CDNs, some preliminary security considerations, and
   also initial guidance for application designers using ALTO.

Status of this Memo

   This Internet-Draft is submitted in full conformance with the
   provisions of BCP 78 and BCP 79.

   Internet-Drafts are working documents of the Internet Engineering
   Task Force (IETF). Note that other groups may also distribute
   working documents as Internet-Drafts. The list of current Internet-
   Drafts is at http://datatracker.ietf.org/drafts/current/.

   Internet-Drafts are draft documents valid for a maximum of six months
   and may be updated, replaced, or obsoleted by other documents at any
   time. It is inappropriate to use Internet-Drafts as reference
   material or to cite them other than as "work in progress."

   This Internet-Draft will expire on July 29, 2011.

Table of Contents

1. Introduction

   Many Internet applications are used to access resources, such as
   pieces of information or server processes, which are available in
   several equivalent replicas on different hosts. This includes, but
   is not limited to, peer-to-peer file sharing applications and Content
   Delivery Networks (CDNs). The goal of Application-Layer Traffic
   Optimization (ALTO) is to provide guidance to applications, which
   have to select one or several hosts from a set of candidates, that
   are able to provide a desired resource. The basic ideas of ALTO are
   described in the problem space of ALTO is described in [RFC5693] and
   the set of requirements is discussed in [I-D.ietf-alto-reqs].

   However, there are no considerations about what operational issues
   are to be expected once ALTO will be deployed. This includes, but is
   not limited to, location of the ALTO server, imposed load to the ALTO
   server, or from whom the queries are performed.

   Comments and discussions about this memo should be directed to the
   ALTO working group: alto@ietf.org.

2. Overview

   The ALTO protocol is a client/server protocol, operating between a
   number of ALTO clients and an ALTO server, as sketched in Figure 1.
   The ALTO working groups defines the ALTO protocol
   [I-D.ietf-alto-protocol].

```
                      +----------+
                      |  ALTO    |
                      |  Server  |
                      +----------+
                           ^
                _.-----|------.
              ,-''      |       `--.
            ,'          |           `.
          (     Network |             )
           `.           |           ,'
            `--.        |        _.-'
              `------|-----''
                          v
      +----------+  +----------+   +----------+
      |  ALTO    |  |  ALTO    |...|  ALTO    |
      |  Client  |  |  Client  |   |  Client  |
      +----------+  +----------+   +----------+
```
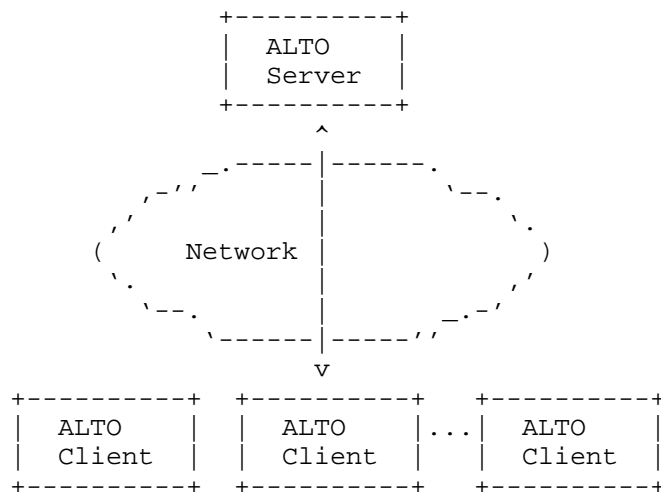
                 Figure 1: Network Overview of ALTO Protocol

2.1. General Placement of ALTO

   The ALTO server and ALTO clients can be situated at various entities
   in a network deployment. The first differentiation is whether the
   ALTO client is located on the actual host that runs the application,
   as shown in Figure 2, (e.g., peer-to-peer filesharing application) or
   if the ALTO client is located on resource directory, as shown in
   Figure 3 (e.g., a tracker in peer-to-peer filesharing).

```
                                              +-----+
                                       =====|     |**
                                   ====      +-----+   *
                                ====                *       *
                             ====                   *       *
      +-----+     +------+=====                +-----+   *
      |     |.....|      |====================|     |   *
      +-----+     +------+=====                +-----+   *
      Source of      ALTO         ====                *       *
      topological    service         ====            *       *
      information                     ====    +-----+   *
                                   =====|     |**
                                        +-----+
```

Legend:
=== ALTO client protocol
*** Application protocol
... Provisioning protocol
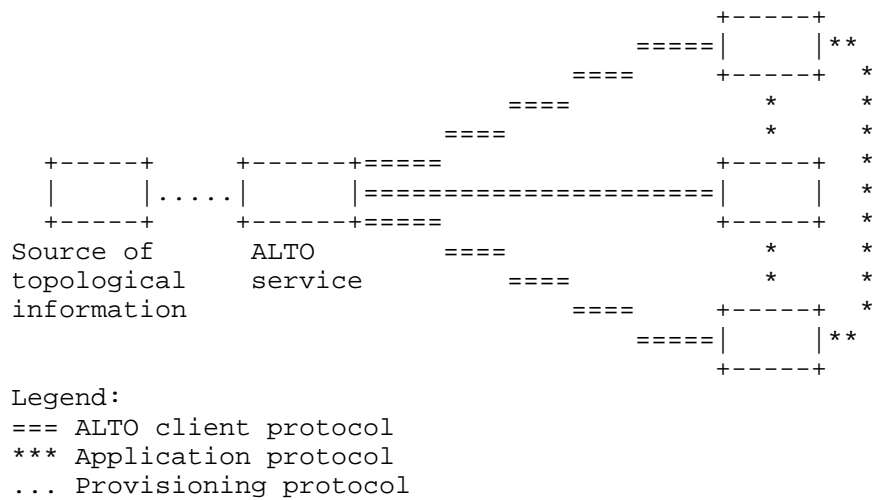
Figure 2: Overview of protocol interaction between ALTO
            elements,scenario without tracker

Figure 2 shows the operational model for applications that do not use
a tracker, such as, edonky, or in if the tracker should be the
querying party. This use case also holds true for CDNs. The ALTO
server can also be queried by CDNs to get a guidance about where the
a particular client accessing data in the CDN is exactly located in
the ISP's network.

```
                                          +-----+
                                    **|     |**
                                  **  +-----+   *
                                **         *     *
                              **           *     *
      +-----+    +------+    +-----+**        +-----+   *
      |     |.....|      |    |     |**********|     |   *
      +-----+    +------+    +-----+**        +-----+   *
      Source of   ALTO        Resource **        *     *
      topological service     directory **        *     *
      information             ("tracker")  **  +-----+   *
                                          **|     |**
                                          +-----+
                                            Peers
      Legend:
      === ALTO client protocol
      *** Application protocol
      ... Provisioning protocol
```

         Figure 3: Overview of protocol interaction between ALTO elements,
                              scenario with tracker

   However, Figure 3 does not denote where the ALTO elements are
   actually located, i.e., if the tracker and the ALTO server are in the
   same ISP's domain, or if the tracker and the ALTO server are managed/
   owned/located in different domains. The latter is the typical use
   case, e.g., taking Pirate Bay as example that serves Bittorrent peers
   world-wide.

2.2. Provided Guidance

   ALTO gives guidance to applications on what IP addresses or IP
   prefixes, and such which hosts are to be preferred according to the
   operator of the ALTO server. The general assumption of the ALTO WG
   is that a network operator would always express to prefer hosts in
   its own network while hosts located outside its own network are to be
   avoided (are undesired to be considered by the applications). This
   might be applicable in some cases but may not be applicable in the
   general case. The ALTO protocol gives only the means to let the ALTO
   server operator to express is preference, whatever this preference
   is. This section explores this space.

2.2.1. Keeping Traffic Local in Network

   ALTO guidance can be used to let applications prefer other peers
   within the same network operator's network instead of randomly
   connecting to other peers which are located in another operator's
   network. Figure 4 shows such a scenario where peers prefer peers in

the same network (e.g., Peer 1 and Peer 2 in ISP1 and Peer 3 and Peer
4 in ISP2).

```
                             ,-------.          +-----------+
              ,---.        ,-'         `-.      |  Peer 1   |
           ,-'     `-.    /    ISP 1   #######|ALTO Client|
          /           \  /             #  \    +-----------+
         /    ISP X    \  |            #  |    +-----------+
        /               \  \          #######|   Peer 2  |
       ;            +-------------------------|ALTO Client|
       |            |   |   `-.        ,-'    +-----------+
       |            |   |     `-------'
       |            |   |   ,-------.          +-----------+
       :            |   ;  ,-'       `#######|   Peer 3  |
       \            |  /  /    ISP 2  #  \    |ALTO Client|
        \           | /  /           #  \    +-----------+
         \          +--------+       #  |    +-----------+
          `-.    ,-'     \   |       #######|   Peer 4  |
            `---'         \  +-----------------|ALTO Client|
             `-.         ,-'  \   +-----------+
                `-------'
```
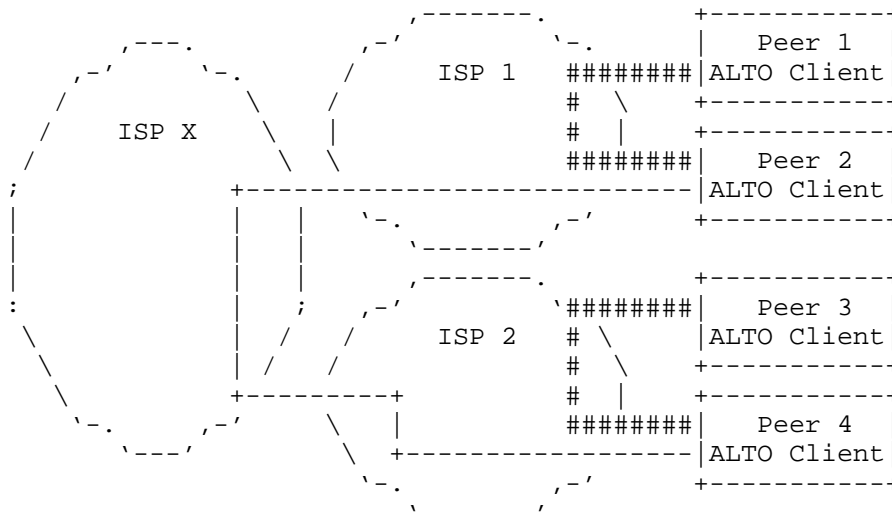
Legend:
preferred "connections"
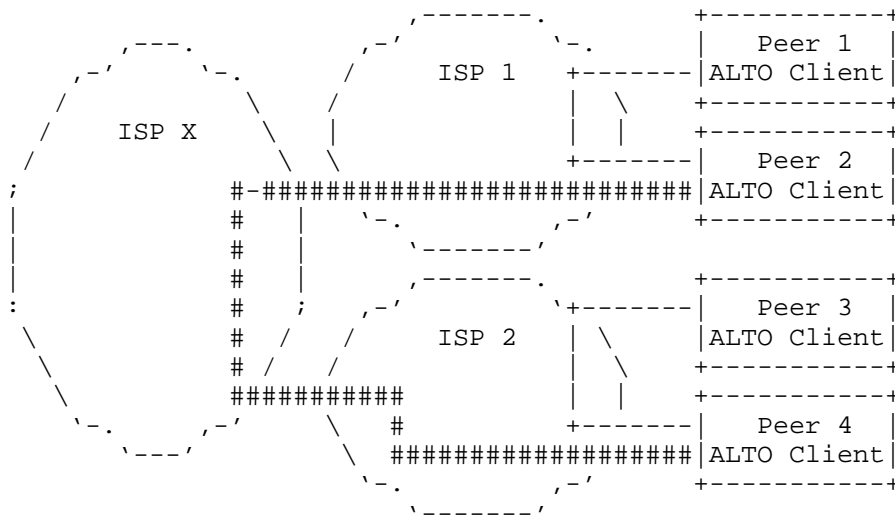--- non-preferred "connections"

              Figure 4: ALTO Traffic Network Localization

   TBD: Describes limits of this approach (e.g., traffic localization
   guidance is of less use if the peers cannot upload); describe how
   maps would look like.

2.2.2. Off-Loading Traffic from Network

   Another scenario where the use of ALTO can be beneficial is in mobile
   broadband networks, e.g., CDMA200 or UMTS, but where the network
   operator may have the desire to guide peers in its own network to use
   peers in remote networks. One reason can be that the wireless
   network is not made for the load cause by, e.g., peer-to-peer
   applications, and the operator has the need that peers fetch their
   data from remote peers in other parts of the Internet.

```
                          ,-------.         +-----------+
              ,---.     ,-'         `-.     |  Peer 1   |
            ,-'     `-. /    ISP 1    +-------|ALTO Client|
           /           \ /            |  \    +-----------+
          /    ISP X    \ |           |   |   +-----------+
         /               \ \          +-------|  Peer 2   |
        ;             #-######################|ALTO Client|
        |             #   |  `-.       ,-'    +-----------+
        |             #   |    `-------'
        |             #   |   ,-------.       +-----------+
        :             #   ;  ,-'       `+-------|  Peer 3   |
        \             #  / /    ISP 2   |  \    |ALTO Client|
         \            # / /            |   \   +-----------+
          \           ##########       |   |   +-----------+
           `-.     ,-'  \    #         +-------|  Peer 4   |
             `---'       \  #################|ALTO Client|
                          \ `-.       ,-'    +-----------+
                           `-.        ,-'
                             `-------'

            Legend:
            === preferred "connections"
            --- non-preferred "connections"

                Figure 5: ALTO Traffic Network De-Localization
```

   Figure 5 shows the result of such a guidance process where Peer 2
   prefers a connection with Peer4 instead of Peer 1, as shown in
   Figure 4.

   TBD: Limits of this approach in general and with respect to p2p.
   describe how maps would look like.

2.2.3. Intra-Network Localization/Bottleneck Off-Loading

   The above sections described the results of the ALTO guidance on an
   inter-network level. However, ALTO can also be used to guide peers
   on which internal peers are to be preferred. For instance, to guide
   Peers on a remote network side to prefer to connect to each other,
   instead of crossing a bottleneck link, a backhaul link to connect the
   side to the network core. Figure 6 shows such a scenario where Peer
   1 and Peer 2 are located in Net 2 of ISP1 and connect via a low
   capacity link to the core (Net 1) of the same ISP1. Peer1 and Peer 2
   would both exchange their data with remote peers, probably clogging
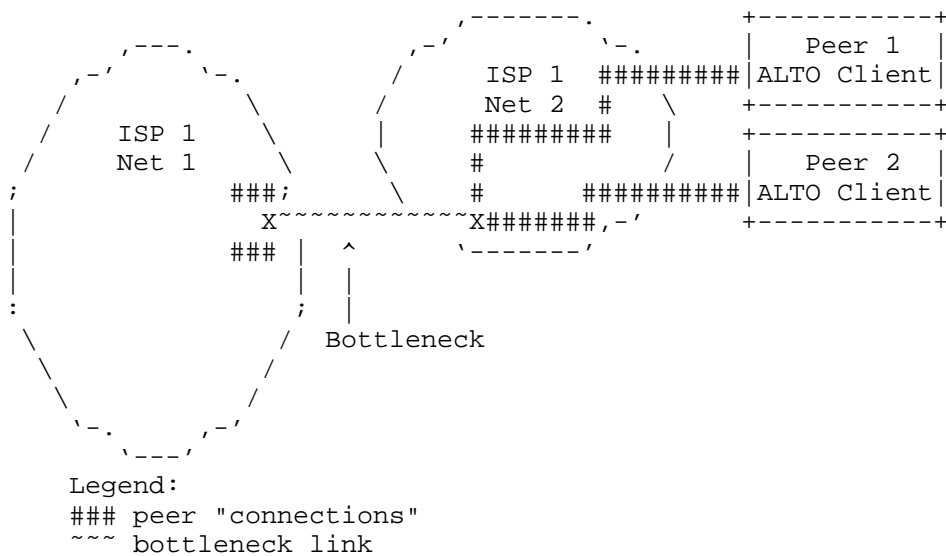   the bottleneck link.

```
                              ,-------.          +-----------+
          ,---.          ,-'          `-.       |  Peer 1   |
       ,-'      `-.      /    ISP 1  ########|ALTO Client|
      /            \    /    Net 2  #   \     +-----------+
     /    ISP 1     \   |     ########    |     +-----------+
    /     Net 1      \   \       #        /     |  Peer 2   |
   ;           ###;   \      #      #########|ALTO Client|
   |            X~~~~~~~~~~X#######,-'       +-----------+
   |           ### |  ^        `-------'
   |               |  |
   :               ;  |
    \             /  Bottleneck
     \           /
      \         /
       `-.    ,-'
          `---'
        Legend:
        ### peer "connections"
        ~~~ bottleneck link
```

          Figure 6: Without Intra-Network ALTO Traffic Localization

The operator can guide the peers in such a situation to try first
local peers in the same network islands, avoiding or at least
lowering the effect on the bottleneck link, as shown in Figure 7.

```
                              ,-------.          +-----------+
          ,---.          ,-'          `-.       |  Peer 1   |
       ,-'      `-.      /    ISP 1  ########|ALTO Client|
      /            \    /    Net 2  #   \     +-----------+
     /    ISP 1     \   |           #    |     +-----------+
    /     Net 1      \   \       #########|  Peer 2   |
   ;            ;    \      #########|ALTO Client|
   |          #~~~~~~~~~~#######,-'       +-----------+
   |           ### |  ^        `-------'
   |               |  |
   :               ;  |
    \             /  Bottleneck
     \           /
      \         /
       `-.    ,-'
          `---'
        Legend:
        ### peer "connections"
        ~~~ bottleneck link
```
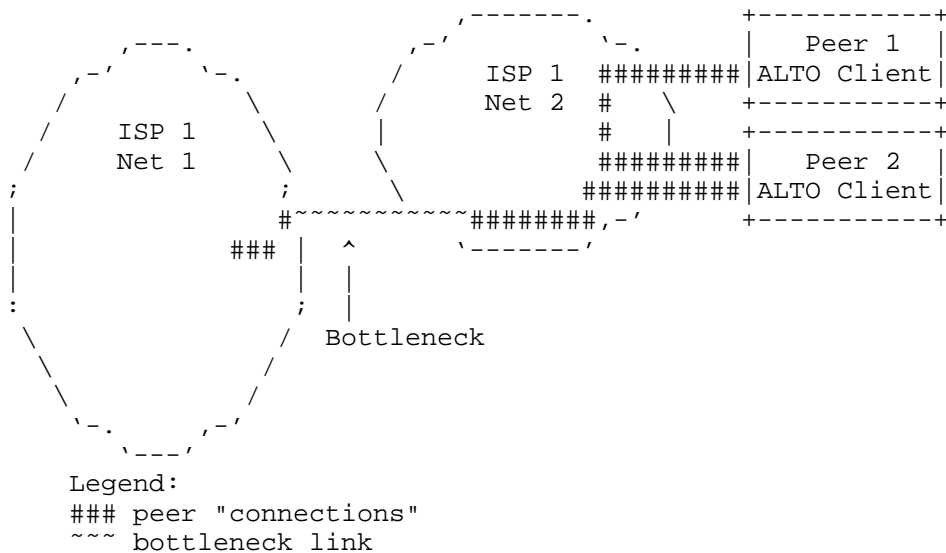
            Figure 7: With Intra-Network ALTO Traffic Localization

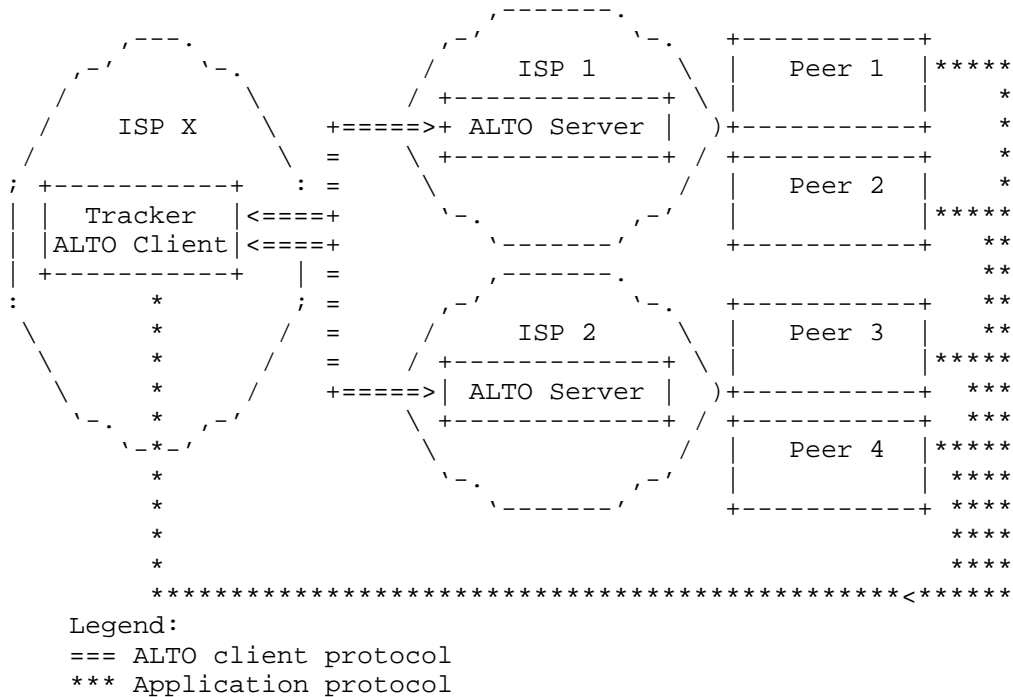TBD: describe how maps would look like.

3. Using ALTO for Peer-to-Peer

```
                              ,-------.
          ,---.            ,-'         `-. +-----------+
        ,-'     `-.       /    ISP 1      \ |  Peer 1   |*****
       /           \     / +-------------+ \|           |   *
      /   ISP X     \   +=====>+ ALTO Server |  )+-----------+   *
     /               \  =   \ +-------------+ / +-----------+   *
    ; +-----------+   :  =    \               / |  Peer 2   |   *
    | | Tracker   |<====+      `-.         ,-' |           |*****
    | |ALTO Client|<====+         `-------'    +-----------+  **
    | +-----------+   | =        ,-------.                    **
    :       *         ; =      ,-'        `-. +-----------+   **
     \      *        /  =     /    ISP 2     \ |  Peer 3   |   **
      \     *       /   =    / +-------------+ \|           |*****
       \    *      /    +=====>| ALTO Server |  )+-----------+  ***
        `-. *   ,-'       \ +-------------+ / +-----------+  ***
         `-*-'             \               / |  Peer 4   |*****
           *                `-.         ,-' |           | ****
           *                   `-------'    +-----------+ ****
           *                                              ****
           *                                              ****
           **************************************************<******
```

        Legend:
        === ALTO client protocol
        *** Application protocol

       Figure 8: Global tracker accessing ALTO server at various ISPs

   Figure 8 depicts a tracker-based system, where the tracker embeds the
   ALTO client. The tracker itself is hosted and operated by an entity
   different than the ISP hosting and operating the ALTO server.
   Initially, the tracker has to look-up the ALTO server in charge for
   each peer where it receives a ALTO query for. Therefore, the ALTO
   server has to discover the handling ALTO server, as described in
   [I-D.kiesel-alto-3pdisc]. However, the peers do not have any way to
   query the server themselves. This setting allows to give the peers a
   better selection of candidate peers for their operation at an initial
   time, but does not consider peers learned through direct peer-to-peer
   knowledge exchange, AKA peer exchange in various peer-to-peer
   protocols.

```
              ,-------.         +-----------+
        ,---.          ,-'         `-. +==>|  Peer 1   |*****
     ,-'     `-.      /    ISP 1      \  = |ALTO Client|    *
    /           \    / +-------------+<=+  +-----------+    *
   /    ISP X    \   | + ALTO Server |<=+  +-----------+    *
  /               \  \ +-------------+ /=  |  Peer 2   |    *
 ;  +---------+    :  \                /  +==>|ALTO Client|*****
 |  | Global  |    |   `-.         ,-'    +-----------+   **
 |  | Tracker |    |      `-------'                       **
 |  +---------+    |     ,-------.         +-----------+   **
 :       *        ;   ,-'         `-. +==>|  Peer 3   |   **
  \      *       /   /    ISP 2      \  = |ALTO Client|*****
   \     *      /   / +-------------+<=+  +-----------+   ***
    \    *     /    | | ALTO Server |<=+  +-----------+   ***
    `-.  *  ,-'     \ +-------------+ /=  |  Peer 4   |*****
      `-*-'          \               / +==>|ALTO Client| ****
        *             `-.         ,-'    +-----------+ ****
        *                `-------'                     ****
        *                                              ****
        ***************************************************<****
```
     Legend:
     === ALTO client protocol
     *** Application protocol
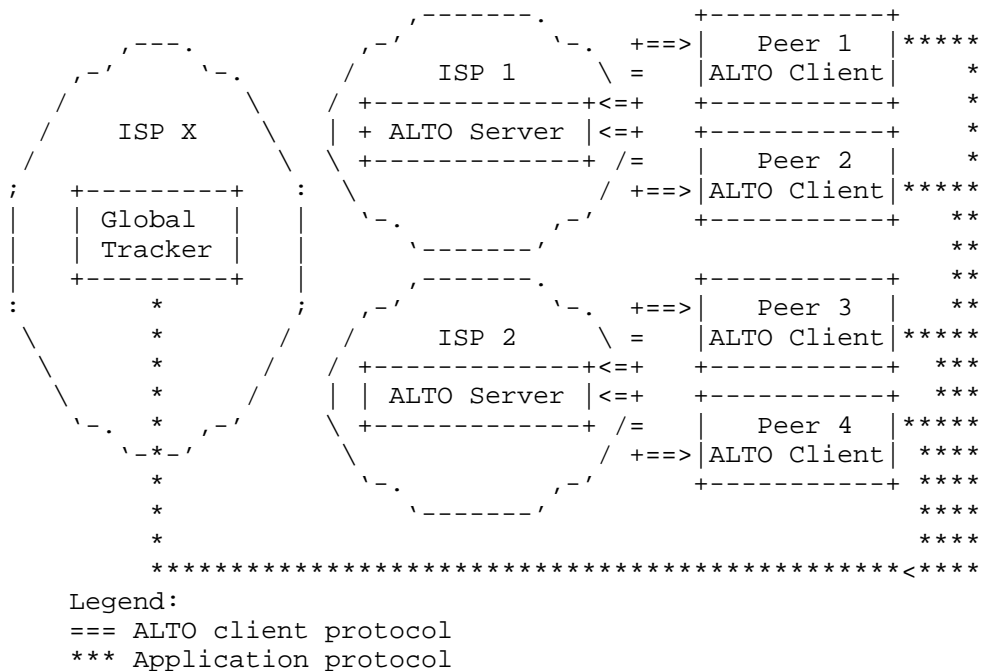
                  Figure 9: Global Tracker - Local ALTO Servers

     The scenario in Figure 9 lets the peers directly communicate with
     their ISP's ALTO server (i.e., ALTO client embedded in the peers),
     giving thus the peers the most control on which information they
     query for, as they can integrate information received from trackers
     and through direct peer-to-peer knowledge exchange.

```
                          ,-------.          +-----------+
           ,---.       ,-' ISP 1  `-. ***>|  Peer 1   |
         ,-'     `-.    /+-------------+\ *   |           |
        /           \  / +   Tracker  |<**   +-----------+
       /   ISP X     \ | +-----==-----+<**   +-----------+
      /               \ \ +-----==-----+ /*   |  Peer 2   |
     ;   +---------+   : \+ ALTO Server |/ ***>|           |
     |   | Global  |   |  +-------------+      +-----------+
     |   | Tracker |   |     `-------'
     |   +---------+   |                       +-----------+
     :        ^        ;    ,-------.          |  Peer 3   |
      \       *       /   ,-' ISP 2  `-. ***>|           |
       \      *      /    /+-------------+\ *   +-----------+
        \     *     /    / +   Tracker  |<**   +-----------+
         `-.  *  ,-'     | +-----==-----+ |   |  Peer 4   |<*
           `---*         \ +-----==-----+ /   |           | *
               *          \+ ALTO Server |/    +-----------+ *
               *           +-------------+                   *
               *              `-------'                      *
               **********************************************
```

         Legend:
         === ALTO client protocol
         *** Application protocol
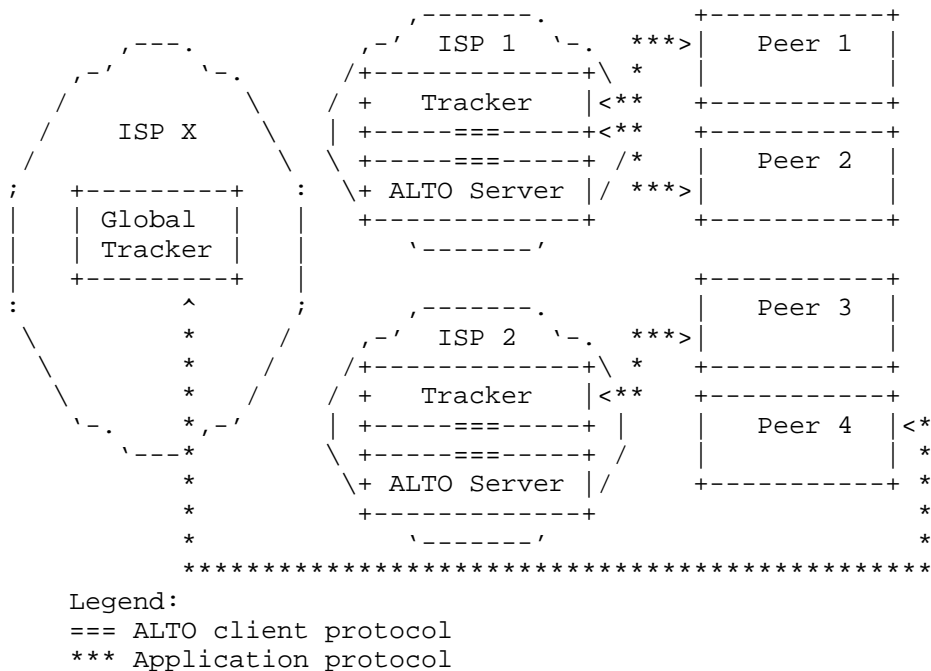
   Figure 10: P4P approach with local tracker and local ALTO server

   There are some attempts to let ISP's to deploy their own trackers, as
   shown in Figure 10. In this case, the client has no chance to get
   guidance from the ALTO server, other than talking to the ISP's
   tracker. However, the peers would have still chance the contact
   other trackers, deployed by entities other than the peer's ISP.

   Figure 10 and Figure 8 ostensibly take peers the possibility to
   directly query the ALTO server, if the communication with the ALTO
   server is not permitted for any reason. However, considering the
   plethora of different applications of ALTO, e.g., multiple tracker
   and non-tracker based P2P systems and or applications searching for
   relays, it seems to be beneficial for all participants to let the
   peers directly query the ALTO server. The peers are also the single
   point having all operational knowledge to decide whether to use the
   ALTO guidance and how to use the ALTO guidance. This is a preference
   for the scenario depicted in Figure Figure 9.

3.1. Using ALTO for Tracker-based Peer-to-Peer Applications

```
.............................         ...............................
: Tracker                    :         : Peer                        :
:   _____                   :         :                             :
: +-_____-+                 :         :             k good          :
: |        |     +--------+  : P2P App. : +--------+ peers +------+  :
: |   N    |     | random |  : Protocol : | ALTO-  |------>| data |  :
: | known  |====>| pre-   |***************>| biased |       | ex-  |  :
: | peers, |     | selec- |  : transmit : | peer   |------>| cha- |  :
: | M good |     | tion   |  : n peer  : | select | n-k   | nge  |  :
: +-_____-+     +--------+  : IDs     : +--------+ bad p.+------+  :
:............................:         :.....^.......................:


        |                                  | ALTO
                                           | client protocol
                                          __|___
                                        +-_____-+
                                        |        |
                                        | ALTO   |
                                        | server |
                                        +-_____-+
```
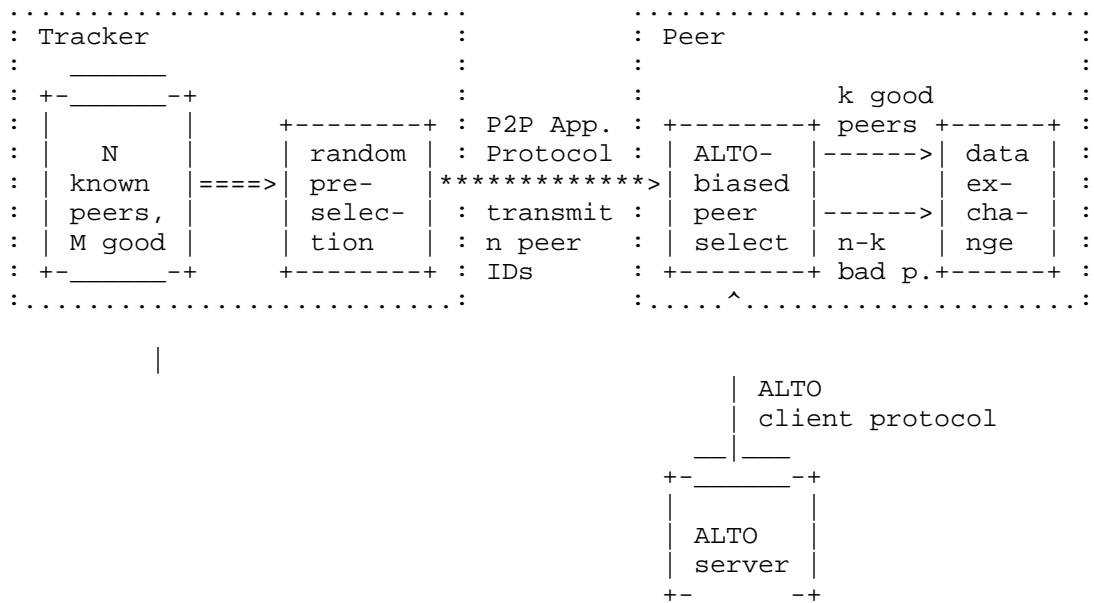
       Figure 11: Tracker-based P2P Application with random peer
                              preselection


```
.............................         ...............................
: Tracker                    :         : Peer                        :
:   _____                   :         :                             :
: +-_____-+                 :         :                             :
: |        |     +--------+  : P2P App. :  k good peers &  +------+  :
: |   N    |     | ALTO-  |  : Protocol :  n-k bad peers   | data |  :
: | known  |====>| biased |*****************************>| ex-  |  :
: | peers, |     | peer   |  : transmit :                 | cha- |  :
: | M good |     | select |  : n peer  :                 | nge  |  :
: +-_____-+     +--------+  : IDs     :                 +------+  :
:......................^.....:         :.............................:

                       |
                       | ALTO
                       | client protocol
                      __|___
                    +-_____-+
                    |        |
                    | ALTO   |
                    | server |
                    +-_____-+
```
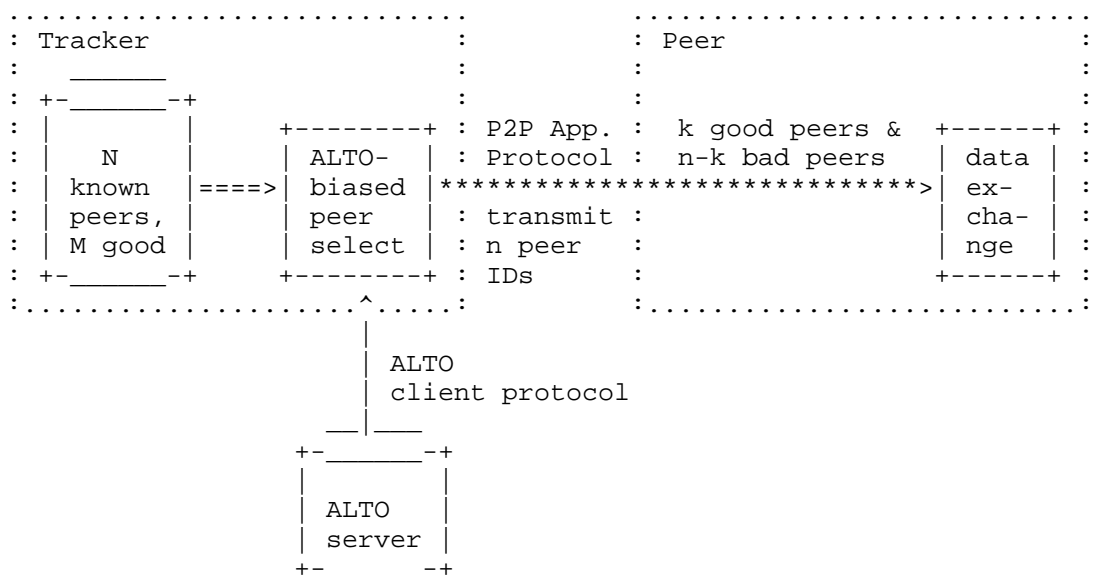
  Figure 12: Tracker-based P2P Application with ALTO client in tracker

   TBD: explain why Figure 12 usually will yield better results wrt.
   peer selection than Figure 11.

3.2. Expectations of ALTO

   This section hints to some recent experiments conducted with ALTO-
   like deployments in Internet Service Provider (ISP) network's. NTT
   performed tests with their HINT server implementation and dummy nodes
   to gain insight on how an ALTO-like service influence a peer-to-peer
   systems [I-D.kamei-p2p-experiments-japan]. The results of an early
   experiment conducted in the Comcast network are documented
   here[RFC5632]

4. Using ALTO for CDNs

   Section 3 discussed the placement and usage of ALTO for P2P systems,
   but not beyond. This section discuss the usage of ALTO for Content
   Delivery Networks (CDNs). CDNs are used to bring a service (e.g., a
   web page, videos, etc) closer to the location of the user - where
   close refers to shorten the distance between the client and the
   server in the IP topology. CDNs use several techniques to decide
   which server is closest to a client requesting a service. One common
   way to do so, is relying on the DNS system, but there are many other
   ways, see [RFC3568].

   The general issue for CDNs, independent of DNS or HTTP Redirect based
   approaches (see, for instance, [I-D.penno-alto-cdn]), is that the CDN
   logic has to match the client's IP address with the closest CDN
   cache. This matching is not trivial, for instance, in DNS based
   approaches, where the IP address of the DNS original requester is
   unknown (see [I-D.vandergaast-edns-client-ip] for a discussion of
   this and a solution approach).

5. Cascading ALTO Servers

   The main assumptions of ALTO seems to be each ISP operates its own
   ALTO server independently, irrespectively of the ISP's situation.
   This may true for most envisioned deployments of ALTO but there are
   certain deployments that may have different settings. Figure 13
   shows such setting, were for example, a university network is
   connected to two upstream providers. ISP2 if the national research
   network and ISP1 is a commercial upstream provider to this university
   network. The university, as well as ISP1, are operating their own
   ALTO server. The ALTO clients, located on the peers will contact the
   ALTO server located at the university.

```
               +-----------+
               |   ISP1    |
               |   ALTO    |
               |  Server   |
               +---------=+
              ,-------=            ,------.
           ,-'        =`-.      ,-'          `-.
          /   Upstream=   \    /   Upstream     \
         (       ISP1 =    )  (       ISP2       )
          \          =    /    \               /
           `-.       =,-'       `-.          ,-'
             `---+---=              `+------'
                 |    =              |
                 |    =====================
                 |,------------.  |       =
               ,-+             `-+  +-----------+
             ,'   University     `. |University |
            (     Network         ) |   ALTO    |
             `. =====================|  Server   |
               `-=             +-'   +-----------+
              =`+-----------'|       |
              = |            |
       +--------+-+      +-+--------+
       |  Peer1   |      |  PeerN   |
       +----------+      +----------+
```
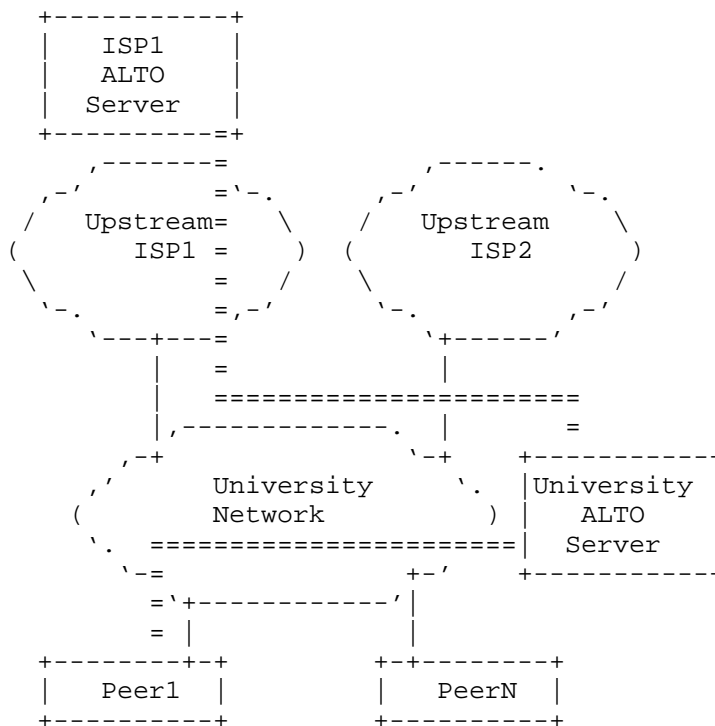
                 Figure 13: Cascaded ALTO Server

   In this setting all "destinations" useful for the peers within ISP2
   are free-of-charge for the peers located in the university network
   (i.e., they are preferred in the rating of the ALTO server).
   However, all traffic that is not towards ISP2 will be handled by the
   ISP1 upstream provider. Therefore, the ALTO server at the university
   has also to include the guidance given by the ISP1 ALTO server in its

replies to the ALTO clients. This can be called cascaded ALTO
servers.

6. Known Limitations of ALTO

   This section describes some known limitations of ALTO in general or
   specific mechanisms in ALTO.

6.1. Limitations of Map-based Approaches

   The specification of the ALTO protocol [I-D.ietf-alto-protocol] uses,
   amongst others mechanism, so-called network maps. The network map
   approach uses Host Group Descriptors that group one or multiple
   subnetworks (i.e., IP prefixes) to a single Host Group Descriptor. A
   set of IP prefixes is called partition and the associated Host Group
   Descriptor is called partition ID. The "costs" between the various
   partition IDs is stored in a second map, the cost map. Map-based
   approaches are chosen as they lower the signaling load on the server,
   as the maps have only to be retrieved if they are changed.

   The main assumption for map-based approaches is that the information
   provided in these maps is static for a longer period of time, where
   this period of time refers to days, but not hours or even minutes.
   This assumption is fine, as long as the network operator does not
   change any parameter, e.g., routing within the network and to the
   upstream peers, IP address assignment stays stable (and thus the
   mapping to the partitions). However, there are several cases where
   this assumption is not valid, as:

   1. ISPs reallocate IPv4 subnets from time to time;

   2. ISPs reallocate IPv4 subnets on short notice;

   3. IP prefix blocks may be assigned to a single DSLAM which serves a
       variety of access networks.

   For 1): ISPs reallocate IPv4 subnets within their infrastructure from
   time to time, partly to ensure the efficient usage of IPv4 addresses
   (a scarce resource), and partly to enable efficient route tables
   within their network routers. The frequency of these "renumbering
   events" depend on the growth in number of subscribers and the
   availability of address space within the ISP. As a result, a
   subscriber's household device could retain an IPv4 address for as
   short as a few minutes, or for months at a time or even longer.

      Some folks have suggested that ISPs providing ALTO services could
      sub-divide their subscribers' devices into different IPv4 subnets
      (or certain IPv4 address ranges) based on the purchased service
      tier, as well as based on the location in the network topology.
      The problem is that this sub-allocation of IPv4 subnets tends to
      decrease the efficiency of IPv4 address allocation. A growing ISP

that needs to maintain high efficiency of IPv4 address utilization
may be reluctant to jeopardize their future acquisition of IPv4
address space.

However, this is not an issue for map-based approaches if changes are
applied in the order of days.

For 2): ISPs can use techniques, such as ODAP (XXX) that allow the
reallocation of IP prefixes on very short notice, i.e., within
minutes. An IP prefix that has no IP address assignment to a host
anymore can be reallocate to areas where there is currently a high
demand for IP addresses.

For 3): In DSL-based access networks, IP prefixes are assigned to
DSLAMs which are the first IP-hop in the access-network between the
CPE and the Internet. The access-network between CPE and DSLAM
(called aggregation network) can have varying characteristics (and
thus associated costs), but still using the same IP prefix. For
instance one IP addresses IP11 out of a IP prefix IP1 can be assigned
to a VDSL (e.g., 2 MBit/s uplink) access-line while the subsequent IP
address IP12 is assigned to a slow ADSL line (e.g., 128 kbit/s
uplink). These IP addresses are assigned on a first come first
served basis, i.e., the a single IP address out of the same IP prefix
can change its associated costs quite fast. This may not be an issue
with respect to the used upstream provider (thus the cross ISP
traffic) but depending on the capacity of the aggregation-network
this may raise to an issue.

6.2. Limitiations of Non-Map-based Approaches

The specification of the ALTO protocol [I-D.ietf-alto-protocol] uses,
amongst others mechanism, a mechanism called Endpoint Cost Service.
ALTO clients can ask guidance for specific IP addresses to the ALTO
server. However, asking for IP addresses, asking with long lists of
IP addresses, and asking quite frequent may overload the ALTO server.
The server has to rank each received IP address which causes load at
the server. This may be amplified by the fact that not only a single
ALTO client is asking for guidance, but a larger number of them.

Caching of IP addresses at the ALTO client or the usage of the H12
approach [I-D.kiesel-alto-h12] in conjunction with caching may lower
the query load on the ALTO server.

6.3. General Challenges

An ALTO server stores information about preferences (e.g., a list of
preferred autonomous systems, IP ranges, etc) and ALTO clients can
retrieve these preferences. However, there are basically two

different approaches on where the preferences are actually processed:

1. The ALTO server has a list of preferences and clients can
   retrieve this list via the ALTO protocol. This preference list
   can be partially updated by the server. The actual processing of
   the data is done on the client and thus there is no data of the
   client's operation revealed to the ALTO server .

2. The ALTO server has a list of preferences or preferences
   calculated during runtime and the ALTO client is sending
   information of its operation (e.g., a list of IP addresses) to
   the server. The server is using this operational information to
   determine its preferences and returns these preferences (e.g., a
   sorted list of the IP addresses) back to the ALTO client.

Approach 1 (we call it H1) has the advantage (seen from the client)
that all operational information stays within the client and is not
revealed to the provider of the server. On the other hand, does
approach 1 require that the provider of the ALTO server, i.e., the
network operator, reveals information about its network structure
(e.g., AS numbers, IP ranges, topology information in general) to the
ALTO client.

Approach 2 (we call it H2) has the advantage (seen from the operator)
that all operational information stays with the ALTO server and is
not revealed to the ALTO client. On the other hand, does approach 2
require that the clients send their operational information to the
server.

Both approaches have their pros and cons and are extensively
discussed on the ALTO mailing list. But there is basically a
dilemma: Approach 1 is seen as the only working solution by peer-to-
peer software vendors and approach 2 is seen as the only working by
the network operators. But neither the software vendors nor the
operators seem to willing to change their position. However, there
is the need to get both sides on board, to come to a solution.

7. API between ALTO Client and Application

    This sections gives some informational guidance on how the interface
    between the actual application using the ALTO guidance and the ALTO
    client can look like.

    This is still TBD.

8. Extensions to the ALTO Protocol

8.1. Host Group Descriptors

   Host group descriptors are used in the ALTO client protocol to
   describe the location of a host in the network topology. The ALTO
   client protocol specification defines a basic set of host group
   descriptor types, which have to be supported by all implementations,
   and an extension procedure for adding new descriptor types . The
   following list gives an overview on further host group descriptor
   types that have been proposed in the past, or which are in use by
   ALTO-related prototype implementations. This list is not intended as
   normative text. Instead, the only purpose of the following list is
   to document the descriptor types that have been proposed so far, and
   to solicit further feedback and discussion:

   o Autonomous System (AS) number

   o Protocol-specific group identifiers, which expand to a set of IP
      address ranges (CIDR) and/or AS numbers. In one specific solution
      proposal, these are called Partition ID (PID).

8.2. Rating Criteria

   Rating criteria are used in the ALTO client protocol to express
   topology- or connectivity-related properties, which are evaluated in
   order to generate the ALTO guidance. The ALTO client protocol
   specification defines a basic set of rating criteria, which have to
   be supported by all implementations, and an extension procedure for
   adding new criteria . The following list gives an overview on
   further rating criteria that have been proposed in the past, or which
   are in use by ALTO-related prototype implementations. This list is
   not intended as normative text. Instead, the only purpose of the
   following list is to document the rating criteria that have been
   proposed so far, and to solicit further feedback and discussion:

8.2.1. Distance-related Rating Criteria

   o Relative topological distance: relative means that a larger
      numerical value means greater distance, but it is up to the ALTO
      service how to compute the values, and the ALTO client will not be
      informed about the nature of the information. One way of
      generating this kind of information MAY be counting AS hops, but
      when querying this parameter, the ALTO client MUST NOT assume that
      the numbers actually are AS hops.

   o Absolute topological distance, expressed in the number of
      traversed autonomous systems (AS).

   o Absolute topological distance, expressed in the number of router
      hops (i.e., how much the TTL value of an IP packet will be
      decreased during transit).

   o Absolute physical distance, based on knowledge of the approximate
      geolocation (continent, country) of an IP address.

8.2.2. Charging-related Rating Criteria

   o Traffic volume caps, in case the Internet access of the resource
      consumer is not charged by "flat rate". For each candidate
      resource provider, the ALTO service could indicate the amount of
      data that may be transferred from/to this resource provider until
      a given point in time, and how much of this amount has already
      been consumed. Furthermore, it would have to be indicated how
      excess traffic would be handled (e.g., blocked, throttled, or
      charged separately at an indicated price). The interaction of
      several applications running on a host, out of which some use this
      criterion while others don't, as well as the evaluation of this
      criterion in resource directories, which issue ALTO queries on
      behalf of other peers, are for further study.

8.2.3. Performance-related Rating Criteria

   The following rating criteria are subject to the remarks below.

   o The minimum achievable throughput between the resource consumer
      and the candidate resource provider, which is considered useful by
      the application (only in ALTO queries), or

   o An arbitrary upper bound for the throughput from/to the candidate
      resource provider (only in ALTO responses). This may be, but is
      not necessarily the provisioned access bandwidth of the candidate
      resource provider.

   o The maximum round-trip time (RTT) between resource consumer and
      the candidate resource provider, which is acceptable for the
      application for useful communication with the candidate resource
      provider (only in ALTO queries), or

   o An arbitrary lower bound for the RTT between resource consumer and
      the candidate resource provider (only in ALTO responses). This
      may be, for example, based on measurements of the propagation
      delay in a completely unloaded network.

   The ALTO client MUST be aware, that with high probability, the actual
   performance values differ significantly from these upper and lower
   bounds. In particular, an ALTO client MUST NOT consider the "upper

bound for throughput" parameter as a permission to send data at the
indicated rate without using congestion control mechanisms.

The discrepancies are due to various reasons, including, but not
limited to the facts that

o the ALTO service is not an admission control system

o the ALTO service may not know the instantaneous congestion status
   of the network

o the ALTO service may not know all link bandwidths, i.e., where the
   bottleneck really is, and there may be shared bottlenecks

o the ALTO service may not know whether the candidate peer itself is
   overloaded

o the ALTO service may not know whether the candidate peer throttles
   the bandwidth it devotes for the considered application

o the ALTO service may not know whether the candidate peer will
   throttle the data it sends to us (e.g., because of some fairness
   algorithm, such as tit-for-tat)

Because of these inaccuracies and the lack of complete, instantaneous
state information, which are inherent to the ALTO service, the
application must use other mechanisms (such as passive measurements
on actual data transmissions) to assess the currently achievable
throughput, and it MUST use appropriate congestion control mechanisms
in order to avoid a congestion collapse. Nevertheless, these rating
criteria may provide a useful shortcut for quickly excluding
candidate resource providers from such probing, if it is known in
advance that connectivity is in any case worse than what is
considered the minimum useful value by the respective application.

8.2.4. Inappropriate Rating Criteria

Rating criteria that SHOULD NOT be defined for and used by the ALTO
service include:

o Performance metrics that are closely related to the instantaneous
   congestion status. The definition of alternate approaches for
   congestion control is explicitly out of the scope of ALTO.
   Instead, other appropriate means, such as using TCP based
   transport, have to be used to avoid congestion.

9. Security Considerations

   The ALTO protocol itself, as well as, the ALTO client and server
   raise new security issues beyond the one mentioned in
   [I-D.ietf-alto-protocol] and issues related to message transport over
   the Internet. For instance, Denial of Service (DoS) is of interest
   for the ALTO server and also for the ALTO client. A server can get
   overloaded if too many TCP requests hit the server, or if the query
   load of the server surpasses the maximum computing capacity. An ALTO
   client can get overloaded if the responses from the sever are, either
   intentionally or due to an implementation mistake, too large to be
   handled by that particular client.

9.1. Information Leakage from the ALTO Server

   The ALTO server will be provisioned with information about the owning
   ISP's network and very likely also with information about neighboring
   ISPs. This information (e.g., network topology, business relations,
   etc) is consider to be confidential to the ISP and must not be
   revealed.

   The ALTO server will naturally reveal parts of that information in
   small doses to peers, as the guidance given will depend on the above
   mentioned information. This is seen beneficial for both parties,
   i.e., the ISP's and the peer's. However, there is the chance that
   one or multiple peers are querying an ALTO server with the goal to
   gather information about network topology or any other data
   considered confidential or at least sensitive. It is unclear whether
   this is a real technical security risk or whether this is more a
   perceived security risk.

9.2. ALTO Server Access

   Depending on the use case of ALTO, several access restrictions to an
   ALTO server may or may not apply. For an ALTO server that is solely
   accessible by peers from the ISP network (as shown in Figure 9), for
   instance, the source IP address can be used to grant only access from
   that ISP network to the server. This will "limit" the number of
   peers able to attack the server to the user's of the ISP (however,
   including botnet computers).

   On the other hand, if the ALTO server has to be accessible by parties
   not located in the ISP's network (see Figure Figure 8), e.g., by a
   third-party tracker or by a CDN system outside the ISP's network, the
   access restrictions have to be more loose. In the extreme case,
   i.e., no access restrictions, each and every host in the Internet can
   access the ALTO server. This might no the intention of the ISP, as
   the server is not only subject to more possible attacks, but also on

the load imposed to the server, i.e., possibly more ALTO clients to
serve and thus more work load.

9.3. Faking ALTO Guidance

   It has not yet been investigated how a faked or wrong ALTO guidance
   by an ALTO server can impact the operation of the network and also
   the peers.

   Here is a list of examples how the ALTO guidance could be faked and
   what possible consequences may arise:

   Sorting An attacker could change to sorting order of the ALTO
      guidance (given that the order is of importance, otherwise the
      ranking mechanism is of interest), i.e., declaring peers located
      outside the ISP as peers to be preferred. This will not pose a
      big risk to the network or peers, as it would mimic the "regular"
      peer operation without traffic localization, apart from the
      communication/processing overhead for ALTO. However, it could
      mean that ALTO is reaching the opposite goal of shuffling more
      data across ISP boundaries, incurring more costs for the ISP.

   Preference of a single peer A single IP address (thus a peer) could
      be marked as to be preferred all over other peers. This peer can
      be located within the local ISP or also in other parts of the
      Internet (e.g., a web server). This could lead to the case that
      quite a number of peers to trying to contact this IP address,
      possibly causing a Denial of Service (DoS) attack.

   This section is solely giving a first shot on security issues related
   to ALTO deployments.

10. Conclusion

   This is the first version of the deployment considerations and for
   sure the considerations are yet incomplete and imprecise.

11. References

11.1. Normative References

   [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate
              Requirement Levels", BCP 14, RFC 2119, March 1997.

   [RFC3568] Barbir, A., Cain, B., Nair, R., and O. Spatscheck, "Known
              Content Network (CN) Request-Routing Mechanisms",
              RFC 3568, July 2003.

11.2. Informative References

   [I-D.ietf-alto-protocol]
              Alimi, R., Penno, R., and Y. Yang, "ALTO Protocol",
              draft-ietf-alto-protocol-06 (work in progress),
              October 2010.

   [I-D.ietf-alto-reqs]
              Kiesel, S., Previdi, S., Stiemerling, M., Woundy, R., and
              Y. Yang, "Application-Layer Traffic Optimization (ALTO)
              Requirements", draft-ietf-alto-reqs-07 (work in progress),
              January 2011.

   [I-D.kamei-p2p-experiments-japan]
              Kamei, S., Momose, T., Inoue, T., and T. Nishitani, "ALTO-
              Like Activities and Experiments in P2P Network Experiment
              Council", draft-kamei-p2p-experiments-japan-04 (work in
              progress), November 2010.

   [I-D.kiesel-alto-3pdisc]
              Kiesel, S., Tomsu, M., Schwan, N., Scharf, M., and M.
              Stiemerling, "ALTO Server Discovery Protocol",
              draft-kiesel-alto-3pdisc-04 (work in progress),
              October 2010.

   [I-D.kiesel-alto-h12]
              Kiesel, S. and M. Stiemerling, "ALTO H12",
              draft-kiesel-alto-h12-02 (work in progress), March 2010.

   [I-D.penno-alto-cdn]
              Penno, R., Raghunath, S., Medved, J., Alimi, R., Yang, R.,
              and S. Previdi, "ALTO and Content Delivery Networks",
              draft-penno-alto-cdn-02 (work in progress), October 2010.

   [I-D.vandergaast-edns-client-ip]
              Contavalli, C., Gaast, W., Leach, S., and D. Rodden,
              "Client IP information in DNS requests",

draft-vandergaast-edns-client-ip-01 (work in progress),
May 2010.

[RFC5632] Griffiths, C., Livingood, J., Popkin, L., Woundy, R., and
Y. Yang, "Comcast's ISP Experiences in a Proactive Network
Provider Participation for P2P (P4P) Technical Trial",
RFC 5632, September 2009.

[RFC5693] Seedorf, J. and E. Burger, "Application-Layer Traffic
Optimization (ALTO) Problem Statement", RFC 5693,
October 2009.

Appendix A. Acknowledgments

   Martin Stiemerling is partially supported by the NAPA-WINE project
   (Network-Aware P2P-TV Application over Wise Networks,
   http://www.napa-wine.org), a research project supported by the
   European Commission under its 7th Framework Program (contract no.
   214412). The views and conclusions contained herein are those of the
   authors and should not be interpreted as necessarily representing the
   official policies or endorsements, either expressed or implied, of
   the NAPA-WINE project or the European Commission.

Authors' Addresses

   Martin Stiemerling
   NEC Laboratories Europe/University of Goettingen
   Kurfuerstenanlage 36
   Heidelberg 69115
   Germany

   Phone: +49 6221 4342 113
   Fax: +49 6221 4342 155
   Email: martin.stiemerling@neclab.eu
   URI: http://ietf.stiemerling.org

   Sebastian Kiesel
   University of Stuttgart, Computing Center
   Allmandring 30
   Stuttgart 70550
   Germany

   Email: ietf-alto@skiesel.de