



Simple Economic Management Approaches of Overlay Traffic in Heterogeneous Internet Topologies

European Seventh Framework Project FP7-2008-ICT-216259-STREP

Deliverable D4.2 Final External Trial and Show Case Set-up and Assessment of Results

The SmoothIT Consortium

University of Zürich, UZH, Switzerland
DOCOMO Communications Laboratories Europe GmbH, DOCOMO, Germany
Technische Universität Darmstadt, TUD, Germany
Athens University of Economics and Business - Research Center, AUEB-RC, Greece
PrimeTel Limited, PrimeTel, Cyprus
Akademia Gorniczo-Hutnicza im. Stanisława Staszica W Krakowie, AGH, Poland
Intracom S.A. Telecom Solutions, ICOM, Greece
Julius-Maximilians Universität Würzburg, UniWue, Germany
Telefónica Investigación y Desarrollo, TID, Spain

© Copyright 2011, the Members of the SmoothIT Consortium

For more information on this document or the SmoothIT project, please contact:

Prof. Dr. Burkhard Stiller
Universität Zürich, CSG@IFI
Binzmühlestrasse 14
CH—8050 Zürich
Switzerland

Phone: +41 44 635 4355
Fax: +41 44 635 6809
E-mail: info-smoothit@smoothit.org

Document Control

Title: Final External Trial and Show Case Set-up and Assessment of Results

Type: Public

Editor(s): Sergey Kuleshov

E-mail: sergeyk@prime-tel.com

Author(s): Rafal Stankiewicz, George Stamoulis, Burkard Stiller, Simon Oechsner, Jerzy Domżał, Robert Wójcik, Ioanna Papafili, Eleni Agiatzidou, Sergey Kuleshov, Konstantin Pussep, Marcin Niemiec, Sergios Souros, Juan Fernández-Palacios, Pedro Aranda Gutierrez, Dirk Staehle

Doc ID: D4.2-v2.1.doc

AMENDMENT HISTORY

Version	Date	Author	Description/Comments
v0.0	10.11.2010	Rafal Stankiewicz, George Stamoulis, Burkard Stiller	ToC proposed with assignment of partner responsibilities for given parts
v.0.1	17.12.2010	Rafal Stankiewicz, Simon Oechsner, Jerzy Domżał, Robert Wójcik, Ioanna Papafili	Section 3 "Simulation results filled in. Content based mainly on WP4-T4.4_win_assessment_of_ETMs-v1.0.doc working document. Description on simulation scenarios and configuration as well as assessment methodology added. Related appendix created.
V0.2	19.01.2011	Eleni Agiatzidou, Ioanna Paopafili, R. Stankiewicz	Contributions from AUEB integrated, additional gaps identified and commented. References provided.
V0.3	26.01.2011	Sergey Kuleshov	Added first draft of trial section. Further work on ToC.
v0.3.1	28.01.2011	Simon Oechsner, Konstantin Pussep, Rafal Stankiewicz	Updates on simulation results to section 3 provided. Editing of section 3, some paragraphs added nad updated
V0.4	15.02.2011	Sergey Kuleshov, Rafal Stankiewicz	Integration of a new version (ToC modified during Zurich meeting) with ion 0.3.2 containing latest updates to section 3 and security section
V0.5	18.02.2011	Fabio Hecht, Martin Waldburger	Added sections 2.5 and 5.1.
V0.9	01.03.2011	All	Added new sections 7, 5.2, 4.5, 4.6, 1 updated sections other sections. Prepared for internal review.
V0.9.1	08.03.2011	George Stamoulis, Rafal Stankiewicz, Konstantin Pussep	Internal review and updates of sections 1-3.
V0.9.2	09.03.2011	Konstantin Pussep, Dirk Staehle, Sergey Kuleshov	Addressed comments in sections 1-3, updated 6.1.1.6 and 4.2.
V1.0-V1.4	10.03.2011	All	Changes. Adding of multiple conclusions and results.
V2.0	14.03.2011/ 22.03.2011	Sergey Kuleshov, Rafal Stankiewicz, Burkhard Stiller	Editorial changes and completions followed by a selected section-based review
V2.1	28.03.2011	Sergey Kuleshov	Minor editorial changes

Legal Notices

The information in this document is subject to change without notice.

The Members of the SmoothIT Consortium make no warranty of any kind with regard to this document, including, but not limited to, the implied warranties of merchantability and fitness for a particular purpose. The Members of the SmoothIT Consortium shall not be held liable for errors contained herein or direct, indirect, special, incidental or consequential damages in connection with the furnishing, performance, or use of this material.

Table of Contents

1	Executive Summary	5
2	Introduction and Overview	6
3	Technical Outcomes	8
3.1	Architecture Overview	8
3.2	Integration into an ISP Network	10
3.3	Performance and Scalability Analysis	11
3.3.1	Methodology	12
3.3.2	Performance test results and conclusions	12
3.3.3	Conclusions of the performance tests	13
4	Simulation Results	14
4.1	“Win” Assessment Methodology	14
4.1.1	Metrics of “win” for the stakeholders	14
4.1.2	The spectrum of simulation scenarios	15
4.2	Assessment of Simulation in Terms of Credibility and Adequacy	16
4.3	Summary of Simulation Results	17
4.3.1	BGP Locality	17
4.3.2	ISP-owned Peer	20
4.3.3	Inter-SIS collaboration based ETM mechanisms	22
4.3.4	QoS-aware ETM	28
4.3.5	Highly Active Peer (HAP)	29
4.3.6	Conclusions	31
5	Trial Results	32
5.1	Promotion of Trial to PrimeTel’s Customers	32
5.2	Trial-specific Modifications	32
5.2.1	Stability and streaming quality	33
5.2.2	Monitoring	34
5.2.3	Configuration, content encryption, and SIS modifications	34
5.3	Data Collection	34
5.3.1	Data aggregation and analysis	36
5.4	Results	38
5.4.1	Qualitative assessment of trial data	38
5.4.2	Performance analysis	43
5.4.3	ETM mechanisms’ comparison	47
5.4.4	External trial results vs. simulation results	48
5.4.5	Assessment of the ETM prototype	48
5.4.6	Socio-economical impact	49
6	SmoothIT in the Current Internet	50
6.1	Regulatory Issues and Network Neutrality	50
6.1.1	Consumer protection	50
6.1.2	Data protection and privacy	51
6.1.3	Intellectual property rights	51
6.1.4	Competition	53
6.2	ETM – Potential Risks	53
6.2.1	Architecture security	54
6.2.2	ETM vulnerabilities	56
6.3	SmoothIT’s Societal Embedding	57

7 Overall Assessment	58
7.1 Assessment of ETMs	61
7.1.1 BGP-Loc	61
7.1.2 Inter-SIS based extensions to locality awareness	63
7.1.3 QoS-aware ETM	67
7.1.4 IoP	68
7.1.5 Combination of BGP-Loc and IoP	69
7.1.6 HAP	69
7.1.7 Overall Triple-Win summary	70
8 Cost-benefit Analysis	73
8.1 Cost-Benefit Analysis of the SmoothIT Prototype	73
8.1.1 ETM mechanisms and cost categories	73
8.1.2 Transit prices	75
8.1.3 Internet and overlay traffic	77
8.1.4 Analysis	77
8.1.5 Relation with existing studies	83
8.2 Global View on Potential Inter-Domain Traffic Savings	83
9 Summary	88
References	89
Acknowledgements	91
Appendix 1. Sample Record for Simulation Experiment	92

1 Executive Summary

The goal of this deliverable is to present the main and major outcomes of the SmoothIT project in terms of the technology, algorithms, assessment of approach and its impact, and its overall applicability in real life environments. It is based on all results of tasks T4.3 Operation of External Trial and T4.4 Assessment of results.

Based on multiple inputs from WP2 and T4.4 this deliverable provides a detailed assessment of simulation results, including the methodology description and assessment of their credibility and adequacy. This is followed by a simulation results summary for a variety of cases, and qualitative assessment of TripleWin. The following ETM (Economic Traffic Management) mechanisms fall under this assessment: BGP-Loc, IoP, Inter-SIS based ETM, QoS-aware ETM and HAP. This constitutes a complete summary of the simulations for all ETM mechanisms studied as part of the project.

In terms of trial results the deliverable describes the modifications that were made to the system developed by WP3 and the trial environment. Definition of raw data collected in the course of the trial is given, followed by aggregation techniques. The aggregated results are presented, interpreted and their qualitative assessment is given. Based on the qualitative analysis and statistical significance, it was deduced that cross-ETM-mechanism comparison of external trial results or their comparison against previous results is not feasible. Conclusions are presented in terms of prototype assessment, which revealed that the prototype deployed in the external trial behaved as expected, providing all required functionality in real life environment without performance issues.

The deliverable also takes a complete high level view on the project and its results. It further positions SmoothIT in the current Internet ecosystem and discusses various regulatory and net-neutrality issues under European perspective. In particular the following areas are addressed: consumer protection, data protection, intellectual property, and competition. Special consideration of potential risks from the security perspective is given on the architecture employed in particular and on the proposed ETM mechanisms in general. A cost-benefit analysis was undertaken for BGP-Loc, IoP, and HAP, assessing their break-even points regarding the relative reduction of the inter-domain overlay traffic attained.

The main results of this deliverable are:

- ETM mechanisms designed in the project are feasible. While some mechanisms are more beneficial to customers and overlay and others to ISPs, different combinations of approaches can lead to a balanced Triple-Win scenario.
- Architecture and ETM realization developed were proven to be usable and stable during external trial. It was also proven that the solution is relatively easy to integrate into real life ISPs environment.
- The Cost-Benefit Analysis reveals that there are feasible break-even points for the adoption of an ETM mechanism by an ISP. More specifically, the break-even point of a small-to-medium ISP when adopting the BGP-Loc, IoP or HAP are within the expected behavior of these mechanisms, as studied in the simulation analysis and partly observed from the internal trial. The situation becomes even more beneficial when larger ISPs are considered.

2 Introduction and Overview

The project SmoothIT aims at defining, developing, and testing Economic Traffic Management (ETM) mechanisms to optimize the traffic impact of overlay applications on underlay networks in order to accomplish a Triple Win situation where network operators overlay providers, and application users' benefit from the approaches undertaken. This deliverable presents the results of Task T4.3 "Operation of External Trial" that had the objective to define the operations for the external trial, to complete the software integration of the internal test-bed, and, finally, to adapt it to the selected external trial platform utilized. The deliverable further includes the results of Task 4.4 "Assessment of Results" that assesses all results achieved with respect to the internal and the external trial. Therefore, the deliverable contains interpretations and technical details based on the outcome of both trials. For the purpose of completeness, the document also contains an overview of the simulation results obtained in WP2 "Theory and Modelling".

This document is organized in the following way:

Section 3 presents the technical outcomes of the project with respect to the prototype and external trial. At first, the developed architecture as used for the internal and external trial is presented. Then, it continues with the integration of the developed ETM architecture into the PrimeTel's network as done for the external trial. Finally, a performance and scalability analysis of the architecture with the focus on the SmoothIT Information Service (SIS) is presented.

Section 4 gives an overview of the simulation studies performed in WP2 from the point-of-view of the overall assessment. At first, the "Win" assessment methodology is presented, including the major metrics in use and the simulation scenarios. Subsequently, the credibility and adequacy of simulation studies is discussed. The section concludes with a summary of the major simulation results grouped by single ETM mechanisms under study or their combination.

Section 5 describes results of the external trial. The section provides details of trial operation, modification, and integration work undertaken, mechanisms used to collect and aggregate the data as well as qualitative analysis of the results. Since the external trial required sufficient participation of users, the actions undertaken to promote the trial to PrimeTel customers is described first. Furthermore, the trial-specific modifications to the SmoothIT prototype developed by WP3 and used in the internal trial are briefly presented. Then, the data collection for the external trial is presented including the aggregation and analysis methodology. The qualitative and quantitative results of the trial are discussed including the comparison of ETM mechanisms, the relation to the simulation results, and the assessment of the SmoothIT prototype as such, and finally the socio-economical impact of the ETM mechanisms based on the external trial results.

In Section 6 regulatory issues and potential risks of the ETM deployment are presented. The discussion includes network neutrality, data protection, and intellectual property rights, issues, as well as security and vulnerability considerations of the ETM architecture.

Section 7 presents the overall assessment of the SmoothIT project. At first, the potential of a "win" situation is presented for each considered ETM mechanism. A concise representation of the assessment results is given using so-called win-win space maps, for single ETM mechanisms and considered combinations.

Section 8 shows a cost-benefit analysis of the SmoothIT prototype for the supported ETM mechanisms. For this purpose, appropriate cost categories are introduced and major influence parameters (transit prices, ratio of overlay traffic etc.) are presented. The resulting cost-benefit analysis is accompanied with the overview of the related studies. Then, a global view of the potential savings that can be achieved by ETM deployment is given.

Finally, Section 9 summarizes the deliverable.

3 Technical Outcomes

This section presents the technical outcomes of SmoothIT. We start with the developed architecture as used for the internal and external trial and continue with the integration of the ETM architecture into the PrimeTel’s network as done for the external trial.

3.1 Architecture Overview

The prototype software used during the external trial was the one designed and developed by the SmoothIT partners, except for the P2P client application which was provided by the P2Pnext [27] FP7 project under LGPL license and was further extended and adapted according to the SmoothIT project needs. In this section we provide an overview of the prototype’s architecture and briefly describe the mechanisms that were implemented and tested throughout the trial. The prototype is the realization of the Economic Traffic Management System (ETMS), which is the framework introduced by SmoothIT in order for next-generation traffic management approaches to consider also the economic implications of the adopted policies. In SmoothIT, only the overlay traffic management is considered.

In deliverable “D3.3 – Documentation of Engineering and Implementation (Final)” [28], the final architecture of the SmoothIT prototype was included. The architecture is modular and extensible and is able to accommodate any future ETM mechanisms, on top of the existing ones. Figure 1 summarizes the architecture and the involved components that were implemented and used in the external trial. The different colors denote parts of the mechanisms that are realized by each component. Note that this is only a subset of the entire SmoothIT architecture, since not all components were implemented and ported to the trial.




Figure 1. The SmoothIT architecture and the involvement of each component to the implemented ETM mechanisms.

The ETM mechanisms implemented and deployed in the external trial are the following:

BGP-Loc (BGP-based Locality Promotion): This mechanism implements a peer ranking algorithm which sorts peer lists provided by local peers, based on network metrics and more specifically on BGP-related metrics. Thus, a peer has at its disposal both the overlay view, via the ranking performed by the overlay application using overlay-specific metrics, and the underlay view, via the ranking provided by the ISP.

IoP (ISP-owned Peer): This mechanism introduces a new peer into the overlay, which belongs and is controlled by the ISP, has enhanced resources (i.e., high access speed and big storage space) and acts as a combination of a super-peer and a smart cache. Depending on the popularity of content among the peers of the local domain, the IoP joins the respective swarms, downloads the content, and starts serving the local peers so as to reduce the unnecessary usage of inter-domain network resources.

HAP (Highly Active Peer): This mechanism is a conceptually similar to the IoP but achieves the same objectives in a distributed manner. Instead of having a central ISP-controlled peer that tries to improve the performance of local peers of a given swarm, the HAP mechanism distributes this responsibility to many regular local peers by assigning them temporarily extra network resources. The decision of which peers should be promoted to HAPs is taken by the ISP and is based by the overlay behaviour of the peers, the swarms they join, the “compatibility” of their actions with the ISP’s goals, etc.

The aforementioned mechanisms are realized through the involvement of certain entities. Some of them already exist in the ecosystem of overlay and underlay networks while other were introduced by the SmoothIT solutions. Some of the existing entities were adapted to host these new mechanisms, some others were left intact. These entities are:

SIS (SmoothIT Information Service): this is the central entity of the SmoothIT architecture. It incorporates a number of components that implement the various functionalities required by the different ETM mechanisms. It is controlled by the ISP and is considered as the means through which the ISP can affect the organization and formation of the overlay networks.

Peer: the standard peer, enhanced with extra functionalities to support the three ETM mechanisms, mainly by providing additional information to the SIS regarding the information about overlay usage.

IoP: the overlay cache that acts as a super peer that acquires quickly and offers continuously the desired content due to its extended network and storage resources.

HAP: a standard peer that is promoted to HAP by being granted with additional network resources, as a result of its ISP-friendly overlay behaviour.

NMS (Network Management System): the interface for the underlay. The NMS provides (mainly to the SIS) all the underlay information required for the different ETM mechanisms to run. This information varies from BGP data to user access profiles.

As already mentioned, the SIS is the central entity of the SmoothIT architecture. It supports all three implemented ETM mechanisms and can support future ones, due to its modular design. It consists of several components that are assigned with different parts of the ETM mechanisms’ functionality. The list and description of all implemented components forming the SIS, are included below:

The Controller determines the core of the system. Its main responsibility is to coordinate each ETM mechanism. Therefore, it receives requests from the overlay application (simple peer, IoP, or HAP), performs calculations based on the ETM mechanism deployed and according to several underlay metrics, such as metering and policy information, and returns to the overlay application the information required. Usually, the information returned is in the form of ranked lists, rating either peers as possible “good” neighbors, or swarms as “popular” ones.

The Metering component collects network information from the underlying NMS in order to support ETM mechanisms implemented by the SIS. This information includes BGP (Border Gateway Protocol) routing tables in order to support locality enforcement algorithms.

The SIS DB (Data Base) is a repository storing all information that may be useful for all modules, such as configuration parameters or business logic.

The Admin component is used by the administrator of the SIS to access and configure the server as well as the deployed ETM mechanisms.

The QoS Manager interacts with the NMS to retrieve access profiles for the ISP’s customers and provide a list of peers that should be promoted to or demoted from HAPs.

The Monitor component is a supportive component used to gather certain overlay and underlay statistics for evaluation purposes.

The Security component provides security services to the SIS, including authentication, access control, and secure communication.

3.2 Integration into an ISP Network

This section describes the integration process into ISP network based on PrimeTel’s experience acquired during the setup of the external trial.

Integration of SmoothIT solution into PrimeTel network proved to be reasonably straightforward thanks to small number of interconnection points and clear APIs. Figure 2 presents how this integration was made.

The SIS to BGP interface, required by the metering component, was implemented using an intermediate BGP server based on Quaga software package [29] that provided full view of the PrimeTel’s main BGP table and exposed it over SNMP protocol. This was required as SNMP protocol generally puts high load due to large amount of BGP full view table. Thus, applying the load on production routers might have had undesired consequences. No other modifications or adaptations were made for this interface.

The deployment of the HAP ETM mechanism required two additional interfaces to be implemented by PrimeTel between SIS Server and Billing & Provisioning System and Traffic shaper. The first interface is a SOAP based web service that exposes information about contracted customer bandwidth profiles and acts as a bridge between Billing System and SIS. It responds to simple queries with one or more IP addresses by returning the bandwidth profile of associated users. The second interface is implemented by SIS to expose the HAP list via SOAP. This list is consumed by traffic shaping configuration solution at regular intervals to alter customer profiles according to HAP ETM approach.




Figure 2. SIS integration — The PrimeTel case.

The deployment of IoP ETM mechanism required no changes neither to PrimeTel's infrastructure nor the SmoothIT release code.

3.3 Performance and Scalability Analysis

The main goal of this section is to analyze the performance capabilities of the SIS. Even though the SW is in a prototype stage, first estimates of its performance are important for the deployment in PrimeTel's network and in other real networks in a near future.

Moreover, in order to assess the results that are reported as part of the internal trial (see [20] for more details), it must be assured that all the tests are carried out while the SIS is working properly.

Considering these premises, the following targets should be covered in the scalability tests:

1. To detect and fix bugs in the system implementation.
2. To detect the presence of possible bottlenecks in the implementation of the SIS system and, if available, to test new versions of the SW modules that solve them.
3. To validate the scalability of the SIS architecture design: to provide an estimation of the current implementation working limits and to validate the usage of the selected protocols.
4. To provide a set of recommendations that should be taken into account for a future commercial implementation/deployment of the SIS system.

3.3.1 Methodology

The core steps in the methodology applied were done as described below.

Scenario description

In order to test the SIS performance capabilities, the scenario shown in Figure 3 has been set up.




Figure 3. Scenario for performance tests.

The following machines can be distinguished:

- *mamen* runs the tool for generating several requests to the SIS.
- *modelnet4* runs an instance of the SIS server that receives all the requests coming from the load generator tool.

Measured parameters

- **Response latency** is defined as the delay from the time a request is sent to the SIS to the time the response is received at the client side. For this response latency, only the capabilities of the SIS server were evaluated. The delay associated to the network itself was ignored. The Average, the 95th Percentile, and the 99th Percentile were calculated from the measured latency values.
- **Blocking Probability:** the SIS should not block any request, so that an occurrence of blocked requests would give an estimation of the failure point of the SIS.

Moreover, in order to monitor the causes of potential performance issues of the SIS server, the CPU usage and the memory usage in the SIS machine were monitored.

3.3.2 Performance test results and conclusions

The performance objective of the final version of the SmoothIT prototype was defined considering the following setup:

- The total number of simultaneous clients was 200.

- According to the current implementation of the Tribler client, every 10 seconds the client sends 1 request for rating the IP addresses. This results in an average rate of 0.1 requests/second per client.
- The system should be able to work for at least 3 days without interruption.

Therefore, the performance objective for the SIS final prototype is to support 20 requests/second without decreasing its performance and run during at least 72 hours without a reboot.

Other important parameters that have been considered are the following:

- The number of IP addresses sent in the request was set to 50, which is the typical size of the list returned by real trackers.
- The SIS manages the information associated to 10 ASes, which are supposed to have any kind of interconnection agreement. Any further ASes is considered as the rest of Internet and their IP addresses will be rated as 0.

The duration of the tests is 120 min, the results can be seen in Figure 4.




Figure 4. SIS response time for a test duration of 120 min.

As can be seen the response time of the system increases linearly when the number of requests per second is above 17 requests/second. This is due to the high usage of the memory, whose usage is almost 100% in a machine which has 2 GB of memory.

3.3.3 Conclusions of the performance tests

Considering that the target performance objectives for the SIS implementation are: (i) to support 20 requests/second and (ii) to run without failures during 72 hours, we conclude that, according to the trends shown in the different tests, the SmoothIT architecture/implementation scales appropriately.

4 Simulation Results

Extensive simulation studies of the proposed ETM mechanisms were performed in WP2 with detailed results to be found in [2]. Each proposed mechanism was verified by simulations, while the most promising solutions were selected and then implemented and deployed in the external trial. The main goal of simulations was to evaluate how all three stakeholders (end users, overlay providers, and ISPs) would benefit from ETM deployment and under which circumstances “TripleWin” arises. To attain this, an innovative methodology of assessment for the ETM mechanisms has been developed and applied. The two main goals for assessing simulation results of ETM were:

- To assess if a “win” situation for each of the three players was achieved (or not). For this purpose considered ETM mechanisms were evaluated under various network and application level conditions to support the selection of ETM mechanisms for the implementation.
- To assess the accuracy of experiments and their credibility.

The following simulators were used:

- Protopeer, an application layer oriented tool for evaluating P2P applications. Based on the existing simulator core, SmoothIT partners developed new application-layer protocols (BitTorrent and Tribler) and appropriate underlay models. Thus, this tool was modified and further enriched with ETM mechanisms for SmoothIT purposes and used for most simulations.
- Eruption.S [5], an AGH proprietary P2P simulator with elements of transport and IP layers implemented. It was necessary for experiments of inter-SIS collaboration under topologies with route asymmetry.
- A custom simulator was used to assess the potential of HAP ETM mechanism [23]. Due to the long-term effects of HAP promotion (where the promotion happens at the range of days) and the requirement to assess the impact on many concurrent distribution swarms a higher level simulation model was required.

4.1 “Win” Assessment Methodology

The assessment methodology did not consider the optimization of a “total cost” metric. Instead, separate objective metrics were used for each player to reflect their diverse requirements and related incentives to employ an ETM mechanism. Several evaluation scenarios have been defined to cover different degree of ETM deployment, popularity of ETM among end-users, various swarm sizes, peer distribution among network domains, network topologies etc. Additionally, a game-theoretic analysis was applied to study interactions of ISPs, regarding decisions whether or not to employ an ETM mechanism and the associated ISPs’ interactions by taking into account the end-user benefit as well.

4.1.1 Metrics of “win” for the stakeholders

From the ISP’s point of view, the ultimate confirmation of “win” is a monetary benefit from ETM deployment. It is, however, not possible to quantify this benefit, since many factors contribute to the overall balance, including deployment and operational costs for an ETM

mechanism, cost of maintenance of IoP or SIS, savings resulting from inter-domain traffic reduction and the structure of interconnection tariffs, business models, marketing factors, etc. Thus, the assessment methodology focused on another quantifiable metric, namely inter-domain traffic reduction, since ISPs benefit mostly thereby. In simulation experiments, the inter-domain traffic reduction was measured directly. Upstream and downstream traffic were evaluated separately, although in certain cases focus was on the upstream only, since there are interconnection tariffs based only on it [19]. A mapping of these benefits in terms of traffic savings to monetary saving is provided in Section 8 where a cost-benefit analysis of different ETM mechanisms is presented.

QoE or QoS metrics are used for measuring the “win” for end-users. For file-sharing P2P applications the most important perceivable parameter is download time (or download speed). It can be strongly influenced by ETM mechanisms, both favorably and adversely. Users will adopt a given mechanism if this improves or, at least, preserves their download times. Ideally, this should be guaranteed on a per-user basis as was done in a certain cases using coupled simulations. However, in most of the cases, it was analyzed by comparing the average values of the metrics with and without an ETM mechanism. Such averages are taken over all peers in a swarm, or over subsets (e.g. those that belong to the same AS). On the other hand, main QoE metrics associated with video streaming applications taken into account in assessment of “win” are the probability of playback stalling, the actual stalling time, and start-up delay. These can be influenced by ETM mechanisms as well.

Assessment of “win” for overlay providers is based on the content availability in the whole overlay (swarm). Another dimension of an overlay provider’s “win” is a decreased traffic volume from its own content servers and reduced load of the servers, as well as an improved performance of the application, which should translate into increased popularity of the service. In reality, these issues often coincide with the objectives of the end-users.. Thus, simulation based evaluation of ETM mechanisms was focused either on assessing win-win for the ISP and the end-users or on the win-win for the ISP and the overlay provider.

4.1.2 The spectrum of simulation scenarios

To obtain a reliable assessment of ETM mechanisms several evaluation scenarios were defined. Depending on the simulation goal simplifications of various extents were assumed. Various network conditions, application behavior etc. were simulated. The simulation scenarios were constructed using:

a) Various network topologies:

- Triangle-shaped: three ASes connected with direct links [2],
- Star-shaped: several ASes connected to a single transit AS (symmetric simulation scenario [3]),
- “Bike”-shaped (complex simulation scenario [4]),
- Hierarchical: with tier 1, tier 2 and tier 3 ASes. Topology mainly used to evaluate inter-SIS collaboration [1],
- Semi-real: a set of ASes with real inter-domain connections and routing was extracted from the Internet topology. This topology reflects also a route asymmetry was used for evaluation of inter-SIS mechanisms [2].

b) Different peer distributions:

- Semi-homogeneous: physical peers belonging to a single swarm are distributed equally between ASes. Since peers switch on and off with some probability, a momentary distribution slightly is semi-homogeneous ,
- Heterogeneous: peers belonging to a single swarm are not equally distributed among ASes. Peer distribution was either selected randomly (e.g., from hyperbolic or Zifp distribution) or arbitrarily, e.g., 10% of physical peers in AS1, 30% in AS2, etc. Heterogeneous distribution allowed for evaluation the impact of ETM on ASes where only a small fraction of peers is located as well as AS having a substantial part of a swarm. Thus, SmoothIT evaluated the potential benefit for peers located in ASes of different sizes as well as ISPs hosting different number of peers.

c) Different swarm sizes: swarms of a few, dozens and hundreds of peers.

d) Varied ETM deployment degree: an interaction between peers located in ASes with ETM deployed and peers located in ASes without ETM was evaluated. The effect on “win” for ISP and end-user was assessed, and in case of partial deployment separately for the adopters and non-adopters of ETM. Three basic scenarios were considered:

- Only one ISP implements ETM
- All ISPs implement ETM
- A subset of ISPs implement ETM

The goal of simulating different deployment degrees was also to answer the question under which conditions it is beneficial for the ISP to implement ETM, and how it is influenced by the actions of other ISPs. This research was also accompanied by game theoretical analysis [2].

e) Varied end-user interest in adoption of ETM: coexistence of users employing ETM and ones declining the usage of ETM within a single swarm

f) Various type of content/application:

- Files (file-sharing application)
- Video (p2p streaming)

4.2 Assessment of Simulation in Terms of Credibility and Adequacy

All simulation experiments were carefully planned and reported. The experiment reports were collected in a unified form. A sample record for a single simulation experiment is provided in Appendix 1.

First, for each simulation scenario a set of objectives, assumptions and expected results were defined. Then, after performing the simulation and processing the results the following questions were answered:

- Were these objectives achieved?
- Were these assumptions correct or reasonable? Was it necessary to modify the assumptions during the work progress?
- Was the scenario representative for a real situation?

- Did the results meet expectations and can the deviations be explained?
- Which improvements are provided by a particular ETM mechanism? Was a win-win situation achieved?

The answers for these questions were the basis for deciding whether to accept or discard the outcome of a certain experiment. In some cases, the decision on modification of assumptions or simulation scenario was made on the basis of the analysis of the results and their credibility. The lessons from each simulation run were used to better plan consecutive experiments. Finally, conclusions on “win” evaluation and ETM performance were drawn.

A careful analysis of statistical credibility of the results was performed. For this reason the methodology of experiments, including the statistical data collection methods, was planned, analyzed, and reported for each experiment. A warm-up period was used in simulations. Most of the measured output values are random variables and the result typically presents a mean value taken from several independent runs, together with confidence intervals.

4.3 Summary of Simulation Results

This section summarizes simulation results. Extended descriptions together with detailed description of results can be found in deliverable D2.6 [4]. This summary focuses on “win” evaluation of ETM mechanisms. In the subsequent subsections results for the following ETMs are presented: BGP Locality, ISP owned Peer, Inter-SIS collaboration-based solutions, QoS-aware ETM, and Highly Active Peer. Also certain combinations of ETM mechanisms deployed together are considered.

Summaries of simulation results are presented in compact tabular form. Table headers describe a basic scenario (evaluated ETM mechanism and some supplementary information specifying the scenario, if applicable). Table rows present specific cases. A brief summary of simulation setup (a topology, peer distribution, etc.) and conclusions related to the evaluation of the “win” situation is provided for each case.

4.3.1 BGP Locality

Table 1. BGP Locality

ETM scenario	BGP-Loc (BNS&BU) [2]
Case	<p>Homogeneous peer distribution, homogeneous access bandwidth, access is bottleneck (16 Mbit/s downlink, 1 Mbit/s uplink).</p> <p>Bike20 and Star20 topologies (20 stub ASes containing the peers in both).</p> <p>Poisson arrival process with a default $\lambda = 1/10s$, average swarm size 120-200 peers.</p> <p>Tribler and BitTorrent overlays under different load conditions (varied seeding time), with different (homogeneous) swarm distributions (number of stub-ASes 10, 20 and 40) and with different shares of locality-promoting peers (0, 25, 50, 75 and 100%) (BitTorrent only).</p>

Results	<p>Win for the ISP (inter-AS traffic reduced), larger relative traffic reduction for</p> <ul style="list-style-type: none"> • shorter seeding times, • a higher share of peers per AS, • a larger share of locality-promoting peers. <p>No-lose for the users (average DL times remain the same as in regBT).</p>
Case	<p>Homogeneous peer distribution, homogeneous access bandwidth, core bottleneck (3 Mbit/s on inter-AS links).</p> <p>Bike20 and Star20 topologies (20 stub ASes containing the peers in both).</p> <p>Poisson arrival process with a default $\lambda = 1/10s$.</p> <p>Tribler and BitTorrent overlays.</p>
Results	<p>Win for the ISP (inter-AS traffic reduced, but less than in previous scenario).</p> <p>Win for the users (average DL and stalling times decrease in comp. to regBT).</p>
Case	<p>Heterogeneous peer distribution (hyperbolic), homogeneous access bandwidth, access is bottleneck (16 Mbit/s downlink, 1 Mbit/s uplink).</p> <p>Star20 topology (20 stub ASes containing the peers).</p> <p>Poisson arrival process with a default $\lambda = 1/10s$.</p> <p>Tribler and BitTorrent overlays, all peers promote locality in BU, BNS and BNSBU cases.</p>
Results	<p>Win for large ISPs (inter-AS traffic reduced), small win/no lose for small ISPs.</p> <p>Win for the users in large ISPs (average DL times decrease, no significant effect on stalling times).</p> <p>Lose for users in small ISPs (average DL times increase, no significant effect on stalling times).</p>
Case	<p>Heterogeneous peer distribution (hyperbolic), homogeneous access bandwidth, access is bottleneck (16 Mbit/s downlink, 1 Mbit/s uplink).</p> <p>Star20 topology (20 stub ASes containing the peers).</p> <p>Poisson arrival process with a default $\lambda = 1/10s$.</p> <p>BitTorrent overlay, only peers in a subset of ASes promote locality (single ASes or top X ASes).</p>
Results	<p>Win for ISPs whose peers promote locality (inter-AS traffic reduced), larger savings for larger ASes and if more ASes promote locality.</p> <p>No-lose for ISPs not promoting locality (no significant change in inter-AS traffic).</p>

	<p>Win for the users promoting locality (average DL times decrease).</p> <p>Lose for users which do not promote locality (average DL times increase), larger increase in DL times if more users/users in larger ASes promote locality.</p>
Case	<p>Homogeneous peer distribution, heterogeneous access bandwidth (each AS one class), access is bottleneck (class 1: 16 Mbit/s downlink, 1 Mbit/s uplink, class 2: 4 Mbit/s downlink, 256 kBit/s uplink).</p> <p>Star20 topology (20 stub ASes containing the peers).</p> <p>Poisson arrival process with a default $\lambda = 1/10$s.</p> <p>BitTorrent overlay.</p>
Results	<p>Win for ISPs (larger traffic savings for fast ISPs).</p> <p>Win for users in fast ISPs (shorter average DL times).</p> <p>Lose for users in slow ISPs (longer average DL times).</p> <p>Lose for swarm as a whole (longer average DL times over all peers).</p>
Case	<p>Homogeneous peer distribution, heterogeneous access bandwidth (each AS all classes), access is bottleneck (class 1: 16 Mbit/s downlink, 1 Mbit/s uplink, class 2: 4 Mbit/s downlink, 256 kBit/s uplink).</p> <p>Star20 topology (20 stub ASes containing the peers).</p> <p>Poisson arrival process with a default $\lambda = 1/10$s.</p> <p>BitTorrent overlay.</p>
Results	<p>Win for ISPs (traffic savings for ISPs).</p> <p>Win for slow-access users (shorter average DL times than in regBT).</p> <p>Lose for fast-access users (longer average DL times than in regBT).</p>
Case	<p>Heterogeneous peer distribution (hyperbolic), homogeneous access bandwidth, access is bottleneck (16 Mbit/s downlink, 1 Mbit/s uplink).</p> <p>Star20 topology (20 stub ASes containing the peers).</p> <p>Poisson arrival process with a default $\lambda = 1/10$s.</p> <p>BitTorrent overlay, locality parameters l_{BNS} and l_{BU} varied ($l = 0.5, 0.9$ (default)).</p>
Results	<p>Win for ISPs smaller for $l_x = 0.5$, this effect is stronger for larger ISPs.</p> <p>Win for the users in large ISPs smaller for $l_x = 0.5$ consequently.</p> <p>Lose for users in small ISPs also smaller, all in all, more fairness in the DL times for $l_x = 0.5$.</p>

4.3.2 ISP-owned Peer

Table 2. IoP (Unchoking policy)

ETM scenario	Insertion of ISP-owned Peer (IoP) and application of the Unchoking Policy - [2]
Case	Topology of 2 ASes inter-connected via a hub AS, homogeneous peer distribution, homogeneous access bandwidth, access is bottleneck, IoP inserted in one AS without policy, single-swarm.
Results	Win for the ISP (significant incoming inter-AS traffic reduction accompanied with outgoing traffic increase compared to the case where no IoP is inserted in any of the ASes). Win for the end-users (average download times reduction compared to the case where no IoP is inserted in any of the ASes).
Case	Homogeneous peer distribution, homogeneous access bandwidth, access is bottleneck, IoP inserted in one AS with policy, Unchoking policy employed (restrictive policy regarding the location of the peers; only local peers are unchoked by the IoP), single-swarm.
Results	Win for the ISP (significant incoming inter-AS traffic reduction as when the IoP does not employ a policy, and also outgoing traffic reduction due to the application of the policy especially compared to the case where no policy applies). Win for the users (significant average download times reduction compared to the case where no IoP is inserted, though slight deterioration compared to the case where no policy applies).

Table 3. IoP (Swarm selection)

ETM scenario	Insertion of ISP-owned Peer (IoP) employing Swarm Selection (SwS) based on overlay factors - [2]
Case	Topology of 2 ASes inter-connected via a hub AS, homogeneous access bandwidth, access is bottleneck, IoP inserted in one AS without any policy, 2-swarms on the overlay, overlay factors on which SwS is based are the file size, the mean inter-arrival time of the leechers, and the mean seeding time of the peers that have become seeds; note that we tune these overlay parameters, one parameter at each time , only for peers of one swarm that are located though in both ASes.
Results	Win for the ISP (higher incoming inter-AS traffic improvements are observed when the IoP selects to join the swarm with either the lower mean inter-arrival time, either the lower mean seeding time, or that serves the largest content file; generally the swarm with the higher capacity needs). Win for the end-users (average download times are always reduced due to the IoP insertion; however reduction is higher for those peers that participate in the swarm that the IoP has joined at each time).

Case	Topology of 2 ASes inter-connected via a hub AS, homogeneous access bandwidth, access is bottleneck, IoP inserted in one AS without any policy, 2-swarms on the overlay, overlay factors on which SwS is based are the file size, the mean inter-arrival time of the leechers, and the mean seeding time of the peers that have become seeds; note that we tune these overlay parameters <i>in pairs</i> only for peers of one swarm that are located though in both ASes.
Results	Win for the ISP (incoming inter-AS traffic is reduced always; however larger improvements are noticed when the IoP selects to join the swarm with the lower mean inter-arrival time,). Win for the end-users (average download times are always reduced due to the IoP insertion; however improvement is larger for those peers that participate in the swarm with the lower mean seeding time).

Table 4. IoP: Bandwidth allocation simulation results

ETM scenario	Insertion of ISP-owned Peer (IoP) employing different Bandwidth Allocation (BwA) strategies - [2]
Case	2 BitTorrent swarms in parallel, distributed over 2 ASes (local As and rest of the world), heterogeneous peer distribution for one swarm (5%, 10%, 15%, 30% in local AS), homogeneous for the other (50% in local AS). Access is bottleneck (16 Mbit/s downlink, 1 MBit/s uplink). Poisson arrival process per swarm with a default $\lambda = 1/10s$, 120 peers on average per swarm. BitTorrent overlay, IoP capacity in total 6 MBit/s, to be distributed among the peers in the local AS in both swarms with different strategies.
Results	Win for ISP, larger traffic savings with a strategy preferring smaller local swarms (Min and Inverse proportional strategies). Win/No-lose for users (DL time reduction for users in preferred swarm).

Table 5. BGP-Loc and IoP

ETM scenario	Combination BGP-Loc + IoP [2]
Case	Heterogeneous peer distribution (hyperbolic), homogeneous access bandwidth, access is bottleneck (16 Mbit/s downlink, 1 MBit/s uplink). Star20 topology (20 stub ASes containing the peers). Poisson arrival process with a default $\lambda = 1/10s$. BitTorrent overlay with and without locality (BGP-Loc), IoP capacity per AS 1 MBit/s and 2 MBit/s.
Results	Win for smaller ISPs (in case IoP costs are lower than cost savings, higher influence of IoP in smaller ASes).

	<p>Lose for larger ISPs in case of combination of BGP-Loc and IoP (incoming traffic increase).</p> <p>Win for swarm as a whole (DL Times are shortened due to cache capacity, higher win for smaller ASes).</p>
Case	<p>Homogeneous peer distribution, homogeneous access bandwidth, access bottleneck (16 Mbit/s downlink, 1 MBit/s uplink).</p> <p>Star20 topology (20 stub ASes containing the peers).</p> <p>Poisson arrival process with a default $\lambda = 1/10s$.</p> <p>BitTorrent overlay with and without locality (BGP-Loc), IoP capacity per AS 1 MBit/s and 2 MBit/s.</p>
	<p>Lose for ISPs in case of BGP-Loc + IoP (higher incoming traffic).</p> <p>Win for users (reduced DL times).</p>

4.3.3 Inter-SIS collaboration based ETM mechanisms

Table 6. Inter-SIS (BNSBU + PeerPop grouping)

ETM scenario	Inter-SIS (BNSBU + PeerPop grouping) [2]
Case	<p>Heterogeneous peer distribution (hyperbolic), homogeneous access bandwidth, access is bottleneck (16 Mbit/s downlink, 1 Mbit/s uplink).</p> <p>Star20 topology (20 stub ASes containing the peers).</p> <p>Poisson arrival process with a default $\lambda = 1/10s$.</p> <p>BitTorrent overlay, PeerPop Grouping Inter-SIS with a group size parameter G varied (G = 10, 30, 50).</p>
Results	<p>Win for ISPs (larger traffic savings for fast ISPs), reduction in traffic for larger ISPs even more pronounced for larger group sizes (in the considered range).</p> <p>No-lose for swarm as a whole (DL Times remain ~the same, fairer distribution for larger group sizes).</p>

Table 7. Inter-SIS (Cost driven peer rating algorithm)

ETM scenario	Inter-SIS – cost driven peer rating algorithm [2]
Case	<p>Heterogeneous peer distribution, topology – semi-real: 29 ASes, in 9 of which peers were located, 20 transit ASes, conservative peer rating algorithm, SIS located in a single domain, contradictory policies of ISPs allowed.</p>

Results	<p>Win for ISP implementing SIS – overall inter-domain traffic decrease. Traffic decreased on peering links (where the traffic is exchanged without cost) as well as provider links (those the ISP has to pay for the traffic volume).</p> <p>No-lose/lose for ISPs not implementing SIS. Overall inter-domain traffic observed in those ASes remains on average the same. The distribution of the traffic on links changes and may lead to increase on costly link and decrease on peering link, finally resulting in higher cost for ISP.</p> <p>Win for peers located in a domain implementing SIS – significant reduction of DL time.</p> <p>No-lose (on average) for peers located in other domains, a little increase or decrease of DL time is observed, but the differences are small.</p>
Case	Heterogeneous peer distribution, topology – semi-real: 29 ASes, in 9 of which peers were located, 20 transit ASes, conservative peer rating algorithm, SISes located in all 9 domains, contradictory policies of ISPs allowed.
Results	<p>Win for all ISPs: overall inter-domain traffic reduced, especially on costly links. The remaining traffic inevitable) is shifted to peering links.</p> <p>Win for users – DL time reduction observed in all domains.</p>
Case	Heterogeneous peer distribution, topology – semi-real: 29 ASes, in 9 of which peers were located, 20 transit ASes, conservative peer rating algorithm, SISes located in all 9 domains, comparison of contradictory and non-contradictory ISPs policies.
Results	<p>If two directly connected have contradictory policies, that is, one tries to decrease the traffic (the one that pay the other for the traffic) while the other promotes peer available through this link it may lead to lose for the one that pay for the traffic. If both agree to decrease the traffic on that link it leads to win situation.</p> <p>Win for users regardless the ISPs policies are contradictory or not.</p>

Table 8. Inter-SIS: (Collaborative BGP Locality and Layered Collaborative BGP Locality)

ETM scenario	Inter-SIS – Collaborative BGP Locality and Layered Collaborative BGP Locality [2]
Case	Homogeneous peer distribution, 8 Tier 3 ASes (1 new peer every 60 sec) 3 Tier 2 ASes (1 new peer every 30 sec), SISes, located in all domains, run BGP Locality except for those that run the mechanism, 2 Tier 2 ISPs (that have a peering agreement) run Collaborative BGP Locality (thus all others run BGP Locality).
Results	<p>Win for Tier 2 ISPs implementing Collaborative BGP Locality – inter-domain traffic (transit traffic that these ISPs pay for) decrease. Traffic increased on peering links (where the traffic is exchanged without extra charging) as well as intra-domain.</p> <p>No-lose for ISPs not implementing Collaborative BGP Locality. Inter-</p>

	<p>domain traffic observed in those ASes remains on average the same. No-loose (on average) for peers located in domains that run Collaborative BGP Locality, a little increase or decrease of download time is observed, but the differences are small. No-loose (on average) for peers located in other domains.</p> <p>Similar results for Layered Collaborative BGP Locality.</p> <p>Tier 2 ISPs do have the incentive to do Collaborative BGP Locality and this will be accepted by their users and by the other ISPs.</p>
Case	<p>Homogeneous peer distribution, 8 Tier 3 ASes (1 new peer every 60 sec) 3 Tier 2 ASes (1 new peer every 30 sec), SISEs, located in all domains, run BGP Locality except for those that run the mechanism, only 2 Tier 3 ISPs (with a peering agreement) run Collaborative BGP Locality (and thus all others run BGP Locality).</p>
Results	<p>No-loose for Tier 3 ISPs implementing Collaborative BGP Locality — overall inter-domain traffic decrease. Traffic increased on peering links (where the traffic is exchanged without cost) as well as intra-domain. No-loose for the upper Tier ISPs (the carriers of Tier 3 ISPs) on the inter-domain links. No-loose for other domains not implementing Collaborative BGP Locality. Inter-domain traffic observed in those ASes remains on average the same. No-loose (on average) for peers located in domains that run Collaborative BGP Locality. No-loose (on average) for peers located in all other domains. Similar results for Layered Collaborative BGP Locality Tier 3 ISPs are indifferent to do Collaborative BGP Locality since they do not gain significantly.</p>
Case	<p>Homogeneous peer distribution, 8 Tier 3 ASes (1 new peer every 60 sec) 3 Tier 2 ASes (1 new peer every 30 sec), SISEs, located in all domains, run BGP Locality except for those that run the mechanism, 2 Tier 3 ISPs and 2 Tier 2 ISPs run Collaborative BGP Locality. Same Tier ISPs have peering agreements with each other. Tier 3 ISPs are customers of Tier 2 ISPs. Also Tier 2 ISPs collaborate with their customers (Tier 3 ISPs).</p>
Results	<p>Win for Tier 2 ISPs implementing Collaborative BGP Locality – overall inter-domain traffic decrease. No-loose for Tier 3 ISPs implementing Collaborative BGP Locality - traffic on the inter-domain links remain the same. No-loose for other domains not implementing Collaborative BGP Locality. Overall inter-domain traffic observed in those ASes remains on average the same. Win (on average) for peers located in Tier 2 ISPs that run Collaborative BGP Locality. No-loose (on average) for peers located in Tier 3 ISPs that run Collaborative BGP Locality. No-loose (on average) for peers located in other domains.</p>

	Tier 2 ISPs do have the incentives to do Collaborative BGP Locality and ask from their customers (Tier 3 ISPs) to do so, thus concluding in a win – win situation between them and their end users.
--	--

Table 9. Inter-SIS (Splitting Chunks)

ETM scenario	Inter-SIS – Splitting Chunks [2]
Case	Homogeneous peer distribution, hierarchical topology: 13 ASes, 8 of which contain peers, 3 Tier 2 and 2 Tier 1 transit ASes, SISes, located in all domains, run BGP Locality, only 2 Tier 3 ISPs, that have a peering agreement, run Splitting Chunks.
Results	<p>Win for ISP implementing Splitting Chunks – overall inter-domain traffic decrease. Traffic increased on peering links (where the traffic is exchanged without cost) as well as intra-domain.</p> <p>No-lose for ISPs not implementing Splitting Chunks. Overall inter-domain traffic observed in those ASes remains on average the same.</p> <p>No-lose (on average) for peers located in domains that run Splitting Chunks, a little increase or decrease of DL time is observed, but the differences are small.</p> <p>No-lose (on average) for peers located in other domains.</p> <p>In case that, ISPs own a large portion of the swarm, such as in this case, they have the incentives to run Splitting Chunks on top of BGP Locality.</p>
Case	Homogeneous peer distribution, hierarchical topology: 13 ASes, 11 of which contain peers (belonging to both Tier 3 and Tier 2), and 2 Tier 1 transit ASes, SISes, located in all domains, run BGP Locality, only 2 ISPs from Tier 2 that have a peering agreement run Splitting Chunks
Results	<p>Win for ISPs implementing Splitting Chunks – overall inter-domain traffic decrease. Traffic increased on peering links (where the traffic is exchanged without cost) as well as intra-domain.</p> <p>No-lose for ISPs not implementing Splitting Chunks. Overall inter-domain traffic observed in those ASes remains on average the same.</p> <p>No-lose (on average) for peers located in domains that run Splitting Chunks, a little increase or decrease of DL time is observed, but the differences are small.</p> <p>No-lose (on average) for peers located in other domains.</p> <p>Thus Tier 2 ISPs have the incentives to implement Splitting Chunks and this will be accepted by their users and by the other ISPs.</p>
Case	Homogeneous peer distribution, hierarchical topology: 13 ASes, 11 of which contain peers (belonging to both Tier 3 and Tier 2), and 2 Tier 1 transit ASes, SISes, located in all domains, run BGP Locality, only 2 Tier 3 ISPs that have a peering agreement run Splitting Chunks.
Results	Win for ISP implementing Splitting Chunks – overall inter-domain traffic decrease. Traffic increased on peering links (where the traffic is ex-

	<p>changed without cost) as well as intra-domain. Win for the upper Tier ISP (Tier 2) on the inter-domain links. Non-lose for other domains not implementing Splitting Chunks. Overall inter-domain traffic observed in those ASes remains on average the same. Lose (on average) for peers located in domains that run Splitting Chunks. No-lose (on average) for peers located in other domains.</p> <p>Despite the win-win situation between Tier 3 ISPs and their carriers, Tier 3 ISPs will not run Splitting Chunks due to the deterioration on the download times of their end users.</p>
--	--

Table 10. Inter-SIS (Combination of collaborative and layered collaborative BGP Locality approaches and splitting chunks)

ETM scenario	Inter-SIS – Combination of Collaborative and Layered Collaborative BGP Locality approaches and Splitting Chunks [2]
Case	Homogeneous peer distribution, 8 Tier 3 ASes (1 new peer every 60 sec) 3 Tier 2 ASes (1 new peer every 30 sec), SISes, located in all domains, run BGP Locality except for those running Collaborative BGP Locality, 2 Tier 2 ISPs (with a peering agreement) run Splitting Chunks on top of Collaborative BGP Locality (and thus all the other ISPs run BGP Locality).
Results	<p>Win for Tier 2 ISPs implementing Splitting Chunks on top of Collaborative BGP Locality – inter-domain traffic decrease. Traffic increased on peering links (where the traffic is exchanged without cost) as well as intra-domain. No-lose for ISPs not implementing Splitting Chunks. Inter-domain traffic observed in those ASes remains on average the same. No-lose (on average) for peers located in domains that run Splitting Chunks on top of Collaborative BGP Locality. No-lose (on average) for peers located in other domains.</p> <p>Similar results for Splitting Chunks on top of Layered Collaborative BGP Locality. Thus Tier 2 ISPs do have the incentive to do Splitting Chunks on top of Collaborative BGP Locality and this will be accepted by their users and by the other ISPs.</p>
Case	Homogeneous peer distribution, 8 Tier 3 ASes (1 new peer every 60 sec) 3 Tier 2 ASes (1 new peer every 30 sec), SISes, located in all domains, run BGP Locality except for those running Collaborative BGP Locality, 2 Tier 2 ISPs (with a peering agreement) run Splitting Chunks on top of Collaborative BGP Locality; (also their Tier 3 customer ISPs run Collaborative BGP Locality while all other ISPs run BGP Locality)
Results	<p>Win for ISPs implementing Splitting Chunks on top of Collaborative BGP Locality – overall inter-domain traffic decrease. No-lose for their customer ISPs on the inter-domain links.</p>

	<p>No-lose for other domains. Inter-domain traffic observed in those ASes remains on average the same.</p> <p>No-lose (on average) for peers located in domains that run Splitting Chunks on top of Collaborative BGP Locality.</p> <p>No-lose (on average) for peers located in the customer ISPs.</p> <p>No-lose (on average) for peers located in other domains.</p> <p>Similar results for Splitting Chunks on top of Layered Collaborative BGP Locality.</p> <p>Thus Tier 2 ISPs do have the incentive to do Splitting Chunks on top of Collaborative BGP Locality when their customer ISPs run Collaborative BGP Locality and this will be accepted by their users and by the other ISPs.</p>
Case	<p>Homogeneous peer distribution, 8 Tier 3 ASes (1 new peer every 60 sec) 3 Tier 2 ASes (1 new peer every 30 sec), SISEs, located in all domains, run BGP Locality except for those running Collaborative BGP Locality, 2 Tier 3 ISPs (with a peering agreement) run Splitting Chunks on top of Collaborative BGP Locality (also their Tier 2 carrier ISPs run Collaborative BGP Locality and all other ISPs run BGP Locality).</p>
Results	<p>Win for Tier 3 ISPs implementing Splitting Chunks on top of Collaborative BGP Locality – overall inter-domain traffic decrease.</p> <p>Loss for Tier 2 ISPs that are carriers of the Tier 3 ISPs.</p> <p>No-lose for other domains not implementing Collaborative BGP Locality. Overall inter-domain traffic observed in those ASes remains on average the same.</p> <p>Loss (on average) for peers located in Tier 3 ISPs that run Splitting Chunks on top of Collaborative BGP Locality.</p> <p>No-lose (on average) for peers located in Tier 2 ISPs that are carriers for Tier 3 ISPs.</p> <p>No-lose (on average) for peers located in other domains.</p> <p>Thus Tier 3 ISPs will not implement Splitting Chunks on top of Collaborative BGP Locality since their end users’ performance deteriorates.</p>
Case	<p>Homogeneous peer distribution, 8 Tier 3 ASes (1 new peer every 60 sec) 3 Tier 2 ASes (1 new peer every 30 sec), SISEs, located in all domains, run BGP Locality except for those running Collaborative BGP Locality, 2 Tier 3 ISPs and 2 Tier 2 ISPs run Splitting Chunks on top of Collaborative BGP Locality (while all others run BGP Locality). Same Tier ISPs have peering agreements with each other. Also Tier 2 ISPs collaborate with their customers (Tier 3 ISPs).</p>
Results	<p>Win for Tier 2 ISPs implementing Splitting Chunks on top of Collaborative BGP Locality – overall inter-domain traffic decrease.</p> <p>No-lose for Tier 3 ISPs implementing Splitting Chunks on top of Collaborative BGP Locality - traffic on the inter-domain links remains the same.</p> <p>No-lose for other domains not implementing Splitting Chunks. Overall inter-domain traffic observed in those ASes remains on average the same.</p>

	<p>No-lose (on average) for peers located in Tier 2 ISPs that run Collaborative BGP Locality.</p> <p>No-lose (on average) for peers located in Tier 3 ISPs that run Collaborative BGP Locality.</p> <p>No-lose (on average) for peers located in other domains.</p> <p>Thus Tier 2 ISPs have the incentives to run Splitting Chunks on top of Collaborative BGP Locality and ask from their customers -Tier 3 ISPs- to collaborate with them.</p>
--	--

4.3.4 QoS-aware ETM

Table 11. QoS awareness

ETM scenario	QoS aware [2]
Case	Homogeneous peer distribution (triangular topology), homogeneous access bandwidth, access is bottleneck
Results	<p>Win for the ISP (inter-AS traffic reduced)</p> <p>Win for the users (average DL times are lower than the same as in regBT)</p>
Case	Heterogeneous peer distribution (triangular topology), homogeneous access bandwidth, access bottleneck
Results	<p>Win for the ISP (inter-AS traffic reduced, but it depends on size of the AS with SIS)</p> <p>Win for the users (average DL times decrease in comp. to regBT in the AS with SIS, but increase in ASes without SIS. The amount of increased/decreased time depends on the size of the AS)</p>
Case	Homogeneous peer distribution (bike topology), homogeneous access bandwidth, access is bottleneck
Results	<p>Win for ISPs (inter-AS traffic reduced, more significant reduction when more ASes use SIS)</p> <p>Win for the users (the average DL time in AS with SIS is lower comparing to regBT; if the number of ASes with SIS grows, the average DL time increases but is still lower than in regBT)</p>

4.3.5 Highly Active Peer (HAP)

Table 12. Static HAP (HAP Selection Metrics and User Behavior)

ETM scenario	Static HAP Case: HAP Selection Metrics and User Behavior [2]
Case	<p>Peer-assisted video-on-demand scenario as above but with the following refinements: heterogeneous topology (5 customer ISPs with 50, 25, 12.5, 6.25 and 6.25% of peers per AS). All ISPs apply HAP promotion. Heterogeneous peer upload capacities (250, 200 and 1000 kbps for 30, 30, and 40% of users, respectively). 50% of users are selfish (do not seed) and 50% seed for 60 minutes after finishing a download. Diurnal usage pattern (3x times higher request probability during the day time). The overlay applies locality-aware peer selection.</p> <p>At first the impact of HAP selection metrics is evaluated, which can be set to <i>random</i>, <i>seeding ratio</i>, or <i>upload amount</i>. Additionally the bandwidth increase per HAP is varied (from 1=no promotion to 5=400% more bandwidth per HAP)</p>
Results	<p>Win for the ISPs (the inter-domain traffic is reduced by up to 40%). The highest benefit is obtained with the upload amount HAP selection metric.</p> <p>The win situation is also observed for the single ISPs (both regarding the in- and outbound traffic, where the relative reduction of inbound traffic is higher for larger ISPs, while the relative outbound traffic reduction is higher for the smaller ISPs).</p> <p>Win for the overlay (the server load is reduced by up to 80%). Again upload amount provides the highest benefit.</p>
Case	<p>In this experiment the ratio of promoted peers is varied from 0 to 1.</p>
Results	<p>Win for the ISPs (biggest benefit if 10% of peers are promoted, corresponding to the amount of heavy altruistic users).</p> <p>Win for the overlay (in particular if 10-20% of peers are promoted corresponding to the amount of heavy users, both altruistic and selfish).</p>
Case	<p>In this experiment the detection of locality-aware users is analyzed. The percentage of locality-aware users is varied from 0 to 100%.</p>
Results	<p>Win for the locality-aware users (over time they are promoted with an increasing probability that is higher than for locality-unaware users).</p>
Case	<p>Same setup as above but 100% of peers are locality-aware. The impact of the promotion interval duration is assessed (between 6 and 36 hours).</p>
Results	<p>Win for the ISPs (the highest benefit for longer promotion durations ≥ 18 hours).</p> <p>Win for the overlay (the highest benefit with the promotion interval of 12 hours. This can be explained with the diurnal traffic pattern where busy and idle phases have the same duration).</p>

Table 13. Dynamic HAP with single BitTorrent swarms

ETM scenario	Dynamic HAP with single BitTorrent swarms [2]
Case	Homogeneous peer distribution, homogeneous access bandwidth, access is bottleneck, triangular topology
Results	<p>Win for the ISP (inter-AS traffic reduced provided that HAP mechanism attract more peers to use SIS). Win for the users (average DL times are lower than the same as in regBT, DL times decrease with the increasing number of HAPs)</p> <p>Win for the users (promotions to HAPs increase peers download and up-load capabilities)</p>

Table 14. Static HAP (Comparison of HAP promotion with and without BGP Locality)

ETM scenario	Comparison of static Highly Active Peer (HAP) promotion with and without locality-awareness (including early adopter study) [2]
Case	<p>Peer-assisted video-on-demand scenario with servers providing the missing resources. Multiple-swarm scenario (10,000 users, 2,000 videos) with long-term duration (4 weeks). Five customer ISPs organized in a star topology with homogeneous peer distribution. Servers reside in a separate domain.</p> <p>Main metrics is the inter-domain traffic (only for the customer ISPs) and the load on the content servers (corresponds to the overlay performance since the servers guarantee stable QoS for users). Also the in- and out-bound traffic of ISPs is considered separately if appropriate.</p> <p>At first the overlay is locality-unaware, while all 5 customer ISPs employ HAP promotion (for 20% of heavy users) with varying bandwidth increase per run.</p>
Results	<p>No-Lose for the ISPs (inter-domain traffic is reduced insignificantly).</p> <p>Win for the overlay (server load is decreased by up to 90%, resulting in lower delivery costs).</p>
Case	Same setup, but now the overlay does not employ HAP promotion, instead the overlay peers behave locality-aware (i.e. local transfers have the highest priority, remote transfers the second priority, while the servers are utilized as a last resort).
Results	<p>Win for the ISPs (both in- and outbound inter-domain traffic can be significantly reduced by up to 30%).</p> <p>No-Lose for the overlay (the server load remains stable).</p>
Case	The two above setups are combined, i.e. the ISPs apply HAP promotion while the overlay behaves locality-aware.
Results	Win for the ISPs (the inbound traffic can be reduced by 30-50%).

	Win for the overlay (server load can be reduced by up to 90%).
Case	Again the same setup, but now only one ISP (A) applies HAP promotion, while the overlay behaves locality-unaware (100% of peers).
Results	Win/No-Lose/Lose for the early adopter ISP A depending on its inter-connection agreements regarding the treatment of in- and outbound traffic (inbound traffic is slightly reduced, but the outbound traffic is increased by up to 50%). No-Lose for other ISPs (that do not apply HAP promotion). Win for the overlay (server load is reduced by up to 50%).
Case	The same scenario with the early adopter ISP A, but now the overlay behaves locality-aware.
Results	Win/No-Lose/Lose for the early adopter ISP A depending on its inter-connection agreements regarding the treatment of in- and outbound traffic (this time the inbound traffic is reduced by up to 50% , while the outbound traffic is again increased by up to 50%). No-Lose for other ISPs (that do not apply HAP promotion). Win for the overlay (server load is reduced by up to 50%).

4.3.6 Conclusions

This section summarizes, on a per ETM mechanism basis, the set of qualitative conclusions drawn for the various simulation scenarios adopted for the evaluation of each such mechanism. Due to the broad spectrum of scenarios simulated and the ample number of experiments run per scenario, the relevant conclusions are considered very reliable.

For most of these mechanisms, a universally applicable conclusion, such that that "Win-Win is always the case", cannot be reached. Indeed, from the outcomes presented here, it is seen that there are a number of important factors, such as homogeneity vs. heterogeneity of the swarm's distribution among ASes that have a considerable impact on whether a mechanism is beneficial for the various players.

Further consolidation of the overall assessment per mechanism is carried out in Section 7, where the outcomes of the trials (and particularly of the external ones, as presented in the next Section) are also taken into account, together with a set of additionally relevant issues. Note that simulations were actually the main means of the basic and detailed evaluation of ETM mechanisms and their behavior, as is most often the case with a new technical approach in the area of networks. Although additionally ETMs have been actually implemented and ported to a real ISP network, the selection of these mechanisms to be deployed, their detailed specification, and the fine tuning of various parameters involved, was crucial and very much supported by the broad set of reliable simulation results.

To this end, the qualitative results of this Section 4.3 already provide a very strong justification that ETM is a promising approach in several general and practically important cases.

5 Trial Results

The external trial had two-fold goal. First, the goal was to show the technical feasibility of the developed solution, check its stability and performance in real life environment and to show that integration into real ISP network and operation in prolonged period of time is possible. The second goal of the trial was to collect detailed usage statistic based on which ETM mechanisms performance under real life conditions could be analyzed with respect to TripleWin goal.

This section provides details of trial operation, modification and integration work undertaken, mechanisms used to collect and aggregate the data as well as qualitative analysis of the results.

5.1 Promotion of Trial to PrimeTel's Customers

A number of steps were taken to attract customers' attention to the new service. In particular good quality content was secured. A few weeks before the trial started SmoothIT project was presented in PrimeTime magazine targeting non-technical audience. Upon trial launch all customers were notified and on each consecutive phase start additional invitations were sent out. The trial environment was designed to be user friendly.

From the early stages of trial planning it was understood that quality of the content would have great impact on level of users' participation. A number of attempts were made to acquire permission to use good quality titles for the trials, but unfortunately most of those were not successful. Finally the content for the trial was provided by one of the PrimeTel partners exclusively for the purposes of the trial and consisted of 15 quality movie titles of average ~30 minutes duration, covering various tourist destinations around the world. The content was offered to the customers on a dedicated website in form of gallery with direct link to torrent files, which once opened would start download and playback. Each movie was accompanied with a screenshot and short description.

The same website contained the information about SmoothIT projects, the software and general usage information targeting non-technical audience. The download link was placed on the home page leading directly to the player setup.

All PrimeTel customers were invited to participate in the trial. The invitation was included in the monthly newsletter and published in PrimeTime magazine, which is delivered to all customers free of charge. The potential audience included ~20K customers. In addition a selected group of ~300 beta testers was invited to join with a separate targeted email. During second and third phases a draw of Sony PlayStation 3 was announced for most active participant to raise people's interest. During the trial planning it was estimated that the response rate would be ~5% which would result in ~1000 real users in the system, while the number of active users was estimated to reach 100 customers. Alas, the trials revealed much lower participation rate, which could be accounted for fact that content provided did not meet customers' expectations.

5.2 Trial-specific Modifications

Certain modifications had to be made in the client software in order to prepare it for the trial as per project requirements or business requirements imposed by the trial environment. The changes mainly affected the following functionality: supporter algorithm for maintain-

ing acceptable QoE, monitoring reports required for results evaluation, and content encoding required by content providers. This section provides a brief account of such changes.

It is important to note that majority of the changes resulted from strict requirement to maintain minimal acceptable level of QoE for customers. Failure to do so would have resulted in no user activity at all. This requirement further led to more complex data analysis, involving MLO metric as defined in [4] being used instead of direct QoE measurements, which required enhanced monitoring component.

5.2.1 Stability and streaming quality

First of all, the preliminary tests performed in December 2009 revealed serious issues regarding the streaming quality of the NextShare version used so far. Even in simple scenarios with few leechers running in an ideal LAN environment the upload utilization of clients was very low and resulted in high stalling delays. This issue was observed in the NextShare release version M12 from December 2008 as provided by the P2P-Next consortium. Therefore, a newer M24 (from December 2009) was obtained, which solved this issue. As a result the whole client-side implementation (including various configuration-related extensions) was migrated to the newer version of NextShare and required significant implementation effort.

Once migration was complete, the software was thoroughly checked to make sure it adheres to all requirements both technical and usability-wise. These checks revealed significant problems with (a) the software's stability (resulting from customer torrent format), (b) background operation, and (c) multiple swarm scenarios. A number of serious usability issues were also noted. The necessary development was undertaken to resolve the problems and bring the client code up to minimal production requirements.

End-user QoE was another issue addressed by the trial preparations. In order to avoid quality degradations due to insufficient supply of video data a well-dimensioned initial seed had to be deployed. However, this raised another issue of having too much data coming from this initial seed, thus, reducing the probability of data uploaded from other peers. As described in [4] a supporter algorithm was designed to address this issue, in such a way that only peers missing data close to the current playback deadlines are connected to and served by the initial seed (supporter algorithm).

Unfortunately, in a real environment the supporter implementation had issues related to the low size of swarms and missing stability of the Tribler tracker that was a crucial component of the supporter algorithm. In order to solve this issue the following steps were performed:

- (1) Instead of the Tribler tracker the official BitTorrent tracker [30] was used for the external trial thus solving the basic tracker-related issues (e.g. departed peers were never deleted from the tracker's contact list and returned to new peers, thus, distorting logs).
- (2) Instead of the original supporter implementation a simplified solution was used. This solution modifies the interest signaling of regular peers to assure that only suffering peers (i.e. missing data in the playout buffers) request data from the initial seed. Only peers that fulfill the suffering condition (in terms of download speed, playout buffer fill state) and are not paused are "interested" in receiving data from the supporter. This means that peers are only receiving data from the supporter seed if they cannot obtain the urgent data from other peers. Normally a seed would dis-

connect any peer from which it receives a not-interested message because it considers such peers as already having downloaded all the data. To overcome this limitation the supporter seed receives an optional parameter “supporter_seed”. If this parameter is set, the seed will not disconnect uninterested peers. To achieve that the modified interested signaling applies only to the connection towards the supporter seed, another configuration parameter “supporter_ips” was introduced and is configurable via the config file. To this end, the modified interest signaling is applied only to connections if the counterpart’s IP address is listed in this parameter value.

Two further issues were identified and fixed: (1) A crucial bug was identified and fixed during trial preparations concerned the implementation of the Give-to-get protocol [31]. Here the mid priority set calculation was incorrect and resulted in chunks being downloaded in a wrong order. (2) Usage of NextShare-specific Merkle Torrents resulted in frequent client crashes when an unfinished download was restarted (hash check of already downloaded chunks failed). Therefore, only regular torrents are used for the trial.

5.2.2 Monitoring

For the purposes of collecting automated client reports a standalone monitoring component was developed. This component implemented server side of the “PhoneHome” protocol available in the NextShare software. It was decided that the component should be implemented as a separate entity as it would only be useful for the purposes of the trial and its separation would ensure greater stability and easier maintenance should any problems with core SIS occur. The component acts as a simple HTTP server accepting XML-formatted reports from the clients. These were then stored in raw files, parsed and recorded into MySQL database. Section 5.2 describes the data structures in detail.

Furthermore, client reports already implemented for the show case were extended to include the information about the data source (IOP_flag) to distinguish IoP and regular peers. Furthermore, in order to compute missed local opportunities statistics about blocks exchanged with other peers were included into statistics. The latter describe which data were downloaded from which peer, including the data offset and block length. This is described in more details in Section 5.3.

5.2.3 Configuration, content encryption, and SIS modifications

Furthermore, various configuration parameters previously available to the headless client (as provided to the internal trial) were ported to the swarmplayer itself. This included upload rate limitation, address of the monitoring component etc.

To obey the restrictions of the copyright-protected content used in the external trial, the content downloaded by clients had to be encrypted. Therefore, a content decoding algorithm based on AES encryption was added as per requirement of content providers.

The SIS server required almost no modification except for the PrimeTel specific customer information component that is used by the HAP algorithm. An adapter specific to PrimeTel was developed and used by the SIS for the trial.

5.3 Data Collection

The data required for assessment of the ETM mechanisms were defined in the [4] and come from two sources — the client itself and the network monitoring system.

Data from the clients are collected by the monitoring subsystem in form of periodic (60s interval) reports sent via HTTP interface in XML format and recorded in the monitoring database. These reports contain raw data about any data transfer performed by the peer on a per-piece basis as well as a description of the client's state at the moment the report is sent. All data are per swarm (swarm id indicated in the report) and aggregated only for the duration of the session.

Table 15. Monitoring report — General section

Key	Description
timestamp	Unix time stamp when report was generated
down_rate	Current download rate of the peer
down_total	Downloaded content during this session in bytes
up_rate	Current upload rate
up_total	Uploaded content during this session in bytes
Listenport	Port on which client software is listening
t_stall	Time spent in stall state
t_prebuf	Time spent for pre-buffering
iop_flag	Indicates if this is an IoP
Filename	Name of the file (for the purposes of the trial there is always one file per swarm)
Infohash	Infohash for the swarm
Status	1 – Playing, 2 – Paused, 3 – Stopped
p_played p_dropped p_late	Number of pieces player, dropped and arrived late respectively
Progress	Progress of download in percent

The peers section of the monitoring report gives information about peers connected at the time report was generated. All statistics in this section are on a per-peer basis.

Table 16. Monitoring report — Peers section

Key	Description
up_total	Total bytes uploaded to this peer
Addr	IP/Port of the peer
g2g	Give-to-get algorithm used (always g2g for the trial' purposes)
Down_str	Coded presentation of Chocked / Interested states. Capitalises either char-

	acter when we're choked (C) and when we're interested (I)
up_str	A string 'cio' which capitalises either character when we're choking the neighbour (C) and when the neighbour is interested (I), and whether this neighbour is optimistically unchoked (O).
Down_total	Total bytes downloaded from the peer
g2g_score	Score of the peer calculated based on give to get algorithm
up_rate	Upload for the peer
Id	PeerID string

The blockstats section contains detailed information about data blocks downloaded by the peer since the last report. The block size can be less than a BitTorrent chunk and as such length/offset is used in reporting. Data are presented as a multidimensional array

Table 17. Monitoring report — Blocks section

Key	Description
0	Piece id (torrent piece)
1	Offset in piece
2	Block length
3	Timestamp when transfer happened
4	Sender IP
5	Sender port
6	Sender peer id

5.3.1 Data aggregation and analysis

As defined in [4] the key aspect of assessing the ETM mechanisms performance in external trial is identifying the Missed Local Opportunities (MLO). Section 6.3.1 of [4] defines MLOs as “*as an event where a local peer is downloading a chunk from a remote location (also referred to as hit) when there is another local peer seeding this chunk wherefrom its downloading was possible according to resource availability*”. The reduction of this metric would have meant respective reduction of incurred inter-domain domain as a result of particular ETM (or their combination) being utilized.

For this purpose a data aggregation tool was created which based on the monitoring data marked the Block transfer records (Table 17. Monitoring report — Blocks section) as MLO. It is important to understand that due to the nature of the reporting and NMS data it is close to impossible to identify whether each specific block transfer could have happened between local peers, i.e. if it is an MLO. Therefore during the analysis only relative comparison will be made.

During data aggregation the following peer classes are defined:

- P — Real peer from within PrimeTel network
- P-A — Automated peer in PrimeTel network
- G — Automated peer in external domain (GLab)
- S — Support peer that is physically located in PrimeTel’s domain but considered as an external peer by all ETM mechanisms.




Figure 5. MLO Identification algorithm

The flowchart in Figure 5 presents the algorithm which was used to process the data to identify the missed local opportunities. The algorithm identifies block transfers that oc-

curred between remote and local peers and then based on the knowledge of the state of other clients attempts to identify if these transfers could have happened between local peers under ideal circumstances.

1. At the entry point the next **potential MLO** is fetched, which is an event of transfer of block **B** from a remote peer **X** to local peer **P**. It is important that all other types of block transfers are ignored as irrelevant.
2. This step builds a list **L** of all potential **local** peers in local domains who have already acquired block **B** by this time. This step imposes only one condition: the block should have been acquired prior to time **T**. All other conditions are imposed at later steps by reducing the list. After each list reduction, the size of the list is checked. If at any point the list is empty algorithm, then the transfer currently examined as a potential MLO is aborted, as this constitutes a condition when no local opportunities exist.
3. The next step reduces list **L** by excluding any peer which is not online at this time. It is impossible to accurately identify when a peer goes offline therefore a heuristic approach is used. If the last monitoring report was received over 5 intervals (i.e. 5 minutes) prior to time **T** the peer is considered offline.
4. Reduce list **L** by excluding those peers who have all upload slots occupied by other local peers.
5. If list **L** is not empty the block **B** is marked as **MLO**.

In addition to MLO identification the data were aggregated into sessions and watch sessions. The session is defined for peer **P** as closed time interval during which the maximum gap between consecutive messages from the peer is less than 10 minutes. The watch session is defined for peer **P** on torrent swarm **S** as a closed time interval during which maximum interval between consecutive messages is less than 10 minutes and during which peer has continuously status = Playing.

5.4 Results

This section presents the results collected during external trial, their qualitative and performance analysis followed by their interpretation from the point of view of ETM comparison and mechanism prototype assessment.

5.4.1 Qualitative assessment of trial data

In this subsection the facts and figures about external trials are presented along discussion of uncertainties arising in data interpretation and various factors which contributed to overall results credibility.

5.4.1.1 Schedule

The trial has been running in three phases:

- Reference phase when no ETM mechanism was enabled
 - November 22nd 2010 — Jan 3rd 2011
- IoP+BGP-Loc phase with two ETM mechanisms running in parallel
 - January 3rd 2011 — 8th February 2011

- HAP phase
 - 8th February 2011 — 22nd of March 2011

The deviation from the initial plan was due to unforeseen technical challenges related to overall stability of the client software code base and problems with supporter algorithm implementation which only manifested themselves under significant load during long period of time. In particular the tests of the client software revealed a number of problems as well as overall lack of stability and robustness minimally required for production launch. As the client code base is a third party product its quality could not be controlled by the project and resolving the technical problems required significant amount of time which lead to the delays.

The decision to launch next phase was taken based on attempt to equalize duration of each phase taking into account practical restrictions due to local holidays which not only affected the operation of the trials but also customers' usage pattern.

The following table 18 shows which content titles were available at each phase. Initially half of the content was made available as it was anticipated that adding new content as trials progress would raise customers' interest. These expectations were only partially fulfilled and as participation level during reference phase was rather low it was decided to re-release all of the content.

Table 18. Content titles availability per phase

torrent id	torrent name	reference phase	IoP+BGP- Loc phase	HAP phase
1	Argentina.mov.tstream	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
2	Rio de Janeiro.mov.tstream	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
3	Thailand.mov.tstream	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
4	Greece.mov.tstream		<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
5	Barcelona.mov.tstream	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
7	SriLanka.mov.tstream	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
10	India.mov.tstream		<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
11	Egypt.mov.tstream	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
12	Chile.mov.tstream		<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
13	Ibiza.mov.tstream	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
16	China.mov.tstream		<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
18	South Africa.mov.tstream		<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
19	Mexico - Yucatan.mov.tstream		<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
20	Australia - East Coast.mov.tstream		<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
21	Australia - Outback.mov.tstream		<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>

5.4.1.2 Users' activity

Figure 6 shows the user activity during the trial. It presents the total number of users active simultaneously. This statistic includes both users actually watching the movie and users having an application running but not watching the movie. The latter users act as seeders (seed the content previously downloaded) but do not download the content. It should be noted that active users sharing different content are summed up. Different colors are used to distinguish phases of the trial.




Figure 6. Users' activity during the trials

Figure 7 presents the statistic of the number of users watching any movie at the same time. The movie watching time is significantly shorter than the time of activity. A user watches a movie for some time (up to 30 min which is the maximum movie length) while she might have a Tribler application active for several hours. Therefore, the situation that two or more peers are watching a movie at the same time is rarer. The spikes on figure 7 are actually time periods related to watching time.

Both figures lead to the conclusion that the user activity in phases with ETM mechanisms deployed was much higher than in the reference phase. The highest activity was in IoP+BGP-Loc phase. This is reflected by a significantly higher number of active users than in the reference phase. A similar conclusion can be derived from the number of users simultaneously watching a movie (maximum 2 in reference phase versus up to 5 in the IoP+BGP-Loc phase). Moreover, only a 4 day long period of high user activity in IoP+BGP-Loc phase can be recognized. The high activity period is indicated by purple color. The user activity in HAP phase is considerably lower. The lowest participation is observed in the reference phase. Additionally, within an approximately a 3 week long period (end of reference phase and beginning of the next) there was no user activity at all. As a result, **the number of data samples** available in each phase significantly differs. The diverse user activity leads also to **uncertainty in the data interpretation** since the results are influenced not only by the ETM mechanism deployment but also by different conditions and comparisons between phases have essentially limited meaning.




Figure 7. Number of users simultaneously watching a movie

5.4.1.3 Number of samples

Due to diverse user participation in individual phases the number of samples collected for each phase also differs. A watch session was selected as a basic data sample.

Few data for the reference phase was collected: total number of watch sessions (with leeching and downloading) is 9.

The movie was watched in its entirety in only one session. In 3 cases a customer downloaded 49% of the content or more. In 4 out of those 9 cases the percentage of content downloaded does not correspond to watch session duration. E.g., 32% of content was downloaded while the movie was watched over 32 minutes, whereas the movie length is 29:55 minutes. Such a situation indicates a Tribler software problem occasionally causing permanent video freezing. Video watching was possible only after restarting the client. However, the Tribler software internal bug could not be identified nor fixed. The existence of the problem could not be validated in a watertight way. In 3 out of 9 cases the watch session lasted longer than the movie size in seconds. In an extreme situation there was a watch session with a duration of more than 2.5 hours. If those suspicious cases are eliminated there are only 5 usable observations in the reference phase.

For IoP+BGP-Loc phase only 3-4 days of high user activity were observed, probably mainly originating from offering a reward for the most active user. Eliminating all probably erro-

neous data the size of obtained a sample equals 31. For HAP phase there are 6 valid samples.

Such a low number of samples leads to a very low level of statistical significance of the results. It is also reflected by confidence intervals provided for average values (see sections 5.4.2.1 and 5.4.2.2).




Figure 8. The number of full seeders per torrent

5.4.1.4 Uncertainty in data interpretation

As shown in Section 5.4.2 some performance improvement was observed in IoP+BGP-Loc and HAP phases. However, there is still a high degree of uncertainty whether the observed effect is really a result of the deployed ETM or originates from other uncontrollable side-effects always present in a real-life experiment and also the generally different situation of the swarms. It is observed that the number of simultaneously active users in phases with ETM deployed was considerably higher as shown in Figures 6 and 7). Additionally, Figures 8 and 9 show the number of active seeders per torrent. Seeders having full (100%) content or partial (less than 100%) content are distinguished, respectively. Seeders are peers that are active and known to the tracker but do not download the video. Leechers become full seeders after they finish the video download and continue watching. Partial and full seeders appear if they start watching and downloading a video but have additionally stored one or more fully or partially downloaded videos on the hard disk and share it to the swarm. Figures 8 and 9 show that the number of seeders was much larger during the high activity period. As a result, the provided upload capacity – independent of IoP or support seed - is also much higher. Thus, a) the situation of a single downloading peer without additional swarms in the reference phase and b) the situation of a few peers downloading a video in

assessed only qualitatively, as an indicator of a positive impact of ETM mechanisms. Confidence intervals are quite large. This is due to small cardinality of samples.

5.4.2.1 Users' benefit

In normal conditions a performance metric that could be used to evaluate "win" for end-users is QoE. For video application such as Tribler one of the main metrics is the probability and amount of playback stalling. If the environment guarantees close to zero stalling probability there is nothing to evaluate, i.e., the introduction of loP and BGP-Loc does not have any effect on perceived video quality. Therefore, the trial setup assumes that playback stalling is possible. However, to avoid users giving up participation in the trial due to low quality, as already explained, the stalling is prevented by adding a support seed functionality (see Section 5.2.2 for more details). In the situation that there is a danger of stalling, the missing block is downloaded from the support seed (that is used only to support starving real peers). Each block downloaded from supporter was counted per movie playback. The higher number of blocks downloaded from supporter during movie playback, the lower QoE is expected. This way, statistics of the number of block downloads from supporter are used indirectly to evaluate impact of ETM deployment on "win" for users.

In the case of the loP+BGP-Loc phase a "win" situation for customers can be shown in two ways.

The first approach is based on the fact that the loP has initially no content. This is due to the assumption of real loP deployment conditions. The first peer downloading a movie (not watched before) creates a new swarm. This swarm is joined also by loP which starts downloading as well. For this reason loP can offer less content for the first downloader and more and more content for subsequently incoming peers. The first user watching a given movie is expected to experience a worse performance than its followers.

It may occur that the loP downloads the torrent file completely during the first download by user. Then, it can act as seed and offer the whole content to next users watching the movie. It may also happen that loP do not manage to download the whole content during the first download by user. Then it will be still a leecher when next users download and watch the movie. The time needed by loP to fully download the file depends on a user behavior and the swarm size. If many real users are interested in the movie at the same time, the swarm is larger and the loP can download the movie faster. If only one real user is active at the beginning the swarm is small (it consists of a real peer and a few automated peers). If a real peer closes the application, also the automated peers stop working and the loP has no longer a source to download the file from.

Summarizing, the loP offers more and more content to subsequent movie downloaders. As a result, it is expected that the first movie watcher experiences lower QoE and downloads more blocks from the support seed than users joining later will do. Such a behavior was observed during the trial.

Figure 10 shows the sample analysis for three torrents. For each watch session we calculated the probability that a block is downloaded from support seed as:

$$P_i^S = \frac{N_i^S}{N_i^{total}}$$

where N_i^S is a number of downloaded from support seed during watch session i while N_i^{total} is the total number of blocks downloaded during watch session i .




Figure 10. The average probability that a block was downloaded from support seed for watch session in loP+BGP-Loc phase

It can be seen that the first download per movie required more blocks to be downloaded from the support seed. It reflects the fact that the first customer downloading a movie experiences a lower QoE than its followers. It should be noted, that the need for downloading a block from a support seed depends not only on the loP share of the content but also on the instantaneous swarm size (the number of seeders and leechers) and the content availability in the whole swarm. Therefore, the probability of downloading a block from the supporter is highest for the first downloader. For consecutive watch sessions it oscillates depending on temporary conditions but, on average, it decreases.

It can be noticed that users activity was highest for torrent id=18. The first downloader watched the whole movie. Then 7 other users started watching the same movie within one hour. At some moment three users were watching (and downloading) simultaneously while the others were seeding previously downloaded content. Thus, many users were active at the same time. As a result, the loP has downloaded the content quickly. In the case of torrent id=5, the first downloader has also downloaded the whole content. However, the user activity was much lower than for torrent 18. Only three users were watching the movie at that day. Simultaneous watching was not observed. Nevertheless, the second and the third watcher experienced better QoE (less probability of downloading a block from support seed) than the first one. The situation is different for torrent id=2. The user activity was considerable lower (one download daily). The first and the second user did not watch the whole movie (they gave up after downloading 55% and 20% of the content, respectively). As mentioned before, after they closed the application, also the automated peers were stopped. As a result, the loP has no opportunity to download the whole movie file. It has

managed to do so while the third customer watched the movie. It downloaded the whole content.




Figure 11. The probability of downloading a block from a support seed during a watch session

The second approach assesses the “win” for users in the case of loP+BGP-Loc and HAP ETM mechanisms. Here, again, the probabilities that a block was downloaded from the support seed are compared. For each watch session the probability of downloading a block from support seed was calculated using the above formula. Afterwards the average values of this probability for each trial phase (reference, loP+BGP-Loc and HAP) was calculated as

$$S[\%] = 100 \frac{1}{\# \text{ watch_sessions}} \sum_{i \in \{\text{watch_sessions}\}} P_i^S$$

and expressed in percents. The first downloads of a torrent during loP phases are omitted in this calculation (Figure 11).

It can be seen that, on the average, blocks are downloaded from support seed with lower probability if ETM is in use.

5.4.2.2 ISPs’ benefit

The assessment of the “win” for an ISP is based on the MLO metric. Due to not enough reliable data for HAP phase the evaluation was restricted to loP+BGP-Loc ETM. We compare the average value of MLO metric for the reference phase and the phase with ETM deployed (the concept of MLO metric is provided in section 5.3.1 and [4]).

For each watch session the probability that a block downloaded during a watch session is recognized as MLO is:

$$P_i^{MLO} = \frac{N_i^{MLO}}{N_i^{total}}$$

where N_i^{MLO} is the number of MLOs observed during watch session i while N_i^{total} is the total number of blocks downloaded during watch session i . Then, the average value from all sessions is calculated and is expressed in percents as:

$$MLO[\%] = 100 \cdot \frac{1}{\#watch_sessions} \cdot \sum_{i \in \{watch_sessions\}} P_i^{MLO}$$




Figure 12. MLO statistics for reference and IoP+BGP-Loc phases

The results for reference and IoP+BGP-Loc phases are presented in Figure 12. A little improvement can be observed for IoP+BGP-Loc case. It should be remembered that the number of samples is not sufficient for statistical credibility (what is also reflected by confidence intervals) and statistics for those two phases are calculated on samples of significantly different sizes. Therefore, the obtained statistic is not a hard evidence but a qualitative positive suggestion that IoP+BGP-Loc offers a benefit to ISP since lower MLO probability reflects lower inter-domain traffic.

Summarizing, on the basis of external trial results safe conclusions about Win for the ISP cannot be derived from the data. There is some positive evidence, put it is not statistically reliable, due to insufficient data, resulting from low user participation in the trial.

5.4.3 ETM mechanisms' comparison

The comparison of performance of ETM mechanisms deployed in the external trial is not possible on the basis of obtained results. As mentioned before the low user participation

made the trial environment not reliable enough to provide generalized conclusions neither on a particular ETM nor for comparing them. Additionally, even the rough qualitative comparison of ETMs is not possible due to significantly different user behavior and swarm conditions in the IoP+BGP-Loc phase and the HAP phase of the trial. It is only possible to provide a rough qualitative assessment of each tested ETM mechanism independently through comparison to the reference phase (without any ETM mechanism deployed) extended by some additional observations, as described above.

5.4.4 External trial results vs. simulation results

A comparison of trial results and simulations results is difficult for several reasons. First, it must be emphasized that regardless how much realistic the simulation scenario was it is never possible to reflect unpredictable real Internet environment. Simulations are always an approximation of reality. They are performed under certain assumptions on network topology, peer distribution and behavior, swarm statistics, file sizes, etc. In the simulations the goal was to emulate conditions as close to reality as possible. Since it is very difficult to evaluate the impact of numerous factors on performance, most variables of a system are assumed to be stable in order to evaluate an impact of a single factor. Otherwise, it would be impossible to distinguish the effect of a single factor if many of them, being highly random, influenced the system at the same time. Instead several simulations scenarios to evaluate ETM mechanisms under various conditions were distinguished. The influence of each particular constituent of the system that was recognized as being significant to ETM assessment was evaluated. Evaluating various factors makes it possible to provide generalized conclusions more credibly.

The trial in a live Internet is different. All factors may be varying and have an impact on the ETM performance at the same time. Additionally, the human factor is crucial. User behavior introduces an almost unpredictable element to the system, reflecting also sociological and psychological aspects of the deployment of new application/system deployment. In general, the trial environment may provide extremely unpredictable and unstable conditions, thus making it difficult to analyze. Such a situation was observed in the external trial. The user participation was highly influenced by the attractiveness of video content and incentives (offered prizes). As a result a very unstable situation was observed, as described before.

It should be also noted that in simulations a long-lasting steady state was assumed, in most cases. This means that network conditions and parameters such as peer arrival rate, average swarm size, etc., were stable during the simulation. In the trial rather bursty behavior of peer activity was observed. Additionally, the swarm size is, on average, quite small in the trial. Such scenarios were not evaluated by simulations.

For the above reasons it was decided not to compare simulation and trial results directly. The only meaningful comparison is putting together the high level conclusions on ETM performance and usability. The simulation and trial results are treated as complementary rather than as comparable results that should confirm one another.

5.4.5 Assessment of the ETM prototype

The goal of the trial was not only to assess win/win for ISP and customer but also to evaluate other aspects of the implementation in a real environment. In this area, the outcome from the external trial is clear.

The trial proved that intended functionality of BGP-Loc, IoP and HAP was achieved. The system was stable. The client software as well as SIS implementation worked without any problems in a live network for three months. In fact, no complaints were indicated by the participants. No problems were observed with the IoP implementation and mechanism promoting peers to HAP status. The network management mechanisms involved worked as specified, and granted additional access bandwidth to customers as expected.

No problems with system scalability were observed either. However, it should be noted that the system load during the external trial was lower than in the internal trial. Therefore, the conclusion on a level of system scalability should be based mainly on the internal trial experiment. The external trial showed that no scalability problems appeared in live network, but very high load was not reached.

5.4.6 Socio-economical impact

In this section, we briefly address the socio-economical impact of the ETM mechanisms evaluated in the external trial that took place in PrimeTel's real network. The SmoothIT external trial is divided into three phases: i) reference, where no ETM was deployed or used, ii) BGP-Loc+IoP, and iii) HAP phase. Based on results obtained during the external trial that have been presented in Section 5.4.2, we aim at assessing the achievement of TripleWin when ETM mechanisms are deployed; in particular we assess whether each player achieves a *lose*, *no-lose* or *win* situation.

ETM employs socio-economic concepts and mechanisms, to deal with the overlay traffic in a way that is incentive compatible for all parties involved, and thus, ETM leads the system to a situation that is mutually beneficial for all end users, overlay providers and ISPs, the so-called "TripleWin" situation. The "TripleWin" translates to traffic savings and therefore cost savings for the ISP and QoE improvements for the overlay and the end-users. In the external trial, "win" for ISP is measured by a metric called MLO (missed local opportunity), therefore the lower the MLO, the higher the gain for the ISP; while "win" for the overlay and the users is measured by the average number of blocks downloaded by the support seed.

As also indicated in Section 5.4.1.2, due to low user participation, we cannot draw secure conclusions from the trial results. However, we can only qualitatively assess the socio-economical impact of the ETM evaluated.

Fig. 10 showed that the probability of blocks downloaded by the support seed has decreased for three swarms during the BGP-Loc+IoP phase (e.g. 21/1 - 26/1). Additionally, Fig. 11 depicts that the average number of blocks downloaded by the support seed during a watch session is lower when the HAP mechanism is deployed. Lower probability of support seed use or lower number of blocks downloaded by the support seed implies higher performance due to the ETM mechanism deployed, and therefore a "win" situation for users. Unfortunately, results on the assessment of "win" for the IoP during this watch session are not secure, since confidence intervals are very broad due to few data samples.

6 SmoothIT in the Current Internet

Additionally, besides the incentive, numerical, and performance investigations, the approach proposed by SmoothIT had been evaluated with respect to the consumer and data protection views as well as the intellectual property rights and under competition views. Therefore, this section concludes with the discussion of potential risks.

6.1 Regulatory Issues and Network Neutrality

Four areas of consideration have been identified from a legal and regulatory point of view for the SmoothIT system and, primarily, the set of ETM mechanisms investigated. From a European perspective, these areas embrace the following domains of law and regulation:

- Consumer protection
- Data protection and privacy
- Intellectual property rights
- Competition issues, including network neutrality

6.1.1 Consumer protection

The European Union (EU) has released 10 principles of consumer protection [11] determining a minimum standard applicable in all Member States. Member States might implement more strict rules on consumer protection, and the specific Member State implementation of the 10 principles might differ in details from Member State to Member State, but the set of 10 principles is binding. These 10 consumer protection principles concern SmoothIT mainly in terms of contractual issues and fair commercial practices. Other concerns such as product safety are of secondary interest only.

A service provider (ISP, in our case) making use of SmoothIT ETMs needs to consider for the relationship with a consumer – *i.e.*, with a service customer that consumes contracted services for private, non-commercial use – the following:

- The terms of an ETM mechanism in use have to be included in the contract the service provider has with the consumer. This has to be done in a visible, understandable (plain language) manner, not in paraphrased legal speak nor in a hidden way (*e.g.*, attachment to a contract). Transparency is especially important in case an ETM is offering economic incentive to a consumer and if this incentive is in direct relation to a service provider's main contractual obligation. For instance, in case of BGP-Loc employed by the user's client in conjunction with a discount, an ISP's main obligation of providing connectivity is (potentially) influenced by the ETM.
- In case an ETM is covered by a contractual agreement between the service provider and consumer, the consumer has to be granted the right to retreat from such a contract within seven working days. If services are provided cross borders, a consumer has to be provided effective means for dispute resolution. The EU Brussels I 0 and Rome I [16] regulations include rules of jurisdiction and applicable law, respectively, for contracts with consumers.

- ETM services shall not be activated automatically, *i.e.*, without consumers having expressed intent to obtain them.

6.1.2 Data protection and privacy

The right to privacy determines one of the basic rights in the EU. It is explicitly addressed in article 8 of the Charter of Fundamental Rights [17]. This basis has been taken up in more detail in two EU directives, the first on the processing of personal data (directive 95/46EC [12]) and the second on the processing of personal data and the protection of privacy in the electronic communications sector (in relation to telecommunications services as opposed to content services; directive 2002/58/EC [15]).

The legislation mentioned is of principle relevance to a service provider making use of ETM as [12] governs “*the processing of personal data wholly or partly by automatic means, and to the processing otherwise than by automatic means of personal data which form part of a filing system or are intended to form part of a filing system.*” This leads to an understanding needed in terms of what is defined as personal data. Personal data is defined in [12] to embrace “[...] *any information relating to an identified or identifiable natural person [...]*”. In case an ETM treats data that can be traced back (directly or indirectly) to an individual, *i.e.*, a service user, such data is personal data. By consequence, the service provider would have to respect the following principles:

1. Any personal data must be collected for an explicit purpose. Hence, the service provider has to be clear about each data item’s relation to the ETM’s purpose. Anything beyond this purpose is to be avoided. Alternatively, a data collection process may be freed from personal data completely.
2. A service provider is required to obtain explicit and unambiguous consent by an individual about whose personal data is collected and processed by an ETM. In case special categories of personal data are processed such as data revealing religious beliefs or concerning health, even more strict requirements may apply. Again, in order to avoid such requirements, a service provider may avoid the collection and treatment of personal data completely – which however would include that an individual cannot even be traced indirectly.

Therefore, depending on a specific ETM, data protection and privacy implications may be considerable. For instance, if an ETM involves caching techniques that allow for (indirect) tracing of individuals, *e.g.*, by combination with IP flow data, a service provider would have to ensure an individual’s rights related to data protection and privacy as stated in [12] and [15]. This is of special concern to the ISP-Owned Peer (IoP) and Highly Active Peer (HAP) ETMs, since the ISP needs to, at least temporarily, store IP addresses or even account for peer activity, in case of the latter. In case of BGP-Loc, providing locality of IP addresses also needs to be done carefully, *e.g.* coarsely enough, not to reveal privacy-sensitive information about users.

6.1.3 Intellectual property rights

Intellectual property rights (IPR) embrace a set of intangible assets in terms of copyright, patents, trademarks, and further related rights (*e.g.*, domain names). For SmoothIT, patents (and less probably trademarks and other neighboring rights) may constitute instruments of interest to protect innovation achieved within the project. With respect to legal or

regulatory impact, however, copyright is seen as the most important area in IPR that a service provider making use of ETM might have to be aware of.

Intellectual right and other legal issues have been identified for the case of P2P networks from the very beginning of their existence [9]. P2P networks employ fully distributed architectures and allow commodity PCs of end-users to cooperate and exchange content capitalizing the bandwidth of end-nodes. The problem lies in the fact that P2P networks have been associated widely with the sharing of illegal and not copyrighted content.

In the EU, there is a row of copyright-related legislation in place. Some of which are dealing with procedural aspects and Member State copyright enforcement harmonization (e.g., Directive 2004/48/EC), while others are concerned with the actual term of protection (e.g., Directive 2006/116/EC). The main EU legislation on copyright in the information society is found in Directive 2001/29/EC [13].

The directive differentiates three types of rights in relation to copyrighted material: reproduction right, right of communication to the public, and the distribution right. In general, these rights are attributed exclusively to the copyright owner, normally the author of a work. Any ETM mechanism making use of caching and/or replicating techniques and/or being involved in the support of content lookup, retrieval, or distribution may risk to inflict with the right to reproduction and the right to distribution.

Especially in cases where a service provider uses an ETM – which it may even promote actively and for which it may even provide financial or technical incentives to users – in a peer-to-peer file sharing context, the service provider may see itself liable for supporting copyright infringements. This may be of particular concern to the IoP ETM, since the ISP would store content that potentially infringes copyright law, but also with BGP Locality (BGP-Loc) and HAP ETM mechanisms, since they give incentives for users to participate in file-sharing, which may include illegal content. In this light, Art. 6 [13] states that “*Member States shall provide adequate legal protection against the circumvention of any effective technological measures, which the person concerned carries out in the knowledge, or with reasonable grounds to know, that he or she is pursuing that objective.*” Even an ETM in support of the distribution and reproduction of copyrighted material might not be perceived to qualify as a means to circumvent a technological measure protecting copyright (e.g., Digital Rights Management, DRM, technology), a service provider might be attributed with a previous act (by a third party) of circumventing DRM.

However, since Content Distribution has been identified as a key revenue opportunity for ISPs and Telecom Operators, research has been done recently towards the development and deployment of authentication and billing systems along with distributed peer-assisted content distribution. Such a capability could allow ISPs to deal with the intellectual property rights issue as well as gain a larger portion of the revenue paid for content. For instance, in [6], an architecture is proposed for helping ISPs and carriers exploit advantages of P2P content distribution networks within a legal framework. Additionally, the use of caches such as IoP to assist the content distribution is also considered in this architecture.

In conclusion, a service provider has to assess the trade-off between economic efficiency by help of the ETM in question and a potential involvement in copyright infringement in certain application context or the deployment cost of an authentication and billing architecture for legal content distribution. The respective risk assessment will reveal whether there is a net benefit in a given scenario. What a service provider, however, will have to consider in any case is that an ISP’s typical defense against copyright claims – that of ignorance

due to a lack of knowledge about potential copyright infringements, since an ISP is providing connectivity and the content-agnostic forwarding of data bits only – will not hold any longer as soon as a service provider is interfering on content level. Nevertheless, in any case, the aforementioned legal issues do not diminish the value of ETM mechanisms and their benefits, because peer-assisted approaches are applied extensively for the distribution of legal content too.

6.1.4 Competition

Competition law in the EU is mainly concerned with avoiding significant market power of dominating players and with ensuring the free flow of services within the internal European market. In this context, an area of specific interest to SmoothIT and its ETMs is in the network neutrality debate. For instance, a promoted preference ranking of internal (ISP-owned) peers may raise anti-competitive concerns with service providers in other autonomous systems yet covering the same customer base, especially when a user is provided by its service provider with economic incentives to consider the preference ranking. This is also relevant to the BGP-Loc ETM, since it is shown to improve performance of P2P applications particularly in larger ISPs while potentially hindering users of smaller ISPs.

At this stage, there is no specific network neutrality act in place in the EU. Political discourse has started, while it remains unclear what its outcome will be. For telecommunication services, the Access Directive [14] ensures non-discriminating access and interconnection. For the specific case of inter-ISP optimizations as foreseen by the SmoothIT system and the respective ETMs, a full regulative assessment cannot be performed unambiguously at this time. Moreover, when an ISP makes an agreement directly with an overlay provider, to mutually apply an ETM mechanism, this could potentially raise network neutrality issues, due to the performance advantage of the ETM-capable application to the other ones. However, this may not be an issue in case ETM is offered as a possibility to all overlay providers. As a general advice, however, a service provider is recommended to assess compatibility with competition law when making use of an ETM related to differentiated pricing, adapted market behavior of an actor, or economic incentives provided to a service user.

6.2 ETM – Potential Risks

A SWOT analysis has been performed by different SmoothIT partners and it was detailed in [21] This analysis has considered main perspectives:

- Scenarios considered for the specification of the SmoothIT solutions and the migration of these solutions to other scenarios and/or applications.
- The results of the different evaluation activities (analytical studies, simulations, and trial).
- The complexity to adapt an application to use the SmoothIT based solutions.
- The way related activities are being done in the rest of the Internet environment, how they are evolving and whether and how they can compete with SmoothIT based solutions.

The main conclusions of this analysis are shown in Table 19.

Table 19. SmoothIT SWOT analysis

Strengths	Weaknesses
<p>A large variety of ETM mechanisms have been investigated.</p> <p>Simulations have been performed in small, medium and large scale scenarios.</p> <p>Flexible architecture.</p> <p>Benefits of ETM mechanisms have been demonstrated in different scenarios.</p> <p>An easily extensible simulator has been developed.</p> <p>BGP-Loc and Inter-SIS are compliant with the ALTO approach and are part of the standardization activities</p>	<p>ETM mechanisms only simulated and also tested in the internal trial but not deployed in large scenarios yet.</p> <p>The SmoothIT system implemented has not yet reached product grade.</p> <p>There is no one-size-fits-all ETM solution.</p>
Opportunities	Threats
<p>Filing patents for ETM mechanisms.</p> <p>New developments could be done based on ETM mechanisms to enhance cache techniques.</p> <p>Potentially easy extension to include non-P2P applications.</p> <p>ETM mechanisms could be considered by an operator in order to reduce its operational costs.</p> <p>High interest in SIS-like solutions due to ALTO.</p>	<p>Competition with other platforms.</p> <p>Legal issues.</p> <p>Decreasing market relevance of P2P traffic.</p> <p>Implementation and maintenance costs can prevent operators to deploy these solutions.</p> <p>Decreasing inter-domain traffic costs.</p>

This analysis has revealed that within the SmoothIT Project not only several ETM mechanisms were evaluated, but also a significant part of them have been simulated, implemented, and developed in short, medium and large scenarios. The embracing insight gained and experience made within these activities is of highest importance for exploitation purposes since the benefits identified of the ETM mechanisms in [19] could be achieved in real networks — an achievement of great value with respect to future commercialization of SmoothIT innovations. Nonetheless, the implemented SmoothIT system constitutes a prototype, which is the contractually envisioned outcome of an FP7 research project. This research prototype, thus, is not an industry-grade product, and so the SmoothIT project has identified the according steps needed towards commercial exploitation. For example, the SmoothIT system will have to deal with some important issues such as security or intellectual rights in order to reach the mentioned commercial product grade.

6.2.1 Architecture security

The security aspects considered during the implementation of the SmoothIT architecture are based on the ITU-T Recommendation X.805. This document “*defines the general security-related architectural elements that when appropriately applied can provide end-to-end network security*” [7].

The recommendation defines vulnerability of the communications systems, threats, attacks and three security layers: infrastructure security, service security and application security. Eight security dimensions can protect a communications system against passive and active threats. They should be considered during a new communications system design. The security dimensions are:

- Authentication,
- Access control,

- Non-repudiation,
- Data confidentiality,
- Communication security,
- Data integrity,
- Availability, and
- Privacy.

In deliverable D3.2 „Initial Documentation of Engineering and Implementation” [34] we presented the security dimensions, also called: requirements and SmoothIT external interfaces (Table 2 of deliverable D3.2). Some possible solutions to protect external interfaces against threats were considered.

With reference to the SmoothIT architecture, four external interfaces were indicated (we assumed that external interfaces are open to attacks). The outside interfaces are:

- Admin interface,
- Inter-SIS interface,
- Network interface, and
- Overlay application interface.

To ensure security of the SmoothIT architecture, the security component was designed and implemented. The main objective of the Security component is protection against active and passive threats. The security component supports all implemented ETM mechanisms. Below, the security dimensions implemented by SmoothIT are collected:

- **Authentication** refers to logins and passwords. The credentials are stored in the database but in a secure format. The MD5 hash algorithm [8] is used to hide plain-text passwords. This algorithm calculates 128-bit hash. Hash function is a one-way function — it is infeasible to find an original message (password) that has a given hash. Therefore, even if an intruder has access to server, it will not be able to obtain the passwords.
- **Authorization** (access control) is based on specific roles and permissions. In particular, an administrator has permissions to read and change internal configuration of the SIS server. Roles and permissions are stored in the database.
- **Non-repudiation** is ensured by means of a dedicated mechanism that assures the accounting service. This service is able to collect all necessary information about used resources and established client-server connections (identity of entity, date, event name, result of event, etc.).
- **Data confidentiality and data integrity** assures secure transfer of sensitive information without the imperceptible change of information. It is integrated by means of encryption and signing provided by the WSsecurity services of JBoss. The asymmetric cipher was considered to protect the data: PKI (*Public Key Infrastructure*)

and X.509 certificates. This solution is based on the mechanisms implemented with the SSL (*Secure Socket Layer*) technique.

- Sometimes an administrator may close the admin web console without logging out. In that situation the session is still open and an intruder can intercept authorized session without authentication. Therefore a special mechanism that protects the system was implemented: the security component does not allow caching passwords in order to avoid unauthorized access to system. This is one of those services, which ensures communication security.

6.2.2 ETM vulnerabilities

To avoid the violation of security, the non-functional requirement R15 was defined (considered in the deliverables D3.1 and D3.4 “Economic Traffic Management Systems Architecture Design” [35], [36]). It defines basic security services which should be implemented to ensure data protection. To meet this requirement, the security component with many security services was designed and implemented.

During the implementation of security the crucial issue is proper balance between security and performance of the system. Too many protection mechanisms implemented in the protocol can completely degrade efficiency. Therefore, we deployed crucial security services. Additionally, a system administrator can turn on or turn off these services.

Security mechanisms and algorithms can protect the system against many attacks but we are not able to create an absolute secure system. Therefore we have to be aware of passive and active threats, which can be preformed against SmoothIT system with ETM mechanisms.

- **Denial of Service (DoS)** and **Distributed DoS (DDoS)** attacks can be performed to prevent or inhibit normal use of server. Such attacks could be performed by flooding the system with unwanted traffic. For example attackers can send a few thousand requests in the same time by means of a network of bots (computers with malicious software). The system will not be able to work properly. Availability assurance, additional network devices (i.e. Intrusion Detection/Prevention Systems) as well as right architecture of system can decrease the probability of successful DoS and DDoS attacks.
- When an intruder pretends to be a trusted entity, we have **masquerade attack** (or spoofing). An attacker can intercept request from client and introduce itself as a trusted entity. Authentication service is the best solution to prevent right entities from these kind of attacks. Nevertheless, users of the system may decrease the security level of this service (i.e. by using the simple and dictionary passwords). Then, an intruder is able to collect some privacy and/or confidential information from system.
- Many different **overflow attacks** exist in network environment. Attackers could try to modify original format of message (i.e. the number of peers which have desirable content) and take control of system or cause abnormal behaviour of system be done. Therefore, very important issue is a high quality implementation process to avoid exceeded data fields in exchanged messages in network.

- An intruder could be interested in **changing content** of exchanged messages in network. The effect of this attack could be ineffective management of overlay network. Good protections from the modification of content are encryption (by means of ciphers) and data integration (by means of one-way hash functions).
- **Eavesdropping** is an attack when an intruder listens to things he/she is not supposed to hear. Even if the content of messages are encrypted, by means of this kind of attack some privacy information about interests and behaviours of entities can be collected.

6.3 SmoothIT's Societal Embedding

In this chapter, the set of major and SmoothIT-relevant legal issues and security consideration was discussed. Four different domains of law and regulation were presented: consumer protection, data protection and privacy, intellectual property rights, and competition issues (including network neutrality). Also a security architecture was presented as well as general ETM vulnerabilities (problems with availability, masquerade, overflow attacks, changing content, and eavesdropping).

7 Overall Assessment

The final Triple-Win assessment of ETM mechanisms proposed in the project was intended to be based on a variety of evaluations, encompassing numerous simulation scenarios, modeling work, as well as the internal and external trials. Different evaluation methods provide a broad view of the system and make it possible to have a complete vision of the system. The general approach was to use simulations and models to evaluate a number of factors influencing the performance of ETM independently. Under the well-defined assumptions (e.g., stable environment, swarm size, or peer behavior) a selected system parameter (or sometimes more of them) was changed in order to evaluate its impact on the system. Although the simulation scenarios were defined carefully to assure they reflect the real system in detail, the trial experiments were designed to verify the system performance in a real environment of the Internet. However, due to the described problems faced with the external trial, the latter provides only rough qualitative assessment, which is considered as supplementary and inadequately reliable. As described above (Section 5.4.4), the simulation and trial results cannot be compared directly.

The metrics of Triple-Win are provided in Section 4.1.1. Metrics of “Win” for an ISP and end-user are well defined. They were used in all evaluations. During the project conduct, the metric of “Win” for an overlay provider appeared to be less obvious. Most of simulation experiments involved BitTorrent file sharing application. In this case, the overlay “Win” was not always directly evaluated. It was assumed that it is mostly aligned with end-user’s “Win”. If the download time for all users is decreased on average, it is a benefit for the whole swarm. As a consequence, the content availability is increased in the overlay. Therefore, “Win” situation for overlay is identified with “Win” for end-user, in most cases.

A more interesting aspect is the assessment of “Win” for the overlay provider in the case of a video application. In addition to the above “Win” metric used for file-sharing application, we may evaluate the load of content provider servers. It was done only for selected simulation scenarios for the HAP mechanism. It should be noted that it raises an interesting research topic related to Content Distribution Networks (CDN). Especially, the combination of CDN and P2P technologies with support of overlay traffic management methods (including such as considered in SmoothIT) are currently on focus. However, these subjects were not explicitly addressed by SmoothIT.

Table 20 presents an overview of the methods used to evaluate ETM mechanisms. It summarizes the most significant ETM mechanisms that were introduced and developed in the project. Three of them, namely BGP-Loc, IoP, and HAP were implemented within the main SmoothIT software release and evaluated in the external trial. Actually, IoP and BGP-Loc were deployed together. These ETM mechanisms were most widely evaluated. It should be noted that during the last year of the project, a lot of attention was devoted to ETM mechanisms based on the Inter-SIS communication and collaboration between ISPs. As a result of the work on this topic (which had started in the second year), three Inter-SIS based mechanisms were recognized as very promising and giving opportunities to improve locality awareness mechanisms: (a) cost and route asymmetry aware peer rating, (b) Peer grouping, and (c) collaborative BGP-Loc taking into account business agreements among ISPs. These were also quite widely evaluated by simulations. An Inter-SIS communication protocol was proposed and specified [33]. This work was also followed by implementation of an Inter-SIS protocol. Also a peer rating functionality that uses cost and route asymmetry information was implemented.

Table 20. A map of ETM mechanisms considered in the SmoothIT project

	BGP-Loc						5	Win	Win/no-lose ¹	
	IoP						5	Win ²	Win	
	HAP							Win	Win	
Inter-SIS	cost/asymmetry aware			3				Win	Win	
	Peer grouping							Win	win ⁴	
	Collaborative BGP-Loc							Win	Win/no-lose	
	Information asymmetry							Win/lose	Lose/win	
	QoS awareness							Win/no-lose	Win/no-lose	
		specification	model	simulation	implementation	internal trial	external trial	ISP	end-user	
		General assessment								

1. "Win" for user depends on a swarm size, AS, and access rate. Especially peers in large ASes win. Performance for peers in small ASes may be improved e.g., by Peer grouping mechanism
2. "Win" is more significant for small ASes, especially if combined with locality awareness
3. Implementation based on SmoothIT release 2.2, developed independently of external trial activities
4. Peers in small ASes win in comparison to BGP-Loc without inter-SIS collaboration (when they may lose). Resulting situation for peers is mostly "no-lose". The same for the whole swarm on the average.
5. BGP-Loc and IoP were deployed together in the external trial

As presented in Table 20, proposed ETMs were evaluated in various ways:

- Simulations
- Analytical models,
- The internal trial, and
- The external trial.

The primary goal of both simulations and modeling was to assess a Win-Win situation associated with an ETM mechanism deployment and its robustness. Numerous simulation scenarios were defined to verify ETM mechanisms under various conditions, as described in Section 4. The only ETM mechanism not evaluated by simulations is DynLoc, which was modeled and verified numerically with a limited set of experiments, involving a simple generator of inter-domain traffic. This partial numerical evaluation dealt with inter-domain traffic reduction in terms of 95th percentile metric. It addressed "win" for the ISP only, for which it provided positive evidence. This is intuitively expected since DynLoc applies locality only when inter-domain traffic is high. For the time intervals that locality is employed, users also are expected to benefit if BGP-Loc is in a "win" for the same set-up. Overall, it can be argued that DynLoc is a "win" for the end-users, whenever BGP-Loc is a "win" too.

Analytical modeling has also significantly supported evaluation and dimensioning of IoP, by means of an approximate Markov model quantifying the transient evolution of a swarm with and without an IoP. Moreover, an innovative unified game-theoretic framework was developed to analyze ISP interactions and dynamics due to ETM.

The primary goal of the internal trial was to evaluate the scalability and stability of a SIS implementation, a Tribler client and an IoP prototype. The results [20] were in general satisfactory and informative. The SIS implementation was proved to be scalable and being

able to provide response up to 20 requests per second. The Tribler client proved to provide reasonable performance and high stability. The internal trial conclusions were taken into account during the preparation of the external trial. Especially, detected problems with the SIS implementation were solved before porting the prototype to the external trial. Additionally, the internal trial evaluated the performance of two ETM mechanisms: BGP-Loc and IoP. Experiments confirmed the performance predicted by simulations and showed that both mechanisms can lead to a Win-Win situation (ISP/end-user). The effect of inter-domain traffic reduction grows if the domain implementing BGP-Loc has more seeds. The usage of the IoP combined with the BGP-Loc ETM has been demonstrated as an effective mechanism to reach a Win-Win scenario [20]. Moreover, an small test-bed was implemented for exploitation purposes, where the efficiency of the IoP and BGP-Loc combinations were evaluated. The test-bed emulated a local ISP with a few peers. However, peers joined a real swarm, so the conditions of the swarm could not be controlled. Even with the small-scale set-up and the not totally controlled environment, the IoP demonstrated good properties since, under certain circumstances, it offered a decrease of up to 60% for the P2P inter-domain traffic and up to 50% of the average download times.

The goal of the external trial was to evaluate not only a Win-Win situation in an operating network but also to verify system’s robustness, scalability and usability. These issues were described in Section 5.4.5, devoted to external trial results.

Table 20 provides also a very general and concise assessment of “Win” provisioning opportunity of ETM mechanisms. Based on latest simulation results, it updates the assessment provided in Section 6.6 of deliverable D2.4 [2].




Figure 13. The idea of a win-win space map

7.1 Assessment of ETMs

The remainder of this section provides a summary of Triple-Win assessment based on input data from all evaluation methods. The assessment of each particular ETM is provided below. In general, a qualitative assessment of ETM mechanisms is presented using traffic savings and user performance as metrics of “win”. In order to present the range of benefits for ISPs and users offered by ETM mechanisms and to compare them, the ETM mechanisms are positioned on a Win-Win space map. The idea of the space map is introduced in Figure 13. The x-axis reflects the situation of a customer while the y-axis presents an ISP’s situation. CAPEX and OPEX provided in y-axis label have an informative meaning. Detailed cost-benefit analysis is provided in Section 8.

7.1.1 BGP-Loc

BGP-Loc ETM was broadly and extensively evaluated, being the basic ETM mechanism. The impact of several factors was taken into account including: peer distribution, ETM deployment by ISPs, impact of access bandwidth bottleneck, etc. The most interesting results are provided in Tables 21 and 22.

It was recognized that in the case of a homogeneous peer distribution end-users perceive a “No-Lose” situation while the ISP “Wins”. An ISP’s benefit grows with its local swarm size. It was also shown that an inter-domain bottleneck has an impact on the users’ situation. Their situation changes from no-lose to win if the inter-domain links become the bottleneck (Figure 14). However, since homogeneous peer distribution is hardly realistic in the Internet, heterogeneous peer distributions were considered.

Table 21. Summary of assessment of "Win" for BGP-Loc: Impact of peer distribution and ETM deployment

	Homogeneous peer distribution; All promote locality	Heterogeneous peer distribution; All promote locality		Heterogeneous peer distribution; Only selected ASes promote locality	
		“large” AS	“small” AS	AS promoting locality	AS not promoting locality
ISP	Win	Win	Win/ No-Lose	Win ¹	No-Lose
peer	No-Lose	Win	Lose	Win	Lose

1. Larger savings for larger ASes and if more ASes promote locality

Table 22. Summary of assessment of "Win" for BGP-Loc: impact of access bandwidth

	Two classes of ISP		Each ISP has two types of access links: slow and fast	
	ISP with slow access links	ISP with fast access links	slow access users	fast access users
ISP	Win	Win	Win	
peer	Lose ¹	Win ¹	Win	Lose

1. Lose for overlay as a whole

In a realistic scenario derived from swarm measurements the ISP deploying BGP-Loc may always expect to benefit. It was also recognized that if some ISPs deploy ETM, then this strengthens the incentive for other ISPs do so as well (to achieve inter-domain traffic reduction and improve performance of its customers). However, the situation of ISP’s peers

is not obvious. Depending on the conditions, their situation after ETM deployment may be either “Win”, “No-Lose” or even “Lose”. The following main factors influencing user performance were recognized: local swarm size, size of the whole swarm, inter-domain bottlenecks, and access bandwidths. The positive effect of BGP-Loc on peers’ performance is stronger if the inter-domain link is the bottleneck. If the bottleneck effect grows, the situation of users moves from “No-Lose” to “Win”. It was also shown that peers in a domain sharing a small fraction of a large swarm may even lose. The impact of the local swarm size on Win-Win is presented in Figure 15. . The impact of the access link speed is summarized in Table 22.




Figure 14. Win-Win space map for BGP-Loc ETM, homogeneous scenario

Summarizing, BGP-Loc offers a clear Win-Win situation for large ISPs with a large local share of a large swarm. The main concern about BGP-Loc is that in some cases (i.e. of a small ISP) the “Win” for end-users might be little or non-existent. Then, users have little incentive to implement locality. If users have better performance while not using ETM, they will not cooperate with their ISP. Therefore, an ISP should deploy additional mechanisms to make the situation beneficial for its customers as well. An example is Peer grouping, an ETM using an Inter-SIS collaboration. It should be also noted that it is very difficult for a user to recognize the “win” since the user is participating not in a single but in a large number of swarms and the perceived “win” is an average.




Figure 15. Win-Win space map for BGP-Loc ETM, heterogeneous scenario

7.1.2 Inter-SIS based extensions to locality awareness

Investigating the mutual influence of deployment of ETM mechanisms in different domains, the project recognized that collaboration between ISPs may provide further benefits for all players. The basis for such solutions is an information exchange between SISes located in different domains. For this reason a generic Inter-SIS communication protocol was proposed. The information exchange between SISes may be used for a variety of mechanisms supporting BGP-Loc. Three solutions were developed by SmoothIT:

- Cost and route asymmetry aware peer rating algorithm ,
- Peer grouping mechanisms,
- Collaborative BGP-Loc according to business agreements.

A brief assessment of those solutions is provided below.

7.1.2.1 Cost and route asymmetry aware peer rating

Even if an ISP applies locality promotion, a portion of the traffic must be exchanged with remote peers. Often an AS is connected to more than one (sometimes several) transit and peering ASes. The cost of the traffic transfer on various links may be different. The aim of this ETM mechanism is to manage the inter-domain traffic in such a way that it is shifted from expensive to cheaper links. For this reason routes to remote peers must be known. Then, the peer rating algorithm triggered by the SIS prefers peers to which a routing path passes a cheaper inter-domain link. It was shown that Win-Win is always achieved by the ISP as well as the peer located in the AS implementing this ETM (Table 23). An ISP benefits from further cost savings while peers obtain reduction of download time (equal or greater than in the scenario with plain BGP-Loc). ASes not implementing this mechanism ex-

perience “No-Lose” situation. It was, however, shown that in some cases the ISP may Lose. For example, it may happen, that two ISPs apply contradictory policies: one tends to reduce the traffic on a given link while the other is interested in increasing it. Such a situation may lead to unintended results and Lose situation for one of them or both. Such scenarios were evaluated as well. Therefore, it is an incentive for cooperation between ISPs and for agreeing on and applying a non-contradictory policy.

Table 23. Summary of "Win" assessments for cost and route asymmetry aware peer rating algorithm (ETM based on an Inter-SIS communication)

	One AS implement Inter-SIS and cost-driven peer ranking		All ASes implement Inter-SIS and cost-driven peer ranking
	AS using ETM	AS not using ETM	
ISP	Win	No-Lose/Lose	Win
Peer	Win	No-Lose	Win

Since route asymmetry is common in the Internet, an Inter-SIS protocol is needed to obtain complete routing information. Especially a full information about downstream path, e.g., the number of inter-AS hops, AS numbers, and the last downstream hop, is usually not available locally. For some of the paths it is possible to get some information from Looking Glass Servers. However, the most reliable information can be provided by communication between cooperating SISes [32].




Figure 16. Win-Win space map for BGP-Loc and Peer grouping

7.1.2.2 Peer grouping

As discussed in the assessment of BGP-Loc mechanism, in case of heterogeneous peer distribution, peers located in ASes with a small fraction of a large swarm Lose if BGP-Loc ETM is deployed. At the same time, peers located in a “large” AS, i.e., an AS containing large number of pees, Win. To avoid this drawback of BGP-Loc, a mechanism called Peer grouping was proposed. Grouping of “small” ASes mitigates this disadvantage and results in at least No-Lose situation for users. In turn, the inter-domain traffic may remain at the same level or be slightly increased but a “small” ISP still Wins and, moreover, its customers are not disappointed (receive at least No-Lose instead of Lose). The latter issue can be, therefore, considered as an additional measure of ISPs Win. By implementing Peer grouping, an ISP gains, since its customers receive an incentive to cooperate and implement locality (in scenario without Peer grouping they lose and are reluctant to implement locality). Benefits introduced by Peer grouping are presented graphically in Figure 16.

7.1.2.3 Collaborative BGP-Loc according to business agreements

Collaborative BGP-Loc distinguishes non-local peers according to the business relationships that the ISPs form with each other. Such a distinction is very important, because it leads to a significant reduction of the transit inter-domain traffic where charging applies. As shown in Table 24, large ISPs (Tier 2 in this case), benefit from elaborating the Collaborative BGP-Loc mechanism. When Tier 2 ISPs having a peering agreement and collaborate with each other exchanging information about their customer ISPs, they achieve a reduction in their costly inter-domain traffic. In this case neither their end-user customers nor their ISP customers (with their users) Lose. In case that Tier 3 ISPs collaborate with each other, there is no effect in the costly inter-domain traffic or in the performance of their end-users. Thus it is unlikely that Tier 3 ISPs will adopt this mechanism. However, if Tier 3 ISPs run Collaborative BGP-Loc, then Tier 2 ISPs have a clear incentive to run the mechanism, too. Indeed, this extra collaboration between them improves the costly transit inter-domain traffic of Tier 2 ISPs and the download times of their end-users. Since in this case Tier 3 ISPs are indifferent in adopting the approach, Tier 2 ISPs should incite it to their customer ISPs by sharing their own benefits with them.

Table 24. Summary of “Win” assessment for Collaborative BGP-Loc ETM (based on an Inter-SIS communication)

Collaborative BGP-Loc between:		Tier 2 – Tier 2	Tier 3 – Tier 3	Tier 2 – Tier 2 & Tier 3 – Tier 3
Tier 2	ISP	Win	No-Lose	Win
	peer	No-Lose	No-Lose	Win
Tier 3	ISP	No-Lose	No-Lose	No-Lose
	peer	No-Lose	No-Lose	No-Lose

In Table 25, a summary of “Win” assessment of the Splitting Chunks ETM is shown. Splitting Chunks is based on Collaborative BGP-Loc (being an add-on to it, in order to reduce traffic redundancy in inter-domain links), since (as shown above) it is beneficial for Tier 2 ISPs to implement the latter mechanism. When Tier 2 ISPs implement Splitting Chunks, comparing to Collaborative BGP-Loc, there is a Win-No-Lose situation, since it reduces the costly transit traffic without affecting the performance of its users. At the same time,

this action of Tier 2 ISPs affects neither the traffic nor the users' performance of the Tier 3 ISPs (who are customers of Tier 2). Tier 3 ISPs will not adopt Splitting Chunks due to slower download times of their users. In case that Tier 2 ISPs already deploy Splitting Chunks on top of Collaborative BNS-BU, Tier 3 ISPs are indifferent in adopting the approach. Again, Win-No Lose situations arise between Tier 2 ISPs and their users and also between Tier 2 ISPs and their customer ISPs. Therefore, Tier 2 ISPs should incite their customer ISPs to follow Splitting Chunks on top of Collaborative BGP-Loc by sharing their benefits with them.

Table 25. Summary of "Win" assessment for the Splitting Chunks ETM (based on Inter-SIS communication)

Splitting Chunks on top of Collaborative BGP Loc between:		Tier 2 – Tier 2	Tier 3 – Tier 3	Tier 2 – Tier 2 & Tier 3 – Tier 3
Tier 2	ISP	Win	No-Lose	Win
	peer	No-Lose	No-Lose	No-Lose
Tier 3	ISP	No-Lose	Win	No-Lose
	peer	No-Lose	Lose	No-Lose

Figure 17 presents the relation between plain BGP-Loc and two inter-SIS based extensions: Cost and route asymmetry aware peer rating and collaborative BGP-Loc.




Figure 17. Win-Win space map for BGP-Loc with inter-SIS extensions: collaborative BGP-Loc and cost and route asymmetry aware peer rating

7.1.3 QoS-aware ETM

The “carrier class overlay services” scenario of the QoS-awareness mechanism was analyzed by means of simulation experiments. In this solution the connections between local peers (peers from the same AS) are prioritized over the others. Priority means that the source node of a local connection receives a higher upload capacity in comparison to other connections. Such a method promotes locality, thus allowing for decreasing the inter-domain traffic in a network. In all experiments it was assumed that access is the bottleneck (homogeneous access bandwidth was assumed). Table 26 summarizes assessment of QoS-aware ETM. It should be noted that ISPs experience more significant traffic reduction when more ASes implement the mechanism. If not all ASes implement the ETM mechanism then download time decreases only in domains with QoS-aware implemented. The user benefit was observed to be lower if the number of ASes deploying this ETM mechanism increases [2]. However, download time is always lower in comparison to a scenario without ETM at all. A win-win space map for QoS aware ETM is provided in Figure 18.

Table 26. Summary of assessment of “win” for QoS-aware

	Heterogeneous peer distribution	Homogeneous peer distribution
ISP	Win	Win
peer	Win	Win

The mechanism generates win-win only if a network experiences congestions. Otherwise, it is neutral, i.e., offers no-lose/no-lose situation.




Figure 18. Win-Win space map for QoS aware mechanism

7.1.4 IoP

The IoP is a resourceful entity aiming to primarily help local peers while localizing traffic too. Thus, in the case of IoP we naturally observe a Win-Win situation in all evaluated simulation scenarios. Especially, it provides a significant improvement of end-users' completion times as well as a considerable inbound inter-domain traffic reduction of the AS that deploys the IoP and, therefore, a reduction of charges for inter-domain traffic. under different charging schemes based on *incoming traffic*.

An important mechanism associated with IoP is the unchoking policy. If applied, only local peers (i.e., located in a domain the IoP is deployed in) may be unchoked by IoP. Under this restriction, the remote peers are not served by the IoP. In consequence, a significant decrease of inbound as well as outbound inter-domain traffic is observed. Therefore, the benefit of ISP is greater than in the case of IoP without unchoking policy. The ISP benefits from the reduction of charges for inter-domain traffic under charging schemes based also on the *outgoing traffic*.




Figure 19. Win-Win space map for IoP

The aforementioned conclusions were drawn by analyzing the impact of the IoP in a single swarm. With multiple swarms in place (as is naturally the case), the basic IoP operation is complemented by the swarm selection and the bandwidth allocation strategies, aiming to optimize the effect of the mechanism. The Swarm Selection strategy (SwS) may be based on various criteria. Selection of swarms with larger content files, low mean inter-arrival times of leechers (i.e., swarms with more leechers), or low mean seeding time (i.e., swarms with less seeders) was done. Evaluation of several scenarios showed that **Win** is achieved by both ISP and local users, including those in the swarms not joined by the IoP. It was also learned that IoP's impact is more significant when it joins the swarm with higher

capacity needs. An additional conclusion is that inter-domain traffic is mostly affected by the mean seeding time, while performance is mostly affected by the mean inter-arrival time.

Furthermore, the Bandwidth Allocation (BwA) strategy can offer a considerable inter-domain traffic reduction. It was learned that IoP capacity allocation to swarms with a small fraction of local peers provides more inter-domain traffic savings.

The IoP Win-Win evaluation is also summarized in Figure 19. . The upper oval reflects relative traffic savings. Since deployment and maintenance cost of IoP are high, a qualitative assessment including CAPEX/OPEX is also provided (lower oval). More details can be found in section 8.

7.1.5 Combination of BGP-Loc and IoP

Additional explanation is needed in the case of combination of IoP and BGP-Loc. If those ETM mechanisms are deployed together, a significant improvement in comparison to no ETM deployment at all was also observed. It is clearly a Triple-Win situation. However, as presented in Section 3.3.2, comparing two scenarios (a) BGP-Loc and (b) IoP+BGP-Loc, the inter-domain traffic reduction achieved by BGP-Loc for large ISPs may be slightly lower in the latter case.

7.1.6 HAP

Another ETM mechanism based on the offering of additional resources is HAP (Highly Active Peer). Instead of deploying IoP, an ISP may use the HAP mechanism. A user cooperating with the ISP is rewarded with *free* additional access bandwidth. Then, this user serves other peers in a similar way to the IoP: it has higher bandwidth resources and seeds the content, thus acting as a cache and offers the most popular as well as less popular content. In this case, legal issues related to IoP are avoided; see also Section 6.1. This solution does not require content-awareness by the ISP.

The HAP mechanism was widely evaluated by simulations (see Section 0) and in the external trial. It should be emphasized that HAP evaluation was done for both a file-sharing application and a video-on-demand application. The latter makes it possible to derive a direct evaluation of “Win” for an overlay provider. It was expressed by the reduction of the load of the overlay provider’s servers resulting in lower delivery costs.

Evaluation of HAP encompassed also HAP promotion and demotion algorithms and their influence on “Win” situation. These algorithms are based on statistics collected by ISP’s metering tools (total upload traffic, local upload traffic, upload bandwidth) and overlay statistics (e.g., seeding ratio).

The experiments showed that HAP offers a **Triple-Win** situation, especially if deployed with locality awareness and preferring locality-aware peers for being promoted to HAP. Then, all players benefit from HAP deployment. The overlay benefit (server load reduction) depends on the scenario: HAP selection mechanisms and locality awareness ratio. However, in all cases, except for the “boundary” scenarios (see below), all players benefit. The “boundary” scenarios were considered in a dedicated study in which HAP promotion was taken with and without BGP-Loc. Results are provided in Table 27. More detailed results are provided in Section 0.

Table 27. Comparison of static HAP promotion with and without locality awareness

Win for:	ISP overlay	Is overlay locality aware (BGP-Loc applied)?	
		Yes	No
Does ISP implement HAP?	Yes	Win	No-Lose
	No	Win	Win

Finally, the case of early adopter ISP was studied, as representative of partial deployment. If only one ISP deploys HAP mechanism it may achieve **Win**, **No-Lose** or **Lose** situation depending on its inter-connection agreements regarding the treatment of in- and outbound traffic. This is because an early adopter decreases inbound traffic but increases outbound traffic. As a result, depending on inter-connection agreement the situation might be as presented in Table 28. Note, that in case when the sum of inbound and outbound traffic determines the payments, a simultaneous increase in outbound traffic (generated by promoted peers serving remote peers in ASes without HAP promotion) negates the inbound traffic reductions (see [2] for more details).

Table 28. HAP early adoption “Win” assessment

Early adopter ISP pays for traffic:	Situation:
<i>Inbound</i>	Win
<i>inbound + outbound</i>	Lose or No-Lose
<i>inbound – outbound</i>	Win

At the same time, it generates **No-Lose** for other ISPs (not applying HAP promotion). However, it is always **Win** for the overlay.

Otherwise, in terms of qualitative assessment of traffic savings, IoP and HAP are hardly comparable since the traffic savings strongly depend on the capacity of the IoP on the one hand side and the number of promoted peers of the other hand side. Including, CAPEX/OPEX, however, HAP achieves traffic savings with less investment than the IoP. This is due to the fact that the extra resources offered by the ISP in the case of HAPs by upgrading their access profiles come at negligible extra costs to the ISP (see Section 8). On the other hand, an IoP requires investment in both hardware and bandwidth.

7.1.7 Overall Triple-Win summary

A space map of all ETM mechanisms proposed by the SmoothIT project is presented in Figure 21. It shows a Win-Win potential of the ETM mechanisms. In general, solutions based on extra resource offering content caching (IoP and HAP) would provide by definition considerable benefit for customers (peer Win) and overlay (especially HAP evaluation was focused on overlay’s “Win” directly). The situation for ISP is not clear, and depends on the inter-domain tariff: those mechanisms reduce inbound traffic, but in some scenarios increase outbound traffic. In turn, locality awareness mechanisms are by definition beneficial for the ISP. They offer significant cost savings related to inter-domain traffic reduction and management. However, the most significant effect is observed for ISPs hosting a large part of a large swarm. The benefit for ISPs having a small local swarm is lower. On the other hand, it was recognized that locality mechanisms not always provide “Win” for

customers. Sometimes they achieve a No-Lose situation, thus an incentive for users should be provided by the ISP or additional supporting mechanisms are needed. For this reason, solutions supporting BGP-Loc were proposed. Some benefit for players is expected if BGP-Loc is supported with inter-SIS based solutions. Inter-SIS ETM deployment will not require additional equipment since only new functionality is to be added to the SIS server. However, the implementation of the ETM is more complex and the OPEX may increase. Finally, the solutions will be beneficial for the ISP if additional cost savings will be higher than costs.




Figure 20. Win-Win space map of IoP and HAP

The solution is also to deploy a combination of locality promotion and caching solution, e.g., IoP+BGP-Loc. The combination will allow to exploit the best features of those two approaches:

- IoP and HAP mechanisms provide clear incentive missing from locality-aware ETM mechanisms.
- IoP and HAP mechanisms are most efficient for small to medium size local swarms.
- Locality-aware ETM mechanisms reduce inbound and outbound inter-domain traffic.
- Locality-aware ETM mechanisms are most efficient for large swarms.




Figure 21. Win-Win space map of ETMs proposed by SmoothIT

8 Cost-benefit Analysis

In the first part of this section, a cost-benefit analysis of the implemented ETM mechanisms is included, so as to provide an insight in monetary terms on how beneficial deploying one of the proposed ETM mechanisms is for an ISP. To do so, certain parameters are considered. These include the costs for the machines to host the system, the unit prices related to the inter-domain traffic, the percentage of the P2P traffic as part of the entire Internet traffic, as well as the effect of the ETM mechanisms in the traffic volumes on the inter-domain links and their 95th percentile. All these parameters are analyzed and certain directions are given of how they should be considered in the calculations of the incurred costs and benefits. Next, the analysis takes place where the break-even points are calculated. Finally, the obtained results are commented and a comparison with actual findings on the effects of the ETM mechanisms from simulative experiments is made.

In the second part, the economic perspective of the potential improvements through ETM, especially locality-awareness based ETM, is analyzed based on the current swarm mixture measured in the Internet. This analysis does not focus on the situation of a single ISP but on the general world-wide optimization potential through ETM based on locality awareness mechanisms like BGP-Loc.

8.1 Cost-Benefit Analysis of the SmoothIT Prototype

Based on a number of assumptions stated and argued for the analysis is performed on a numerical basis.

8.1.1 ETM mechanisms and cost categories

First, a breakdown of costs related to the ETM mechanism deployed takes place in this subsection. There are three ETM mechanisms currently implemented and studied: the BGP-Loc, the IoP and the HAP. For all the mechanisms, the existence of a server hosting the SIS is required. Moreover, for the IoP mechanism, additional software, hosted in the same but more powerful machine and implementing the actual IoP (i.e. the smart overlay cache/super-peer), is required.

For the SIS, the following cost categories are identified:

- *Server equipment*: the hardware required to host the SIS.
- *Installation*: the effort required by a technician to install and configure the SIS.
- *Operation and Maintenance*: the costs that are related to the everyday operation and periodic maintenance of the SIS.
- *Amortization period*: the period over which the respective costs are amortized.

The aforementioned cost categories also apply for the IoP. Equipment, Installation, Operation and Maintenance costs are applicable to the installation and operation of a smart overlay cache, most probably hosted to a different machine from the one hosting the SIS, mainly for scalability issues. For the case of the IoP, one additional cost is considered, the cost of the over-provisioned resources of the IoP, i.e. the increased link capacity of the

IoP. Note that this expense does not actually burden the ISP, since he owns the link. However, such a link is still considered as an expense, in the sense that the ISP “loses” the income he could attain from selling such a link to an actual customer.

Regarding the third ETM mechanism, the HAP, no additional costs were indentified. The main reason was that the underlying NMS mechanisms to realize the dynamic adaptation of access profiles of certain customers are considered as common and available. Thus, no extra costs for their realization are incurred.

Apart from the above costs, each mechanism introduces some benefits that need to be translated into monetary gains. In all three cases, the benefits appear in terms of inter-domain traffic reduction, which result in lower inter-domain costs. These savings are monetized in order to be comparable with the incurred costs.

Before proceeding with the benefits, a table is provided summarizing the different costs categories per ETM and also including an estimation of the actual values as an aggregation of the proposed costs suggested by the ISPs belonging to the consortium. From TID, the following input was received:

- SIS:
 - Equipment: 900€
 - Installation: 200€
 - Operation: 300€/year
 - Amortization period: 4 years
- IoP:
 - Equipment: 2000€
 - Installation: 200€
 - Operation: 250€/year
 - Amortization period: 4 years
 - IoP link (10 Mbps): 28€/month
- HAP:
 - No extra costs from the BGP-Loc case

At the same time, PrimeTel has provided the following:

- SIS:
 - Equipment: 2500€
 - Installation: 3000€
 - Operation: -
 - Amortization period: 5 years
- IoP:
 - Equipment: 2500€
 - Installation: 3000€

- Operation: -
- Amortization period: 5 years
- IoP link (10 Mbps): -
- HAP:
 - No extra costs from the BGP-Loc case

SmoothIT aggregated the proposed figures by considering the fact that in the IoP case, one machine could host both the SIS and the IoP, with higher however technical specifications. These prices are included in the following table. Note that the amortization period for the above equipment is considered to be 4 years.

Table 28. Cost categories per implemented ETM mechanism

ETM mechanism	Cost category	Cost
BGP-Loc	equipment	1000 €
	installation	300 €
	operation and maintenance (per year)	300 €
IoP	equipment	1800 €
	installation	300 €
	operation and maintenance (per year)	400 €
	access link (monthly price for a 10 Mbps symmetric link, averaged over 4 years)	28 €
HAP	(same as BGP-Loc case, no additional costs)	-

8.1.2 Transit prices

SmoothIT now comes to analyze the prices related to the inter-domain traffic so as to calculate, in monetary terms, the savings from the deployment of an ETM mechanism. Inter-domain expenses are tightly related with the contracts that define the pricing policies between an ISP and its higher-tier connectivity provider. These contracts differ from case to case, but all define a unit price for the traffic volume (average, maximum, 95th-percentile, or other) transferred in the inter-domain link (in the downlink or uplink direction, or the difference between these two).

According to DrPeering.net [23], the transit unit prices follow a continuously decreasing trend. This is depicted in the following table, regarding the Internet Transit Prices (1998-2014) at the U.S. Internet Region:

Table 19. Internet Transit Prices (1998-2014) at the U.S. Internet Region
(source: DrPeering.net)

Year	Unit Price (per Mbps per month)	Decline
1998	\$1200	
1999	\$800	33%
2000	\$675	16%
2001	\$400	40%
2002	\$200	50%
2003	\$120	40%

2004	\$90	25%
2005	\$75	17%
2006	\$50	33%
2007	\$25	50%
2008	\$12	52%
2009	\$9	25%
2010	\$5	44%
2011	\$3.25	35%
2012	\$2.34	28%
2013	\$1,57	33%
2014	\$0.94	40%
2015	\$0.63	33%

Additionally, the transit unit cost depends on the link capacity that is leased. If all the related fees are excluded and SmoothIT focuses only on the usage costs, then the following table reveals the structure of the tariffs in the inter-domain link market, considering the 2009 prices:

Table 20. Tiered Pricing Structure in 2009 (source: DrPeering.net)

Link Capacity	Unit price (per Mbps)	Difference (%) from base price
100 Mbps	\$ 9	
1 Gbps	\$ 8	11,11%
2 Gbps	\$ 7	22,22%
3 Gbps	\$ 7	22,22%
4 Gbps	\$ 7	22,22%
5 Gbps	\$ 7	22,22%
6 Gbps	\$ 6	33,33%
7 Gbps	\$ 6	33,33%
8 Gbps	\$ 6	33,33%
9 Gbps	\$ 6	33,33%
10 Gbps	\$ 6	33,33%

The previous tables reveal a lot about the transit pricing schemes, even though one cannot take for granted the above estimates and assumptions. In the following subsections, the calculations taking place are based on these tables, in order to calculate the gains from deploying one of the ETM mechanisms, during a 4-year period. Note that the prices are here expressed in US dollars but for the purposes of the analysis, they are converted to Euros.

Another important remark is that, though not stated explicitly, the above prices are related to the 95th percentile of the traffic flowing on the inter-domain link. The 95th percentile is the most common metric for traditional inter-domain charging schemes. SmoothIT has conducted some experiments to check the relation between an average decrease on the inter-domain traffic and the resulting reduction on the 95th percentile. Preliminary results have shown that there is no great divergence between these two. Hence, for simplicity reasons, SmoothIT considers only the average reduction on inter-domain traffic, achieved by the deployment of an ETM mechanism.

8.1.3 Internet and overlay traffic

Having commented on the inter-domain unit prices, it is important for the upcoming analysis to consider realistic values regarding the percentage of traffic that is generated by P2P applications as well as the overall trends of Internet traffic for the following years.

In the study provided by [25], the following figures are expected to characterize the global IP traffic and the P2P file sharing traffic within the next few years. Note that though the percentage of P2P traffic decreases, the absolute volumes increase with a compound annual growth rate (CAGR) of 16%.

Table 31. The evolution of P2P traffic percentage (source: CISCO)

Year	Global IP traffic (PB/ month)	P2P File Sharing traffic (PB/month)	Percentage
2009	10942	3510	32,08%
2010	15205	4157	27,34%
2011	21181	4839	22,85%
2012	28232	5531	19,59%
2013	36709	6305	17,18%
2014	47176	7303	15,48%

It is also worth noticing that global Internet traffic increases with a CAGR of 34%, and hence it is safe to assume that the inter-domain link capacities are increasing with the same rate. Note also that here only the P2P file sharing traffic is considered. Having in mind that P2P video streaming is constantly increasing during the last years (with a CAGR of 47% but with small absolute volumes at the moment), the portion of total P2P traffic will be higher than what is depicted in the above table. However, this traffic is not currently considered.

8.1.4 Analysis

In order to proceed with the cost-benefit analysis, the setup of the network environment is described. The objective is to have a realistic scenario, while using the figures and their projections as described in the previous subsections so as to arrive at useful conclusions. Then, the cost-benefit analysis takes place, considering the next 4 years, from to 2011 to 2014, a period equal to the amortization period of the required equipment.

8.1.4.1 Assumptions

First, the dimensioning of the network domain to be studied is considered. As a basis, SmoothIT uses the size of PrimeTel's network (see Section 5.1), but SmoothIT expands it so as to consider a bigger-sized ISP. During the external trial, one server was used to host the SIS and one machine to host the IoP. PrimeTel's inter-domain link is of capacity 1.5 Gbps. Considering that the capacity of an ISP's inter-domain link is indicative of the size of its customer base (since all the Internet traffic needs to be accommodated through that link), SmoothIT assumes that one machine is adequate to serve a population of peers for a domain with an inter-domain link up to 2.5 Gbps. Above that, a new machine will be required for every 2.5 Gbps of inter-domain link capacity.

The complete justification of the above assumption is as follows: during the internal trial conducted within the project, it was observed that the prototype SIS maintained a stable performance when it received requests of a rate up to 17 requests per second. PrimeTel's

inter-domain link is 1.2 Gbps. If SmoothIT assumes that the link is utilized 80% and that 40% of the traffic is P2P traffic (SmoothIT takes the “worst” case so as to provide a relatively conservative approach) then the rate of P2P traffic flowing into the inter-domain link is 480Mbps. The average access speed of a PrimeTel’s customer is 2 Mbps, hence SmoothIT assumes that at most 240 P2P transactions may take place. Considering that a standard P2P client contacts every 30 seconds the tracker to request new overlay information (new peers to contact), SmoothIT assumes that this is the same rate that a peer contacts the SIS. Hence, the average transactions/requests that the SIS receives is 8 requests per second. Thus, SmoothIT assumes (from the previously observed limit) that the SIS can serve up to twice the number of peers, hence a domain with twice the capacity on the inter-domain link. This results in a link of 3 Gbps, but adopt the more conservative limit of 2.5 Gbps.

Having this as a basis, SmoothIT considers a “generic” ISP with a 4 Gbps inter-domain link. Keeping in mind the relation that was considered earlier between the evolution of global Internet traffic and the inter-domain link capacities, a CAGR of approximately 34% and assuming that inter-domain links are around 75% utilized (avoiding congestion conditions), the following table presents the evolution of the ISP’s network throughout the 4 years under study.

Table 32. Network and inter-domain link dimensioning

Year	Link Capacity (Mbps)	Utilized Capacity (Mbps)	Hardware Multiplication Factor (HMF)
2011	4000	3000	2
2012	5000	3750	2
2013	7000	5250	3
2014	10000	7500	3

8.1.4.2 Per Month Analysis

Let’s now present how the monthly costs and savings are estimated. To begin with, year 2011 is considered. The same analysis holds for the next years, by adjusting the prices and volumes according to the respective estimates presented in the previous subsections.

Regarding the monthly costs, if SmoothIT amortizes the capital expenses over the 4 year period then the monthly costs are calculated to be 52,08€ for the BGP-Loc/HAP cases while for the IoP case the costs are 105,08€ if one machine is used. But due to the capacity of the inter-domain link (4 Gbps), SmoothIT assumes that 2 machines will be required to serve the number of peers expected. Hence, the above cost figures must be doubled for 2011.

SmoothIT needs to estimate the savings realized by the respective average reduction introduced by an ETM mechanism, considering the transit prices of Table 1 and the tariff structure of Table 2. It turns out that for 2011 the price unit is 1.86€, if SmoothIT considers that the base price for 100Mbps is \$3.25 and that the difference at the prices of 4Gbps from 100Mbps is approximately 22%. The following table summarizes the savings realized in one month of 2011.

Table 33. Estimated monthly savings in inter-domain traffic (volume and euros) for 2011 by the introduction of an ETM mechanism

Reduction on P2P part of inter-domain traffic	Savings (Mbps)	Savings (€)
2%	13,71	25,49 €
4%	27,42	50,97 €
6%	41,13	76,46 €
8%	54,84	101,94 €
10%	68,55	127,43 €
12%	82,26	152,92 €
14%	95,97	178,40 €
16%	109,68	203,89 €
18%	123,39	229,37 €
20%	137,10	254,86 €
22%	150,81	280,34 €
24%	164,52	305,83 €
26%	178,23	331,32 €
28%	191,94	356,80 €
30%	205,65	382,29 €

To provide a better understanding on how these values were calculated, let's consider the case of a 2% reduction on the P2P share of the traffic. The link's capacity is 4000 Mbps, but only 3000Mbps are used. The P2P share of traffic for 2011 is 22.85%. Hence, the savings are $3000 * 22.85\% * 2\% = 13.71$ Mbps, and if multiplied by 1.86€ SmoothIT gets 25.49€.




Figure 22. Monthly costs and savings for the three ETM mechanisms for year 2011

Having calculated the costs and savings, the break-even point is easy to be estimated. The following diagram depicts this point for the BGP-Loc/HAP case and for the IoP case. One can observe that the deployment of the BGP-Loc/HAP mechanism is beneficial even for an average reduction of 8%, while for the IoP deployment, the break-even point is achieved with a 16% average reduction.

The same analysis is repeated for the next three years (2012, 2013 and 2014). The following table summarizes the assumptions and findings for every year.

Table 34. Assumptions and break-even points per month per year of the amortization period

Year	Link capacity (Mbps)	Link utilization (Mbps)	Unit price	HMF	Break-even point	
					BGP-Loc/HAP	IoP
2011	4000	3000	1.86€	2	~ 8%	~ 16%
2012	5000	3750	1.34€	2	~ 10%	~ 21%
2013	7000	5250	0.77€	3	~ 22%	~ 46%
2014	10000	7500	0.46€	3	~ 29%	~57%

8.1.4.3 Overall results (4-year aggregated analysis)

The analysis is concluded by providing the total costs and savings during the 4 year amortization period. The total costs and savings and compared and the break-even point is calculated on the total numbers. The following table summarizes this information.

Table 35. Total costs and savings for the 4 year amortization period (2011-2014)

Average Reduction	Savings (Mbps)	Savings (€)	Total Costs (€)	
			BGP-Loc / HAP	IoP
2%	835,94	836,84 €	6.250,00 €	12.610,00 €
4%	1671,88	1.673,68 €	6.250,00 €	12.610,00 €
6%	2507,81	2.510,52 €	6.250,00 €	12.610,00 €
8%	3343,75	3.347,36 €	6.250,00 €	12.610,00 €
10%	4179,69	4.184,20 €	6.250,00 €	12.610,00 €
12%	5015,63	5.021,04 €	6.250,00 €	12.610,00 €
14%	5851,57	5.857,88 €	6.250,00 €	12.610,00 €
16%	6687,50	6.694,72 €	6.250,00 €	12.610,00 €
18%	7523,44	7.531,56 €	6.250,00 €	12.610,00 €
20%	8359,38	8.368,40 €	6.250,00 €	12.610,00 €
22%	9195,32	9.205,24 €	6.250,00 €	12.610,00 €
24%	10031,26	10.042,08 €	6.250,00 €	12.610,00 €
26%	10867,19	10.878,92 €	6.250,00 €	12.610,00 €
28%	11703,13	11.715,76 €	6.250,00 €	12.610,00 €
30%	12539,07	12.552,60 €	6.250,00 €	12.610,00 €

Observe that for the BGP-Loc/HAP case, the break-even point is achieved with a traffic reduction of approximately 15%. In the case of the IoP, the break-even point is achieved at an average reduction of around 30%.

An important remark for the analysis so far, is that the resulting numbers for costs, savings and break-even points heavily depend on the future development of absolute P2P traffic and transit prices. Here, SmoothIT has shown the case where transit prices fall much faster than P2P traffic grows, as revealed by the studies from Cisco and DrPeering summarized in Sections 0 and 8.1.3.




Figure 23. Total costs and savings for the three ETM mechanisms for years 2011-2014

8.1.4.4 Impact of server capacity and ISP size

The previous cost-benefit analysis was done for a scenario similar to the external trial, an ISP the size of PrimeTel and a server with a capacity according to the SmoothIT prototype. In this section, the cost-benefit analysis is extended to study the impact of server capacity and ISP size.

The server capacity is chosen relative to the capacity of the prototype and specified as the amount of inter-domain P2P traffic the server is able to handle. One prototype server is able to handle 400 Mbps of inter-domain P2P traffic. In 2011 with 660 Mbps of inter-domain P2P traffic 2 prototype servers are required. In 2014 with 1125 Mbps of inter-domain P2P traffic 3 prototype servers are required.

The size of an ISP is specified by the link capacity of the ISP in 2011 according to the previous cost-benefit analysis. ISPs with 2011 Link capacities of 4Gbps, 8 Gbps, and 16 Gbps are considered. The size of the ISPs is expressed by their link capacities, while total traffic is taken to grow by 34% per year. In contrast to the previous analysis, the link capacities are not rounded but taken as continuous values. As an example, the 2014-link-capacity of a 4Gbps-link-capacity-in-2011 ISP is 9624 Mbps and not rounded up to 10 Gbps as previously done. The inter-domain traffic costs per Mbps and month are determined as describe above, the price reduction due to larger link capacities is truncated at 10 Gbps according to Table 2.

Since “larger” servers are more cost-efficient, let us consider a server with 500% capacity of the prototype that is able to handle 2 Gbps of inter-domain P2P traffic. In 2014 for a 16Gbps-in-2011 ISP, the total amount of inter-domain traffic is 38.5 Gbps and a 15.5% P2P traffic share in 2014 leads to a little less than 6 Gbps inter-domain P2P traffic. Consequently, three 500%-prototype-capacity-servers need to be installed for this amount of traffic. Hence, the monthly costs in 2014 are 156€ for the 3 servers when considering BGP-Loc. Note that since this analysis aims at revealing the impact of server capacity and ISP size in a representative case of an ETM mechanism, BGP-Loc is considered only. IoP and HAP are not considered in this analysis since they scale in another way with the amount of P2P traffic. Indeed, while BGP-Loc is applied to the large swarms and thus needs to handle requests by most peers, IoP and HAP profit from more P2P traffic through a larger diversity in selecting the most efficient swarms for caching. Supporting the total amount of P2P traffic becomes the more ineffective the larger the ISP. Indeed, HAP and in particular IoP may focus on the most efficient swarms and handle only a part of the P2P traffic. BGP-Loc, however, is not swarm aware and needs to respond to all requests it receives. Going back to the example, the monthly costs per Mbps of inter-domain traffic is 0.46€ as in Table . Consequently, the required ETM efficiency to achieve the break-even point is 5.6% since 340 Mbps of inter-domain traffic have to be saved which is 5.6% of the total 6 Gbps of P2P traffic.




Figure 24. Impact of server capacity

Figure 24 shows this required ETM efficiency for different server capacities and ISP sizes. Years 2011 to 2014 are not considered as a whole but the development from 2011 to 2014 is shown. The circular markers show the relation of ETM efficiency and server capacity. The color of the markers indicates the year, the size of the marker the ISP size, larger

markers for larger ISPs. As mentioned earlier, SmoothIT relates the size of the ISP to the size of the inter-domain link, thus SmoothIT considers the cases of 4Gbps, 8Gbps and 16Gbps links related to the respective marker sizes. The curves show that assuming a more realistic server capacity the required ETM efficiency decreases. Even with a server capacity three times the capacity of the prototype, an ETM efficiency of 10% is sufficient to achieve the break-even point, even in the last year (2014), which is the most unfavorable one. This means that BGP-Loc becomes interesting also for medium-sized ISPs that do not host larger local swarm partitions and thus do not achieve an inter-domain traffic reduction of 50% to 90% but rather 10% to 20%. Further, the required ETM efficiency also decreases with the size of the ISPs, however only to a certain degree. The gain only results from the fact that small ISPs are not able to fully utilize the capacity of the server(s) while for larger ISPs the percentage of the unutilized server capacity diminishes.

8.1.5 Relation with existing studies

What is provided up to now is a calculation of the break-even points that render the adoption of ETM mechanisms beneficial for the ISP, in terms of monetary gains. What remains to be checked is whether such decrease in the inter-domain traffic is feasible to happen by any of the studied ETM mechanisms.

To do so, first SmoothIT briefly refers to some of the results obtained by simulative studies and are included in previous deliverable [2]. Indeed, in Section 4.1 of [2], it is stated that the traffic savings in the case of the BGP-Loc ETM mechanism can reach a maximum of 60% of average reduction, depending on the swarm sizes, peer distribution and heterogeneity (or not) of domains, with more realistic being the figures of 30%. In the case of the loP evaluation (cf. Section 4.3 of [2]), depending on the case under study, the decrease in the inter-domain traffic has been observed to vary from 10% to 30%. Finally, for the HAP case (cf. Section 4.5 of [2]) savings are observed to reach 30-50% of the original traffic.

To complement the above findings, SmoothIT refers to the findings of the internal trial, where only the BGP-Loc and the loP ETM mechanisms were evaluated [20]. In this deliverable, it is mentioned that the BGP-Loc mechanism achieved an average reduction of 4-5% while the loP mechanism achieved an average reduction of 20% approximately. Due to the fact that in the internal trial, a small-sized swarm was emulated, the reduction on inter-domain traffic is not expected to be large enough, especially in the BGP-Loc case.

Hence, it is obvious, both by simulative studies and by internal trial that in most cases the ETM mechanisms perform equal or even better than what is required for the achievement of the break-even point. It is safe to assume that there is no major monetary concern from the ISP's point of view in deploying one of the ETM mechanisms, or even combination of mechanisms. For the latter case, the cost-benefit effects have not been assessed in detailed, but it is expected that a combination of ETM mechanisms performs better than a single ETM mechanism, under certain circumstances while the associated cost does not exceed by far that of the most expensive mechanism. Thus, their deployment still remains beneficial for the ISP.

8.2 Global View on Potential Inter-Domain Traffic Savings

The previous section presented a detailed cost-benefit analysis from the perspective of a small to medium sized ISP and revealed the required efficiency of the ETM in terms of inter-domain traffic reduction in order to achieve the break-even-point. In this section the potential inter-domain traffic reduction for P2P traffic is estimated on a global perspective.

This estimation is based on swarm measurements performed in WP1 and published in [26]. More details of the analysis described in this section are also found in that paper. Please note that the approach taken in this section and the previous cost-benefit analysis are complementary. This section focuses on the potential global impact of a locality promoting ETM, which as the analysis will show is mainly achieved by major ISPs applying the ETM to the largest swarms. Small ISPs will not significantly contribute to the global inter-domain traffic reduction but as shown in the previous section they may nevertheless achieve a monetary gain by deploying the ETM.

The idea of the analysis is to derive an upper bound for the inter-domain traffic savings based on music and movie swarm measurements that reveal the swarm size and the partitioning of the swarms into ASes. In order to quantify the potential inter-domain traffic savings, the number of peers in a swarm is used as an indicator of how much traffic a swarm produces in relation to the other swarms. Further, the assumption is made that peer access capacities are not correlated with swarm sizes such that these are neglected in the analysis.

Figure 25 presents the cumulative estimates ('music traffic T_0 ', 'movie traffic T_0 ') for the top $x\%$ of the largest swarms normalized by the total amount of traffic. The figure reveals that 10% of the swarms of the movies data set contain 80% of the peers and are consequently responsible for the same fraction of the total traffic according to the aforementioned assumptions. Weighting the number of peers in a swarm with the size of the exchanged file ('traffic T_f ', legend: 'with file sizes') to estimate the amount of traffic, almost the same results are obtained as when using just the number of peers ('w/o file sizes'). This is in particular true for the movie traffic. For the music files the difference is small.




Figure 25. Total and potentially local traffic of the top $x\%$ of the swarms with and w/o considering the size of the exchanged files

Next, a very simple and optimistic approximation is developed for estimating the potential of locality awareness. For each swarm, δ_2 is the fraction of peers in the swarm which are not the sole peer in their AS. Assuming an ideal locality algorithm those peers produce no

inter-domain traffic. In doing so it is neglected which peers are seeders and leechers and possible performance degradations for simplicity reasons. Then, δ_2 is the fraction of 'potentially local traffic' of that swarm. Figure 25 shows this value weighted by the total traffic of the swarm for the music and movie files. The figure confirms that, on global scale, locality awareness can produce a recognizable effect in a small subset of all swarms. However, it also shows that the potential savings of inter-domain traffic are quite larger in the big swarms which are responsible for the vast majority of BitTorrent traffic. Therefore, the overall optimization potential of locality awareness is about 65% for the movie files and about 40% for the music files. In other words, around 35% (60%) of the overall movie (music) traffic is produced by peers which are alone in their AS. Therefore, no locality awareness mechanism can avoid this part of the inter-domain traffic though caching-type ETM mechanisms may still work on this part of the traffic though the cost-benefit-ratio will probably be low. A larger additional benefit can be expected from BGP-Loc and the Inter-SIS extensions since they do not only focus on minimizing inter-domain traffic but also distinguish inter-domain and peering traffic which is not covered in this analysis. Emphasizing the positive outcome of this analysis, locality awareness has to be deployed in only 10% of the swarms in order to achieve 90% of the theoretically possible reduction of inter-domain traffic.




Figure 26. Calculation of the optimization potential of ETM mechanisms

Having dealt with a theoretical upper bound on the savings so far, an estimation of actual inter-domain traffic savings is derived. This is illustrated in Figure 26 for the movie swarms. From the measurements SmoothIT has seen that when considering only the top 10% swarms 60% of the total traffic is potentially local traffic while 40% of the traffic is inter-domain traffic in any case. Assuming that 30% of the potentially local traffic is already kept local by current BitTorrent, the optimization potential of ETM mechanisms applies to 70% of the potentially local traffic. Actually, 30% is already a high estimate of the local traffic. For a local (e.g., German or French) movie the largest AS has a share of around 35% and

in an international (English) movie the largest AS has a share of about 20% of the traffic. Assuming all peers to be equal in regular BitTorrent, in the largest ASes of the swarms we have between 20% and 35% of local traffic, and in the other ASes the share of local traffic is even lower. Thus, 30% is actually a high number and the actual local share is expected to be between 10% and 30%.




Figure 27. Example values for inter-domain traffic reductions achieved by ETM mechanisms




Figure 28. Dependency of ETM efficiency and inter-domain traffic savings

Considering further the efficiency of a given ETM approach, which varies between 0 and 1, leads to the percentage of inter-domain traffic which was actually saved. The efficiency of the ETM algorithms corresponds to the fraction of the optimization potential achieved by

the ETM in the end. This is shown in the last row of Figure 26. This gain of a given ETM mechanism is further depicted in Figure 27 which focuses on the third and fourth column of Figure 26. In particular, SmoothIT sees that an ETM mechanism with a reasonable efficiency of 70% can influence 30% of the total BitTorrent traffic. Finally, this leads to a reduction of the inter-domain traffic to 35%.

This dependency of inter-domain traffic reduction and efficiency of an ETM mechanism is further studied. Figure 28 shows the inter-domain traffic savings depending on the ETM efficiency. The different curves indicate an absolute amount of already local BitTorrent traffic of 6%, 12% and 18% corresponding to 10%, 20%, and 30% of the potentially local traffic as introduced in Figure 26. The figure shows that up to 50% of the inter-domain traffic can be influenced by ETM mechanisms having a high efficiency (70% to 90%) like BGP-Loc in large swarms, for example. The current share of local traffic has only a minor impact.

9 Summary

The project SmoothIT has defined, developed, and evaluated several mechanisms to perform economic management of the traffic generated by overlay applications, with the objective to accomplish TripleWin; i.e. the situations where network operators, overlay providers, and users' benefit from the employed approaches. This deliverable presents the final evaluation assessment of those ETM mechanisms that were fully specified, simulated, implemented and tested (Sections 4 and 7), including an account of work undertaken as part of external trial operation together with detailed description of data collected, methods in which these were aggregated, analyzed and interpreted (Section 5). In addition project results are analyzed from legal, security and cost benefit perspectives (sections 6 and 8 respectively).

Based on the studies made as part of this deliverable it can be concluded that the ETM mechanisms designed are in general a technically feasible, performance-wise promising solution to the problem of reducing inter-domain traffic while at the same time improving users' QoE, or at least maintaining it. The thorough investigations undertaken have also revealed the circumstances under which each ETM mechanism (or combination thereof) does attain the goal of a balanced Triple-Win, and thus can be considered as a viable solution to be imposed by ISPs and to be embraced by users and/or overlay application providers. Therefore, while there is no generally applicable ETM mechanism, there are available appropriate mechanisms for most of the realistic cases analyzed, leading to stable situations that are beneficial for all players even if it is necessary for some of them (e.g. small ISPs) to react to the initially harmful actions of the others. The analysis of economic costs and benefits resulting from application of these mechanisms was performed. This analysis revealed that there do exist break-even points for ETM adoption by ISPs of various sizes both small-to-medium and large ones, with the later being in more beneficial circumstances. These break-even points correspond to performance benefits that were already proved by the evaluation to be indeed achievable by the mechanisms.

Moreover, the experience from external trial revealed that the developed prototype ETM implementation can be integrated and operated in real life ISP network, performing as defined by requirements with performance within expected ranges. Finally, the study performed regarding security and legal/regulatory risks and related issues that should be dealt with in such a case.

Overall, ETM is technically feasible and promising both performance-wise and economically, and can serve as a practically applicable solution to a problem faced in the Internet nowadays as well as the basis for developing new mechanisms for the economic and incentives-based management of other emerging applications.

References

- [1] The SmoothIT project: Deliverable 2.3 — ETM Models and Components and Theoretical Foundations (Final), October 2009.
- [2] The SmoothIT project: Deliverable 2.4 — Performance, Reliability, and Scalability Investigations of ETM Mechanisms, September 2010.
- [3] The SmoothIT project: Deliverable 2.5 — Comprehensive Test-bed and Trial Parameter Set Definition (Part I), April 2009.
- [4] The SmoothIT project: Deliverable 2.6 — Comprehensive Test-bed and Trial Parameter Set Definition (Part II), April 2010.
- [5] W. Krzysztofek, P. Wydrych, R. Stankiewicz, Z. Duliński, M. Kantor, P. Chołda, B. Polaczyk, M. Kandyba, "Eruption.S," <http://www.kt.agh.edu.pl/~wydrych/Research/Eruption.S>
- [6] B. Falchuk, D. Gordon, D. Marples, Content Distribution for Telecom Carriers, *Journal of Multimedia*, vol. 1, no. 5, August 2006
- [7] ITU-T Recommendation X.805: *Security architecture for systems providing end-to-end communications*, 2003
- [8] R. Rivest: *The MD5 Message-Digest Algorithm*, IETF RFC 1321, 1992
- [9] P. Rodriguez, S. M. Tan, C. Gkantzidis, On the feasibility of Commercial, Legal P2P Content Distribution, 2005
- [10] Council of the European Union: *Council Regulation (EC) No 44/2001 of 22 December 2000 on Jurisdiction and the Recognition and Enforcement of Judgments in Civil and Commercial Matters*. Official Journal of the European Communities, L12:1–23, January 2001.
- [11] European Commission, Directorate-General for Health and Consumer Protection: *Consumer Protection in the European Union: Ten Basic Principles*. Brussels, Belgium, 2005. http://ec.europa.eu/consumers/cons_info/10principles/en.pdf
- [12] European Parliament and the Council of the European Union: *Directive 95/46/EC of the European Parliament and of the Council of 24 October 1995 on the Protection of Individuals with Regard to the Processing of Personal Data and on the Free Movement of such Data*. Official Journal of the European Communities, L281:31–56, November 1995.
- [13] European Parliament and the Council of the European Union: *Directive 2001/29/EC of the European Parliament and of the Council of 22 May 2001 on the Harmonization of Certain Aspects of Copyright and Related Rights in the Information Society*. Official Journal of the European Communities, L167:10–19, June 2001.
- [14] European Parliament and the Council of the European Union: *Directive 2002/19/EC of the European Parliament and of the Council of 7 March 2002 on Access to, and Interconnection of, Electronic Communications Networks and Associated Facilities (Access Directive)*. Official Journal of the European Communities, L108:7–20, April 2002.
- [15] European Parliament and the Council of the European Union: *Directive 2002/58/EC of the European Parliament and of the Council of July 2002 Concerning the Processing*

- of Personal Data and the Protection of Privacy in the Electronic Communications Sector (Directive on Privacy and Electronic Communications)*. Official Journal of the European Communities, L201:37–47, July 2002.
- [16] European Parliament and the Council of the European Union: *Regulation (EC) No 593/2008 of the European Parliament and of the Council of 17 June 2008 on the Law Applicable to Contractual Obligations (Rome I)*. Official Journal of the European Union, L177:6–16, July 2008.
- [17] European Parliament, the Council of the European Union, and the European Commission: *Charter of Fundamental Rights of the European Union (2000/C 364/01)*. Official Journal of the European Communities, C364:1–22, December 2000.
- [18] F. Hecht, T. Bocek, R. G. Clegg, R. Landa, D. Hausheer, and B. Stiller: *LiveShift: Mesh-Pull P2P Live And Time-Shifted Video Streaming*. Technical Report IFI-2010-0009, University of Zurich, Department of Informatics, September 2010.
- [19] The SmoothIT project: Deliverable 2.2 — ETM Model and Components (Initial Version), December 2008.
- [20] The SmoothIT project: Deliverable 4.1 — Internal Test-bed Trial Set-up Specification and Results, November 2009.
- [21] The SmoothIT project: Deliverable 5.4 — Exploitation Plan (Final), February 2011.
- [22] K. Pussep, S. Kuleshov, C. Gross, S. Soursos, An Incentive-Based Approach to Traffic Management for Peer-to-Peer Overlays, Incentives, Overlays, and Economic Traffic Control, Springer, 2010.
- [23] Internet Transit Prices - Historical and Projected, DrPeering.net
- [24] W. Willinger, V. Paxson, and M. Taqqu, “*Self-similarity and Heavy Tails: Structural Modeling of Network Traffic*”, In *A Practical Guide to Heavy Tails: Statistical Techniques and Applications*, Adler, R., Feldman, R., and Taqqu, M.S. (Edts), Birkhauser, 1998.
- [25] Cisco Visual Networking Index: Forecast and Methodology, 2009–2014, June 2010
- [26] T. Hoßfeld, F. Lehrieder, D. Hock, S. Oechsner, Z. Despotovic, W. Kellerer, M. Michel. *Characterization of BitTorrent Swarms and their Distribution in the Internet*. Computer Networks, <http://dx.doi.org/10.1016/j.comnet.2010.11.011>, 2010.
- [27] P2P Next EU/ICT project, <http://www.p2p-next.org/>
- [28] The SmoothIT project: Deliverable 3.3 — Documentation of Engineering and Implementation (Final), July 2010
- [29] Quagga Routing Suite, <http://www.quagga.net/>
- [30] BitTorrent software source code, <http://download.bittorrent.com/dl/archive/>
- [31] J.D.D. Mol, J.A. Pouwelse, M. Meulpolder, D.H.J. Epema, and H.J. Sips, Give-to-get: Free-riding-resilient video-on-demand in p2p systems, Proceeding of the 15th SPIE/ACM Multimedia Computing and Networking (MMCN’08), 2008
- [32] Z. Duliński, M. Kantor, W. Krzysztofek, R. Stankiewicz, and P. Chołda, “Optimal Choice of Peers based on BGP Information,” in Proc. 2010 IEEE International Conference on Communications ICC 2010, Cape Town, South Africa, May, 2010.

- [33] Z. Dulinski, R. Stankiewicz, P. Cholda, P. Wydrych, B. Stiller, Inter-ALTO communication protocol, IETF Internet draft, draft-dulinski-alto-inter-alto-protocol-00, Jun. 2010
- [34] The SmoothIT project: Deliverable 3.2 — Initial Documentation of Engineering and Implementation (Initial Version)
- [35] The SmoothIT project: Deliverable 3.1 — Economic Traffic Management Systems Architecture Design (Initial Version)
- [36] The SmoothIT project: Deliverable 3.4 — Economic Traffic Management Systems Architecture Design (Final)

Acknowledgements

This deliverable was made possible due to the very large, constructive, and open help of the WP4 team, the Assessment Task Force of the SmoothIT team within this STREP, and all individuals, which includes besides the deliverable authors the internal reviewers. Very many thanks to all of them!

Appendix 1. Sample Record for Simulation Experiment

This appendix presents a sample record for simulation.

For consistency, each objective, assumption and expected result received a unique number. The following notation was used:

X-SSn.m

where

- for X we use: O for objective, A for assumption, or E for expected result.
- SS stands for simulation scenario
- n is a number of the scenario
- m is a serial number of an objective, assumption or expected result within a scenario.

For example, O-SS2.1 denotes objective number 1 for simulation scenario number 2.

Records for simulation experiments' were collected in a separate internal document not published entirely in deliverables. Below we present a sample record.

Simulation Scenario 2:

Symmetric BitTorrent Simulation Scenario for BGP-based ETM without SIS cooperation

This scenario was used to evaluate the effectiveness of the BGP-Loc ETM for BitTorrent under different conditions as well as to compare client-side mechanisms to use the underlay information provided by the SIS.

Objectives

- O-SS2.1 To study the impact of BGP-Loc-based SIS mechanism on one BitTorrent swarms when the total swarm is subdivided into K local swarms of equal size
- O-SS2.2 To study the impact of the Biased Neighbor Selection and Biased Unchoking mechanisms as described in D2.3, in comparison to the regular BT implementation
- O-SS2.3 To study the impact of ETM mechanism w.r.t. peering, inter-AS and transit links

Basic Assumptions

- A-SS2.1 Homogeneous peer bandwidths, parameters and capabilities
- A-SS2.2 Exponential peer arrival process, equally distributed on the available ASes
- A-SS2.3 Bandwidth bottlenecks of connections are either in the access network or in the inter-AS links (experiments were done for both)
- A-SS2.4 Bike topology, as described in D2.3. As a result, inter-domain traffic is symmetric.

A-SS2.5 One global SIS server with global information; acts as SIS for all peers promoting locality.

Expected results

E-SS2.1 The BGP-Loc ETM leads to inter-domain traffic reduction and prefers intra-AS, peering, inter-AS and transit links, in that order.

E-SS2.2 Recommendation of one client-side algorithm as the best alternative for implementation

E-SS2.3 Download time reduction only possible in scenarios with inter-AS bottlenecks (also see related work)

Were the objectives achieved?

O-SS2.1 Yes. A large saving potential regarding inter-domain and transit traffic and download times was identified for the considered topology, depending on the experiment.

O-SS2.2 Yes. A combination of the client-side mechanisms BNS and BU achieved the highest efficiency in the experiments. The inner workings of the mechanisms were understood.

O-SS2.3 Yes. A clear preference for intra-AS peers before peers in peering ASes before remote peers was observed.

Discussion on assumptions correctness, adequacy, limitations

A-SS2.1 Simplification of real peer characteristics. Will be relaxed in later scenarios with heterogeneous bandwidths.

A-SS2.2 Adequate assumption for steady-state swarms.

A-SS2.3 Adequate assumption.

A-SS2.4 Simplification of real-world topologies. May be diversified (larger number of possible hops, asymmetry...) in future scenarios.

A-SS2.5 Adequate assumption for scenarios with no inter-SIS communication. Information availability at peers identical to case with one SIS per AS.

Discussion on expected vs. obtained results

E-SS2.1 Large traffic savings observed with BGP-Loc, obtained result fits expected result.

E-SS2.2 Combination of BNS and BU identified as best solution, obtained result fits expected result

E-SS2.3 Obtained result fits expected result. In experiments with the bottlenecks just in the access network, download times were not influenced by locality promotion.

Discussion on statistical credibility of the results

To evaluate the swarm behavior in its steady state, we simulated 6.5 hours and discarded the initial 1.5 hours in which the swarm grows to its steady-state size.

As the main performance indicators in the experiments, the used inter-AS bandwidth and

the download times of the file were reported. The bandwidth was measured by logging the amount of data flowing over each link of the topology every 60 seconds. Then for each run, the mean value of these one minute intervals was computed.

The download times experienced by all peers in one run are averaged. The steady state of the simulation is 5 hours long, so that we gather data from at least 1500 peers per run. If not stated otherwise, we conducted at least 10 simulation runs per parameter setting in order to generate statistically reliable results. Confidence intervals were smaller than 10% of the mean value for all data points (to be updated when more results/runs are available).

Summary

From the observed results, we draw the conclusion that the usage of the SIS-generated locality information in the client is of high importance. Even if the SIS provides good data, it is of no worth to the provider or the user if it is not used effectively. In this regard, the combination of Biased Neighbor Selection and Biased Unchoking works better than BNS or BU alone.

The reason for this is that BNS ensures that there are more local and close neighbors in the neighbor set of a peer than in regular BitTorrent. It achieves inter-AS traffic reduction simply by trusting that a higher fraction of neighbors in the same or peering ASes indirectly leads to a higher fraction of locally forwarded traffic.

In contrast, BU prefers local and close neighbors in the most crucial mechanism for deciding the direction of traffic flows, the unchoking process. It can directly affect the fraction of local and remote traffic. However, this does only work if a) there are actually local interested neighbors and b) there are enough interested neighbors that a preference of these has an effect. In low load scenarios and in swarms with a low number of peers in the same AS, these conditions are not always fulfilled.

In combination, these two mechanisms combine their advantages, leading to a synergy. BNS provides the local neighbors necessary for BU to function effectively, while BU ensures that the change in the neighbor set composition created by BNS is used to maximum effect.

With regard to scenarios, the locality promoting strategies work best in swarms with high load, since there the number of interested peers in the neighbor set is generally larger. Therefore, a preference of local or close peers in the unchoking process has a larger effect here. Also, larger fractions of the swarm per AS increase the efficiency of locality promotion, since there are more local and close neighbors that can be utilized.

Not surprisingly, if only a fraction of peers in the swarm support locality, the performance improvement increases with the fraction of locality-promoting peers. In general, locality-promoting peers perform better than their counterparts that use the regular implementation if the user performance is limited by inter-AS bottlenecks.

The quality of the improvement achieved by locality promotion depends on the topology evaluated. We could corroborate known results from related studies that in scenarios with connection bottlenecks only in the access networks, download times remain largely unaffected by locality promotion, while savings in utilized inter-AS bandwidth can be achieved. In contrast, if tight bottlenecks between ASes exist, still a little inter-AS bandwidth can be saved, but the largest effect here is on the download times, since these are affected strongly by the bandwidth limitation. Download times can be improved significantly by locality promotion in these scenarios.