ALTO S. Kiesel
Internet-Draft University of Stuttgart
Intended status: Standards Track M. Stiemerling
Expires: September 9, 2010 NEC Europe Ltd.
March 8, 2010

ALTO H12
draft-kiesel-alto-h12-02

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 applications, which have to select one or several hosts from a set of candidates, that are able to provide a desired resource. This memo proposes the Simple ALTO (H12) protocol.

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. 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. This memo proposes the Simple ALTO (H12) protocol. The H12 protocol is a client/server protocol between ALTO clients and ALTO servers, where ALTO clients can be either peer-to-peer applications residing on end hosts or peer-to-peer tracker servers.

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.kiesel-alto-reqs].

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

2. Protocol Framework

The ALTO protocol is a client/server protocol, operating between a number of ALTO clients and an ALTO server, as sketched in Figure 1

```
                 +----------+
                 |  ALTO    |
                 |  Server  |
                 +----------+
                       ^
                _.-----|------.
            ,-''       |       `--.
          ,'           |           `.
         (     Network |             )
          `.           |           ,'
            `--.       |       _.-'
                `------|-----''
                       v
    +----------+  +----------+   +----------+
    |  ALTO    |  |  ALTO    |...|  ALTO    |
    |  Client  |  |  Client  |   |  Client  |
    +----------+  +----------+   +----------+
```

Figure 1: Network Overview of ALTO Protocol

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 . This approach has been proposed by [I-D.shalunov-alto-infoexport].

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. This approach has been initially described in [ACM.ispp2p], but never been described on the protocol level.

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.

Therefore, this does memo proposes to integrate both approaches in one protocol and offer a way for clients and servers to learn each preferred way of operating.

3. H12 Operational Model

The P4P protocol proposal [I-D.penno-alto-protocol] assumes that the ALTO server maintains two different databases that store the server-side information necessary to guide ALTO clients. There is the network map that provides a mapping between network prefixes and a macro called partition ID (PID) and the cost map that relates PIDs to costs.

H12 assumes also that the H12 server internally maintains two maps, one for the network partitioning and the other for the associated costs, but does not need that the information stored in these maps is also conveyed to the H12 client in one piece. However, this memo is not specifying how the server is implemented, it is only specifying the ALTO protocol.

The client puts one or several host location attributes, about which it wants to receive a rating, in the query message.

The server replies with a list of network location attributes, in the same format as in the query, and the respective ratings for the requested attributes. However, the number of lines in this list may be shorter or longer than in the query, and the prefix lengths may be different:

- o The server may decide not to give any rating for a specific location attribute. In this case, a default value applies.

- o Instead of rating several location attributes with long prefix lengths (in particular: individual IP addresses) individually, the server may decide to give only one rating for a broader address range (i.e., prefix length is shorter).

- o Instead of giving one rating for a large address range, the server may decide to give several ratings for smaller ranges (i.e., i.e., each returned entry has a prefix length that is longer that requested).

The actual rating is given for each rating criterion as a signed integer value. A value of zero (0) means "default value". This value is to be used if the server has no information regarding this (network location attribute, rating criteria) tuple, or if it does not want to disclose it. Positive values mean that this location is "better" than default and therefore should be preferred for peer selection, while negative values indicate the location to be "worse" than default and therefore that it should be avoided. The meaning of "better" and "worse", as well as the scale has to be defined individually for each rating criterion.

This approach gives both sides, i.e., server and clients, to still exchange their desired information and level of precision, but also gives the chance to hide information if necessary and desired.

4. Proposed Protocol Semantics

H12 uses HTTP/1.1 and TCP as transport protocol between H12 clients and H12 servers. The encoding of the message body is done with XML. The usage of HTTP is similar to [I-D.penno-alto-protocol] also with the intention to reuse existing HTTP software and deployments.

H12 is aiming at keeping the level of involvement of the application that is using ALTO as low as possible. I.e., requiring an application, such as p2p file sharing, to use ALTO is already a considerable step. The implementers of the application must be able to use ALTO with a very low effort. It is assumed that the complexity of ALTO, in terms of implementation and operational effort, is mainly handled at the server.

Unlike the H1H2 protocol[I-D.stiemerling-alto-h1h2-protocol] the H12 protocol does not have several modes of operation, which have to be negotiated at the startup. Instead it allows the client and the server some flexibility in the requests and the responses while using only on mode of operation.

4.1. Locating the H12 Server Capabilities

H12 clients initially need to locate the right H12 server that is in charge of serving them. This step and the technical solution to locate such ALTO server is currently discussed within the ALTO working group. This memo does not yet define such H12 server discovery.

4.2. Learning the H12 Server Capabilities

This section describes how an ALTO client can learn about the capabilities of the ALTO server.

H12 clients initially need to locate the right H12 server that is in charge of serving them. This step and the technical solution to locate such ALTO server is currently discussed within the ALTO working group. This memo does not yet define such H12 server discovery.

The first step for a H12 client, before it can start querying for ALTO guidance, is to request the H12 server capabilities. The server capabilities are, e.g., administrative information (operator of the server, contact addresses, etc), the supported host location attributes (IP addresses or IP prefixes), the supported rating criteria, and the URIs to query for ALTO guidance. The H12 protocol uses only a single static URI path for retrieving the capability information. All other query URIs are announced by the server during

the capability retrieval.

4.3. Redirection

There are basically two cases where a H12 server has to redirect request to other locations:

a. the queried H12 server is overload and can tell about other H12 server;

b. the queried H12 server is overload and cannot tell about other H12 server;

c. the queried H12 server is solely used as entry point and redirects the actual H12 server;

d. the querying host in not allowed to use this ALTO server (e.g., host in ISP1 is querying ALTO server in ISP2) (which is a sub case of (a)).

4.4. Querying the ALTO Server

An ALTO client can query on its own or on behalf of other peers (e.g., a tracker). This is indicated in the resource consumer host location attribute rc_hla in the ALTO query. The query body itself contains the list of IP addresses or IP prefixes the ALTO client is asking guidance for. This shows a example list Figure 2 of IP addresses queried for

```
195.37.70.39/32         # mito.netlab.nec.de
193.141.139.237/32      # www.nec.de
58.89.210.171/32        # www.nec.co.jp
122.224.8.143/32        # www.huawei.cn
202.103.147.132/32      # www.zte.com.cn
135.245.1.29/32         # www.alcatel.de
139.15.248.12/32        # www.bosch.de
141.113.97.34/32        # www.daimler.de
129.206.0.0/16          # university of heidelberg
129.13.0.0/16           # university of karlsruhe
129.69.0.0/16           # university of stuttgart
130.83.0.0/16           # university of darmstadt
130.149.0.0/16          # university of berlin (TU)
171.67.0.0/16           # stanford university
129.78.64.24            # university of sidney
12.110.110.204/32       # www.nsa.gov
85.180.57.61/32         # some random residential DSL user (ALICE)
84.56.180.139/32        # some random residential DSL user (Arcor)
62.227.16.206/32        # some random residential DSL user (DTAG)
80.238.206.25/32        # some random residential DSL user in .ch
```

Figure 2: Example Candidate IPs for Query

The query is constructed as show in the below exampleFigure 3. The client requests guidance for the IP prefixes out of Figure 2 for its own IP address (prefix='195.37.70.39/32') stated in the rc_hla.

```
 <?xml version="1.0" encoding="UTF-8"?>
<alto xmlns='urn:ietf:params:xml:ns:p2p:alto'>
    <group_rating_request db_version='1234'>
        <pri_ratcrit crit='pref'/>
        <rc_hla><ipprefix version='4' prefix='195.37.70.39/32'/></rc_hla>
        <cnd_hla>
            <ipprefix version='4' prefix='195.37.70.39/32'/>
            <ipprefix version='4' prefix='193.141.139.237/32'/>
            <ipprefix version='4' prefix='58.89.210.171/32'/>
            <ipprefix version='4' prefix='122.224.8.143/32'/>
            <ipprefix version='4' prefix='202.103.147.132/32'/>
            <ipprefix version='4' prefix='135.245.1.29/32'/>
            <ipprefix version='4' prefix='139.15.248.12/32'/>
            <ipprefix version='4' prefix='141.113.97.34/32'/>
            <ipprefix version='4' prefix='129.206.0.0/16'/>
            <ipprefix version='4' prefix='129.13.0.0/16'/>
            <ipprefix version='4' prefix='129.69.0.0/16'/>
            <ipprefix version='4' prefix='130.83.0.0/16'/>
            <ipprefix version='4' prefix='130.149.0.0/16'/>
            <ipprefix version='4' prefix='171.67.0.0/16'/>
            <ipprefix version='4' prefix='129.78.64.24'/>
            <ipprefix version='4' prefix='12.110.110.204/32'/>
            <ipprefix version='4' prefix='85.180.57.61/32'/>
            <ipprefix version='4' prefix='84.56.180.139/32'/>
            <ipprefix version='4' prefix='62.227.16.206/32'/>
            <ipprefix version='4' prefix='80.238.206.25/32'/>
        </cnd_hla>
    </group_rating_request>
</alto>
```


Figure 3: XML encoded Query

This ipprefix tag carries a full IP address or an IP address prefix, leaving the client the choice how much of an IP address it wants to reveal to the server. That is, the client can request information for one or several specific IP addresses (prefix length equal 32 or 128), for address ranges, or for "the whole Internet" (prefix length equal 0). However, the "whole Internet" is not really referring to the whole Internet as such, as no single entity can have such a big knowledge, but to whatever broader scope the server can give guidance about. This scope can include, for instance, its own complete network.

Furthermore, the client specifies one or several rating criteria, such as operator preference, lower bound for delay, etc. Here is a work-in-progress list of such rating criteria, consisting of two levels of rating criteria offered to the client are:

- o Primary rating criterion
- o Further rating criteria

The offered rating criteria are:

- o operator's relative preference
- o Topological distance (Number of AS hops)
- o minimum boundary for upload bandwidth

4.5. ALTO Server Response

This section discussions at this of point of time only a positive reply. All other cases are TBD in this write-up. The listed response is shortened, see Section 1 for the full answer. The examplatory answer is listed for the IP address 193.141.139.237/32 and 202.103.147.132/32, and for the IP prefix 129.13.0.0/16.

The rating response given in the candidate host location attributes (cnd_hla) is different for the single requests, depending on what information can be delivered by the server. For 193.141.139.237/32, the server replies with two prefixes belonging to the same ISP. For 202.103.147.132/32, the server replies with even more details about other prefixes belonging to the same operator. The ensures that the client automatically learns even more prefixes the operator gives the same guidance for. A simple response is shown for the query about 129.13.0.0/16, where the response contains only the same prefixes as in the request.

```
<alto xmlns="urn:ietf:params:xml:ns:p2p:alto">
  <group_rating_reply statuscode="200">

    <cnd_hla overall_rating="1">
      <info type="country" unit="ISO-3166-1" value="DE" />
      <info type="X-NEC-map_of_internet" unit="areacode" value="2" />
      <ipprefix prefix="193.141.139.0/24" version="4" />
      <ipprefix prefix="193.141.140.0/22" version="4" />
    </cnd_hla>

        <cnd_hla overall_rating="3">
      <info type="country" unit="ISO-3166-1" value="CN" />
      <info type="X-NEC-map_of_internet" unit="areacode" value="3" />
      <ipprefix prefix="202.95.252.0/22" version="4" />
      <ipprefix prefix="202.120.24.0/25" version="4" />
      <ipprefix prefix="202.120.24.128/26" version="4" />
      <ipprefix prefix="202.120.24.192/27" version="4" />
      <ipprefix prefix="202.120.0.0/20" version="4" />
      <ipprefix prefix="202.120.16.0/21" version="4" />
      <ipprefix prefix="202.96.0.0/12" version="4" />
      <ipprefix prefix="202.112.0.0/13" version="4" />
    </cnd_hla>

    <cnd_hla overall_rating="1">
      <info type="country" unit="ISO-3166-1" value="DE" />
      <info type="X-NEC-map_of_internet" unit="areacode" value="2" />
      <ipprefix prefix="129.13.0.0/16" version="4" />
    </cnd_hla>

   <pri_ratcrit crit="pref" />
    <rc_hla>
      <info type="country" unit="ISO-3166-1" value="DE" />
      <info type="X-NEC-map_of_internet" unit="areacode" value="2" />
      <ipprefix prefix="195.37.0.0/16" version="4" />
    </rc_hla>
  </group_rating_reply>
</alto>
```


Figure 4: XML encoded Query

The response contains also a resource consumer host location attribute (rc_hla). This rc_hla echos partially the information from the request, but gives actually guidance to the ALTO client in what scope this information can be distributed amongst other peers. In this response, the server allows the redistribution of the received guidance to peers with the IP prefix 195.37.0.0/16.

5. Security Considerations

This initial version of this memo does not yet a full security considerations, but they will be added in future revision.

minimum boundary for upload bandwidth (AKA provisioned upload bandwidth): criminal suspects can easily re-use the geographical coordinates of an IP address (taken from whois) and google maps to correlate IP addresses and wealth of subscribers of that IP address.

6. Conclusion

This memo presents a very basic protocol, for sure work in progress, and is requesting feedback from the ALTO working group. Sebastian Kiesel is implementing the herein proposed protocol.

7. References

7.1. Normative References

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

7.2. Informative References

[ACM.ispp2p]
Aggarwal, V., Feldmann, A., and C. Scheideler, "Can ISPs and P2P systems co-operate for improved performance?", In ACM SIGCOMM Computer Communications Review (CCR), 37:3, pp. 29-40.

[I-D.kiesel-alto-reqs]
Kiesel, S., Popkin, L., Previdi, S., Woundy, R., and Y. Yang, "Application-Layer Traffic Optimization (ALTO) Requirements", draft-kiesel-alto-reqs-02 (work in progress), March 2009.

[I-D.penno-alto-protocol]
Penno, R. and Y. Yang, "ALTO Protocol", draft-penno-alto-protocol-04 (work in progress), October 2009.

[I-D.shalunov-alto-infoexport]
Shalunov, S., Penno, R., and R. Woundy, "ALTO Information Export Service", draft-shalunov-alto-infoexport-00 (work in progress), October 2008.

[I-D.stiemerling-alto-h1h2-protocol]
Stiemerling, M. and S. Kiesel, "ALTO H1/H2 Protocol", draft-stiemerling-alto-h1h2-protocol-00 (work in progress), March 2009.

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

1. Full XML-Response


```
<alto xmlns="urn:ietf:params:xml:ns:p2p:alto">
  <group_rating_reply statuscode="200">
    <cnd_hla overall_rating="1">
      <info type="country" unit="ISO-3166-1" value="DE" />
      <info type="X-NEC-map_of_internet" unit="areacode" value="2" />
      <ipprefix prefix="195.37.0.0/16" version="4" />
    </cnd_hla>
    <cnd_hla overall_rating="1">
      <info type="country" unit="ISO-3166-1" value="DE" />
      <info type="X-NEC-map_of_internet" unit="areacode" value="2" />
      <ipprefix prefix="193.141.139.0/24" version="4" />
      <ipprefix prefix="193.141.140.0/22" version="4" />
    </cnd_hla>
    <cnd_hla overall_rating="3">
      <info type="country" unit="ISO-3166-1" value="JP" />
      <info type="X-NEC-map_of_internet" unit="areacode" value="3" />
      <ipprefix prefix="58.87.128.0/17" version="4" />
      <ipprefix prefix="58.88.0.0/13" version="4" />
    </cnd_hla>
    <cnd_hla overall_rating="3">
      <info type="country" unit="ISO-3166-1" value="CN" />
      <info type="X-NEC-map_of_internet" unit="areacode" value="3" />
      <ipprefix prefix="122.224.0.0/12" version="4" />
      <ipprefix prefix="122.240.0.0/13" version="4" />
    </cnd_hla>
    <cnd_hla overall_rating="3">
      <info type="country" unit="ISO-3166-1" value="CN" />
      <info type="X-NEC-map_of_internet" unit="areacode" value="3" />
      <ipprefix prefix="202.95.252.0/22" version="4" />
      <ipprefix prefix="202.120.24.0/25" version="4" />
      <ipprefix prefix="202.120.24.128/26" version="4" />
      <ipprefix prefix="202.120.24.192/27" version="4" />
      <ipprefix prefix="202.120.0.0/20" version="4" />
      <ipprefix prefix="202.120.16.0/21" version="4" />
      <ipprefix prefix="202.96.0.0/12" version="4" />
      <ipprefix prefix="202.112.0.0/13" version="4" />
    </cnd_hla>
    <cnd_hla overall_rating="3">
      <info type="country" unit="ISO-3166-1" value="US" />
      <info type="X-NEC-map_of_internet" unit="areacode" value="1" />
      <ipprefix prefix="135.197.0.0/16" version="4" />
      <ipprefix prefix="135.198.0.0/15" version="4" />
      <ipprefix prefix="135.200.0.0/13" version="4" />
      <ipprefix prefix="135.208.0.0/12" version="4" />
      <ipprefix prefix="135.224.0.0/11" version="4" />
```

```
         <ipprefix prefix="136.0.0.0/9" version="4" />
         <ipprefix prefix="136.128.0.0/12" version="4" />
         <ipprefix prefix="136.144.0.0/16" version="4" />
       </cnd_hla>
       <cnd_hla overall_rating="1">
         <info type="country" unit="ISO-3166-1" value="DE" />
         <info type="X-NEC-map_of_internet" unit="areacode" value="2" />
         <ipprefix prefix="139.11.0.0/16" version="4" />
         <ipprefix prefix="139.12.0.0/14" version="4" />
         <ipprefix prefix="139.16.0.0/13" version="4" />
         <ipprefix prefix="139.24.0.0/14" version="4" />
         <ipprefix prefix="139.28.0.0/15" version="4" />
         <ipprefix prefix="139.30.0.0/16" version="4" />
       </cnd_hla>
       <cnd_hla overall_rating="1">
         <info type="country" unit="ISO-3166-1" value="DE" />
         <info type="X-NEC-map_of_internet" unit="areacode" value="2" />
         <ipprefix prefix="141.113.0.0/16" version="4" />
       </cnd_hla>
       <cnd_hla overall_rating="1">
         <info type="country" unit="ISO-3166-1" value="DE" />
         <info type="X-NEC-map_of_internet" unit="areacode" value="2" />
         <ipprefix prefix="129.206.0.0/16" version="4" />
       </cnd_hla>
       <cnd_hla overall_rating="1">
         <info type="country" unit="ISO-3166-1" value="DE" />
         <info type="X-NEC-map_of_internet" unit="areacode" value="2" />
         <ipprefix prefix="129.13.0.0/16" version="4" />
       </cnd_hla>
       <cnd_hla overall_rating="1">
         <info type="country" unit="ISO-3166-1" value="DE" />
         <info type="X-NEC-map_of_internet" unit="areacode" value="2" />
         <ipprefix prefix="129.69.0.0/16" version="4" />
         <ipprefix prefix="129.70.0.0/16" version="4" />
       </cnd_hla>
       <cnd_hla overall_rating="1">
         <info type="country" unit="ISO-3166-1" value="DE" />
         <info type="X-NEC-map_of_internet" unit="areacode" value="2" />
         <ipprefix prefix="130.83.0.0/16" version="4" />
       </cnd_hla>
       <cnd_hla overall_rating="1">
         <info type="country" unit="ISO-3166-1" value="DE" />
         <info type="X-NEC-map_of_internet" unit="areacode" value="2" />
         <ipprefix prefix="130.149.0.0/16" version="4" />
       </cnd_hla>
       <cnd_hla overall_rating="3">
         <info type="country" unit="ISO-3166-1" value="US" />
         <info type="X-NEC-map_of_internet" unit="areacode" value="1" />
```

```
      <ipprefix prefix="171.64.0.0/12" version="4" />
    </cnd_hla>
    <cnd_hla overall_rating="3">
      <info type="country" unit="ISO-3166-1" value="AU" />
      <info type="X-NEC-map_of_internet" unit="areacode" value="3" />
      <ipprefix prefix="129.78.0.0/16" version="4" />
    </cnd_hla>
    <cnd_hla overall_rating="3">
      <info type="country" unit="ISO-3166-1" value="US" />
      <info type="X-NEC-map_of_internet" unit="areacode" value="1" />
      <ipprefix prefix="12.109.8.120/29" version="4" />
      <ipprefix prefix="12.109.8.128/25" version="4" />
      <ipprefix prefix="12.109.9.0/24" version="4" />
      <ipprefix prefix="12.109.10.0/23" version="4" />
      <ipprefix prefix="12.109.12.0/22" version="4" />
      <ipprefix prefix="12.109.16.0/20" version="4" />
      <ipprefix prefix="12.109.32.0/19" version="4" />
      <ipprefix prefix="12.109.64.0/18" version="4" />
      <ipprefix prefix="12.109.128.0/17" version="4" />
      <ipprefix prefix="12.129.72.0/27" version="4" />
      <ipprefix prefix="12.129.0.0/18" version="4" />
      <ipprefix prefix="12.129.64.0/21" version="4" />
      <ipprefix prefix="12.110.0.0/15" version="4" />
      <ipprefix prefix="12.112.0.0/12" version="4" />
      <ipprefix prefix="12.128.0.0/16" version="4" />
    </cnd_hla>
    <cnd_hla overall_rating="1">
      <info type="country" unit="ISO-3166-1" value="DE" />
      <info type="X-NEC-map_of_internet" unit="areacode" value="2" />
      <ipprefix prefix="85.176.0.0/13" version="4" />
    </cnd_hla>
    <cnd_hla overall_rating="1">
      <info type="country" unit="ISO-3166-1" value="DE" />
      <info type="X-NEC-map_of_internet" unit="areacode" value="2" />
      <ipprefix prefix="84.56.0.0/13" version="4" />
    </cnd_hla>
    <cnd_hla overall_rating="1">
      <info type="country" unit="ISO-3166-1" value="DE" />
      <info type="X-NEC-map_of_internet" unit="areacode" value="2" />
      <ipprefix prefix="62.225.180.0/22" version="4" />
      <ipprefix prefix="62.225.184.0/21" version="4" />
      <ipprefix prefix="62.225.192.0/18" version="4" />
      <ipprefix prefix="62.226.0.0/15" version="4" />
    </cnd_hla>
    <cnd_hla overall_rating="2">
      <info type="country" unit="ISO-3166-1" value="CH" />
      <info type="X-NEC-map_of_internet" unit="areacode" value="2" />
      <ipprefix prefix="80.238.128.0/17" version="4" />
```

```
        </cnd_hla>
        <pri_ratcrit crit="pref" />
        <rc_hla>
          <info type="country" unit="ISO-3166-1" value="DE" />
          <info type="X-NEC-map_of_internet" unit="areacode" value="2" />
          <ipprefix prefix="195.37.0.0/16" version="4" />
        </rc_hla>
      </group_rating_reply>
    </alto>
```


Figure 5: XML encoded Query

2. 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

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

Email: ietf-alto@skiesel.de

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

Phone: +49 6221 4342 113
Fax: +49 6221 4342 155
Email: stiemerling@cs.uni-goettingen.de