Application-Layer Traffic Optimization (ALTO) Requirements

draft-kiesel-alto-reqs-02.txt

Sebastian Kiesel (Editor) <kiesel@nw.neclab.eu>

IETF-74, ALTO WG, San Francisco, 2009-03-26

draft-kiesel-alto-reqs-02.txt

Authors

- S. Kiesel, Editor, NEC
- L. Popkin, Pando Networks, Inc.
- S. Previdi, Cisco Systems, Inc.
- R. Woundy, Comcast Corp.
- Y. R. Yang, Yale University

Further contributors

- Z. Despotovic, DOCOMO
- V. Gurbani, Alcatel Lucent
- J. Livingood, Comcast Corp.
- E. Marocco, Telecom Italia
- S. Niccolini, NEC
- M. Stiemerling, NEC
- Many other people on the p2pi & alto mailing lists

Related research projects

- NAPA-WINE (EC 7th FP) [NEC people]
- P4P working group [L. Popkin, Y.R. Yang]

Discussions before ALTO WG mailing list was established:

- http://ubiq.tilab.com/pipermail/alto-reqs/ (archives)

draft-kiesel-alto-reqs-02.txt

- Overview: changes
- Overview and discussion: Requirements
- Overview and discussion: Rating Criteria
- Conclusion & next steps

draft-kiesel-alto-reqs-02: changes since -01

- Enhanced terminology, add (non-normative) description of framework architecture to better explain scope
- Most of the reqs. are (and were) about the “ALTO Client Protocol”
  - stated explicitly
  - add sections for reqs. wrt. other interfaces (few reqs added so far)
- Remove some requirements that were
  - too closely related to the “Sorting Oracle” approach
  - de-facto implementation specs (e.g., overload control)
  - related to very specific corner cases
- Split remaining reqs. in more clauses -> number increases
- Labels RQv02-?? to avoid ambiguities after renumbering
- Discussion of possible Host Location Attributes & Rating Criteria

Table of Contents

Terminology & framework architecture

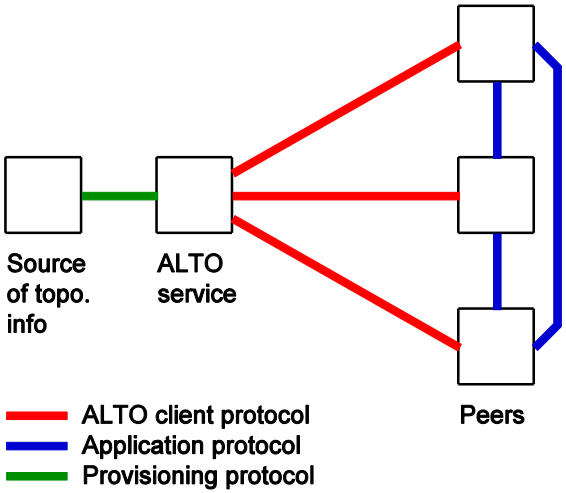

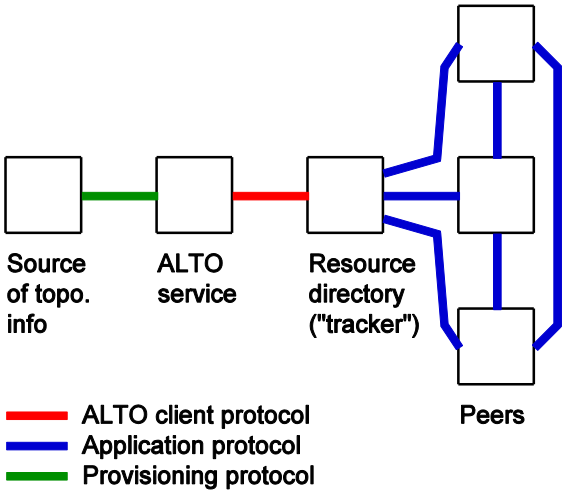

Scenario without resource directory

- Large number of associations
- Peer's IP addresses visible as matter of principle

Scenario with resource directory (e.g., "tracker")

- ResDirectory needs not only 1:m guidance matrix but m:n
- ResDirectory may cache all info

TBD: add discovery and Inter-ALTO server proto to the picture

Is this non-normative illustration useful?

Requirements: ALTO Client Protocol

General Requirements (“Compliance”)

- ALTO Server MUST implement the ALTO client protocol (RQv02-1)
- ALTO Clients MUST implement the ALTO client protocol (RQv02-2)

- *REQ. RQv02-3: The detailed specification of a protocol is out of the scope of this document. However, any protocol specification that claims to implement the ALTO client protocol MUST be compliant to the requirements itemized in this document.*

  Should the wording assume that exactly one ALTO Client protocol specification will be the official outcome of this working group?

Requirements: ALTO Client Protocol

Protocol Semantics

- Client MUST be able to solicit guidance (RQv02-4)
- Client SHOULD be able to specify rating criteria (RQv02-5)
- Server MUST be able to express guidance (RQv02-6)
- Client Protocol MUST support Client placement both in regular P2P node (RQv02-7) and in Resource Directory (RQv02-8)
- Target-aware query mode: resource & candidate providers known
  Target-independent query mode: guidance will be evaluated later
  CP MUST support one (RQv02-9) & SHOULD support both (RQv02-10)
- CP SHOULD have lifetimes (RQv02-11) & aging mechanism (RQv02-12)
- CP MUST be designed in a way that parties other than the network operator can provide the ALTO Service (RQv02-13)
- Different ALTO services MUST be able to coexist (RQv02-14)
- CP MUST be extensible (RQv02-15) & MUST have versions (RQv02-16)

Requirements: ALTO Client Protocol

Error handling and overload protection

- Application MUST also work without ALTO guidance (RQv02-17)
- ALTO Client Protocol MUST use TCP transport (RQv02-18)
- An overloaded ALTO Server MUST be able to inform clients, and
  - Ask clients to throttle query rate (RQv02-19)
  - Redirect clients to another server (RQv02-20)
  - Terminate conversation (RQv02-21)
  - Reject new conversation attempts (RQv02-22)

Requirements: ALTO discovery

- ALTO clients MUST be able to use the server discovery mechanism, in order to find out where to send queries (RQv02-23)
- The server discovery mechanism SHOULD be able to return the respective contact information for several servers. (RQv02-24)
- The ALTO server discovery mechanism SHOULD be able to indicate preferences for each returned ALTO server contact information. (RQv02-25)
- The ALTO server discovery mechanism SHOULD be independent of specific link-layer protocols or access network arch's. (RQv02-26)

- Missing: If a resource directory wants to do third-party queries on behalf of a "distant" resource consumer, how to find the ALTO server with the respective knowledge?
- Relation to other drafts assessing ALTO server discovery?

Requirements: Security

- Client Protocol MUST support mutual authentication (RQv02-27)
- Client protocol MUST support different levels of detail in queries and responses, to protect privacy of operators (e.g., network topology) (RQv02-28) and users (RQv02-29)
- The ALTO client protocol SHOULD be defined in a way, that the operator of one ALTO server cannot easily deduce the resource identifier (e.g., file name in P2P file sharing) which the resource consumer seeking ALTO guidance wants to access. (RQv02-30)
- The ALTO protocol MUST include appropriate mechanisms to protect the ALTO service against DoS attacks. (RQv02-31)

Host Location Attributes

Proposal for attribute definition procedure

- Define initial set of mandatory and optional attributes
- Establish registry (IANA) for adding further (optional) attributes

Attributes that have been discussed or are used in other drafts:

- IP address or IP address ranges (CIDR notation)
- Autonomous System Numbers
- Group IDs which expand to a set of other identifiers (IP, AS No.)

- Is this process proposal reasonable
  (how likely is it that additional useful attributes are defined later)?
- Is it reasonable to include AS Numbers and require the client to have a mapping mechanism? Better map in ALTO server?

Rating Criteria

Proposal for attribute definition procedure

- Define initial set of mandatory and optional attributes
- Establish registry (IANA) for adding further (optional) attributes

Attributes that have been discussed or are used in other drafts:

- Relative topological distance
- Absolute topological distance (AS hops, router hops)
- Absolute physical distance based on geolocation of IP address
- Relative operator's preference (e.g., based on peering costs)
- Charging / volume caps (not well understood so far)
- Upper/lower bounds for bandwidth and RTT – goal is to quickly exclude candidate peers from further P2P measurements, if it is clear that the candidate cannot meet performance requirements
- Inappropriate: if the primary goal is congestion control – use TCP!

Conclusion

- Cleanup of requirements
  - Clearer terminology
  - Better structured in subsections regarding different interfaces
  - Remove req's implicitly assuming the "sorting oracle" approach
  - Incorporated feedback from various discussions
- Added discussion of possible Host Location Attributes
- Added discussion of possible Rating Criteria
  - Probably further discussion needed here

- Next steps?

draft-kiesel-alto-reqs-02.txt

Appendix B. Acknowledgments

The authors would like to thank

- **Vijay K. Gurbani <vkg@alcatel-lucent.com>**
- **Enrico Marocco <enrico.marocco@telecomitalia.it>**

for fostering discussions that lead to the creation of this document, and for giving valuable comments on it.

Sebastian Kiesel, Saverio Niccolini, Jan Seedorf, and Martin Stiemerling are 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.

Laird Popkin and Y. Richard Yang are grateful to the many contributions made by the members of the P4P working group and Yale Laboratory of Networked Systems. The P4P working group is hosted by DCIA.

Backup slides

draft-kiesel-alto-reqs-02.txt

3. ALTO requirements

3.1. ALTO client protocol

3.1.1. General requirements

REQ. RQv02-1: The ALTO service is provided by one or more ALTO servers. ALTO servers MUST implement the ALTO client protocol, for receiving ALTO queries from ALTO clients and for sending the corresponding ALTO replies.

REQ. RQv02-2: ALTO clients MUST implement the ALTO client protocol, for sending ALTO queries to ALTO servers and for receiving the corresponding ALTO replies.

REQ. RQv02-3: The detailed specification of a protocol is out of the scope of this document. However, any protocol specification that claims to implement the ALTO client protocol MUST be compliant to the requirements itemized in this document.

draft-kiesel-alto-reqs-02.txt

3.1.2. Protocol semantics

REQ. RQv02-4: The format of the ALTO query message MUST allow the ALTO client to solicit guidance for selecting appropriate resource providers.

REQ. RQv02-5: The ALTO guidance is be based on the evaluation of one or several rating criteria (see Section 5). The ALTO query message SHOULD allow the ALTO client to express which rating criteria should be considered, as well as their relative relevance for the specific application that will eventually make use of the guidance.

REQ. RQv02-6: The format of the ALTO reply message MUST allow the ALTO server to express his guidance for selecting appropriate resource providers.

draft-kiesel-alto-reqs-02.txt

With respect to the placement of ALTO clients, several modes of operation exist:

- One mode of ALTO operation is that ALTO clients may be embedded directly in the resource consumer (e.g., peer of a DHT-based P2P application), which wants to access a resource.
- Another mode of operation is to perform ALTO queries indirectly, via resource directories (e.g., tracker of a P2P application), which may issue ALTO queries to solicit preference on potential resource providers, considering the respective resource consumer.

REQ. RQv02-7: The ALTO client protocol MUST support the mode of operation, in which the ALTO client is directly embedded in the resource consumer.

REQ. RQv02-8: The ALTO client protocol MUST support the mode of operation, in which the ALTO client is embedded in the resource directory.

draft-kiesel-alto-reqs-02.txt

With respect to the timing of ALTO queries, several modes of operation exist:

- In target-aware query mode, an ALTO client performs the ALTO query when the desired resource and a set of candidate resource providers are already known, i. e., after DHT lookups, queries to the resource directory, etc.
- In target-independent query mode, ALTO queries are performed in advance or periodically, in order to receive "target-independent" guidance, which will be cached locally and evaluated later, when a resource is to be accessed.

REQ. RQv02-9: The ALTO client protocol MUST support at least one of these two modes, either the target-aware or the target-independent query mode.

REQ. RQv02-10: The ALTO client protocol SHOULD support both the target-aware and the target-independent query mode.

draft-kiesel-alto-reqs-02.txt

REQ. RQv02-11: The ALTO client protocol SHOULD support lifetime attributes, to enable caching of recommendations at ALTO clients.

REQ. RQv02-12: The ALTO client protocol SHOULD specify an aging mechanism, which allows to give newer recommendations precedence over older ones.

REQ. RQv02-13: The ALTO client protocol MUST be designed in a way that the ALTO service can be provided by an operator which is not the operator of the IP access network.

REQ. RQv02-14: The ALTO client protocol MUST be designed in a way that different instances of the ALTO service operated by different providers can coexist.

draft-kiesel-alto-reqs-02.txt

REQ. RQv02-15: The ALTO client protocol MUST include support for adding protocol extensions in a non-disruptive, backward-compatible way.

REQ. RQv02-16: The ALTO client protocol MUST include protocol versioning support, in order to clearly distinguish between incompatible major versions of the protocol.

draft-kiesel-alto-reqs-02.txt

3.1.3. Error handling and overload protection

REQ. RQv02-17: Any application designed to use ALTO MUST also work if no ALTO servers can be found or if no responses to ALTO queries are received, e.g., due to connectivity problems or overload situation.

REQ. RQv02-18: The ALTO client protocol MUST use TCP based transport.

draft-kiesel-alto-reqs-02.txt

REQ. RQv02-19: An ALTO server, which is operating close to its capacity limit, MUST be able to inform clients about its impending overload situation, and require them to throttle their query rate.

REQ. RQv02-20: An ALTO server, which is operating close to its capacity limit, MUST be able to inform clients about its impending overload situation, and redirect them to another ALTO server.

REQ. RQv02-21: An ALTO server, which is operating close to its capacity limit, MUST be able to inform clients about its impending overload situation, and terminate the conversation with the ALTO client.

REQ. RQv02-22: An ALTO server, which is operating close to its capacity limit, MUST be able to inform clients about its impending overload situation, and reject new conversation attempts.

draft-kiesel-alto-reqs-02.txt

3.2. ALTO server discovery

REQ. RQv02-23: ALTO clients MUST be able to use the ALTO server discovery mechanism, in order to find out where to send ALTO queries.

REQ. RQv02-24: The ALTO server discovery mechanism SHOULD be able to return the respective contact information for several ALTO servers.

REQ. RQv02-25: The ALTO server discovery mechanism SHOULD be able to indicate preferences for each returned ALTO server contact information.

REQ. RQv02-26: The ALTO server discovery mechanism SHOULD be independent of specific link-layer protocols or access network architectures. For example, many broadband access networks use DHCP for configuration, while others use PPPoE. In contrast, DNS is available in virtually all Internet access networks.

draft-kiesel-alto-reqs-02.txt

3.3. Security and privacy

REQ. RQv02-27: The ALTO client protocol MUST support mechanisms for mutual authentication and authorization of ALTO clients and servers.

REQ. RQv02-28: The ALTO client protocol MUST support different levels of detail in queries and responses, in order for the operator of an ALTO service to be able to control how much information (e.g., about the network topology) is disclosed.

REQ. RQv02-29: The ALTO client protocol MUST support different levels of detail in queries and responses, in order to protect the privacy of users, to ensure that the operators of ALTO servers and other users of the same application cannot derive sensitive information.

draft-kiesel-alto-reqs-02.txt

REQ. RQv02-30: The ALTO client protocol SHOULD be defined in a way, that the operator of one ALTO server cannot easily deduce the resource identifier (e.g., file name in P2P file sharing) which the resource consumer seeking ALTO guidance wants to access.

REQ. RQv02-31: The ALTO protocol MUST include appropriate mechanisms to protect the ALTO service against DoS attacks.

draft-kiesel-alto-reqs-02.txt

4. Host location attributes are used in the ALTO client protocol to describe the location of a host in the network topology. The following list gives an overview on such attributes that have been proposed in the past, or which are in use by by ALTO-related prototype implementations.

   One possible way forward is to define the syntax and semantics of a mandatory set of attributes, which have to be understood by all entities that implement the ALTO client protocol. Furthermore, defining a set of optional attributes, as well as a procedure for allocating new attributes (e.g., an IANA registry) may be required. However, there was no broad discussion of this issue so far and no consensus has been reached. Therefore, the only purpose of the following list is to document the attributes that have been proposed so far, and to solicit further feedback and discussion:

draft-kiesel-alto-reqs-02.txt

- IP address or range of IP addresses (in CIDR notation)
- Autonomous System (AS) number
- 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).

draft-kiesel-alto-reqs-02.txt

5. 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 following list gives an overview on such rating criteria that have been proposed in the past, or which are in use by by ALTO-related prototype implementations.

   One possible way forward is to define the syntax and semantics of a mandatory set of criteria, which have to be understood by all entities that implement the ALTO client protocol. Furthermore, defining a set of optional criteria, as well as a procedure for allocating new criteria (e.g., an IANA registry) may be required. However, there was no broad discussion of this issue so far and no consensus has been reached. Therefore, the only purpose of the following list is to document the attributes that have been proposed so far, and to solicit further feedback and discussion.

draft-kiesel-alto-reqs-02.txt

5.1. Distance-related rating criteria

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

draft-kiesel-alto-reqs-02.txt

- Absolute physical distance, based on knowledge of the approximate geolocation (continent, country) of an IP address.
- Relative operator's preference: higher numerical value indicates that the application should prefer this candidate resource provider over others with lower values (if no other reasons speak against it, such as probed throughput). Again, as this is a relative measure, the ALTO service does not have to indicate how the values have been computed. Examples could be: cost for peering or transit traffic, traffic engineering inside the own network, and other policies.

draft-kiesel-alto-reqs-02.txt

5.2. Charging-related rating criteria

- 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, throtteld, or charged separately at an indicated price). The interaction of several applications running on a host, out of which some use this attribute while others don't, as well as the evaluation of this attribute in resource directories, which issue ALTO queries on behalf of other peers, are for further study.

draft-kiesel-alto-reqs-02.txt

5.3. Performance-related rating criteria

The following rating criteria are subject to the remarks below.

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

draft-kiesel-alto-reqs-02.txt

- An arbitrary lower bound for the RTT between resource consumer and the candidate resource provider (only in ALTO replies). 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.

draft-kiesel-alto-reqs-02.txt

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

- the ALTO service is not an admission control system
- the ALTO service may not know the instantaneous congestion status of the network
- the ALTO service may not know all link bandwidths, i.e., where the bottleneck really is, and there may be shared bottlenecks
- the ALTO service may not know whether the candidate peer itself is overloaded
- the ALTO service may not know whether the candidate peer throttles the bandwidth it devotes for the considered application

draft-kiesel-alto-reqs-02.txt

- 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.

draft-kiesel-alto-reqs-02.txt

- 5.4. Inappropriate rating criteria

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

How ALTO could work in a tracker-based P2P network

Physical topology

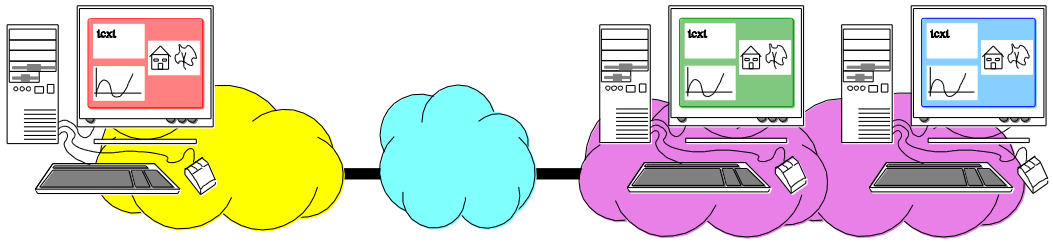

How ALTO could work in a tracker-based P2P network

Client Application
(out of scope)

Application overlay topology

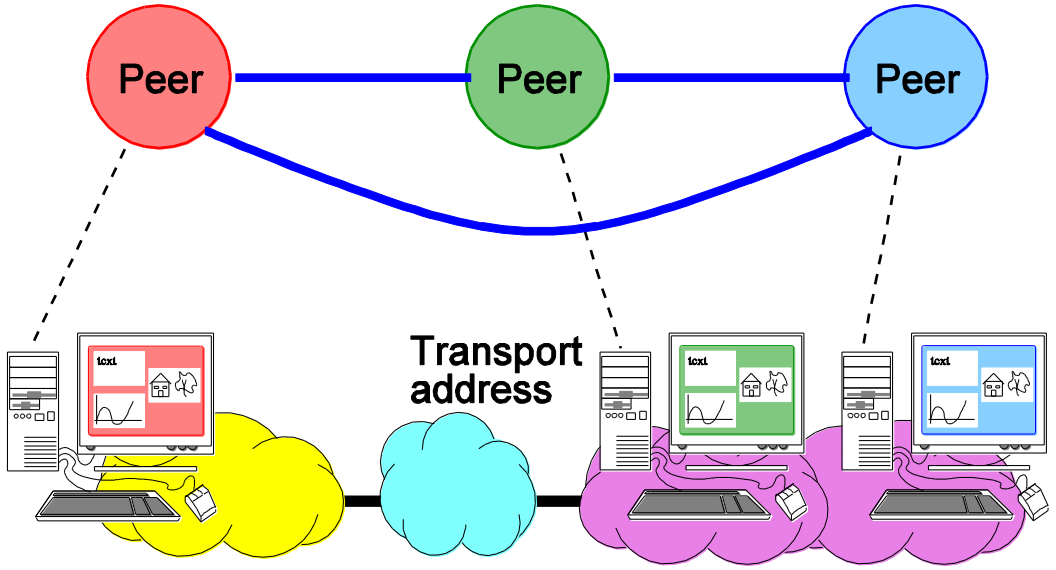

Physical topology

How ALTO could work in a tracker-based P2P network

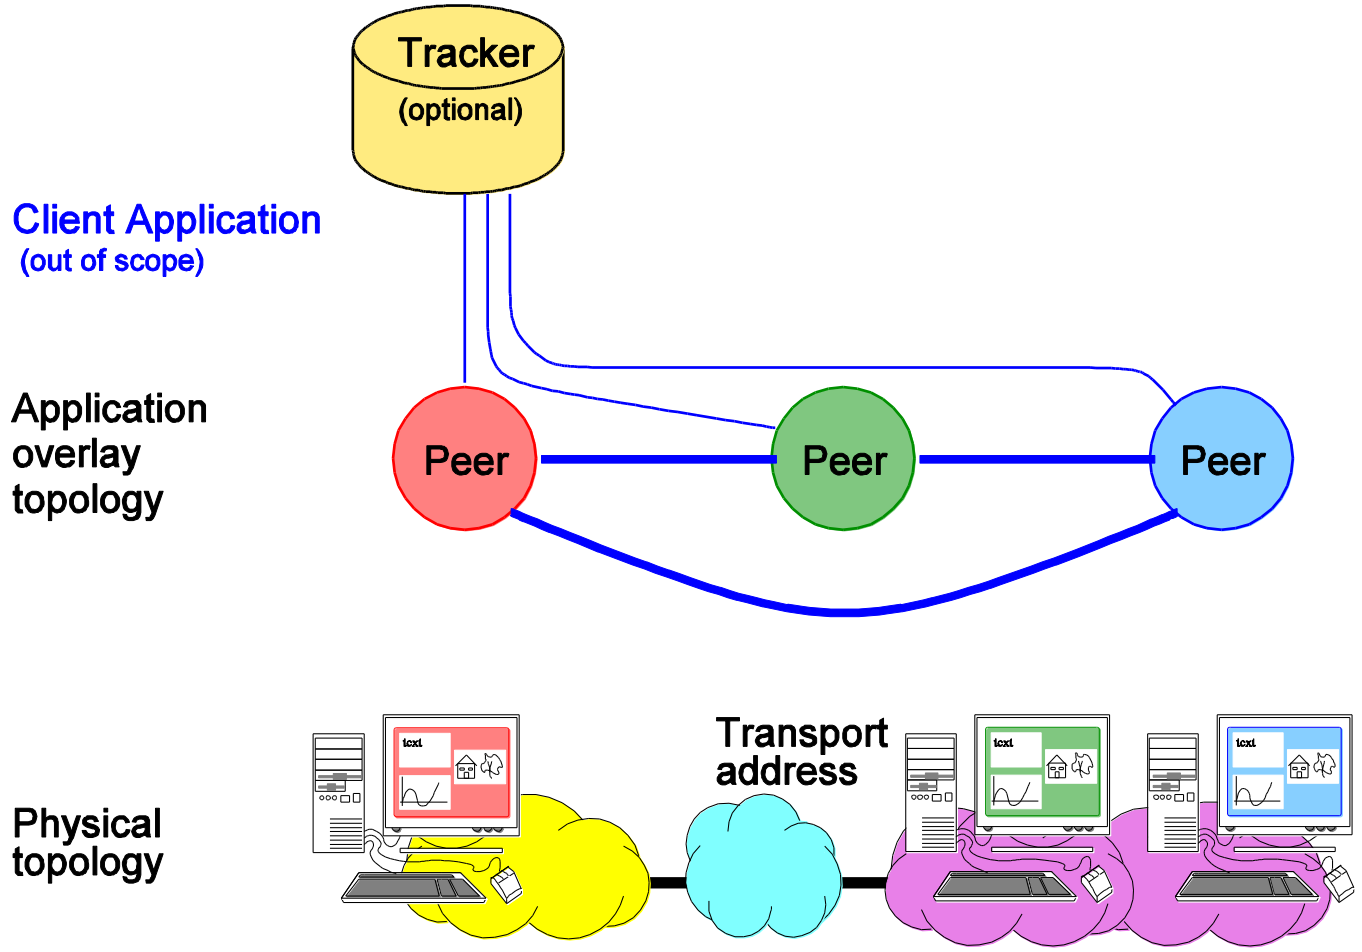

How ALTO could work in a tracker-based P2P network

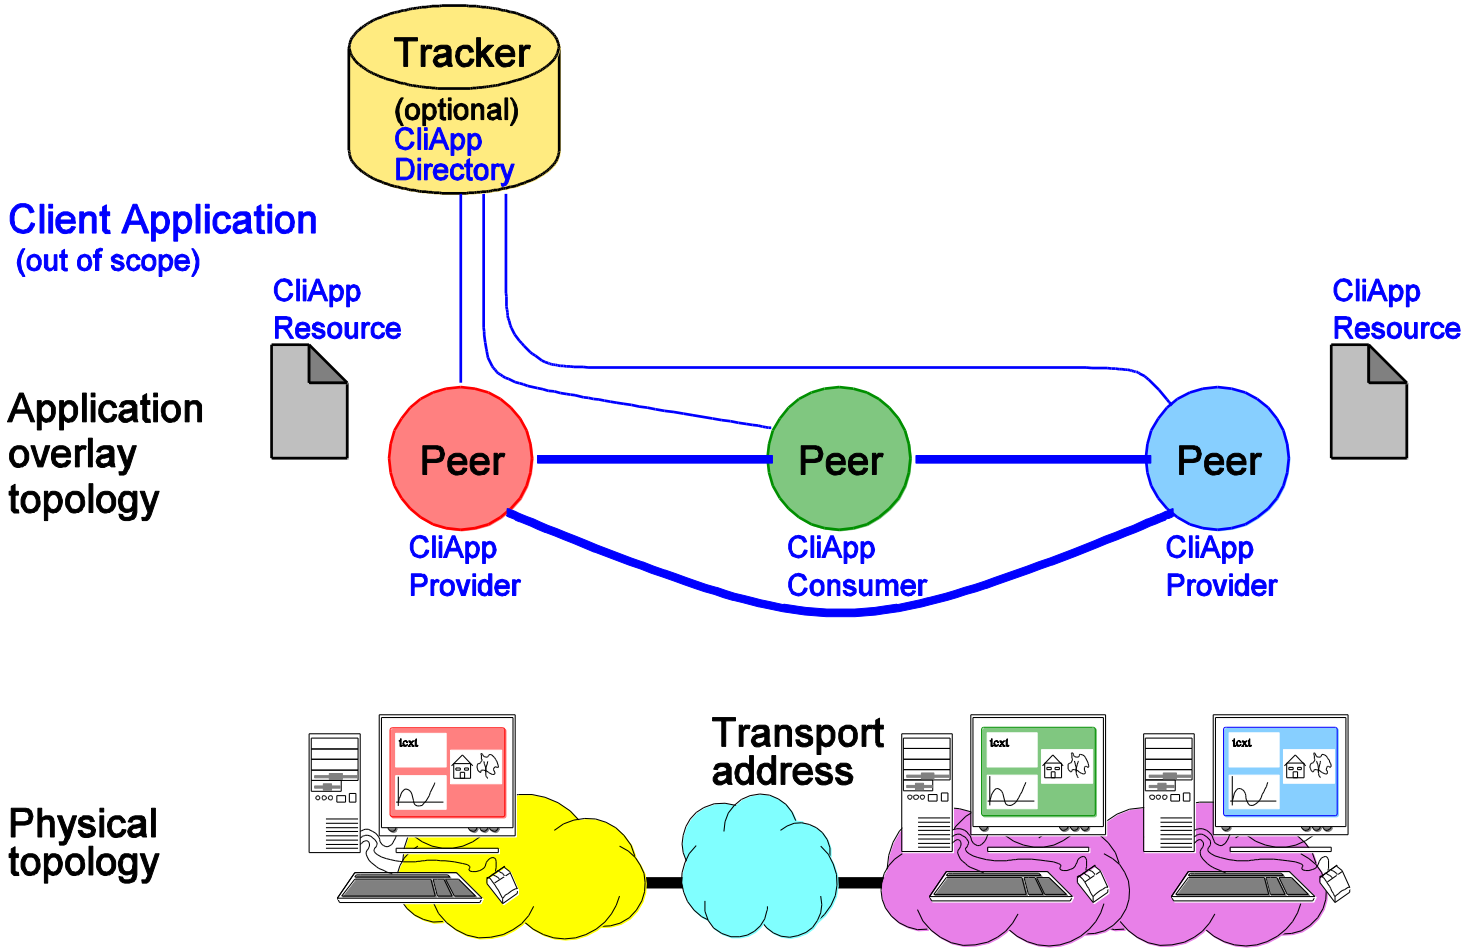

How ALTO could work in a tracker-based P2P network

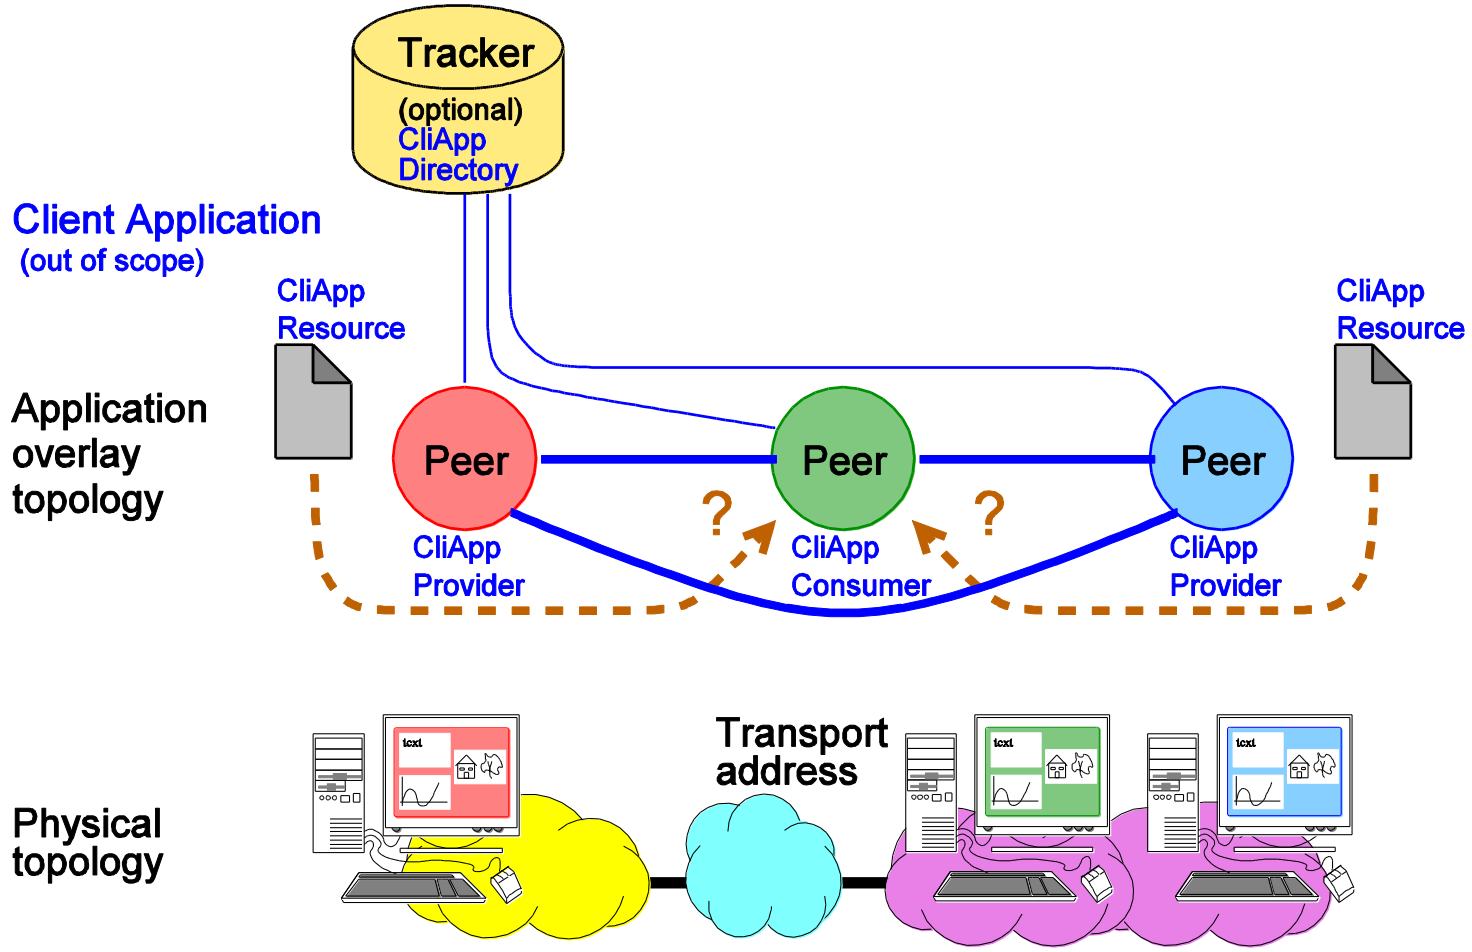

How ALTO could work in a tracker-based P2P network

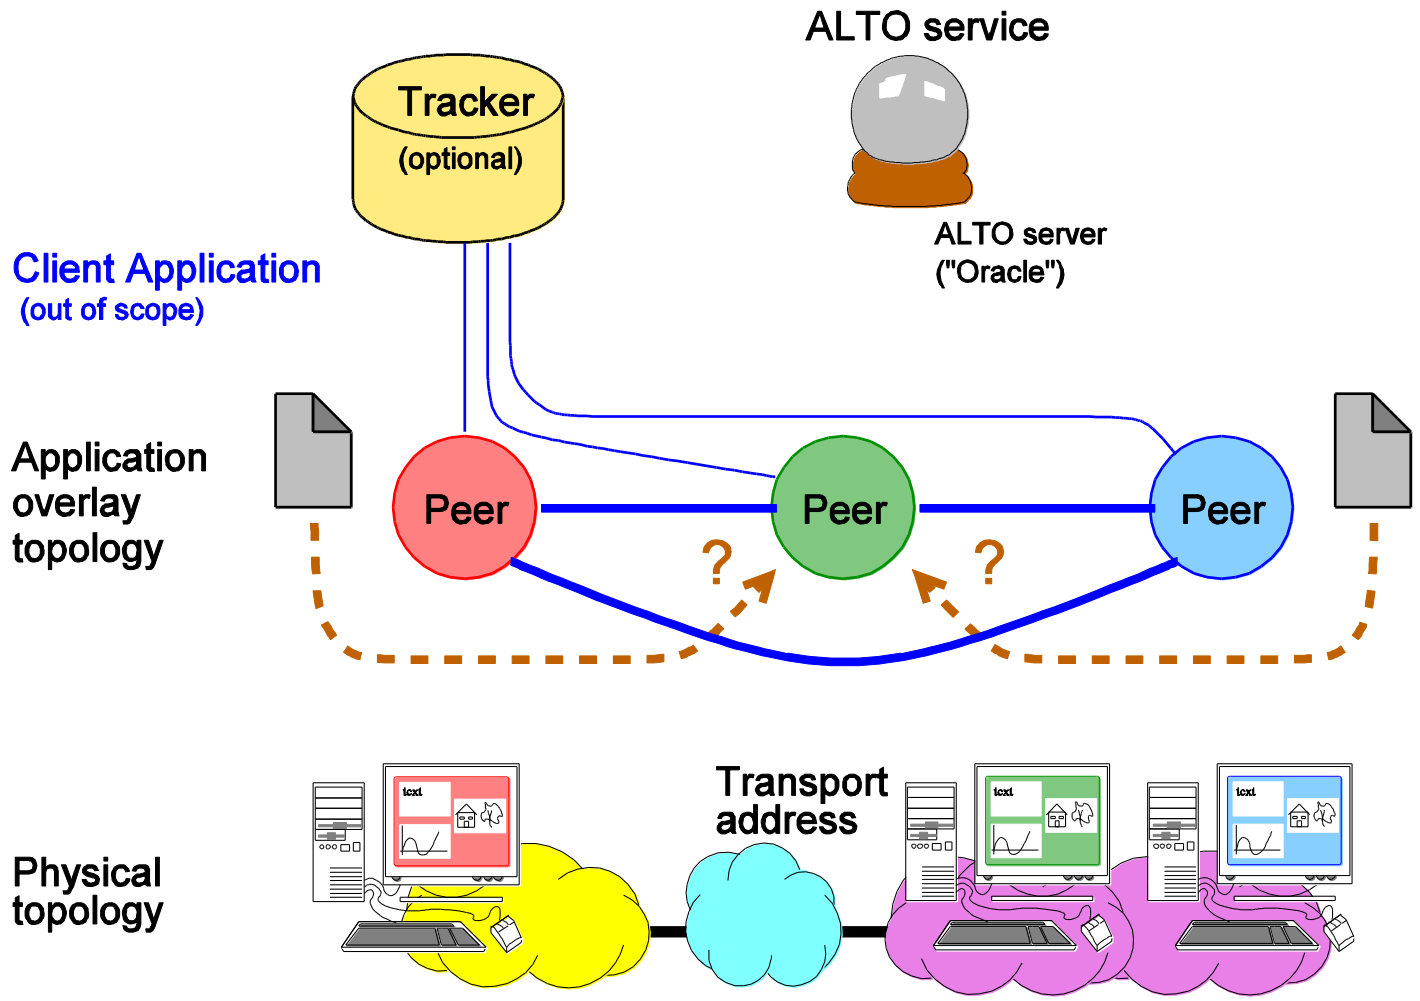

How ALTO could work in a tracker-based P2P network

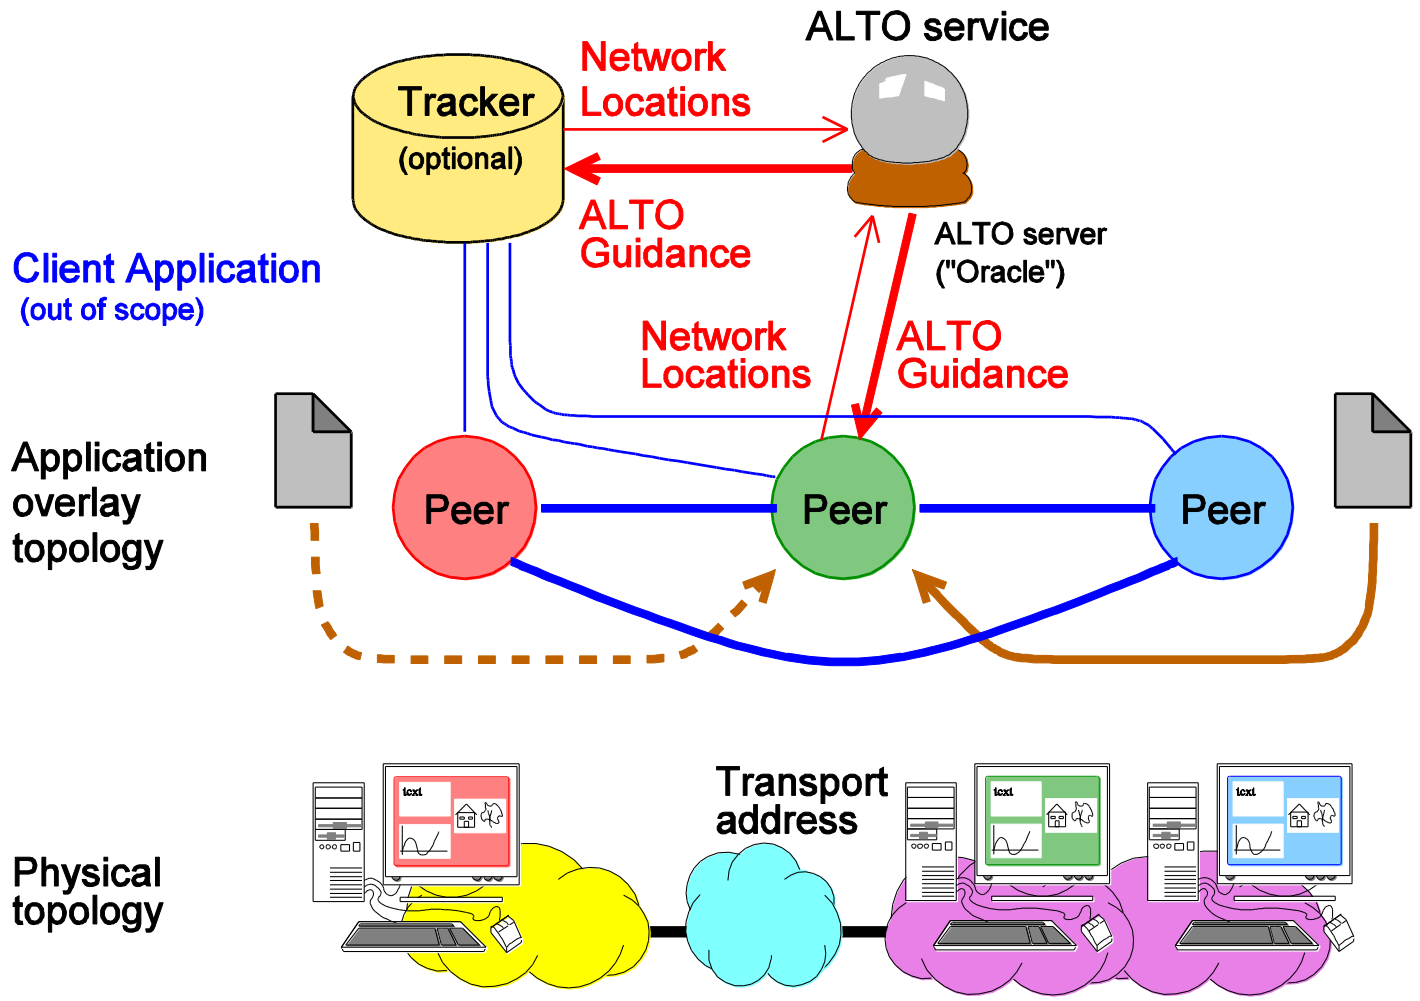

How ALTO could work in a tracker-based P2P network

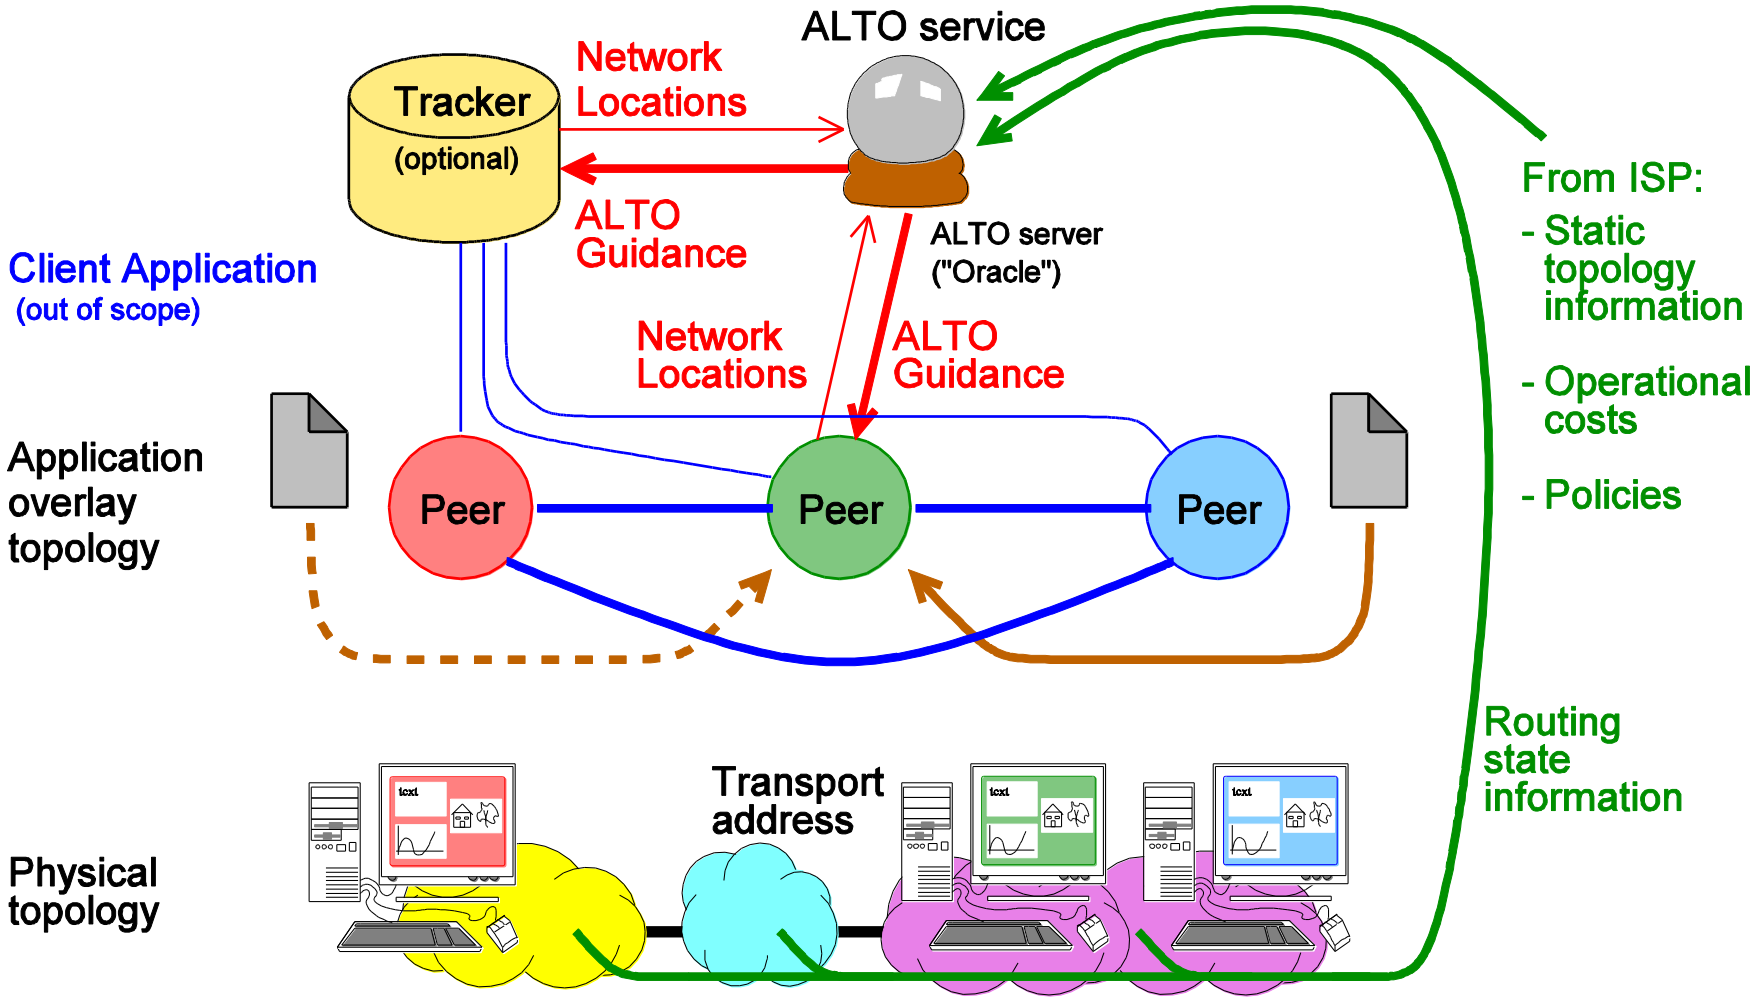

How ALTO could work in a tracker-based P2P network

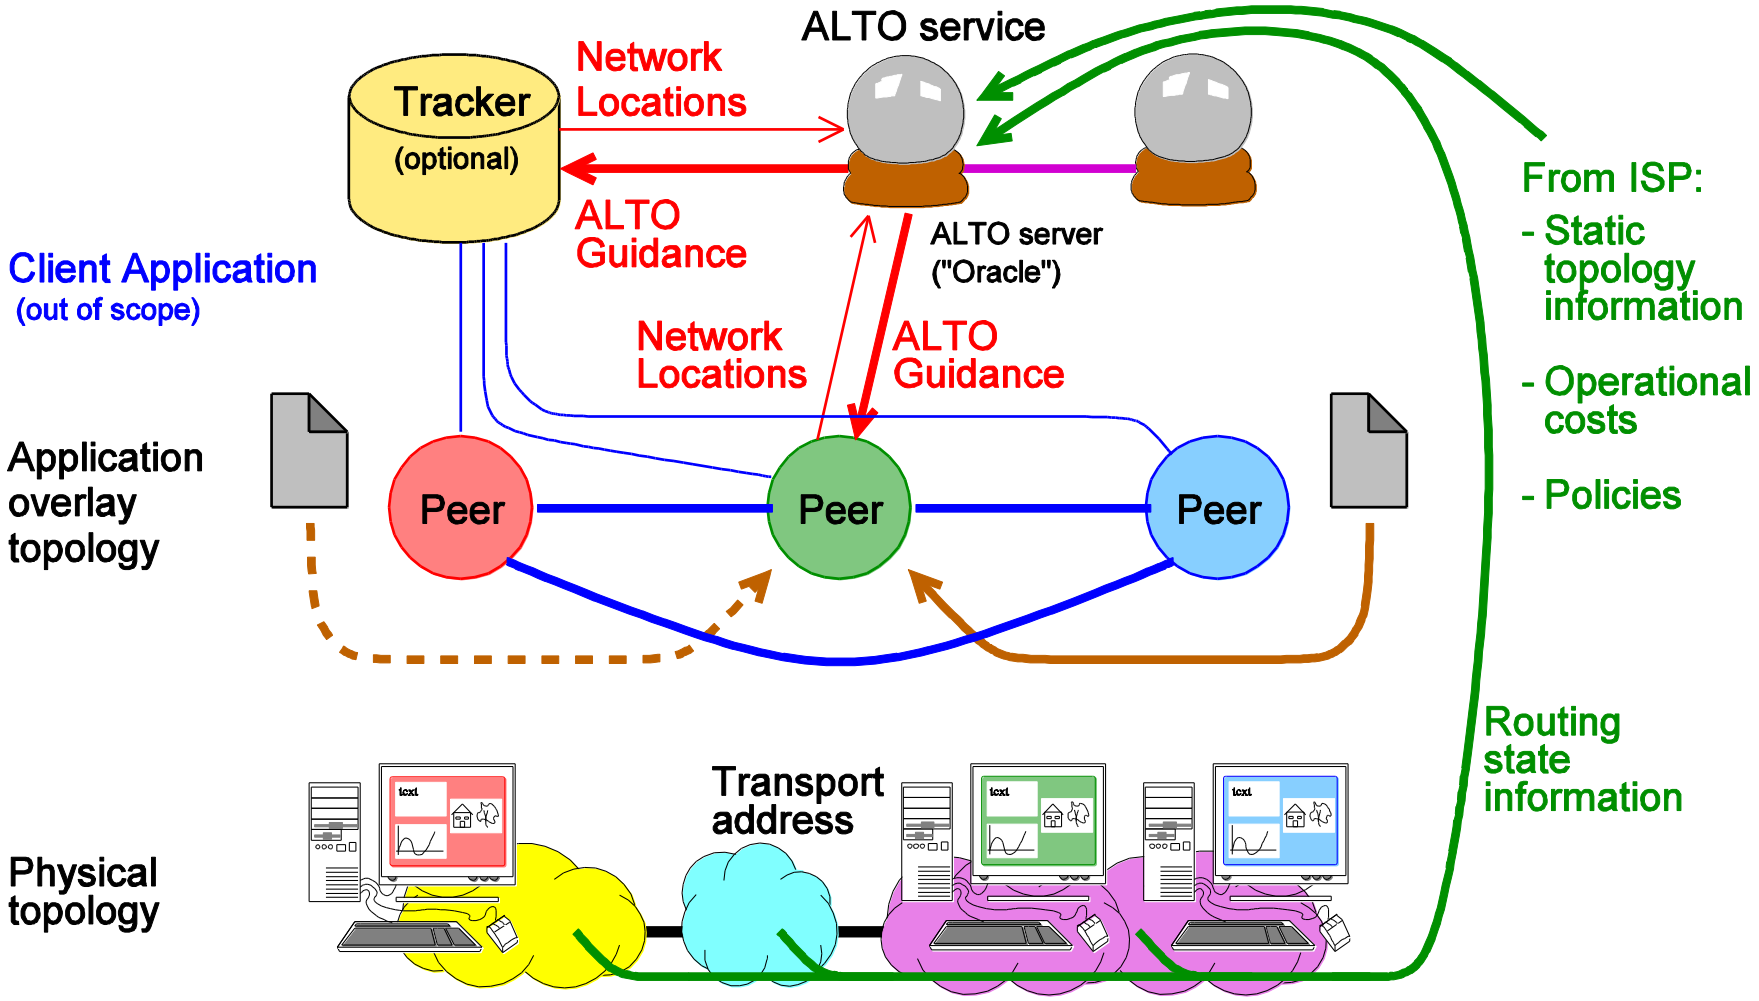

draft-kiesel-alto-reqs-02.txt

Appendix B. Acknowledgments

The authors would like to thank

- **Vijay K. Gurbani <vkg@alcatel-lucent.com>**
- **Enrico Marocco <enrico.marocco@telecomitalia.it>**

for fostering discussions that lead to the creation of this document, and for giving valuable comments on it.

Sebastian Kiesel, Saverio Niccolini, Jan Seedorf, and Martin Stiemerling are 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.

Laird Popkin and Y. Richard Yang are grateful to the many contributions made by the members of the P4P working group and Yale Laboratory of Networked Systems. The P4P working group is hosted by DCIA.