

Application-Layer Traffic Optimization (ALTO) Requirements
draft-kiesel-alto-reqs-00.txt

Sebastian Kiesel (Editor) <kiesel@nw.neclab.eu>
IETF-72, ALTO BoF, Dublin, 2008-07-29

draft-kiesel-alto-reqs-00.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

- V. Gurbani, Alcatel Lucent
- J. Livingood, Comcast Corp.
- E. Marocco, Telecom Italia
- S. Niccolini, NEC
- J. Seedorf, NEC
- M. Stiemerling, NEC

Related research projects

- NAPA-WINE: Network-Aware P2P-TV Application over Wise Networks (EC 7th FP) [NEC people]
- P4P working group (hosted by DCIA) and Yale Laboratory of Networked Systems [L. Popkin, Y.R. Yang]

Outline

draft-kiesel-alto-reqs-00

- Terminology
 - Requirements
 - Open issues
-
- Further steps

Initial ALTO Terminology

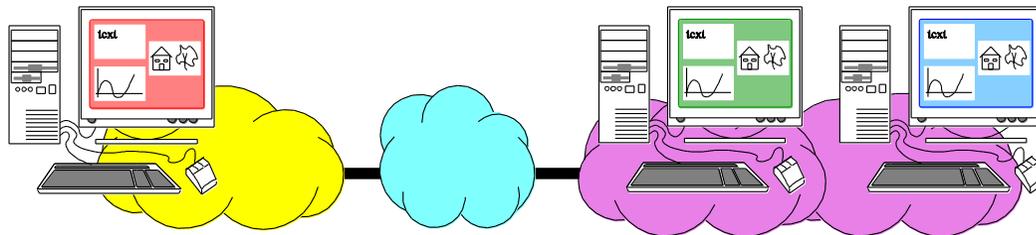
- Note: This is a requirements draft
 - ➔ Don't think too much about solutions
- However: It's hard to talk about requirements without having any common view at the problem, the terminology, and the framework
 - ➔ So ...

The following gives an **example** how ALTO could look like,
in a tracker-based P2P network

It is **NOT** an elaborated proposal for a framework architecture!
Eventually, ALTO must be more general and flexible as shown here

Terminology: How ALTO **could** work in a Tracker-based P2P network

Physical topology

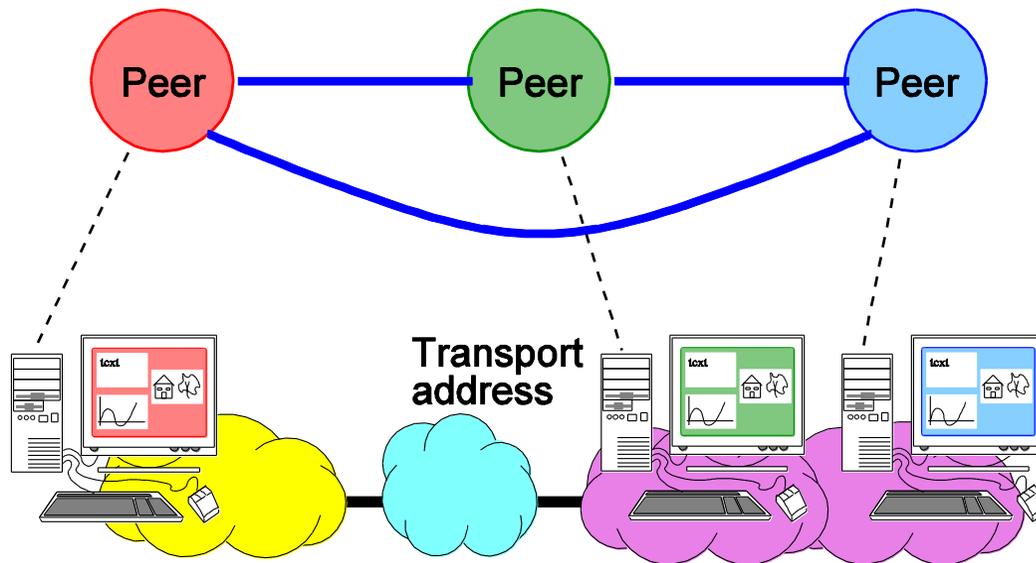


Terminology: How ALTO **could** work in a Tracker-based P2P network

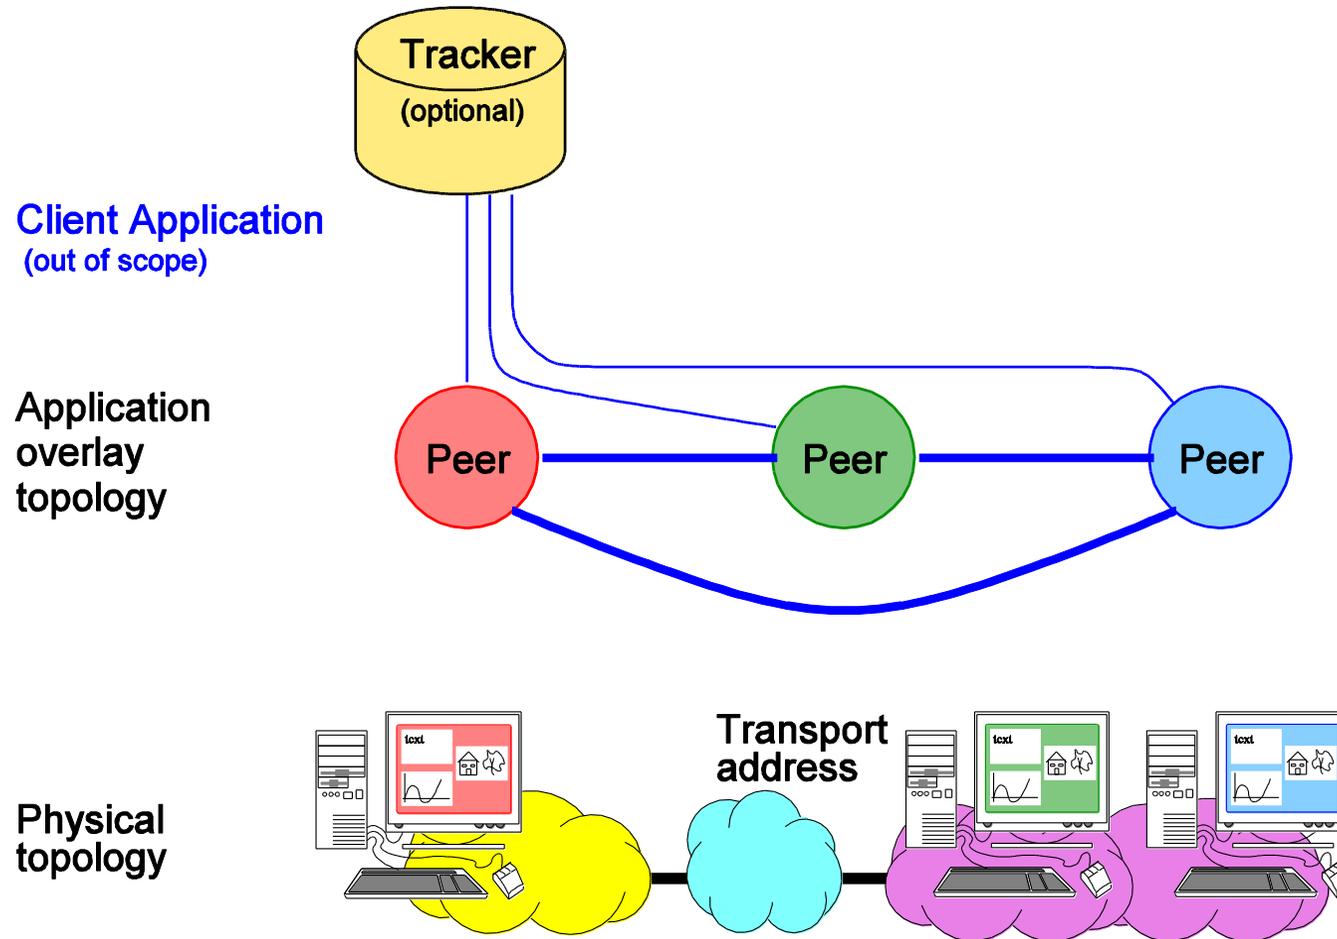
Client Application
(out of scope)

Application
overlay
topology

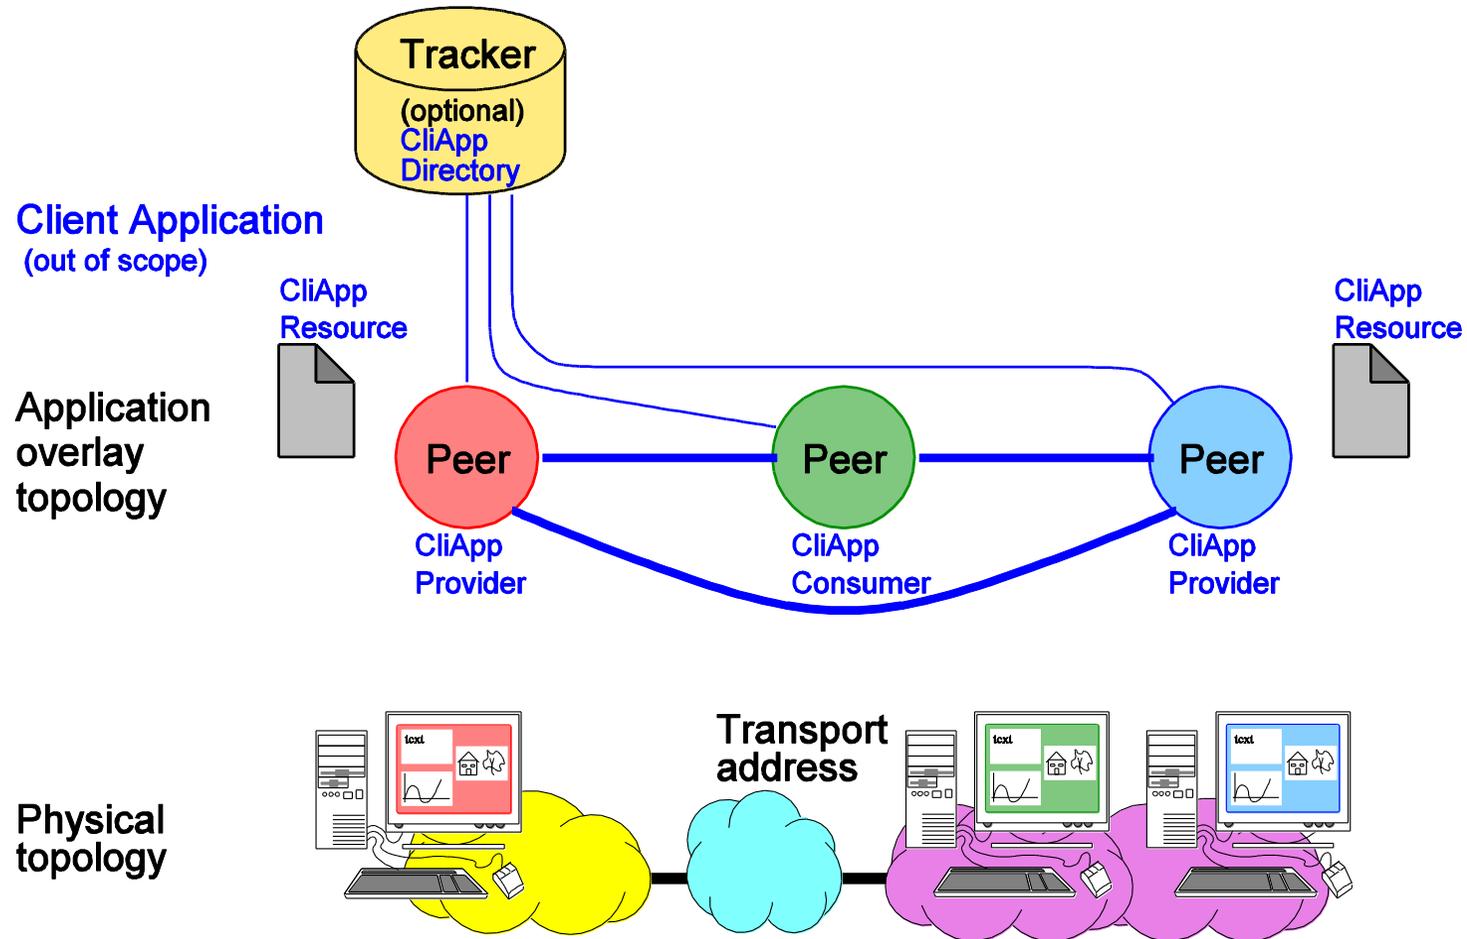
Physical
topology



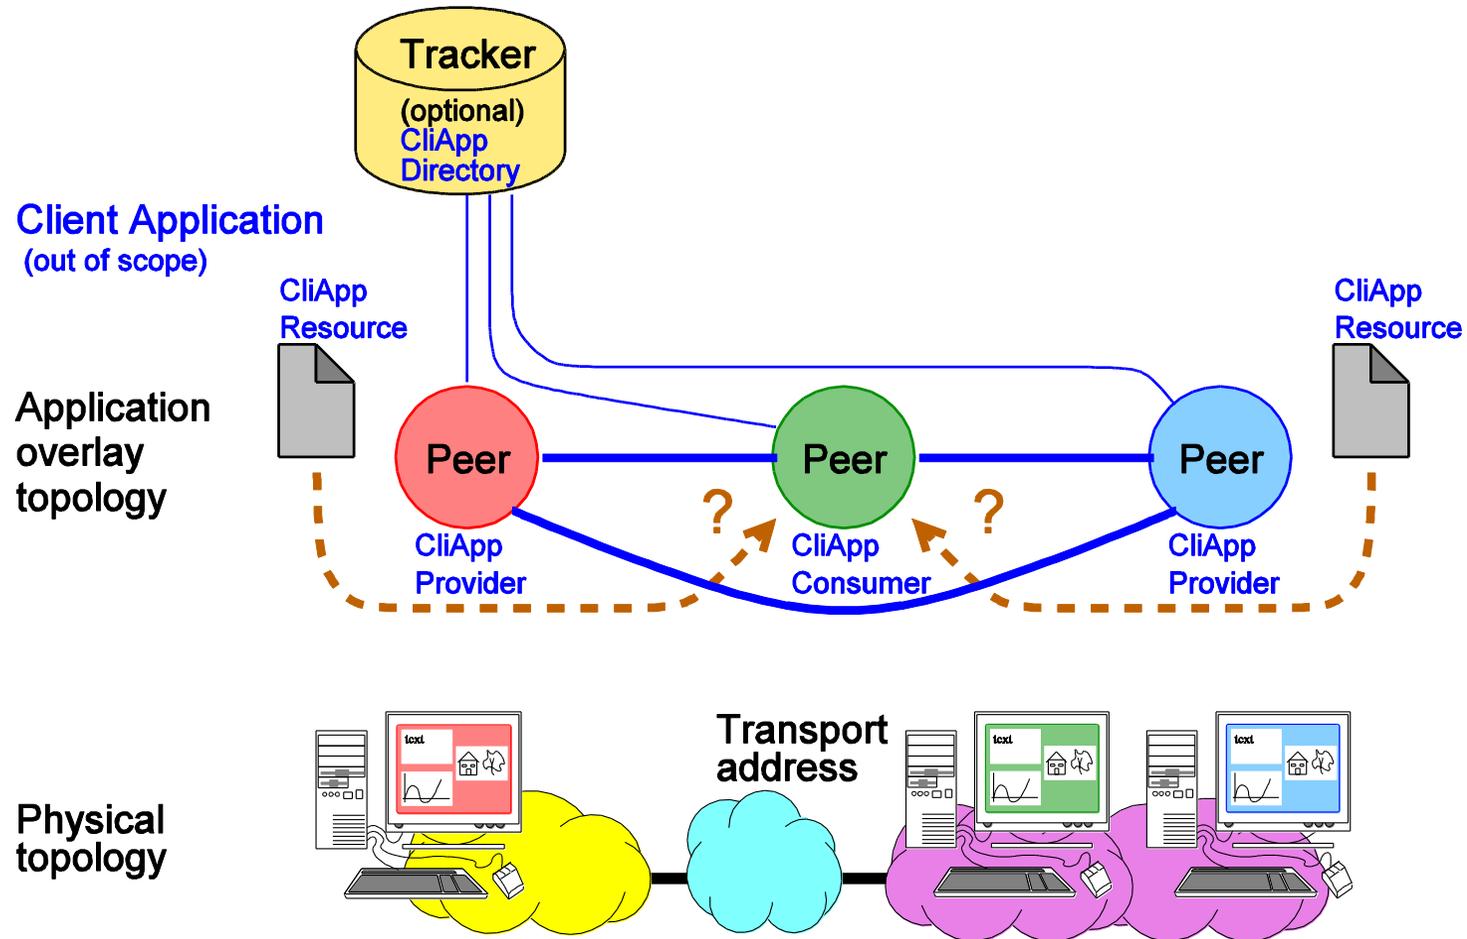
Terminology: How ALTO **could** work in a Tracker-based P2P network



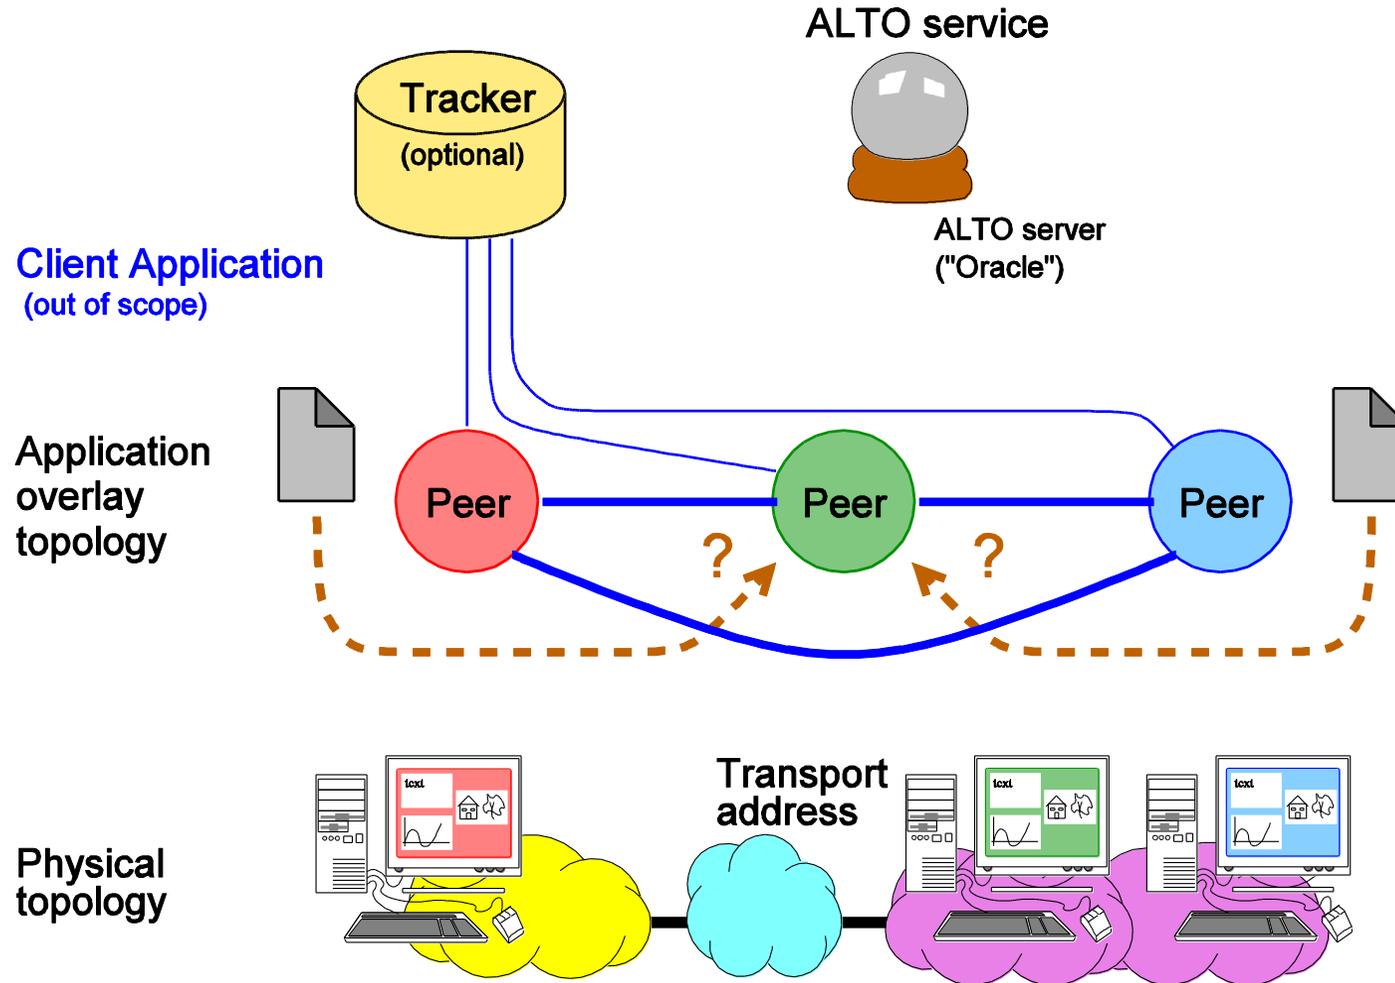
Terminology: How ALTO **could** work in a Tracker-based P2P network



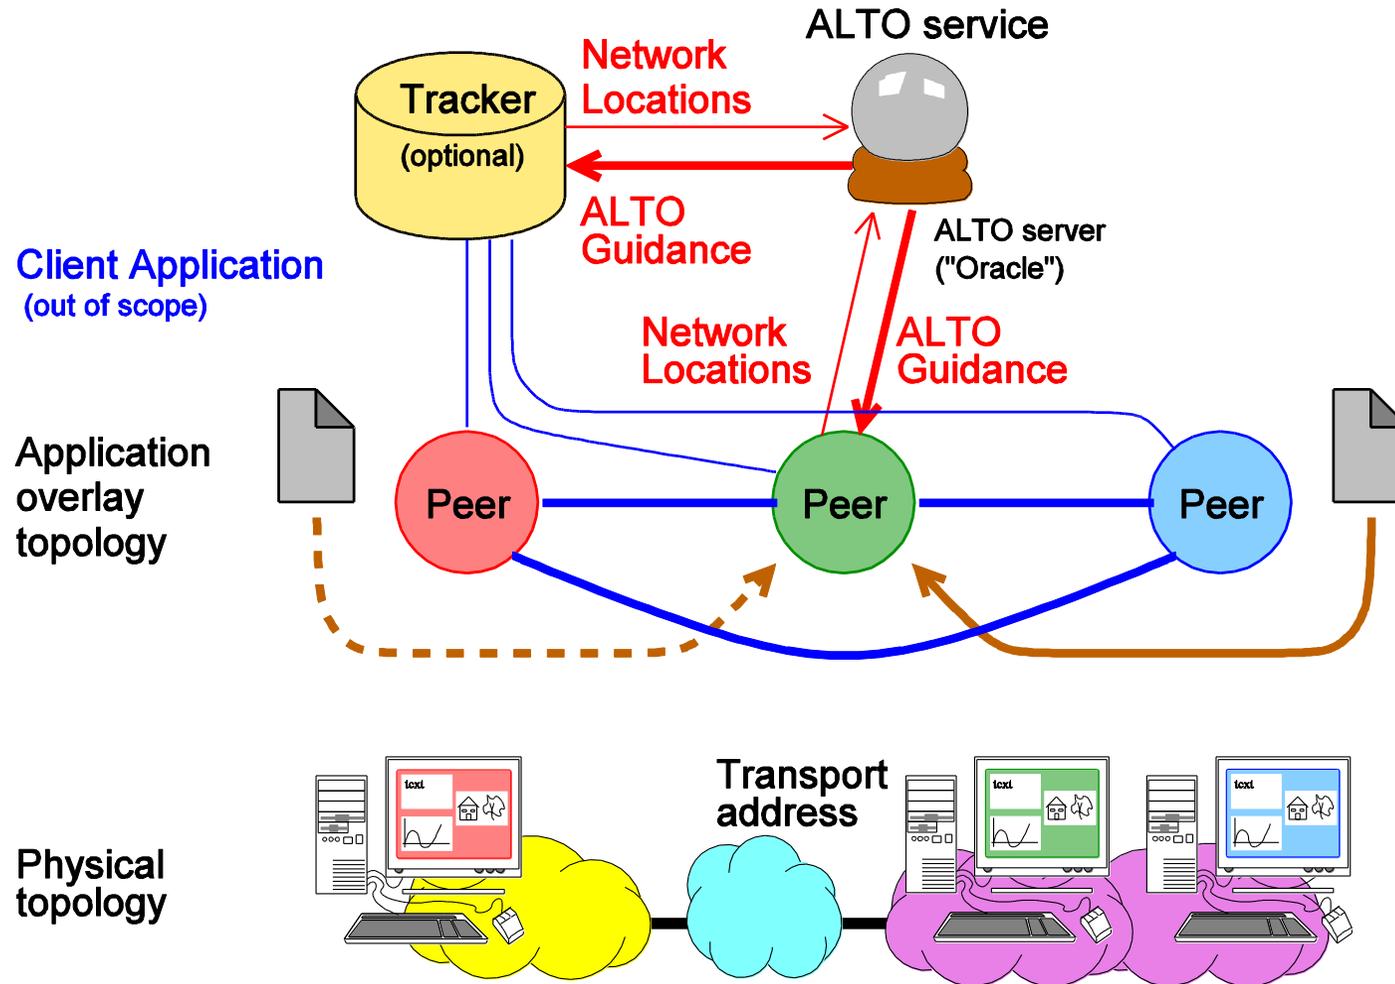
Terminology: How ALTO **could** work in a Tracker-based P2P network



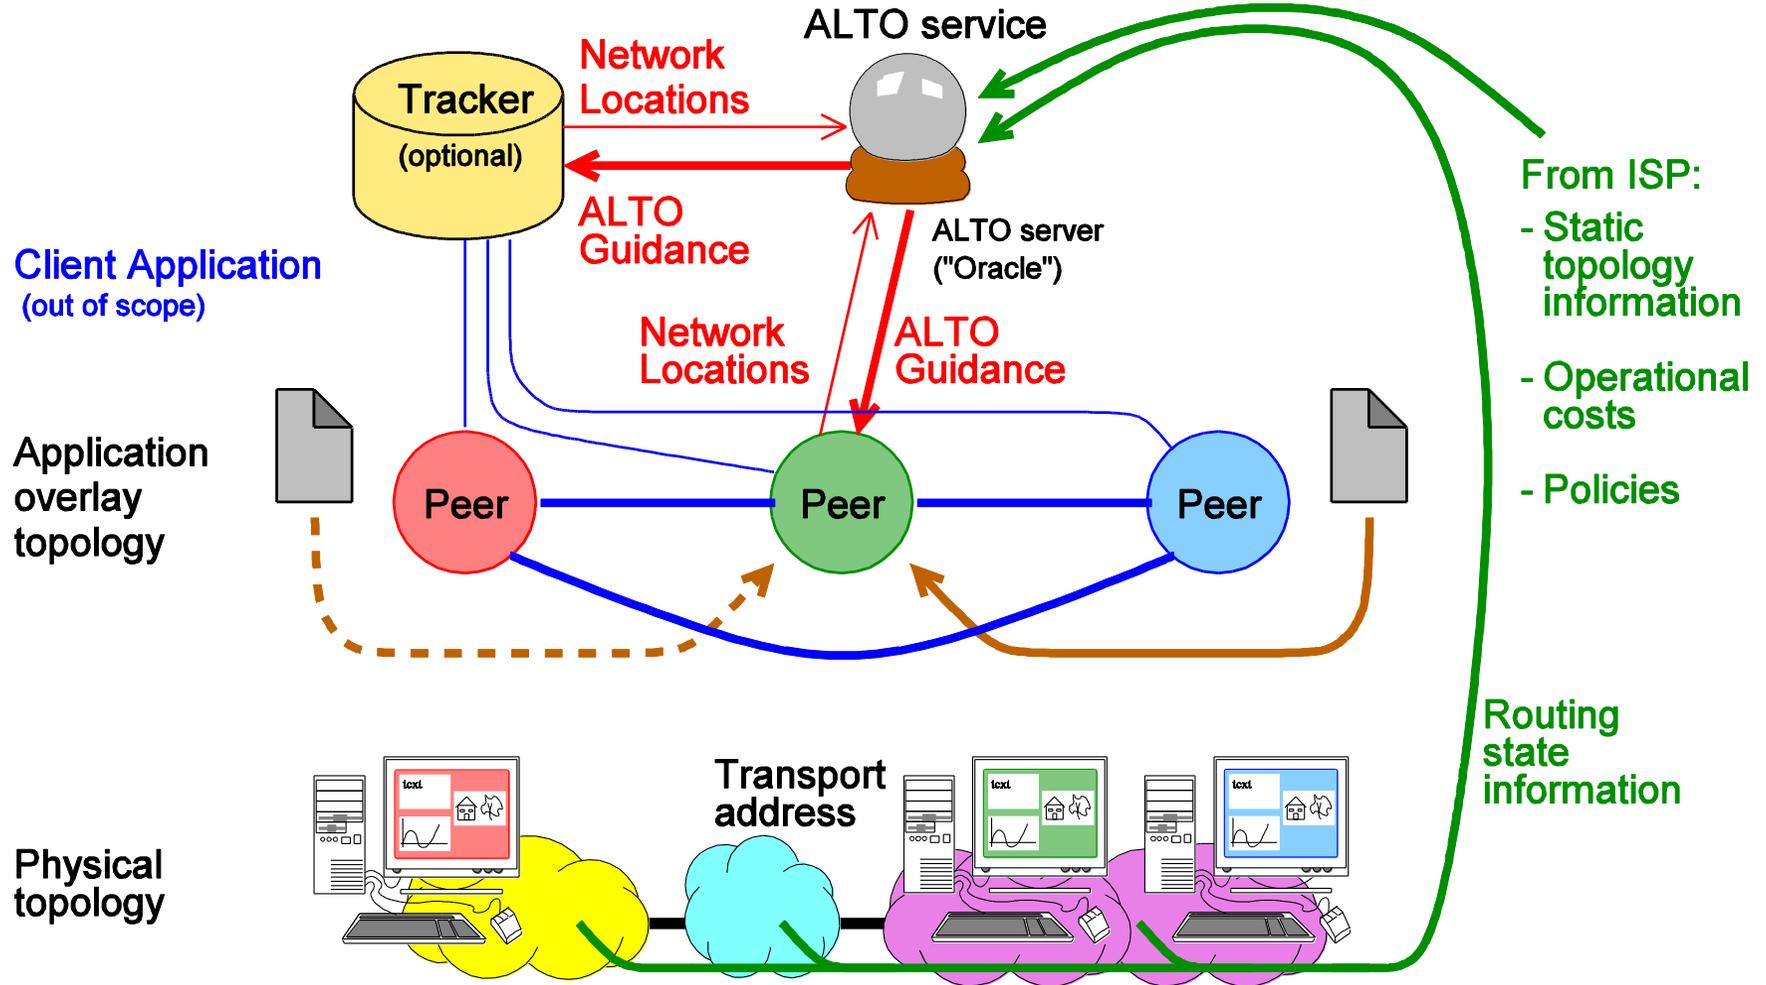
Terminology: How ALTO **could** work in a Tracker-based P2P network



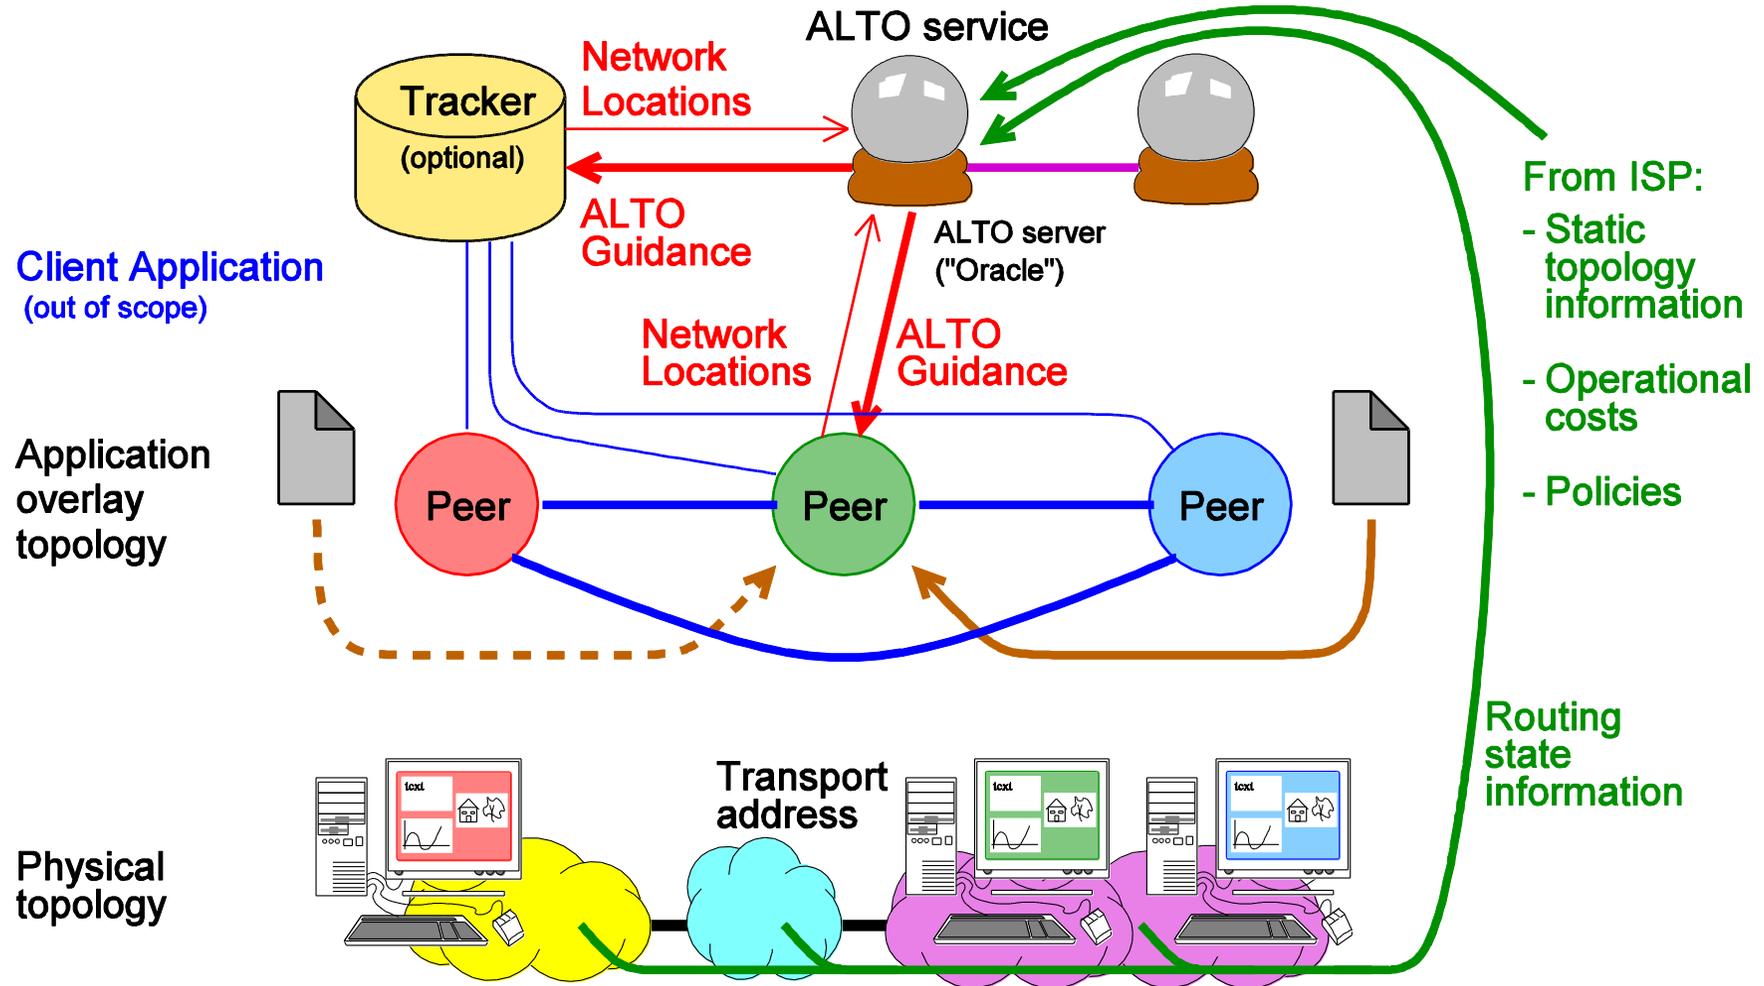
Terminology: How ALTO **could** work in a Tracker-based P2P network



Terminology: How ALTO **could** work in a Tracker-based P2P network



Terminology: How ALTO **could** work in a Tracker-based P2P network



Terminology

Still at a very early stage

- TBD: consider non-P2P use cases
- TBD: move to a separate draft
- Client Application
- Client Application Resource
- Client Application Resource Identifier
- Client Application Provider
- Client Application Consumer
- Client Application Directory
- Transport Address
- Host Location Attribute
- ALTO Service
- ALTO Server
- ALTO Client

Requirements Overview

- ALTO Interface
- Error handling and overload protection for ALTO
- Security and Privacy

Requirements: ALTO Interface

An ALTO server may have to implement several interfaces, e.g., for

- Feeding topology related information into the ALTO server
 - Maybe, we don't need a new protocol:
 - Import BGP, OSPF, etc.
 - Static configuration via config file, etc. (→ implementation specific)
- Coordination, data exchange, or referral between ALTO servers
 - Hard to decide at this stage of the discussions
- Querying the ALTO service (by clients)
 - **Let's focus on this interface / protocol for now!**

Requirements: ALTO Interface

- The ALTO protocol is a query/response protocol
- The ALTO server rates different providers for the same resource, based on criteria related to connectivity between client & provider
- It must be known that the resource is actually available at the provider (discovery), or that it is willing to fetch it for us (cache)!
- Client doesn't have to give resource provider's address, instead, a location attribute (e.g., prefix, AS number) may be sufficient
- **TBD:** is reply a sorted list of locations/addresses, or will the server return values for every location?

Example rating criteria

- Rating criteria are not specified, only examples
- ➔ Move to a separate document in the future
- ➔ Protocol must be extensible

Example criteria

- Topological distance, e.g.,
 - AS hop count
- Expected cost for transport

Criteria that **SHOULD NOT** be used for rating

- Performance metrics related to instantaneous congestion status. This has to be probed and adapted to continuously, e.g., using TCP

Requirements:

Relation to Congestion Control

REQ. 7: The information available from the ALTO service is not a replacement for congestion control mechanisms. Applications using ALTO MUST ensure that they do not cause congestion in the network, e.g., by using TCP transport, which includes congestion control mechanisms.

- This requirement will most likely not result in a protocol feature
- However, consensus among ALTO designers with respect to this issue is crucial!

Error handling and overload protection

- Client applications **MUST** also work reasonably without ALTO
- Retransmission and load control strategy
 - Simple transaction based query protocol assumed
 - Inspired by TCP's CWND
 - Explicit server overload flag
- Client may specify required level of detail
 - Resources spent at ALTO server for rating alternatives should reflect the amount of data to be exchanged or session duration
- Lifetime attributes for caching

Security and Privacy

- An ALTO provider must be able to control information disclosure (e.g., about the network topology, policies, operational costs, etc.)
- An ALTO user must be able to control information disclosure
REQ. 22: One ALTO interface should be defined in a way, that the operator of one ALTO server cannot easily deduce the resource (e.g., filename in P2P file sharing) which the requester wants to obtain.
- **TBD:** The ALTO service provider might want to inform the client, that he does not want to be told content IDs, even if that could lead to better recommendations. Do we need to support this?
- Several independent ALTO instances must be able to coexist
- Mutual authentication & authorization, encryption, DoS protection

Open Issues

- Is the ALTO service offered by servers, or is ALTO just a data format, e.g., for sending topology info from routers to “trackers”?
- Relation between resource provider discovery and selection
 - Combined or separate?
 - Query/response or asynchronous pushing of topology information
 - Which models to work on? One combined or separate protocols?
- How to avoid additional exposure of privacy sensitive or legally relevant information to ALTO operator (e.g., P2P file sharing IDs)
- Do we need special protocol elements and handling for caches, or are they just altruistic peers with good connectivity?
- How to generalize ALTO for non-P2P usage?

...AND MANY MORE!

Further steps

- Updates reflecting discussions here and on mailing list
- Adjust requirement levels (SHOULD ↔ MUST)
- Separate framework architecture and terminology draft(s)
- Separate specification and assessment of rating criteria
- Start the actual design

draft-kiesel-alto-reqs-00.txt

Appendix B. Acknowledgments

The authors would like to thank

- o Vijay K. Gurbani <vkg@alcatel-lucent.com>
- o 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 Stiernerling 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.