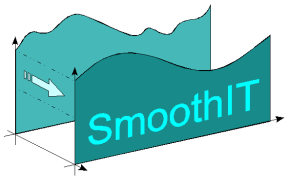

Simple Economic Management Approaches of Overlay Traffic in Heterogeneous Internet Topologies

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

Deliverable D5.3
Standardization Survey

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. Stanislawa Staszica W Krakowie, AGH, Poland
Intracom S.A. Telecom Solutions, ICOM, Greece
Julius-Maximilians Universität Würzburg, UniWue, Germany
Telefónica Investigación y Desarollo, TID, Spain

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: Standardization Survey

Type: Public

Editor: Krzysztof Wajda

E-mail: wajda@kt.agh.edu.pl

Authors: Maria Angeles Callejo, Zoran Despotovic, Juan Fernández-Palacios, Fabio Hecht, Michael Makidis, Konstantin Pussep, Sergios Souros, Dirk Staehle, Krzysztof Wajda, Haris Neophytou

Reviewers: Maria Angeles Callejo, Zoran Despotovic, Dirk Staehle

Doc ID: D5.3-v1.1.doc

AMENDMENT HISTORY

Version	Date	Author	Description/Comments
V0.1	July 15, 2008	Krzysztof Wajda	Initial version, proposing table of content
V0.2	October 1, 2008	Krzysztof Wajda	ALTO standardization effort added
V0.3	November 30	Krzysztof Wajda	OGF and ITU standardization efforts added, IETF ALTO group updated
V0.3a	December 2	Michael Makidis	Additional contribution to the OpenGrid forum section
V0.3b	December 8,2008	Fabio Hecht	Added ALTO chapter
V0,3c	December 9,	Konstantin Pussep	Added description of SIP, NSIS, P2PSIP
V0.3d	December 10	Sergios Souros	Added contribution on IPsphere.
V0.33	December 12	Dirk Staehle	Added description of IRTF P2P and LEDBAT
V0.3f	December 12	María Ángeles Callejo Rodríguez, Juan Fernandez Palacios	
V0.3g	December 12	Zoran Despotovic	Added additional description of ALTO and LEDBAT
V0.4	December 18	Michael Makidis, Fabio Hecht, Konstantin Pussep, Sergios Souros, Juan Fernández-Palacios, Maria Angeles Callejo, Dirk Staehle, Zoran Despotovic,	Contributions from partners added (AUEB, UZH, TUD, TID, ICOM, UniWue),
V0.5	December 30	Fabio Hecht, Zoran Despotovic, Krzysztof Wajda	Incorporating IETF 72 and 72 meetings' reports, merging DoCoMo contribution, IRTF NMRG subsection added, checking cross-references,
V0.5a	January 11, 2009	Maria Angeles Callejo	Review of whole document, additions to Summary section, included TID involvement in standardization.
V0.5b	January 13	Fabio Hecht	Provided.an update on UZH involvement in standardization
V0.6	January 13	Dirk Staehle	Review of whole document, provided an update on UniWue involvement in standardization.
V1.0	January 14	Krzysztof Wajda	Addressing comments of reviewers, recent contributions, sending comments to the authors, merging feedback from the authors.
V1.1	April 7, 2009	Sergios Souros	Update with reviewer comments

indirect, special, incidental or consequential damages in connection with the furnishing, performance, or use of this material.

Table of Contents

(This page is left blank intentionally.)

1 Executive Summary

This deliverable is aiming at reflecting the internal discussion carried out within SmoothIT consortium related to the relevant standardization bodies and forums relevant for our project goals. The purpose of this document is twofold:

- A brief overview of ongoing efforts towards architectures, control and management plane issues and definition of services being in the core of interest for SmoothIT consortium;
- To lay a foundation for original proposal coming from our project (rather natural target organization is IETF ALTO WG (Internet Research Task Force Application-Layer Traffic Optimization Working Group), less relevant seems to be the recently established IETF LEDBAT WG (Internet Research Task Force Low Extra Delay Background Transport Working Group)).

Extensive standardization efforts relevant to overlay networks are currently being carried out. They reflect the complexity of telecommunications concepts, architectures, protocols, algorithms, etc. At least the following questions arise: What can or should be standardized? Whether it is possible or relevant to control P2P systems and influence users' behaviour?

According to the first discussions, the primary sources of standards and incentives relevant for SmoothIT are IETF and IRTF (Internet Research Task Force), however other standardization bodies like ITU (International Telecommunications Union), IPSphere Forum or OGF (Open Grid Forum) have been also recognized as important. The goals of these bodies, how SmoothIT can benefit from them and how the project results can contribute to them are analysed along the document.

The discussion about a choice of relevant standardization institutions and documents carried out since the beginning of SmoothIT project is now completed by new IETF approach to start ALTO working group. Within the timeframe of standardization discussion, starting from preliminary document [13], dated July 2008, we are now at the position to witness running of this new IETF activity.

2 Introduction

The Internet is today's ubiquitous network and platform to offer flexible services to individual as well as corporate users. Its autonomic and decentralized behaviour is not only based on IP as routing mechanism but is supported by many justified and agreed concepts; many of them can be called standards.

The term "*standard*" has rather broad meaning and concerning its technical aspects, it denotes a formal document that establishes uniform engineering criteria, methods, processes and (not so strictly) also practices. If defining a specification of standard it means giving an explicit set of requirements and features of component, system, service (these three concepts being of interest in telecommunications), and additionally an item or material (used usually in engineering).

2.1 A Position of Contemporary Operator, ISP, and Network User

In the 70s and 80s, when voice service was the dominating offer, traditional, "incumbent" or "national" operators have had monopolistic positions. Telecommunications market was static and shared by two types of players: operators and customers (which, moreover, did not have the chance to select between different operators). In the 90s, along with development of the Internet, technological progress in computer science, market liberalization and globalization, one could observe an evolution of entities sharing the network resources and now we can consider four types of players: operators (being owners of infrastructure), ISPs (Internet Service Provider), content and service providers and the end users with the capability to choose service and network providers. Obviously such simple classification can be further complicated since single entity can eg. play two roles.

The SmoothIT framework is focused towards ISP – service providers and end user relationships. In particular, due its important contribution to the traffic amount with overlay applications as dominant part of generated and sent traffic.

2.2 Purpose of the Document D5.3

One of the SmoothIT WP5 objectives is to outline and potentially support relevant standardizations in field of relevance to SmoothIT. This implies two main activities: the survey of standardization bodies, organizations, and activities and the preparation of input(s) to them.

In this deliverable, a brief overview of standardization activities mostly relevant to the SmoothIT framework is given and first discussions about what can or should be standardized are also presented. This is the general idea also presented in the SmoothIT Description of Work: "to carefully follow standardization (in particular coming from IRFT, IETF and IPsphere Forum), which is relevant for network management and economic support mechanisms of SmoothIT."

Further work within this task will lead towards a final document: "D5.5: Final Report on Liaisons, Standardizations, and Dissemination", where basic achievements of SmoothIT consortium will be presented, with the the aim to contribute to the standards, in one way or another.

2.3 Document Outline

This document serves either as a repository of standards and concepts, which are of agreed value for Partners of SmoothIT consortium, and as a foundation for future efforts to make the achievements of this project visible to the external audience. The scope includes related materials, methods, targets, partner involvement, and schedule.

The reminder of the document is structured as follows:

- Chapter 3, presents the actual and planned involvement of the SmoothIT partners into standardization efforts (e.g., research projects and consortia).
- Chapter 4, which is the most essential for this deliverable, describes six standardization efforts chosen by SmoothIT consortium as important and relevant to the project. These efforts are important for further evolution of the Internet and for new position/vision of business players in networks and for new generation services.
- Recent advances and possible standardization efforts with involvement of SmoothIT are described in Chapter 5.
- Chapter 6 outlines the main conclusions of the document. In particular, after the analysis of the different fora, we identify the targets for SmoothIT regarding standardisation efforts.
- Two reports from recent IETF meetings are included in Annex A and B, showing the results of BoF (Bird of Feather) meetings of ALTO and LEDBAT WGs.

3 Partners' Involvement in Standardization

The goal of this section is to identify initially and provide information about the involvement of SmoothIT partners (both past and current) into standardization efforts and also to list the standardization institutions or activities, which are of importance to the SmoothIT consortium. The most relevant standardization efforts are described in Section 4.

When combining contributions from partners describing their involvement in standardization bodies it shows that these activities are complementary.

UZH

UZH is participating in the IETF at the ALTO WG and, to a lesser extent, at the LEDBAT WG.

DoCoMo

IETF: SIP/P2PSIP

- Service discovery for SIP-based distributed services (draft)

IETF and 3GPP (Third Generation Partnership Project): SIP/SIPPING (SIP INvestiGation)

- SIP session mobility (Request for Comments (RFC) with the IETF)

- Is standardized in 3GPP Rel 9 for IMS (IP Multimedia Subsystem) Service Continuity

TUD

No activity in any of the discussed standardization bodies.

AUEB

Some involvement undertaken in IPsphere and Open Grid Forum activities.

PrimeTel

No activity in any of the discussed standardization bodies.

ICOM

- is a member of the DSL (now Broadband) and WiMAX forums,
- has contributed to IETF IPPM (IP Performance Metrics), TV-Anytime, and IEEE (Institute of Electrical and Electronics Engineers) 802.16 work-groups,
- has been a member of ETSI (Telecommunications Standards Institute) and ITU-T (International Telecommunications Union – Telecommunications Sector), and
- has been monitoring IETF AVT (Audio Video Transport) and IETF MMUSIC (Multiparty Multimedia Session Control) workgroups

UniWue

Currently UniWue is not involved in any standardization bodies. In the future UniWue intends to contribute to the IETF ALTO WG in the scope of the SmoothIT project.

TID

TID has some involvement in ITU and IPSphere. TID is also planning to contribute to the IETF ALTO WG with the rest of SmoothIT partners

AGH

Thanks to participation in IST LION [37] and IST NOBEL [38] projects, AGH had an opportunity to contribute to ITU efforts towards defining an intelligent, modern transport network. The main, long-lasting achievement coming from these efforts is a definition of ASON platform (being the counterpart of IETF-based GMPLS concept).

Contributions to ITU-T SG17dealing with Identity Management Systems for NGN (Next Generation Networks) are planned.

4 Most Relevant Standards and Their Description

On the basis of declarations of Partners' related to standardization activities, the possible focus of SmoothIT in T5.3, still open for discussion, is presented below.

4.1 ITU

From the perspective of current rapidly changing IT standards and development of Internet, based on the decentralized approach, the ITU (International Telecommunication Union, [41]) seems to be the organization with rather slow and limited progress without proper relevance to modern networking. However, there are a few arguments to analyze past achievements of ITU and to follow current works. These arguments are mentioned below:

1) ITU is the oldest international organization dedicated to standardize and regulate international radio and telecommunications (founded as the International Telegraph Union in Paris in 1865).
2) ITU, as a specialized agency of United Nations for agency for information and communication technologies, has wide reputation and support both from official administrations and companies. Currently (November 2008), there are 191 Member States and more than 700 Sector Members and Associates.
3) ITU has the responsibility standardization, allocation of the radio spectrum, and organizing interconnection arrangements between different countries to allow international communications, previously only phone calls, currently any kind.
4) Till mid 80s, ITU was basic standardization organization with broad range of interest. As a consequence, all standards dealing with telephone and data networks, transport systems, radio and TV, signalling and network management, also formal methods and specialized programming languages, were made as ITU Recommendations (prior to 1992 ITU was known as the International Telephone and Telegraph Consultative Committee or CCITT). Then the way and speed of standardization was not acceptable for telecom industry, so practical impact of ITU decreased. Currently, ITU is gaining its momentum by standardizing IP (Internet Protocol) networks, Internet, NGN and other novel technologies.

Recommendations proposed by ITU are split into groups indicated by single letters, each relevant to specific group of problems, areas of networking, applications or administration issues. For SmoothIT consortium 2 groups are of special interest:

H: Audiovisual and multimedia systems

Y: Global information infrastructure, Internet protocol aspects and next-generation networks

ITU took also the responsibility for promotion and regulations of Internet and Information Society, approaching this problem rather from administrative and political positions then technical ones [42]. In the section Internet Governance [42] there are many aspects mentioned and briefly defined. Paragraph 58 states: “Internet governance includes more than Internet naming and addressing. It also includes other significant public policy issues such as, inter alia, critical Internet resources, the security and safety of the Internet, and developmental aspects and issues pertaining to the use of the Internet”. Then, in paragraph 59 one can find: “Internet governance includes social, economic and technical issues including affordability, reliability and quality of service”. These two paragraphs express briefly similar targets as for SmoothIT project.

4.2 IRTF P2PRG

The IRTF peer-to-peer research group (P2PRG) [35] was initiated in 2003 with the major goal to serve as a bridge between short-term interests of the commercial P2P world and the long-term research on fundamental P2P issues mainly driven by academia. In doing so the group offers a forum for researchers to explore a broad range of fundamental P2P issues. Additionally, the group intends to fill the gap that there is no foundation upon which one can build a unified, P2P network on the Internet. To this end, the research group also emphasizes the near-term goals like a classification of the P2P problem space into problems with and without existing solutions, or the development of descriptive model(s) of

peer-node organization. The group intends to offer input to the IETF as a starting place for possible groups standardizing new protocols that are useful in building P2P applications.

The original charter of the working group contains some interesting definitions and statements on P2P networks:

“Peer-to-Peer (P2P) is a way of structuring distributed applications such that the individual nodes have symmetric roles. In P2P applications a node may act as both a client and a server. A key concept for P2P systems is to permit any two peers to communicate with one another in such a way that either ought to be able to initiate the contact. As such, P2P is a powerful tool for organizing cooperative communities - both in the research and commercial domains - with common goals. In addition, as commercial P2P deployment on the Internet has raced ahead of research and standards, issues as basic as interoperable, scalable, P2P communication protocols have been set aside. There is no foundation upon which one can build a unified, P2P network on the Internet.”

After some activities in early phase of the group it can be stated now that the group is currently inactive. One of the main achievements relevant for SmoothIT is the Informational RFC 4981 [15] which gives an overview of search methods in robust P2P networks.

Currently, an attempt was made by CISCO and Alcatel Lucent to reactivate the group by proposing a new charter emphasizing the cooperation with the new established IETF working groups ALTO, LEDBAT, and P2PSIP and provide a forum to discuss research topics that exceed the working group charter. Initial research topics include taxonomy of the uses of P2P applications, an empirical study of peer distributions in BitTorrent swarms, survey of peer selection algorithms, in particular with respect to the impact of peer selection strategies on cross-ISP fairness and with respect to P2P. The topics of the new charter are definitely of high relevance for SmoothIT and are partly also issues actively pursued by SmoothIT. However, the success of the reactivation of P2PRG is questionable since there was little response to the charter proposal. SmoothIT should definitely observe what is going on in the research group and a future active research group might be a forum for some results of SmoothIT.

Table 1: Main characteristics and achievements of IRTF P2PRG

Relevance for IT community, P2P, Future Internet initiative	Currently none since de facto inactive, may change with the new charter
Relevance for SmoothIT project	Topics of new charter are partly overlapping with SmoothIT problem statement. High relevance if the group becomes active again
Description of main goals and achievements of standard (standardization activities), scope of its usage	Referring to the new charter: underlay awareness, interaction of different P2P networks, search in DHTs, security with protection against malicious nodes, performance and monitoring of P2P overlays
Description of details, if studied/known	None

4.3 IRTF NMRG

Network management is a very complex and multilevel framework aiming at planning, monitoring, supervising and controlling of network elements and network itself. More specifically, the network management refers to the activities, methods, procedures, and tools that pertain to the operation, administration, maintenance, and provisioning of networked systems – using the nomenclature related directly to activities run within networks.

During the evolution of network management concepts two mainstreams has emerged: TMN (Telecommunication Management Networks), oriented towards traditional, voice-like networks and SNMP-based (Simple Network Management Protocol) packet-oriented ones. Internet community in natural way deals with packet networks. The most important issues here are already mentioned SNMP and MIB (Management Information Base).

The IRTF Network Management Research Group (NMRG) was initiated in March 2000 and its primary goal was expressed by charter description: "[It] … provides a forum for researchers to explore new technologies for the management of the Internet. In particular, the NMRG will work on solutions for problems that are not yet considered well understood enough for engineering work within the IETF. The initial focus of the NMRG will be on higher-layer management services that interface with the current Internet management framework. This includes communication services between management systems, which may belong to different management domains, as well as customer-oriented management services." [34]. Main achievements of IRTF NMRG are described in RFC3430, RFC3444, RFC3780 and RFC3781.

Recent NMRG meeting was held in October 2008 in Munich [35] and was devoted to experiences and ideas in the usage of Netflow/IPFIX (IP Flow Information Export) in network management. These specific subjects reflect the importance of flow-based approaches which are currently used in various areas of network management, e.g., accounting, dependency discovery, and intrusion detection.

Table 2: Main characteristics and achievements of IRTF NMRG

Relevance for IT community, P2P, Future Internet initiative	Not direct, but the awareness of network management concepts and methods is of great importance.
Relevance for SmoothIT project	The evolution of NMRG towards management of flow-based approaches gives an interesting framework, also for P2P systems.
Description of main goals and achievements of standard (standardization activities), scope of its usage	Four RFCs available, the most interesting seems to be RFC3780, defining SMIng - Next Generation Structure of Management Information, however the work towards this concept was stopped.
Description of details, if studied/known	Tbd. Not studied.

4.4 SIP, NSIS, and P2PSIP

This section covers three groups of session initiation protocols: SIP, NSIS, and P2PSIP.

4.4.1 Session Initiation Protocol (SIP)

The SIP (Session Initiation Protocol) protocol is designed to provide signalling capabilities to various applications, since originally it was mainly designed for Voice-over-IP. It tries to be independent of the underlying transport layer, e.g., TCP (Transmission Control Protocol), UDP (User Datagram Protocol), or SCTP (Stream Control Transmission Protocol). The protocol is text-based, allowing for humans to read and analyze SIP messages. SIP is standardized by the IETF SIP Working Group [25]. Its charter states:

"SIP is a text-based protocol, similar to HTTP (Hyper-text Transport Protocol) and SMTP, for initiating interactive communication sessions between users. Such sessions include voice, video, chat, interactive games, and virtual reality." [23]

There are numerous extensions proposed to extend the protocol, e.g. a locality conveyance capability where clients can exchange their location information and servers route clients' messages based on clients' location [24].

Since SIP was designed to support PSTN-like features for IP networks, it is server-based with different kinds of network infrastructure entities: proxy servers, registrars, and redirect servers. Additionally, gateways can be used to establish connections with other protocols and even networks (e.g., PSTN (Public Switched Telephony Network)).

SIP can be used for Instant Messaging and presence information, or to support streaming applications (video and audio) where the data itself is transported via other protocols, e.g., RTP (Real-time Transport Protocol). A Java-based reference implementation of SIP is JAIN-SIP [26].

Table 3: Main characteristics of SIP.

Relevance for IT community, P2P, Future Internet initiative	SIP can be considered as a general purpose signalling protocol for the Internet that can be used supplementary to other (transport) protocols. Currently, it is mostly used in VoIP/IM applications and it is not clear whether it will gain significant importance outside of this field.
Relevance for SmoothIT project	Since there are no heavy-hitters overlays that rely on SIP, it can be only interesting for SIS-related communication. But seem to be other, more lightweight, solutions based on DNS-extensions and protocols such as REST or XML-RPC.
Description of main goals and achievements of standard (standardization activities), scope of its usage	Transport protocol independent session signalling.
Description of details, if studied/known	Some of SIP limitations are the restriction to text-only messages and the necessity of

	servers (since they can communicate in a decentralized way with each other). Therefore, applications that require a lot of signalling traffic (such as P2P overlays) might require other protocols (e.g. P2P-SIP).

4.4.2 NSIS: Next Steps in Signalling

"The Next Steps in Signalling Working Group is responsible for standardizing an IP signalling protocol with QoS signalling as the first use case"[20]. It further considers the handling of NAT (Network Address Translation) devices that cause problems to SIP applications. NSIS should extend the capabilities of SIP with QoS guarantees.

Table 4: Main characteristics of NSIS.

Relevance for IT community, P2P, Future Internet initiative	NATs, lack of QoS (Quality-of-Service) are common issues with demanding applications. Here NSIS tries to complement SIP.
Relevance for SmoothIT project	Currently, NSIS seems to be not very widely used, since it could be applied for QoS signalling in overlay applications.
Description of main goals and achievements of standard (standardization activities), scope of its usage	QoS signalling and NAT traversal seems to be the main difference to SIP. The NSIS WG position NSIS as a replacement to currently used RSVP.
Description of details, if studied/known	One of their main activities is the GIST (General Internet Signalling Transport) protocol [19]. There are also QoS NSLP, NAT/FW NSLP and a Diagnostics NSLPNSIS protocols. One available implementation of NSIS protocols is FreeNSIS (formerly OpenNSIS) [18].

4.4.3 P2PSIP

P2PSIP is considered by another IETF working group that tries to achieve functionality similar to SIP without server-based infrastructure [21]. Main solutions use Distributed Hash Tables (DHT) to allow participants to find each other and to establish sessions. Since not all peers have to be DHT nodes. Different approaches exist, including those that try to reuse existing DHTs (e.g., OpenDHT), build their own DHTs or to provide interoperability among several DHTs. There are several (mostly industry driven) implementation of P2PSIP DHTs [22].

Ideally, P2PSIP efforts could result in a general-purpose DHT that can be used to discover and locate nodes or peers even if they are constantly moving and changing their IP addresses.

4.5 OpenGrid Forum

OpenGrid Forum (OGF) is the international community dedicated to accelerating grid adoption to enable business value and scientific discovery by providing an open forum for grid innovation and developing open standards for grid software interoperability [39]. OGF is supported by 40 organizational members, classified by level of their involvement. The only one Platinum member is Microsoft, among Gold members are IBM, Oracle, HP Invent, Fujitsu, Shell, Intel, Nortel. At present there are 142 documents available from OGF, many of them seem to be of some interest for the SmoothIT project.

Technical activity of OGF is split into three areas called “functions”:

- eScience,
- Enterprise,
- Standards.

Within Enterprise function there is a working group Telecomm Community Group with general target being relevant to SmoothIT: “This group focuses on how network providers and the Grid community can seek to create common approaches in dealing with significant changes in the offering of services and applications, as well as using infrastructure that will support grids and their evolution. It will seek advice and opinions from customers, network providers, network equipment vendors, Grid middleware and applications vendors and technologists.”

Moreover, there are several Working Groups within the OGF that might be of interest to SmoothIT. The Network Mark-up Language Working Group (NML-WG) aims to create a network description ontology in order to facilitate interoperability between projects. This scheme would be used to create inter-domain maps of network elements at various abstraction levels. This aligns with the need for SmoothIT’s ETM (Economic Traffic Management) components to interface, understand and make decisions based on an AS’s topology.

The Network Measurement and Control WG is developing schemata for network measurement and control based on the work done by the perfSONAR project. This project developed XML-based (eXtended Mark-up Language) protocols in order to monitor and control the network infrastructure. This aligns with SmoothIT’s need to make the SIS monitor and/or control various network elements in order to provide its service to the end users. Similarly, the OGF Network Measurements Working Group (NM-WG) identifies network metrics that are relevant to distributed applications.

Finally, the Network Service Interface WG (NSI-WG) defines an open interface to network services that can be used by external entities such as users, applications and middleware. This WG also covers service interoperability in multi-domain environments as well as configuration, monitoring and orchestration of network resources under particular agreements and policies for service setup. Again, this aligns to SmoothIT’s SIS (SmoothIT Information Service)/ETM component functions.

Table 5: Main characteristics and achievements of OGF

Relevance for IT community, P2P, Future Internet initiative	Increasing importance since this organization migrates from academic-oriented towards more open architectures and applications.

Relevance for SmoothIT project	Generally OGF is oriented towards all aspects of distributed and high performance computing and such targets have some similarities to SmootIT project.
Description of main goals and achievements of standard (standardization activities), scope of its usage	Wide scope of standardization for high performance grid networking.
Description of details, if studied/known	In [10] there is an effort to define and tailor functionality of OGSA (Open Grid Services Architecture), developed by OGF, with fulfilled requirements of P2P systems. Despite the fact that P2P systems have different properties than grid ones, such as low level of security and trust, special connectivity characteristics (flexible IP networking) and usage models (mainly file sharing and collaboration), defining and developing a set of requirements for OGSA with the focus on P2P, is important and recognized by OGF.

4.6 NGN Architectures

The NGN concept is described in ITU Y series Recommendations [41].

As stated in [43], the term NGN (Next Generation Networks) is commonly used to give a name to the changes that are being carried out in the service provision infrastructures. It is more or less clear that the main focus of the NGN could be summarized in the following key topics:

(i) to provide better access (high bandwidth, QoS),

(ii) to be able to efficiently carry different services, and

(iii) to integrate mobile and fixed architectures and services.

Furthermore, all these features should be implemented in the networks in such a way that simplicity for the end users and operators' management is a given. Therefore, taking into account the foreseen users' requirement for improved QoE (Quality-of-Experience) and taking into account the role of the NGN for building frameworks to provide QoS, one of the major challenges is to guarantee the users' QoS requirements between the end points involved in the communication. As stated in [43], a new architecture must be designed in order to address this goal.

In [44], a draft picture of the NGN Architecture is depicted and described. In this architecture the following strata can be distinguished:

- The Transport Stratum includes:
    - The transport functions (in the access, metro and core network). The evolution of all these functionalities are related to the evolution of the network technologies itself (e.g., new optical solutions for the core networks, FTTH (Fiber to the Home), new wireless mechanisms).

- The transport control functions (resource and admission control functions and network attachment control functions). Currently there two standardization bodies in charge of leading the evolution of this control plane: TISPAN (Telecommunications and Internet converged Services and Protocols for Advanced Networking) and ITU-T.

- The Service Stratum which includes the service control functions including service user profile functions; and the application support functions and service support functions. In principle, any service stratum can use the transport stratum capabilities but, maybe the most clear standardization (and also commercial) initiative is the IMS (IP Multimedia Subsystem), triggered by the 3GPP (Third Generation Partnership Project), which specifies an environment where the network operator is in charge of providing the services.

Taking into account these strata, the only way to tune the transport capabilities for any service will be to use the transport control services. Therefore, it is important to identify which are the reference points that can be used. In the Figures 1 and 2 the ITU-T ([45]) and ETSI/TISPAN ([46]) approaches for the NGN Architectures are shown. In these figures, the functionalities associated to the Resource and Admission Control (RACF in ITU-T and RACS in TISPAN) and the Network Attachment functions (NACF and NASS) are part of the Transport Control Functions.

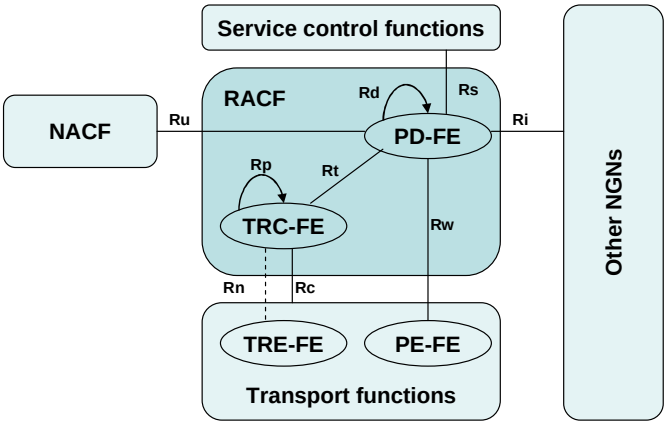

Figure 1: ITU-T NGN Architecture

As shown in Figures 1 and 2 the reference interfaces to access the transport control capabilities are the Rs in the case of the ITU-T approach and the Gq' in the case of the ETSI/TISPAN interface.

Since an important target of the SmoothIT project is the specification of mechanisms that could be used to tune the network to the overlay requirements, these interfaces can be used with this purpose. In this way, the SmoothIT solution can be considered a part of the Service Stratum that interacts with the Transport Functionalities through the Control Transport Functionalities. This interaction can be used to address the following two scenarios:

- The overlay end user requests guaranteed capabilities for specific flows. In this case the SmoothIT solution will adapt this request to the NGN Transport Control expected request.

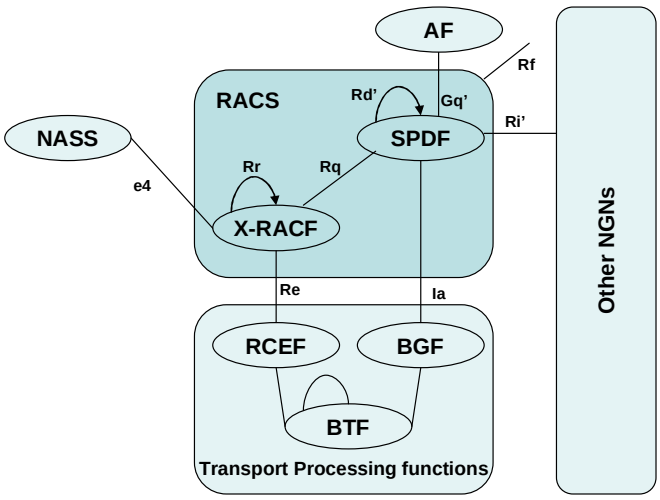

Figure 2: ETSI/TISPAN NGN Architecture

- The ISP fulfils an agreement with the Overlay Service Provider (following a model similar to the IPSphere) and uses the NGN Transport Control capabilities to effectively enforce this agreement, as it is shown in the next figure.

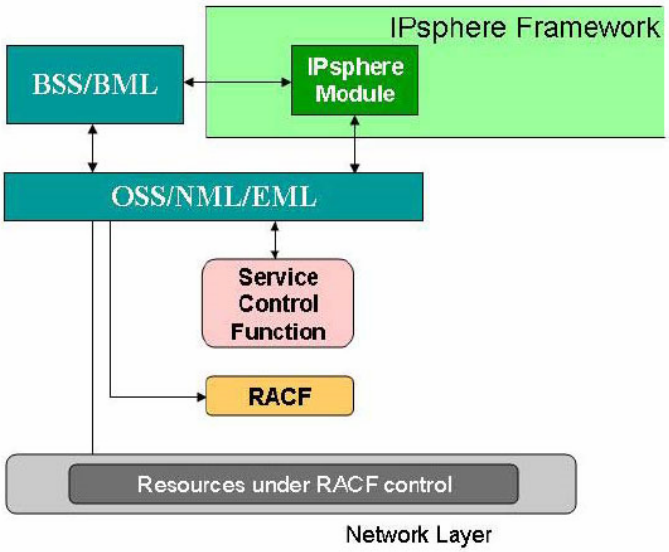

Figure 3: IPSphere and NGN relationship

Table 6: Main characteristics of NGN concept

Relevance for IT community, P2P, Future Internet initiative	NGN is used to name the network evolution. Therefore, any advance on the network evolution would have an important impact on the IT community, the P2P and the Future Initiative.

	In particular for the Future Internet Initiative, the AKARI Project (the Japanese Future Internet Initiative) is focused on the evolution of the NGN as the started point to build the Future Internet.
Relevance for SmoothIT project	As explained above, SmoothIT can use and contribute to the NGN evolution. • SmoothIT will take advance of the interface to the Transport Control Functions in order to tune the Transport Capabilities. • SmoothIT solution will represent a Service Stratum that can be standardized in the future.
Description of main goals and achievements of standard (standardization activities), scope of its usage	• 3 main capabilities are distinguished: service, control and transport. • Data formats are defined to common flow description at the transport control functionalities. • Important achievements in the protocols to manage the user location and authorization.
Description of details, if studied/known	The Gq' interface of the ETSI/TISPAN mainly provides a service that allows the initiation and finalization of a policy request that can be related to a specific user flow or to the user profile.

4.7 IPsphere Framework

The scope of the IPsphere Framework [47] is to offer the ability to abstract and compose telecommunications and IT resources into services via a standardized messaging structure. To do so, IPsphere uses the principles of a service-oriented architecture (SOA), and delivers business flexibility. Service creation utilizes a standardized registration and discovery process. It also supports payments for resource usage, quality assurance, and regulatory compliance. With complete service autonomy and control, the IPsphere Framework meets the needs of Service Providers, Network Operators, Content Providers, Application Developers and System Integrators.

The Release 1 Technical Specification details three key areas:

- Service Abstraction describes the business and technical characteristics of a service and its constituent service elements.
- Service Composition identifies and selects elements that satisfy these technical and business requirements.

- The IPsphere Service Structuring Stratum (SSS) is where the abstraction and composition takes place and it provides support for structuring, executing and assuring these services.

In the context of a typical network service provider's operating environment, the following figure illustrates the broad nature of the interaction with operational/business support subsystems and element/network management systems that are required to manage these pan-provider and multi-stakeholder services. Although the diagram presents the IPsphere in the context of a bilateral relationship between providers/stakeholders, the IPsphere framework is fully extensible to support any configuration of multilateral relationships.

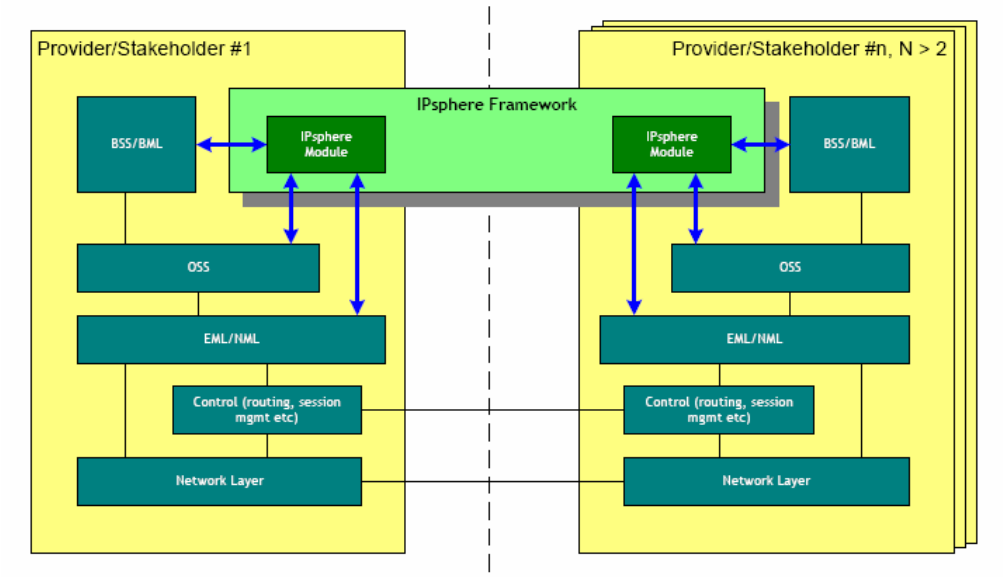

Figure 4: IPsphere Framework in the context of typical network provider's operating environment

A fundamental objective of this framework is to define interfaces for existing systems to participate in the 'meet in the middle' interaction between service-based views and resource-based views, as described above. It also defines how vendors of management, billing and publishing systems can participate in the framework. The diagram in Figure 5 below, illustrates the IPsphere Framework context in more detail and shows how a Service Provider offers a service to a Customer.

In September, 2008, IPsphere Forum became an integral part of TM Forum through the TM Forum IPsphere Transaction. As already mentioned, Release 1 Technical Specification is already available, and now the Forum works for the 2nd release which will focus more on the interaction between Service Provider and Third-party providers.

Table 7: Main characteristics and achievements of IPsphere Framework

Relevance for IT community, P2P, Future Internet initiative	Highly relevant with all-IP convergence and NGN networks, IPsphere offers a way to compose services and built new markets and business relationships.
Relevance for SmoothIT project	IPSphere is defining the SSS (Service

	Structuring Stratum) as the stratum that will be in charge of interfacing the control functionalities of the NGN networks in order to implement new business relationships. On the other hand, the QoS based ETM mechanism proposes a solution that interfaces the transport control functionalities of the NGN which can be seen as a specific implementation of the Service Stratum (following ITU-T NGN terminology) and which, in fact, can be seen a specific implementation of the SSS (since in the QoS Based ETM, SmoothIT is also considering the SLA agreement with a content provider). Up to now, SmoothIT has started to disseminate this proposal in ITU-T and, if the feedback is positive, SmoothIT will consider the presentation of these scenarios in the IPSphere.
Description of main goals and achievements of standard (standardization activities), scope of its usage	• Defines an information model to capture the requirements of a structured service. • Defines the flows for distributing the configuration information to the participating elements so they can be configured to deliver the service. • Receives fault and performance information from the underlying resources so that service owner can take preventative action or inform the customer of a violation of contracted service level agreement. • Generates events to support auditing, notification, billing and settlement, etc.
Description of details, if studied/known	Not studied.

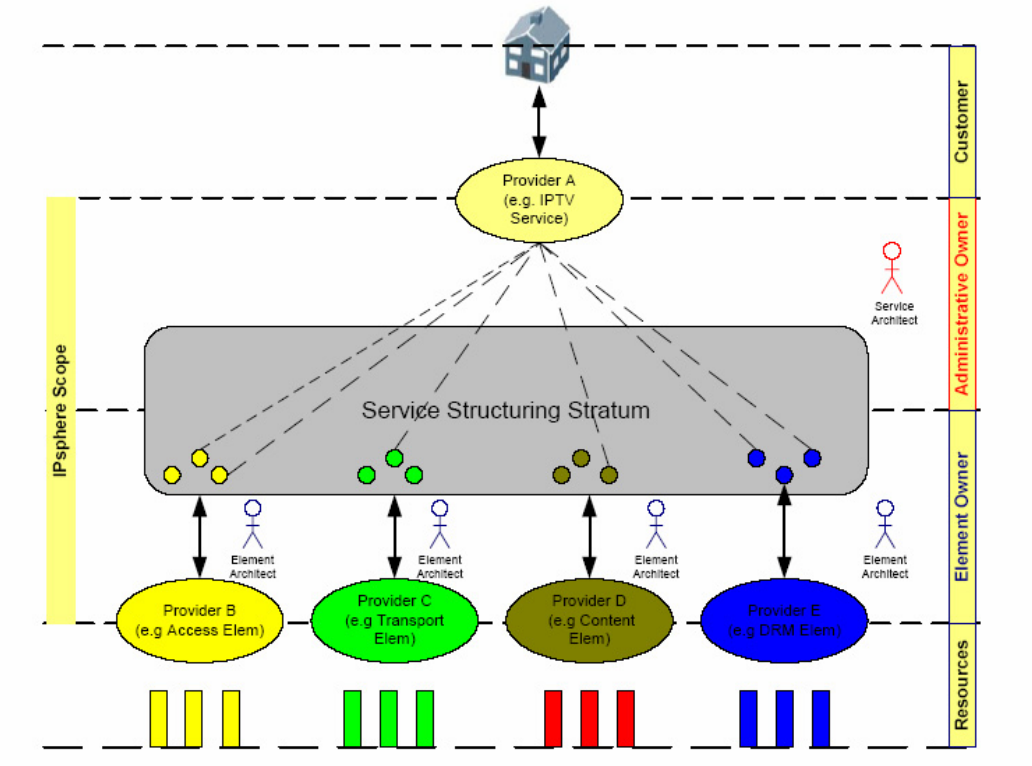

Figure 5: IPsphere Framework context diagram

5 Recent IETF Standardization Efforts

Just recently there was an ongoing work towards establishment of ALTO (Application-Layer Traffic Optimization) group within IETF. This ALTO initiative has been approved as a WG in the Applications area with issued the introductory document [13], (previous version of this draft was dated July 2008). The main motivation of ALTO WG is to define a management framework for peer-to-peer applications. For further information and updated repository of documents, reader should refer to [30].

Traditional client/server architectures use rather stable and well defined architectures while P2P applications use resources distributed across the Internet and exchange large amounts of data with other nodes. Preliminary vision of ALTO problem reduces to providing topology information regarding the underlying network to P2P applications and completing P2P applications' functionalities with selection of the best endpoints among those that are available for the connections they are going to establish.

5.1 IETF ALTO

The ALTO (Application-Layer Traffic Optimization) [32] has been approved as an IETF working group (WG) in November 2008. It aims specifically at managing P2P traffic to lower costs for ISPs and improve performance for users. ALTO explores the fact that, in a P2P application, content is replicated in multiple peers and applications have the freedom to choose which peer to request content from – this process is called peer selection. Most

applications currently perform a random peer selection, to explore path heterogeneity and attempt to achieve higher robustness.

ALTO specifically aims at peer selection, to provide a service (the ALTO service) to feed P2P applications with locality information so they can achieve “better-than-random” peer selection. Although not clearly defined, the word “better” has interpretation similar to the SmoothIT Triple-Win.

Some P2P applications already perform measurements to select peers with which it considers it will achieve higher performance. Common methods are to use IP prefix match, GeoIP database, WHOIS database, AS number matching, or to look at other available services with topology information. However, cost information from the ISP is not currently available to P2P applications.

Replication mechanisms, such as Tit-For-Tat (TFT) are part of many P2P applications and also play a role in peer selection. They enforce fairness by providing incentives for peers to provide resources in order to consume. The use of reputation mechanisms have the effect that peers with higher capacity and willingness to provide also receive better performance. It can be then considered that reputation mechanisms up to a certain point already attempt to optimize peer selection to maximize performance. ALTO would, in one hand, provide a better starting point for reputation to be established between peers with higher potential. Reputation mechanisms are, however, an important factor that is not within the scope of ALTO.

The use cases targeted by ALTO are file sharing applications, cache/mirror selection, live media streaming, real-time communications, and Distributed Hash Tables (DHT). The objective in all use cases is the same – reduce cost for the ISP and increase performance by exploring localization of traffic.

The ALTO service is foreseen to be provided in a client/server fashion by ISPs, third parties and/or user communities. Communication between those instances has not been discussed yet by the WG.

Service discovery is an important aspect that the group intends to target. It is not clear yet whether a current IETF standard will be used for this purpose or if there is necessity for the development of another service discovery service.

User privacy is also mentioned as a concern. However, at this point, it is not fully defined what type of information is considered sensitive and how the design of the protocol will take privacy into consideration. Most likely it will be addressed by not requiring the user to disclose detailed information to the ALTO server, but that alone would not address the case when a user does not want his or her IP to be ranked by an ALTO server at all; or the case he or she is not aware of the implications of having its IP address ranked. ISP privacy is also foreseen to be addressed by the WG, pointing that they must not necessarily disclose detailed topology information that might be considered strategically sensitive. The WG has not yet specified how much information will need to be disclosed though. The trade-off is not clear.

The ALTO WG also aims to consider security issues, although at this point very superficial at the moment. The current draft points the optional character of ALTO as a security mechanism, which is awkward. It reads that in case ISPs are malicious, or there is an asymmetry of interest between ISPs and P2P applications, the latter would be able to detect such abuse and stop using the ALTO service. This is not a solution to the security problems that may arise with the introduction of a new service.

There are several drafts [32] showing research work that could serve to model an ALTO service. P4P [28] and the work developed at TU Berlin [1] are two candidates that fit into the proposed framework. Both suggest a client/service approach in which the client (a P2P application) constantly queries the server, sending a list of peers that offer the same resource, to request advice on peer selection. There are also suggestions to build the ALTO service in such a way that the server exports all information to the clients so they can perform all processing and decisions. In this case, there is no need for peers to send list of potentially selected peers to the ALTO server.

A standardized open protocol helps to achieve interoperability among different services. However, research must be done in order to guide applications, ISPs and service providers on how to use such a service in order to successfully achieve a triple-win situation. This is not the scope of IETF.

With regard to future work, the ALTO WG intends to release its main protocol as a standard by March 2010. The charter is foreseen to be dissolved again or re-chartered in August 2010.

5.2 IETF LEDBAT

The LEDBAT (Low Extra Delay Background Transport) [28] has been approved as an IETF working group (WG) in November 2008 based on the Internet Draft [17] resulting from the TANA (Techniques for Advanced Networking Applications) BoF.

The main goal of the working group is to standardize a congestion control mechanism that should saturate the bottleneck, maintain low delay, and yield to standard TCP. Concretely, the new working group is driven by two major tasks which should both result in RFC submissions in October 2009. The first one is to design an experimental congestion control algorithm for less-than-best-effort "background" transmissions, i.e., an algorithm that attempts to scavenge otherwise idle bandwidth for its transmissions in a way that minimizes interference with regular best-effort traffic. The second one is to produce a document that clarifies the current practices of application design and reasons behind them and discusses the tradeoffs surrounding the use of many concurrent TCP connections to one destination and/or to different destinations.

The first work item is driven by the fact that when transferring large data volumes TCP fills the potentially large bottleneck buffer and thus causes large delays for concurrent interactive applications like VoIP or web browsing. In particular home uplinks that carry mainly P2P traffic are seen as the worst case where filled buffers may lead to RTTs (Round-Trip Time) in the order of seconds. The idea is to solve this problem by a transport/application layer solution which may be an extension or an alternative to TCP. Requirements for the solution are to saturate the bottleneck link while keeping the delay for other traffic as low as possible and to yield to normal TCP traffic. The solution should work well for simple FIFO (First-in First-out) drop tail queues but also be able to utilize the potential of AQM (Active Queue Management), DiffServ (Differentiated Services), and ECN (Explicit Congestion Notification) when present. Current proposals are a variation of TCP with an adaptive receive window, called “Low Priority TCP: Receive Window control” by Microsoft and a congestion control mechanism deployed in BitTorrent, i.e. on application layer.

The second work item is to give an overview of how applications make use of multiple parallel TCP connections. The reasons for a P2P application and a web browser are

different. While a web browser wants to achieve a higher bandwidth and more robustness, a P2P client mainly intends to keep a meshed overlay with long-lived TCP connections. The motivation behind studying multiple TCP connections is that TCP tries to achieve per connection fair sharing of the bottleneck bandwidth and an application with multiple parallel TCP connections consequently receives too much bandwidth. Currently, a first draft [14] still under the TANA label is present. The goal is to obtain an overview which applications use multiple TCP connections and to understand the intentions of doing so and the advantages and consequences on the performance.

The scope of the LEDBAT does not overlap with SmoothIT since it proposes transport layer and congestion control solutions. However, the results of LEDBAT might have an impact how the data transfer in P2P applications takes place in the future. The possible effects of such a congestion control should be considered in the design of SmoothIT's ETM and possible chances for exploiting these mechanisms should be investigated. Accordingly, the achievements of the working group will be further monitored.

Table 8: Main characteristics and achievements of IETF LEDBAT WG

Relevance for IT community, P2P, Future Internet initiative	Definitely relevant. A new congestion control mechanism might change the traffic control on lower layers. The performance of overlay applications might also be affected.
Relevance for SmoothIT project	No direct relevance. SmoothIT cannot contribute to LEDBAT. However, introducing new congestion control could have an impact on SmoothIT performance. Investigations about multiple TCP connections could be interesting.
Description of main goals and achievements of standard (standardization activities), scope of its usage	1. Develop congestion control for less than best effort applications that yields to standard TCP 2. Gain understanding around multiple TCP connections.
Description of details, if studied/known	None

6 Summary and Conclusions

In this deliverable, theSmoothIT project reviewed ongoing standardization efforts relevant to our project. Based on the relevance and the potential impact of these efforts on the future operation of the Internet, we selected two standardization bodies to actively participate in: ITU and IETF ALTO working group.

SmoothIT can take advantage and also contribute to the current specification of NGN Transport Control Functions. In particular, as described in the main document, using the Gq interface of the ITU, SmoothIT can provide a service that allows the initiation and finalization of specific policy requests (being the way to implement QoS enforcement in the

networks). In this way, the SmoothIT solutions will act as the Service Stratum being defined in the ITU. The exchanged data formats could constitute a good input for the NGN standardisation bodies.

IETF ALTO working group aims at locality-based peer selection, leading to what SmoothIT calls “triple-win” In a much wider context Thus it is the most natural choice for standardization efforts of the SmoothIT project. Another good reason for the decision to actively participate in the ALTO WG is the fact that ALTO has the potential of defining a protocol to be widely deployed and used in the Internet. As such, it already generates enormous interest in the Internet community and SmoothIT expects that this interest will even grow in the future.

7 References

[1] V. Aggarval, A. Feldman, Ch, Scheideler: Can ISPs and P2 Users Cooperate for Improved Performance? ACM SIGCOMM Computer Communication Review, Volume 37, Number 3, July 2007.

[2] S. Le Blond, A. Legout, W. Dabbous: *Pushing BitTorrent Locality to the Limit*, available at http://arxiv.org/abs/0812.0581.

[3] P. Calhoun, J. Loughney, E. Guttman, G. Zorn, J. Arkko: *Diameter Base Protocol*; Internet Engineering Task Force (IETF) RFC 3588, September 2003.

[4] K. Bhatia: *Peer-To-Peer Requirements On The Open Grid Services Architecture Framework*, GFD.49, Open Grid Forum, July 12, 2005.

[5] O. Bonaventure, D. Saucez, B. Donnet: *The case for an informed path selection service, draft-bonaventure-informed-path-selection-00.htm*, February 18, 2008.

[6] S. Bradner: RFC2026, The Internet Standards Process, October 1996.

[7] A. Clemm: Network Management Fundamentals, Cisco Press, 2006.

[8] C. Griffiths, J. Livingood, R. Woundy, *Comcast's ISP Experiences In a Recent P4P Technical Trial*, draft-livingood-woundy-p4p-experiences-02, October 28, 2008.

[9] V. Hilt, I. Rimac, M. Tomsu, V. K. Gurbani, E. Marocco: *A Survey on Research on the Application-Layer Traffic Optimization (ALTO) Problem*, draft-hilt-alto-survey-00.htm, July 3, 2008.

[10] S. Jha, A. Mertzky: *A Requirements Analysis for a simple API for Grid applications*, GFD.71, Open Grid Forum, May 9, 2006.

[11] J. Ni, H. Xie, S. Tatikonda, Y. R. Yang: *Network Routing Topology Inference from End-to-End Measurements*, *INFOCOM 2008*. March 2008.

[12] S. Kiesel, L. Popkin, S. Previdi, R. Woundy, Y R. Yang: *Application-Layer Traffic Optimization (ALTO) Requirements*, draft-kiesel-alto-reqs-01.txt, November 3, 2008.

[13] E. Marocco, V. Gurbani: *Application-Layer Traffic Optimization (ALTO) Problem Statement*, draft-marocco-alto-problem-statement-03.txt, November 2, 2008 (it obsoletes version 2, issued 10 July 2008).

[14] R. Penno and J. Iyengar: TANA *Practices and Recommendations*, draft-penno-tana-app-practices-recommendation-01.txt, Internet Draft, 3-Nov-08.

[15] J. Risson, T. Moors: Survey of Research towards Robust Peer-to-Peer Networks: Search Methods, RFC 4981, September 2007.

[16] S. Shalunov, R. Penno, R. Woundy: *ALTO Information Export Service*, draft-shalunov-alto-infoexport-00.txt, October 27, 2008.

[17] S. Shalunov, *Transport for Advanced Networking Applications (TANA) Problem Statement*,draft-shalunov-tana-problem-statement-01.txt, Internet draft, 14-Jul-08.

[18] FreeNSIS: http://user.informatik.uni-goettingen.de/~nsis/home.html.

[19] GIST draft: http://tools.ietf.org/html/draft-ietf-nsis-ntlp-16.

[20] NSIS charter http://www.ietf.org/html.charters/nsis-charter.html

[21] P2PSIP WG http://www.p2psip.org/

[22] P2PSIP implementations: http://www.p2psip.org/implementations.php.

[23] SIP WG charter: http://www.ietf.org/html.charters/sip-charter.html.

[24] SIP status page: http://tools.ietf.org/wg/sip/.

[25] SIP WG supplementary page: http://www.softarmor.com/sipwg/.

[26] Jain-SIP, http://jain-sip.dev.java.net/.

[27] L. Wen, Y. Zhang: *P2P Traffic Localization by Traceroute and 2-Means Classification*, draft-zhang-alto-traceroute-02.txt, October 23, 2008.

[28] H. Xie, A. Krishnanmurphy, Y. N. Yang, A. Silberschatz: *P4P: Proactive Provider Participation for P2P*, YALEU/DCS/TR-1377, January 2007.

[29] Y. Zhang, H. Liao, Naibao.Zhou: *P2P Traffic Localization by Alias Tracker for Tracker-based P2P*, draft-zhang-alto-attp-02.txt, October 23, 2008.

[30] URL: http://alto.tilab.com/docs/charter.txt.

[31] URL: https://www.ietf.org/html.charters/alto-charter.html.

[32] URL: http://tools.ietf.org/wg/alto/.

[33] URL: http://tools.ietf.org/wg/ledbat/agenda, November 18, 2008.

[34] URL: http://www.irtf.org/charter?gtype=rg&group=nmrg

[35] URL: http://www.ibr.cs.tu-bs.de/projects/nmrg/meetings/2008/munich/.

[36] URL: http://www.irtf.org/charter?gtype=rg&group=p2prg.

[37] URL: http:www.telecom.ntua.gr/lion/ IST–LION web page..

[38] URL: http:// www.ist-nobel.org, IST-NOBEL Web page.

[39] URL: http://www.ogf.org/.

[40] URL: http://www.ogf.org/gf/group_info/charter.php?review&group=telco-cg.

[41] URL: http://www.itu.int/.

[42] ITU World Summit on information society*: Tunis agenda for Information Society*, 18 November 2005, http://www.itu.int/wsis/docs2/tunis/off/6rev1.html.

[43] ITU-T Y.2000, *General overview of NGN*.

[44] ITU-T Y.2012, *Functional requirements and architecture of the NGN release 1*.

[45] ITU-T Y.2111, *Resource and Admission Control functions in Next Generation Networks*.

[46] ETSI ES 282 003, Telecommunications *and Internet converged Services and Protocols for Advanced Networking (TISPAN); Resource and Admission Control Sub-System (RACS): Functional Architecture.*

[47] TeleManagement (TM) Forum, The IPsphere Framework, http://www.tmforum.org/IPsphereFramework/6214/home.html.

8 Abbreviations

3GPP	Third Generation Partnership Project
ALTO	Application-Layer Traffic Optimization
AQM	Active Queue Management
AVT	Audio Video Transport
BoF	Bird of Feather
CCITT	International Telephone and Telegraph Consultative Committee
DHT	Dynamic Hash Table
DiffServ	Differentiated Services
DSL	Digital Subscriber Line
ECN	Explicit Congestion Notification
ETM	Economic Traffic Management
ETSI	European Telecommunications Standards Institute
FIFO	First-in First-out
FTTH	Fiber to the Home
GMPLS	Generalized Multiprotocol label Switching
IETF	Internet Engineering Task Force
IMS	IP Multimedia Subsystem
IRTF	Internet Research Task Force
LEDBAT	Low Extra Delay Background Transport
GIST	General Internet Signalling Transport
HTTP	Hyper-text Transport Protocol
IEEE	Institute of Electrical and Electronics Engineers
IP	Internet Protocol
IPFIX	IP Flow Information Export
IPPM	IP Performance Metrics
IRTF	Internet Research Task Force
ISP	Internet Service Provider
ITU	International Telecommunications Union
MMUSIC	Multiparty Multimedia Session Control
NACF	Network Attachment functions
NAT	Network Address Translation
NGN	Next Generation Networks
NM-WG	Network Measurements Working Group

NML-WG	Network Mark-up Language Working Group
NMRG	(IRTF) Network Management Research Group
NSI-WG	Network Service Interface Working Group
NSIS	Next Steps in Signalling
OGF	Open Grid Forum
OGSA	Open Grid Services Architecture
P2P	Peer-to-peer
P4P	Provider Portal (concept) for P2P
PSTN	Public Switched Telephony Network
QoE	Quality-of-Experience
QoS	Quality-of-Service
RACF	Resource and Admission Control
RFC	Request for Comments
RTP	Real-time Transport Protocol
RTT	Round-Trip Time
SCTP	Stream Control Transmission Protocol
SIP	Session Initiation Protocol
SIPPING	Session Initiation Protocol INvestiGation
SIS	SmoothIT Information Service
SmoothIT	Simple Economic Management Approaches of Overlay Traffic in Heterogeneous Internet Topologies
SNMP	Simple Network Management Protocol
STREP	Specific Targeted Research Project
TANA	Techniques for Advanced Networking Applications
TCP	Transmission Control Protocol
TFT	Tit-For-Tat
TISPAN	Telecommunications and Internet converged Services and Protocols for Advanced Networking
TM Forum	Telecommunications Management Forum
TMN	Telecommunication Management Networks
UDP	User Datagram Protocol
WG	Working Group
XML	eXtended Mark-up Language

9 Acknowledgements

This deliverable was made possible due to the large and open help of the WP5 team of the SmoothIT project within this STREP, which includes besides the deliverable authors as indicated in the document control, Robert Wójcik (AGH) and Jerzy Domżał (AGH) and Jan Derkacz (AGH) as well. Many thanks to all of them. A final review was performed by Burkhard Stiller.

Annex A (For Information Only):

IETF 72 – ALTO Meeting

72nd Internet Engineering Task Force Meeting

Author: Fabio Hecht

Facts

Date: July 27 - August 1, 2008

Location: Citywest Hotel, Dublin, Ireland

Participants: 1300+

Participant from UZH: Fabio Hecht

ALTO BoF

The 72nd IETF Meeting took place in a hotel in the far west of Dublin, with more than 1300 attendees. Seven meetings took place in parallel, in a total of three sessions per day, in six days. Attendance in the ALTO BoF meeting is estimated at about 15x2x10=300 participants.

Enrico Marocco, Vijay K. Gurbani – Problem Statement

The meeting began with a presentation about the problem statement. The focus was on P2P traffic. Some approaches were presented, and a line was drawn a line between peer-to-peer (resource-based) and Client/Server (host-based), although the line is not so clear. The problem statement presented was about the same as we have on SmoothIT. He argued that peers performing measures are expensive and not accurate. He mentions throughput as a special case – I wonder whether he will provide a solution for that. The purpose of this BoF is to try and become a WG in the IETF to define only the protocol – not caring about metrics, for example. He mentioned privacy and security aspects but simplified the solution by stating that the service is optional. I think it can be better than that. The current draft can be obtained from http://tools.ietf.org/html/draft-marocco-alto-problem-statement, which contains what has been discussed on the mailing list.

After the presentation, many questions popped up form the audience. Reducing inter-domain domain traffic is THE problem to solve? Not the only, but if this is about problem statement, it must be very clear what the main problem is. Another point is about this specific routing information (intradomain vs. interdomain): is it already available to overlay applications? Some say it can be figured, but still there should be a standard way to fetch it, not a hack. Another person points out that optimization is a balance, and maximizing intradomain traffic may not be ideal – what are they trying to optimize? A more experienced engineer from Verizon mentioned that access network is usually the

bottleneck. Also, application scenarios are a must. Although the protocol should not focus on the metrics, they are important to be foreseen in order to justify the effort to define it.

Volker Hilt - ALTO Survey

The next presentation dealt with a survey, available on http://tools.ietf.org/html/draft-hilt-alto-survey. He started by showing examples on how to select a good peer. He presented several solutions, such as IDMaps, Vivaldi, iPlane, Ono Project, and P4P.

After the presentation, there were many questions, some of the challenging the whole mechanism. One of them stated that nowadays the ISP don't necessarily own its facilities, they rent it. The ISP does not own its domain. Also, a provider may cover several countries – intradomain won't necessarily mean low cost and high quality. Another interesting question was: would all of those presented approaches fit into whatever will be defined? Since there are many metrics, how shall semantic information be kept consistent?

Henning Schulzrinne – Distance-related Network Costs

The presentation of Henning Schulzrinne was mainly about costs and pricing. He showed the current cost of bandwidth to be about 0.15 US$/Mbps/month. Downloading a DVD would cost approximately 1.05 US$ to an ISP. He presented a discrete step graph, picturing cost proportional to distance. Then he mentioned alternative charging schemes, for example, a client losing priority after hitting a certain limit, or having free local traffic (potentially limited), or varying according to time of day (at nights faster), and others. The bottom line was that the economic decisions should be taken by the customers.

As usual, there were many highly critical questions. One participant stated that the highest cost is backbone cost, doesn't matter if traffic goes long or short, if uses the backbone, it costs. Then, a discussion about what a backbone is took place. Another interesting question challenged whether the price should be proportional to the cost.

Sebastian Kiesel – Requirements

Kiesel's presentation focus on how the system could work in a tracker- and DHT-based P2P system. An ALTO server may have to implement several interfaces: feeding topology info, providing coordination between ALTO servers, besides the client interface. He presented the requirements for the protocol that can be viewed in detailed in http://tools.ietf.org/html/draft-kiesel-alto-reqs. Finally, he showed some examples of criteria: distance and cost, but it is still unclear what metrics shall be used.

Questions regarded caching and lifetime of results, peer selection process that can be bidirectional (providing peer also performs selection), and issues related to ISPs caching content.

Reinaldo Penno – Caching and Peer Selection

The presentation of Reinaldo Penno was about cache discovery – not strategies. He presented some techniques, for example using DNS, or using the ALTO server. This is the point on which the meeting got controversial. Should the ALTO server help the peer-to-peer application to find content or caches? If that is the case, there are already many resource discovery protocols, why not reuse an existing protocol instead of creating a new one for this proposal?

Dimitri Papadimitriou – Similar Problems in Multi-Homed Networks

The last presentation targeted the assumption that, in the future, more and more content will be available through multiple sources – be them in peer-to-peer networks, CDNs.

Multiple sources mean multiple paths, and even multiple protocols can be available (e.g. IPv4 and IPv6). Dimitri showed that those applications would need to select the best paths to reach a particular content, and the ALTO server can play a role. He showed current solutions and proposed one of his own. A detailed overview of his presentation can be obtained from http://www.ietf.org/internet-drafts/draft-bonaventure-informed-path-selection-00.txt.

Charter discussion

After all the presentations, there was a general discussion about diverse subjects. One of the most discussed points regarded service discovery – some people were meaning discovering the ALTO server and some people were thinking about using the ALTO server to obtain alternative sources, such as caches or additional peers. It was more or less agreed that they shouldn't reinvent the wheel, and use one of the many standards already in place for such a role. Also, why put together provision and discovery of service? Generally speaking, separation is a good thing.

There was also criticism about the vagueness of the problem statement. It should be short and to the point. Also, the requirements proposed already state the solution. That is very biased.

At the end, the chairs requested the audience to hum if they were in favor of continuing the work creating an ALTO WG. Many people agreed, shyly. Then, people were requested to hum if they were against it. This time, less people hummed, but they did it vehemently. There were several arguments against the standardization of the protocol. Everyone agrees that this is a problem to be solved. However, there are people who are not convinced this is the best solution, since it is not clear that peers get a better answer by trusting the ISPs than they can figure out by themselves; or doubt the approach is the best possible. On the other hand, there are many separate technologies popping up to set up ISP-assisted peer selection, and a standard would be desirable, otherwise there will be many incompatible solutions. The fact is that there is a lot of research work involved and that is not the purpose of the IETF. It might be too early to define a standard, and research – like SmoothIT – must be done to prove the point to all the community.

The discussion is still ongoing in the mailing list p2pi@ietf.org.

Annex B (For Information Only):

IETF 73

73rd Internet Engineering Task Force

Author: Fabio Hecht

Facts

Date: November 16-21, 2008

Location: Hilton Hotel, Minneapolis, MN, USA

Participants: 1115 registered

Participant from UZH: Fabio Hecht

ALTO – Application-Layer Traffic Optimization WG

The newly created ALTO WG is attracting great interest from the IETF. An estimated number of 240 persons attended the meeting.
The first talk was a short introduction, by the chairs, that presented the last important modifications in the charter, stressing important changes. The most important ones are the following:

- Focus is peer selection only;
- Goal is to perform better than random peer selection – not optimal.

After this talk, Aaron Falk talked about the IRTF p2prg (Peer-to-Peer Research Group). They are creating a new charter and are looking for chairs.

The next talk, by Enrico Marocco (chair, Telecom Italia), presented the latest updates in the problem statement draft, reviewing discussions since Dublin. The document already points the solution: “a topology information (...) will allow applications to improve their performance and will help ISPs make a better use of their network resources”. It is very clear that the objective of the WG is to improve performance and reduce inter-domain traffic, focusing on localization of traffic. They say it can also be something other than localization, but it is not clear what can that be. The original document contained the word “oracle”, referring to the solution of Feldmann et al., but due to requests the term has been changed. The first comment from the audience was that solutions should not be part of the problem statement – an analysis of the solution space must be done in the first place. Another concern regards privacy. The document states that “the application does not have to disclose information it may consider sensible”, but it is actually very difficult to say the least to determine what is private and what is acceptable.

Sebastian Kiesel (USTUTT) presented his draft on requirements. He showed changes in the document since the Dublin meeting, the most important ones being the following:

- Not seeking the optimal solution, just better than random one;
- Removed the suggestion of a sorting oracle, that would be not appropriate for requirements;

- Define core set of attributes for expressing preference, extensible to other ones.

The document reads much like the one on problem statement, and a long discussion about this was started, including the theme of why a requirements document is useful after all. In the end, it was more or less agreed that the requirements should be revisited after the problem statement document is refined.

The next talk, by Richard Yang (P4P), was entitled “P4P Design and Implementation”. They are working with two services: location and pDistance. Location is stable and returns a PID for each peer. One PID groups peers that are close together, and the granularity can be played with (AS level is suggested). The pDistance service returns a distance between two peers. PIDs can be interdomain or inter-domain, and the return value can be ordinal, or numeric (he prefers the latter). Types of metrics can be: hopcount, air-mile, cost (which is the default). A person in the audience asked. why not to use one of the existing Internet maps (“looking glasses”) that are available, and how different would it be to just use them? This is a point that is still controversial.

Richard Woundy presented then the talk Comcast's Experiences In a P4P Technical Trial, in which he showed some details about the draft with same title. They are working directly with P4P and ran a test with their customers. The results were improved download speed, and localization (less interdomain traffic). The experiment involved a single, 21MB, file. It was a Pando client update, so users were forced to download it. Criticisms to this experiment include:

- The file is small in comparison to what people have been trading in file sharing applications;
- It was a forced download, so the gains are maximized due to a large swarm, which is not always the case.

The next presentation was about the draft “ALTO Information Export Service”, which suggests that the clients download a table containing full preference information from the ISP so they do not have to keep on constantly querying the ALTO service. He argues that the P2P applications can already do a pretty good job figuring out routing information. What is missing is only the ISP preference, which can be described in a small enough table to be downloaded completely. The presentation included the format of the file, as contained in the draft, with three fields per record: designator (“asn” or “cidr”), AS number or IP prefix (CIDR), and a priority. In his example, the application would sort the peers in three sets: preferred, default, and to be avoided. The size of the table in these tests was 1.5 MByte (compressed). Possible issues recognized by the author are the redistribution of information by peers (could produce outdated information) and, service discovery. This work is being carried on by BitTorrent.

Stefano Previdi (Cisco) presented the next talk, entitled Routing Proximity. His goal is to establish metrics to be classify peers with regard to proximity. Routing databases (ISIS/OSPF/BGP) have already proximity metrics, so they can be used for the purpose of calculating the distance between peers. He believes everything needed to achieve localization is already available. ALTO should be just an interface to support routing (e.g. BGP) information to overlay. An important question is the addition of other metrics, such as cost, link capacity, and congestion – term that appears as a prominent motivation in the charter text.

The last presentation, “A Multi Dimensional Peer Selection Problem”, by Saumitra Das, discussed the fact that many different factors influence peer selection. Some information

might come from ISPs and but ultimately the peers make the selection using information such as reputation. He suggests that different types of ALTO servers (e.g., P4P) can coexist.

The meeting ended with a word from the chairs to keep up current work and discussion on the mailing list.

LEDBAT – LEss Than Best Effort Transport WG

LEDBAT is the newly created WG that stems from the TANA BoF. The name was changed since some people considered the term "Advanced Network Applications" too unspecific. Just like ALTO, it aims at coping with P2P traffic, but works at the transport layer. The idea is to design a protocol that does not interfere with regular (best-effort) traffic, utilizing unused bandwidth. Moreover, the protocol would be able to “scavenge” network resources that would otherwise be unused.

Stanislaw Shalunov (BitTorrent) opened the session with a charter recapitulation. The objective is “to standardize a congestion control mechanism that should saturate the bottleneck, mantain low delay, and yield to standard TCP”. It originated at the P2PI workshop at MIT in May/2008, creating the ALTO and TANA BoFs in Dublin, which lead to the ALTO and LEDBAT WGs in Minneapolis. The problem being solved is that TCP fills router buffers if congested, and the buffer can be large, the likely worst case is the home uplink. Since most traffic on home uplink is P2P-related, it leads to a delay in other applications that users might be using. There are two work items: experimental congestion control and current practices of applications (using multiple connections). Applications (BitTorrent, web browsers, e-mail servers) create multiple connections to try and maximize throughput and add stability. The practice is common, but considered more a poorly documented hack. The idea is to research and document how applications use multiple connections. Bob Briscoe mentioned that they should take into consideration how to respond to the congestion. Another person asked whether the protocol should "yield to TCP" or to be something better, more modern than TCP.

Satish Raghunath (Juniper.net) presented the next talk, entitled “LEDBAT - App Practices and Recommendations”. He mostly talked about why current P2P applications open multiple connections. Mentions reliability that comes with diversity, but comes with overhead, and impact delay-sensitive traffic, due to more state needed within TCP termination devices and middleboxes. The applications try and find a “right” number of connections – too few or too many can cause problems – the objective is to maximize download speed, for example, in BitTorrent. The objective is to elaborate a document with recommendations to the number of connections. The audience asked, whether he will/did look also at UDP or only TCP? At the moment, only TCP, someone (!) could contribute with a draft. Another person pointed out that Firefox opens up to 16 connections to each host, is that a good or a bad number? People also pointed out that is not true that opening multiple connections always makes downloads faster. They experimented with the iPhone over 3G and stated that handshake and congestion control does not work, especially for small images it may not be worth to open another connection. Next step: do the research!

Murari Sridharan (Microsoft) presented next “Low priority TCP: Receive-Window Control”. It was a paper presentation about BATS (background transfer service). They adapt the receiver window to create a low priority service, and he explains how. The algorithm has 2 modes: rate limiting mode (1) and window scaling mode (2). Mode 1: gets accurate RTT samples, mode 2: uses binary search to drive towards target window, assuming the value lies between Wmin and Wmax; depending on whether there is congestion, window size is

adjusted. Bottom line: it maintains low delay and yields to TCP. It requires no support from the network, although additional information helps it adapt quicker. He suggests this work as a starting point of how LED could look like. Question 1: has he analyzed if it works with competition between several connections? The presented answered he is working on getting these numbers. Other important questions regarded RTT independence and whether the background flow should be starved if necessary and respective answers are that it “tries” to be independent from RTT (whatever this means) and that starvation can be controlled.

The next presentation was very short – only 5 minutes – by Nick Weaver (Berkeley), entitled “A Couple Academic Thoughts on LEDBAT”. He affirmed that there are two separate problems without use of packet marking or AQM (active queue management): detect buffer occupancy problems and detect and yield in common congestion to other types of traffic. In his opinion, LEDBAT should be defined as a TCP operating mode. Only one side should need to use the defined congestion control policy (à la 4CP) in order to ease deployment issues. He raises the question on whether DiffServ marking should be used. Bob Briscoe thinks that marking would not work, because it doesn't matter what they put there the operator won't believe. The point is that they would be marking them to be low priority, not to get higher priority.

Stanislaw Shalunov (BitTorrent, chair) presented “Low Extra-Delay Background Transport”, his idea of what could be standardized by LEDBAT. In his view, the main problem is that TCP fills the buffer and it can be large, introducing high delay in case of congestion. This delay breaks real-time applications like VoIP when a P2P application (like BitTorrent) is running. It also slows down considerably traffic that is not real time, like web browsing. He raises the question of how large should the buffer be, but does not have the answer. This raises a long discussion, and Stanislaw mentions measuring one-way delay, which is deemed impossible by at least some of the audience. He presents details of his approach, which includes using smaller packets and estimating queue delay in order to reduce window size before packet loss occurs. The presenter states though that he has done it and tested in BitTorrent DNA by 7M active users. The audience asks whether he has any numbers to show and confirm his statements, but he has not. Part of the audience did not really agree that using small packets should help, the reality is that packets are of a small size for a long time. The author says that he uses smaller packets in order to minimize serialization time and be able to obtain faster speed in a slow link.

The meeting ended with a word from the chairs pointing the WG to future work. They envision researching how applications use multiple connections in order to maximize their download speed and plan to further refine current studied approaches.