ETSI GS MOI 010 V1.1.1 (2010-05)

Group Specification

Measurement Ontology for IP traffic (MOI); Report on information models for IP traffic measurement

Reference

DGS/MOI-0001

Keywords

IP, traffic, ontology, information model

ETSI

650 Route des Lucioles
F-06921 Sophia Antipolis Cedex - FRANCE

Tel.: +33 4 92 94 42 00 Fax: +33 4 93 65 47 16

Siret N° 348 623 562 00017 - NAF 742 C
Association à but non lucratif enregistrée à la
Sous-Préfecture de Grasse (06) N° 7803/88

Important notice

Individual copies of the present document can be downloaded from:
http://www.etsi.org

The present document may be made available in more than one electronic version or in print. In any case of existing or perceived difference in contents between such versions, the reference version is the Portable Document Format (PDF). In case of dispute, the reference shall be the printing on ETSI printers of the PDF version kept on a specific network drive within ETSI Secretariat.

Users of the present document should be aware that the document may be subject to revision or change of status. Information on the current status of this and other ETSI documents is available at
http://portal.etsi.org/tb/status/status.asp

If you find errors in the present document, please send your comment to one of the following services:
http://portal.etsi.org/chaircor/ETSI_support.asp

Copyright Notification

Contents

Intellectual Property Rights

IPRs essential or potentially essential to the present document may have been declared to ETSI. The information pertaining to these essential IPRs, if any, is publicly available for **ETSI members and non-members**, and can be found in ETSI SR 000 314: *"Intellectual Property Rights (IPRs); Essential, or potentially Essential, IPRs notified to ETSI in respect of ETSI standards"*, which is available from the ETSI Secretariat. Latest updates are available on the ETSI Web server (http://webapp.etsi.org/IPR/home.asp).

Pursuant to the ETSI IPR Policy, no investigation, including IPR searches, has been carried out by ETSI. No guarantee can be given as to the existence of other IPRs not referenced in ETSI SR 000 314 (or the updates on the ETSI Web server) which are, or may be, or may become, essential to the present document.

Foreword

This Group Specification (GS) has been produced by ETSI Industry Specification (ISG) Measurement Ontology for IP traffic (MOI).

Introduction

A number of different models, information schemas and best-pratices have recently been proposed in order to cope with a lack of *de-facto* standards for interoperability and the extreme heterogeneity of tools' operational modes and repositories' internal organization of data that one can observe, as of now, in the Traffic Measurement and Analysis domain. To accomplish any standardization goal in the TMA field, a detailed analysis of such proposals, and of existing information models for IP traffic measurement is mandatory. A further step is then needed in order to unify the existing models into a set of well-defined ontological models, which will fully describe the domain of Internet traffic measurements and will tackle the most problematic aspects such as legally-compliant privacy protection and support for widely accepted QoS/QoE parameters.

1 Scope

The present document constitutes an analysis of information models for IP traffic measurement. This will include the basic definitions and state-of-the-art study, as well as the main guidelines to specify a complete set of vocabulary of classes and relations to describe Internet measurements, supporting QoS parameters and offering privacy protection, by studying existing schemas that are currently used to describe such information.

The present document is to give an initial focus and guide the process of the MOI ISG. The focus is on the key QoS parameters and the key approaches in privacy protection when manipulating, analysing and distributing IP traffic measurements.

2 References

References are either specific (identified by date of publication and/or edition number or version number) or non-specific. For specific references, only the cited version applies. For non-specific references, the latest version of the reference document (including any amendments) applies.

Referenced documents which are not found to be publicly available in the expected location might be found at http://docbox.etsi.org/Reference.

NOTE: While any hyperlinks included in this clause were valid at the time of publication ETSI cannot guarantee their long term validity.

2.1 Normative references

The following referenced documents are necessary for the application of the present document.

Not applicable

2.2 Informative references

The following referenced documents are not necessary for the application of the present document but they assist the user with regard to a particular subject area.

[i.1] "IP Flow Information Export (ipfix)".

NOTE: See http://www.ietf.org/dyn/wg/charter/ipfix-charter.html.

[i.2] "IP Flow Anonymisation Support", IETF Internet Draft , November 2009, E. Boschi and B. Trammel.

NOTE: See http://tools.ietf.org/html/draft-ietf-ipfix-anon-03.

[i.3] "Packet Sampling (psamp)" (concluded WG).

NOTE: See http://www.ietf.org/dyn/wg/charter/psamp-charter.html.

[i.4] "Benchmarking Methodology (bmwg)".

NOTE: See http://www.ietf.org/dyn/wg/charter/bmwg-charter.html.

[i.5] "Performance Metrics for Other Layers (pmol)".

NOTE: See http://www.ietf.org/dyn/wg/charter/pmol-charter.html.

[i.6] "Common Control and Measurement Plane (ccamp)".

NOTE: See http://www.ietf.org/dyn/wg/charter/ccamp-charter.html.

[i.7] "IP Performance Metrics (ippm)".

NOTE: See http://www.ietf.org/dyn/wg/charter/ippm-charter.html.

[i.8] "Relax-NG".

NOTE: See http://en.wikipedia.org/wiki/RELAX_NG.

[i.9] Revision 387.

NOTE: See http://anonsvn.internet2.edu/svn/nmwg/.

[i.10] "Internet Measurement Data Catalog".

NOTE: See http://www.datcat.org/.

[i.11] "Traffic Measurements and Models in Multi-Service Networks (2007 - 2009). Winner of the Celtic gold award 2009".

NOTE: See http://projects.celtic-initiative.org/tramms/.

[i.12] "What is perfSONAR?".

NOTE: See http://www.perfsonar.net/.

[i.13] P. Ohm, D. Sicker, and D. Grunwald: "Legal issues surrounding monitoring during network research", in Proceedings of the 7th ACM SIGCOMM Conference on Internet Measurement (IMC "07), San Diego, USA, October 24 - 26, 2007, pp. 141 - 148.

[i.14] M. Barbaro and T. Zeller Jr.: "A face is exposed for AOL searcher No. 4417749", The New York Times, August 9, 2006.

[i.15] G. D. Bissias, M. Liberatore, D. Jensen, and B. N. Levine: "Privacy vulnerabilities in encrypted HTTP streams", in Proceedings of the 5th Workshop on Privacy Enhancing Technologies (PET 2005), Cavtat, Croatia, May 30 - June 1, 2005.

[i.16] S. Bellovin: "A technique for counting NATted hosts", in Proceedings of the 2nd ACM SIGCOMM Workshop on Internet Measurement (IMW '02), Marseille, France, 6 - 8 November 2002, pp. 267 - 272.

[i.17] D. Koukis, S. Antonatos, D. Antoniades, P. Trimintzios, and E.P. Markatos: "A generic anonymization framework for network traffic", in Proceedings of the 2006 IEEE International Conference on Communications (IEEE ICC 2006), Istanbul, Turkey, June 11 - 15, 2006.

[i.18] R. Pang, M. Allman, V. Paxson, and J. Lee: "The devil and packet trace anonymization", ACM SIGCOMM Computer Communication Review, Vol. 36, No. 1, pp. 29 - 38, January 2006.

[i.19] Y. Lindell and B. Pinkas: "Privacy preserving data mining. In Advances in Cryptology - CRYPTO '00", volume 1880 of Lecture Notes in Computer Science, pages 36--54. Springer-Verlag, 2000.

[i.20] Privacy Preserving Data Mining Bibliography.

NOTE: See http://www.cs.umbc.edu/~kunliu1/research/privacy_review.html.

[i.21] References to Privacy-Preserving Data Mining Literature.

NOTE: See http://privacy.cs.cmu.edu/dataprivacy/papers/ppdm/.

[i.22] Privacy Preserving Data Mining Publications.

NOTE: See http://www.cs.ualberta.ca/%7Eoliveira/psdm/pub-by-year.html.

[i.23] R. Canetti, Y. Ishai, R. Kumar, M. K. Reiter, R. Rubinfeld, and R. N. Wright: "Selective private function evaluation with applications to private statistics", Proc. of the 20th ACM Symposium on Principles of Distributed Computing (PODC), 2001.

[i.24] Matthew Roughan and Yin Zhang: "Secure distributed data-mining and its application to large-scale network measurements", ACM SIGCOMM Computer Communication Review, Volume 36, Issue 1 (January 2006).

[i.25] Mitra, P., Pan, C., Liu, P., and Atluri, V. 2006: "Privacy-preserving semantic interoperation and access control of heterogeneous databases", in Proceedings of the 2006 ACM Symposium on information, Computer and Communications Security (Taipei, Taiwan, March 21 - 24, 2006). ASIACCS '06. ACM, New York, NY, 66-77.

[i.26] T. Finin, A. Joshi, L. Kagal, J. Niu, R. Sandhu, W. Winsborough and B. Thuraisingham: "ROWLBAC: representing Role Based Access Control in OWL", in Proceedings of the 13th ACM Symposium on Access Control Models and Technologies (SACMAT'08), Estes Park, CO, USA, June 11 - 13, 2008.

[i.27] A. Noorollahi Ravari, M. Amini, R. Jalili: "A Semantic Aware Access Control Model with Real Time Constraints on History of Accesses", in Proceedings of the 3rd International Workshop on Secure Information Systems (SIS'08), Wisla, Poland, 20 - 22 October 2008.

[i.28] G. V. Lioudakis, E. A. Koutsoloukas, N. Dellas, G. M. Kapitsaki, D. I. Kaklamani, I. S. Venieris: "A Semantic Framework for Privacy-Aware Access Control", in Proceedings of the 3rd International Workshop on Secure Information Systems (SIS'08), Wisla, Poland, 20 - 22 October 2008.

[i.29] Organization for the Advancement of Structured Information Standards (OASIS): "OASIS eXtensible Access Control Markup Language (XACML) TC", 2004.

NOTE: See http://www.oasis-open.org/committees/xacml/.

[i.30] T. Moses: "OASIS Privacy Policy Profile of XACML v2.0", OASIS Standard, February 2005.

[i.31] FP7 ICT project PRISM (PRIvacy-aware Secure Monitoring).

NOTE: See http://fp7-prism.eu/.

[i.32] FP7 ICT project MOMENT (Monitoring and Measurement in the Next Generation Technologies).

NOTE: See http://fp7-moment.eu/.

[i.33] A. Salvador, J. E. López de Vergara, G. Tropea, N. Blefari-Melazzi, Á. Ferreiro, Á. Katsu: "A Semantically Distributed Approach to Map IP Traffic Measurements to a Standardized Ontology", International Journal of Computer Networks & Communications (IJCNC), vol. 2, Issue 1, pp 13-31, January 2010.

[i.34] ITU-T Recommendation X.509: "Information technology - Open Systems Interconnection - The Directory: Public-key and Attribute Certificate Frameworks", August 2005.

[i.35] G. V. Lioudakis, F. Gogoulos, A. Antonakopoulou, D. I. Kaklamani, I. S. Venieris: "An Access Control Approach for Privacy-Preserving Passive Network Monitoring" in Proceedings of the 4th International Conference for Internet Technology and Secured Transactions (ICITST-2009), London, UK, November 9 - 12, 2009.

[i.36] IETF RFC 5101: "Specification of the IP Flow Information Export (IPFIX) Protocol for the Exchange of IP Traffic Flow Information", B. Claise, Ed.

[i.37] M. Casassa Mont: "Dealing with Privacy Obligations: Important Aspects and Technical Approaches," in Proceedings of the International Workshop on Trust and Privacy in Digital Business (TrustBus 2004), Zaragoza, Spain, August 30-September 3, 2004.

[i.38] Project, Strohmeier, F., et al: "D03 - MOME Final Project Report", IST MoMe Project, 2006.

NOTE: See http://www.ist-mome.org/deliverables/mome-wp0-0603-d03-update_final_report.pdf.

[i.39] LOBSTER IST Project.

NOTE: See http://www.ist-lobster.org

[i.40] D. Antoniades, M. Polychronakis, A. Papadogiannakis, P. Trimintzios, S. Ubik, V. Smotlacha, A. Øslebø and E. P. Markatos. LOBSTER: "A European Platform for Passive Network Traffic Monitoring".

NOTE: In Proceedings of the 4th International Conference on Testbeds and Research Infrastructures for the Development of Networks & Communities (TRIDENTCOM), March 2008, Innsbruck, Austria.

[i.41] D. Antoniades, M. Polychronakis, S. Antonatos, E. P. Markatos, S. Ubik, and A. Øslebø. Appmon: "An Application for Accurate per Application Network Traffic Characterization".

NOTE: In Proceedings of the IST Broadband Europe 2006 Conference, December 2006, Geneva, Switzerland.

[i.42] DIOR Project.

NOTE: See http://arantxa.ii.uam.es/~networking/projects/DIOR/index.htm. Also in José L. García-Dorado, José Alberto Hernández, Javier Aracil, Jorge E. López de Vergara, Francisco Montserrat, Esther Robles and Tomás de Miguel, "On the Duration and Spatial Characteristics of Internet Traffic Measurement Experiments", IEEE Communications Magazine, vol. 46, issue 11, November 2008.

[i.43] RIPE Document Store.

NOTE: See http://www.ripe.net/ripe/docs.

[i.44] "RIPE Routing Working Group Recommendations on Route-flap Dam".

NOTE: See http://www.ripe.net/ripe/docs/routeflap-damping.html.

[i.45] ETOMIC project.

NOTE: See www.etomic.org.

[i.46] P. Mátray, I. Csabai, P. Hága, J. Stéger, L. Dobos, G. Vattay: "Building a Prototype for Network Measurement Virtual Observatory" Proceedings of ACM SIGMETRICS - MineNet 2007, 12 June 2007, San Diego, CA, USA (2007).

[i.47] DIMES project.

NOTE: See www.netDimes.org.

[i.48] MINER, Salzburg Research: "MINER - Measurement Infrastructure for network Research", 2008.

NOTE: See http://miner.salzburgresearch.at.

[i.49] Quality of experience.

NOTE: See http://en.wikipedia.org/w/index.php?title=Quality-of-experience

[i.50] ETSI ETR 003: "Network Aspects (NA); General aspects of Quality of Service (QoS) and Network Performance (NP)".

[i.51] ITU-T Recommendation E.800: "Terms and Definitions Related to Quality of Service and Network Performance Including Dependability".

[i.52] ITU-T Recommendation X.641: "Quality of Service: Framework", Geneva, Switzerland, December 1997.

[i.53] ITU-R Recommendation BT.500-11: "Methodology for the subjective assessment of the quality of television pictures".

[i.54] Rubino, G. and Varela M.: "A new approach for the prediction of end-to-end performance of multimedia streams", First International Conference on the Quantitative Evaluation of Systems, 2004. QEST 2004. Proceedings. September 2004.

[i.55] Rodríguez-Bocca, P., Cancela, H., and Rubino, G. 2007. Video quality assurance in multi-source streaming techniques. In Proceedings of the 4th international IFIP/ACM Latin American Conference on Networking (San José, Costa Rica, October 10 - 11, 2007). LANC '07. ACM, New York, NY, 83-93.

[i.56] Alfonso Sánchez-Macián, Jorge E. López de Vergara, Encarna Pastor, Luis Bellido: "A System for Monitoring, Assessing and Certifying Quality of Service in Telematic Services". Knowledge-Based Systems, Vol. 21, Issue 2, March 2008, Elsevier, ISSN 0950-7051.

[i.57] Alfonso Sánchez-Macián, David López Berzosa, Jorge E López de Vergara, Encarna Pastor Martín: "A Framework for the Automatic Calculation of Quality of Experience in Telematic Services", Proceedings of the 13th HP-OVUA Workshop, Côte d'Azur, France, 21-24 May 2006. ISBN 3000187804.

[i.58] EFIPSANS project.

NOTE: See http://www.efipsans.org/.

[i.59] R. Chaparadza: "Requirements for a Generic Autonomic Network Architecture Suitable Requirements for Autonomic Behavior Specifications of Decision-Making-Elements for Diverse Networking Environments", International Engineering Consortium (IEC) Annual Review in Communications, vol. 61, December 2008.

[i.60] R. Natale: "Converting SNMP MIBs to SOA/Web Services Management Artifact; draft-natale-snmp-mibs-to-ontology-00", IETF Network Working Group Internet-Draft, August 2007.

NOTE: See http://tools.ietf.org/html/draft-natale-snmp-mibs-to-ontology-00.

[i.61] Kun Liu, Hillol Kargupta, Jessica Ryan: "Random Projection-Based Multiplicative Data Perturbation for Privacy Preserving Distributed Data Mining", IEEE Transactions on Knowledge and Data Engineering, vol. 18, no. 1, pp. 92-106, Jan., 2006.

[i.62] S. Agrawal and J.R. Haritsa: "A Framework for High-Accuracy Privacy-Preserving Mining", Proc. 21st Int'l Conf. Data Eng. (ICDE'05), pp. 193-204, Apr. 2005.

[i.63] IETF RFC 3577: "Introduction to the Remote Monitoring (RMON) Family of MIB Modules".

[i.64] IETF RFC 2819: "Remote Network Monitoring Management Information Base".

[i.65] IETF RFC 1513: "Token Ring Extensions to the Remote Network Monitoring MIB".

[i.66] IETF RFC 2613: "Remote Network Monitoring MIB Extensions for Switched Networks Version 1.0".

[i.67] IETF RFC 3144: "Remote Monitoring MIB Extensions for Interface Parameters Monitoring".

[i.68] IETF RFC 3273: "Remote Network Monitoring Management Information Base for High Capacity Networks".

[i.69] IETF RFC 3434: "Remote Monitoring MIB Extensions for High Capacity Alarms".

[i.70] IETF RFC 2021: "Remote Network Monitoring Management Information Base Version 2 using SMIv2".

[i.71] IETF RFC 2895: "Remote Network Monitoring MIB Protocol Identifier Reference".

[i.72] IETF RFC 3395: "Remote Network Monitoring MIB Protocol Identifier Reference Extensions".

[i.73] IETF RFC 2896: "Remote Network Monitoring MIB Protocol Identifier Macros".

[i.74] IETF RFC 3287: "Remote Monitoring MIB Extensions for Differentiated Services".

[i.75] IETF RFC 3729: "Application Performance Measurement MIB".

[i.76] IETF RFC 4150: "Transport Performance Metrics MIB".

[i.77] IETF RFC 4711: "Real-time Application Quality-of-Service Monitoring (RAQMON) MIB".

[i.78] IETF RFC 4149: "Definition of Managed Objects for Synthetic Sources for Performance Monitoring Algorithms".

[i.79] ITU-T Recommendation Y.1540: "Internet protocol data communication service - IP packet transfer and availability performance parameters".

3 Abbreviations

For the purposes of the present document, the following abbreviations apply:

CCAMP	Common Control and Measurement Plane
IMDC	Internet Measurement Data Catalog
IPFIX	IP Flow Information export
IPPM	IP performance metrics
IS-IS	Intermediate System to Intermediate System
MIB	Management Information Base
NIC	Network Information Centre
NMWG	Network Measurement Working Group
NOC	Network Operations Centre
OSPF	Open Shortest Path First
OWL	W3C Web Ontology Language
PMOL	Performance Metrics for Other Layers
PQoS	Perceived Quality of Service
QoS/QoE	Quality of Service/Quality of Experience
RELAX NG	Regular Language for XML Next Generation
RRD	Round-Robin Database
SOA	Service Oriented Architecture
TMA	Traffic Monitoring and Analysis

4 Working Groups and Metrics for Network Measurements

We start by reviewing the most common external data representation efforts from IETF. Then this clause includes an outline of the Open Grid Forum work on measurement data standardization, and continues with PerfSonar, a mediator for monitoring services developed by GEANT and other partners. A detailed overview of the MIBS RMON working-group activities and of the DatCat meta-data repository is finally given.

4.1 Network Data Representation Models from IETF

There are a few attempts to standardise network measurement protocols in order to provide a common understanding of measures across different networks and organisations, promoted by IETF. This clause provides a short list of a few of them, their characteristics, progress, future and implementation, if any.

4.1.1 IPFIX (IP Flow Information Export Charter)

The IETF IPFIX [i.1] working group defined a protocol to transmit information about captured flows. It has specified both the Information Model (to describe IP flows) and the IPFIX protocol to transfer IP flow data from IPFIX exporters to collectors, which is used to transmit to a collector the captured information flows. It considers a flow as a group of packets sent from the same source to the same destination through the same protocol.

The work towards standardisation is quite advanced and several drafts have been published as RFCs so far, and some others are about to be reviewed. However, the group activities are still ongoing and new issues are raised. Ongoing activities are now focusing on anonymization, with an initial draft on this specific issue [i.2] made available since November 2009.

4.1.2 PSAMP (Packet Sampling)

PSAMP [i.3] is an IETF group whose target is to define a standard set of capabilities for network elements to sample subsets of packets by statistical and other methods. They should be simple as they are supposed to work at maximal line rate ubiquitously.

The standard will specify a set of selection operations by which packets are sampled, will specify the information that is to be made available for reporting on sampled packets, describe protocols by which information on sampled packets is reported to applications; describe protocols by which packet selection and reporting can be configured. Unreliable transport is permitted to allow ubiquitous deployment.

The standard will specify:

- Selectors for packet solving.
- Packet information available for reporting.
- Sampled packet reports format.
- Report Streams format for a stream of packet reports.
- Multiple Report Streams requirements for parallel packet samplers in one network element.
- Configuration and Management packet format.
- Presentation, Export and Transport of Packet Reports.

4.1.3 BMWG (Benchmarking Methodology)

The BMWG [i.4] group tries to make a series of recommendations concerning the measurement of the performance characteristics of various internetworking technologies focusing on:

- The systems or services that are built from these technologies, describing the class of equipment, system, or service being addressed.
- The performance characteristics that are pertinent to that class.
- The set of metrics that aid in the description of those characteristics.
- The methodologies required to collect said metrics.
- The requirements for the common, unambiguous reporting of benchmarking results.

These standards will be limited to technology characterization using simulated stimuli in a laboratory environment, and will not be prepared for live, operational networks. The most interesting goal of the WG is to produce benchmarks which strive to be vendor independent and have universal applicability to a given technology class, but not to deal with acceptance criteria or performance requirements. Works are well developed, with most of the methodologies taken to AD Review and around 30 RFC have been published so far.

4.1.4 PMOL (Performance Metrics for Other Layers)

This group has finished its activities in November 2008, with the publication of two Internet Drafts, available at [i.5]. The first one is a framework and guidelines memo which describes any necessary elements of performance metrics of protocols and applications transported over IETF-specified protocols (such as the formal definition, purpose, and units of measure) and the various types of metrics that characterize traffic on live networks (such as metrics derived from other metrics, possibly on lower layers).

Most important for other standardization initiatives can be the enclosed guidelines for a performance metric development process that includes entry criteria for new proposals (how a proposal might be evaluated for possible endorsement by a protocol development working group), and how an successful proposal will be developed.

The second Internet Draft is a proof-of-concept, defining performance metrics for SIP, based on the previous draft-malas-performance-metrics. This memo serves as an example of the framework and the PMOL development process in the IETF.

4.1.5 CCAMP (Common Control and Measurement Plane)

The CCAMP [i.6] has been a very active and prolific WG, which focused on defining measurement methods on the physical path and core tunnelling technologies. Its results included the definition of protocol-independent metrics for describing links and paths required for routing, signalling, protocols and extensions to them required for link and path attribute measurement (LMP among others).

CCAMP WG has been working on the following most interesting tasks:

- Define how the properties of network resources gathered by a measurement protocol can be distributed in existing routing protocols, such as OSPF and IS-IS.
- Define abstract link and path properties needed for link and path protection.

Functional specification of extensions for routing and signalling required for path establishment is also under the scope of the WG.

Collaborations with other IETF groups have been very fluid. The standards are quite developed with around 60 RFCs published. This group seems to have closed its activities in June 2009.

4.1.6 IPPM (IP performance metrics)

The IPPM [i.7] WG is developing a set of quantitative unbiased standard metrics that can be applied to the quality, performance, and reliability of Internet data delivery services. These metrics will be designed such that they can be performed by network operators, end users, or independent testing groups. NOC/NIC services are not included.

Metrics to be standardised are connectivity, one-way delay and loss, round-trip delay and loss, delay variation, loss patterns, packet reordering, bulk transport capacity and link bandwidth capacity. Some of them have already been completed and published.

This working group intends to produce also a protocol to enable communication among test equipment that implements the one-way metrics. The intent is to create a protocol that provides a base level of functionality that will allow different manufacturer's equipment that implements the metrics according to a standard to interoperate.

The WG will also produce a MIB to retrieve the results of IPPM metrics, such as one-way delay and loss, to facilitate the communication of metrics to existing network management systems. Around twenty RFCs have been published on some of the metrics.

As a concluding remark of this clause, it shall be noticed that an Internet Draft about converting SNMP MIBs to corresponding ontology models has been proposed in 2007 by R. Natale [i.60], but the term "ontology" was loosely employed in the document to mean SOA management artifacts in general. The initiative failed to generate a specific working group, and was considered to be an extension of the "MIB-to-XML" work.

4.2 Open Grid Forum Network Measurement Working Group

This clause analyses Open Grid Forum Network Measurement Working Group (NMWG) work on sharing knowledge about measurement tools and metrics. The NMWG developed an infrastructure in order to communicate among different systems knowledge in relation to network measurement. Basically, the works of NMWG focus on a common vocabulary used to provide information about different measuring tools. Actually, a particular XML Schema is defined for each of these tools, in a particular language, RELAX NG (Regular Language for XML Next Generation) [i.8], and the information is sent in XML code defined by that schema.

NMWG has published a public SVN repository of their work, which can be accessed at [i.9]. There is one compact syntax RELAX NG file (.rnc) for each tool, and a few more with information used by all of them. There are also some examples of how to code in XML from a tool wrapper, which prepares XML messages to be sent following the XML schema of the tool.

The .rnc schema files can be categorised as:

- Description of tools.
- Description of output data.
- Description of topological features, which are included by other schemas and defined for several network levels.
- Description of errors and other system outcomes, included by other schemas.
- General structure for the messages, including different statistics, times and other measurements, and also interface utilisation description.
- Specific elements for Perfsonar implementation.

The advantage of Relax NG files is that they can be written in a compact syntax, more similar to Perl or Python, rather than XML files, and at the same time, these files can be automatically mapped to their XML Schemas. In addition, this language is widely used and there is plenty of information about it.

The disadvantage of this specification is that it does not include any semantic information which could be used in combination with inference rules to make decisions.

Unfortunately, it seems that NMWG have not met for long time, and its website was last time updated on August, 2007, so it is hard to know whether they keep improving the project or not. However they keep updating a subversion repository with small changes, around once a month.

4.3 MIBS RMON Working Group

The family of RMON MIBs is described in RFC 3577 [i.63]. This family starts with the RMON-1 MIB (RFC 2819 [i.64]), which provides statistics focused on the data-link level. It also provides alarms, as well as traffic capture capabilities. This MIB, defined for Ethernet networks, is completed with Token Ring extensions in RFC 1513 [i.65]. At the same time, SMON MIB (RFC 2613 [i.66]) and IfTopN MIB (RFC 3144 [i.67]) are defined for switched networks. Finally, two additional MIBs have been defined (RFC 3273 [i.68] and RFC 3434 [i.69]) for high capacity networks, where 32 bit counters are too small to measure the traffic.

RMON-2 MIB (RFC 2021 [i.70]) extends RMON-1 by providing an analysis of the monitored traffic in the network, transport and application levels. For this, it is necessary to define a set of protocol identifiers (RFC 2895 [i.71]), protocol operations identifiers (RFC 3395 [i.72]), as well as a macro to define new identifiers (RFC 2896 [i.73]). RMON-2 is complemented with DSMON MIB (RFC 3287 [i.74]), to monitor differentiated services traffic, and with APM MIB (RFC 3729 [i.75]), TPM MIB (RFC 4150 [i.76]) and RAQMON MIB (RFC 4711 [i.77]), to measure application performance at several levels and perspectives. This set of MIBs is completed with the SSPM MIB (RFC 4149 [i.78]), useful to generate traffic from a synthetic source as a way to perform active measurements.

All these MIBs have in common that they define a set of control tables to define the monitoring tasks, and another set of tables to read the information obtained as a result of those tasks.

4.4 Internet Measurement Data Catalog (DatCat)

DatCat [i.10], developed and run by CAIDA, is an Internet Measurement Data Catalog (IMDC), a searchable registry of information about network measurement datasets. It serves the global network research community by allowing anyone to find, annotate, and cite data contributed by others, and soon by allowing anyone to contribute new data.

The goals of DatCat IMDC are:

- To facilitate searching for and sharing of data among researchers. Finding data to use in network research has historically been difficult. By serving as a shared global resource where anyone can find the data needed for network analysis, DatCat mitigates a significant barrier to research.
- To enhance documentation of datasets via a public annotation system. Instead of relying on the data contributor alone to document the data, DatCat allows any researcher to annotate datasets with problems, features, or missing information they discover in the data, thereby increasing the utility of the datasets.

- To advance network science by promoting reproducible research. Reproducibility of results is a cornerstone of good science, but requires that the researcher's data is available to others. Similarly, to get the most meaningful comparison of analysis methodologies and algorithms, researchers must test them against the same data. By putting their data in DatCat or using data already in DatCat, and then citing the IMDC Handle in their published results, researchers can make it easier for others to obtain their data and validate their results or perform alternate analyses on the same data.

Note that IMDC does not store the data (or tools) itself, but only metadata, that is, descriptions of the data and instructions for obtaining it. The storage of the data itself remains in the hands of the contributor. As such, it may or may not be freely available; it might, for example, reside on a password-protected server, or require asking the owner of the data. IMDC does not dictate the terms of availability of the data, it just helps you with the first step of *finding* the data.

Information in IMDC is organized as Objects, each of which describes a real-world object or idea. For the purposes of finding and obtaining data, the most important types of objects are:

- Data. The core of IMDC is the Data object. A Data object describes a dataset in a single file in its most natural working form, even if the data is not made available directly in that form.

- Data Collection. A set of Data objects with a common purpose and/or collected as part of a single effort. When searching, it is often most convenient to search for Collections as a unit rather than searching through thousands of individual Data objects.

- Package. A Package object describes a collection of one or more data files, in a form that can be downloaded or otherwise made available. Package objects usually represent compressed archives of data files, but can be as simple as a single uncompressed data file, if that file is the downloadable form.

- Location. Location objects represent the method for obtaining a package. Often this will be a URL linked directly to the package (external to IMDC), but it can also be text instructions (e.g. for packages that require human approval or agreement to an AUP).

IMDC is designed to work with any browser that supports standard HTML. IMDC does not require graphics, cookies, JavaScript, or CSS, but it will take advantage of those features if available to make the interface more convenient, faster, and generally more pleasant, so we recommend using a browser with those features enabled. All infrastructural IMDC text is in ASCII English, although user-contributed text may contain other languages and character sets.

Interoperability with the DatCat repository is, as of now, achievable through the Import and Export key functionalities: for importing IMDC has an XML based meta-data format and a set of tools to help in creating the XML needed for submitting meta-data to the catalog. The tools are a programmatic Perl API and a declarative interface called 'subcat' based on the YAML data format (www.yaml.org). IMDC does not have an official data export format or method, but it has an undocumented prototype of a data export facility, that outputs in YAML format.

4.5 Projects related to Traffic Monitoring and Analysis

4.5.1 TRAMMS

The main objective of TRAMMS [i.11] was to model traffic in multi-service IP networks, and to develop tools for monitoring of QoS and bottlenecks in networks. The project has closed in December 2009 and its models are built upon data acquired in different parts of Europe and combined with new tools developed within the project, they aim at bringing significant new insight into network traffic, bottleneck analysis, user behaviour and QoS monitoring.

The main focus in TRAMMS has been to increase the knowledge of traffic patterns and possibilities for traffic management and QoS monitoring of IP networks. This has been realized through development of low cost tools for analyzing QoS parameters, available bandwidth on end-to-end links, routing events, as well as using available state of-the art deep packet inspection devices to analyse actual traffic in live access networks.

Parameters such as applications used, trends in application usage, penetration of applications, peak hours, peak rates, service specific user behaviour have been analysed, and typical user types have been defined. The influence on the user behaviour from different first mile technologies has been studied as well as the difference in user behaviour between different regions in Europe.

Measurements from the application to the packet level per household were collected in real networks located in different countries (Sweden and Spain) covering different types of access (FTTH, xDSL, CMTS, GGSN, university network). Measurements from a large amount of users were gathered for long periods of time (close to 3 000 TiB of traffic volume was analysed in Spanish and Swedish networks from 2007 to 2009 in periods ranging from several days to several years).

The project has promoted standardization of active end-to-end capacity measurement methods in the International Telecommunication Union (ITU). The main result so far is the acceptance and inclusion of the "IP-layer capacity framework" in ITU-T Recommendation Y.1540 [i.79].

4.5.2 EFIPSANS

The EFIPSANS-IP-Project [i.58] that started in January 2008, aims at exposing the features in IPv6 protocols that can be exploited or extended for the purposes of designing or building autonomic networks and services. A study of the emerging research areas that target desirable user behaviours, terminal behaviours, service mobility, e-mobility, context-aware communications, self-ware, autonomic communication/computing/networking will be carried out. Out of these areas desirable autonomic(self-*) behaviours in diverse environments e.g. end systems, access networks, wireless versus fixed network environments will be captured and specified. Appropriate IPv6 protocol or architectural extensions that enable the implementation of the captured desirable autonomic behaviours will be sought and specified.

The work conducted in EFIPSANS is based on the recently proposed Generic Autonomic Network Architecture (GANA) [i.59]. GANA sets the principles and guidelines that need to be followed according to EFIPSANS's vision of the Future Internet design. GANA is supporting context-awareness through combining information models with a set of ontologies. Since information and data models are not capable of representing the detailed semantics required to reason about behaviour, GANA augments the use of knowledge extracted from information and data models with ontologies. Information about the current state is collected according to the information model and - through the use of a domain-specific ontology - relationships are added between context data and the behaviours of the autonomic nodes. Consequently, the nodes are able to reason about the modelled information and interact in a way that accurately reflects the overall network behaviour.

The GANA ontology is based on the GANA meta-model that is the information model designed in EFIPSANS. The GANA meta-model, defines via a semantically precise meta-model the relationships among the identified elements. It deals with and formalises each and every aspect of GANA up to details that are needed in order to be able to analyse, verify, build and evaluate autonomic behaviours. The GANA ontology is designed in order to enrich this model with semantics, enable information gathered from the network to be analyzed and ensure that the model accurately reflects the current operational status.

The GANA Ontology adds semantics to the autonomic elements that are defined in the GANA architecture. It describes the autonomic entities and the relationships among them, through the definition of a number of basic classes and properties. Part of the ontology is built upon the MOMENT one [i.32], and extends it by adding descriptions for the supported services in an autonomic network, the characteristics and the available interfaces of the GANA defined elements and the possible interactions among them.

The following basic classes are defined (in alphabetical order) in order to describe general concepts that are present in an autonomic network: *AutonomicBehaviour, AutonomicNode, Capabilities, CommunicationInterface, ControlLoop, Element, Event, GanaPlane, Goal, Mechanism, MonitoringData, Policy, Profile* and *Protocol.*

4.5.3 PerfSonar

PerfSonar [i.12] is a service for network performance monitoring, making it easier to solve end-to-end performance problems on paths crossing several networks. These services act as an intermediate layer, between the performance measurement tools and the diagnostic or visualization applications. This layer is aimed at making and exchanging performance measurements between networks, using well-defined protocols.

PerfSonar is a services-oriented architecture. That means that the set of elementary functions have been isolated and can be provided by different entities called services. All those services communicate with each other using well-defined protocols. These services are implemented using actual tools for building web services.

PerfSonar's main services are:

- Measurement Point Service: Creates and/or publishes monitoring information related to active and passive measurements.

- Measurement Archive Service: Stores and publishes monitoring information retrieved from Measurement Point Services.
- Lookup Service: Registers all participating services and their capabilities.
- Authentication Service: Manages domain-level access to services via tokens.
- Transformation Service: Offers custom data manipulation of existing archived measurements.
- Resource Protector Service: Manages granular details regarding system resource consumption.
- Topology Service: Offers topological information of networks

5 Network Monitoring and Personal Data Protection

Network monitoring not only may lead to privacy violations but it is also surrounded by legal implications [i.13]. Privacy-sensitive information is not at all limited to the payload of the network packets, i.e. the content of the monitored communications. In fact, this case could be even considered as a trivial one from a privacy protection point of view, since the confidentiality of the content can be adequately guaranteed by using strong end-to-end encryption. Personal data can be extracted from the various protocols' headers (e.g. a visited web-site or the peers of a VoIP call), from data that are supposed to be anonymized (e.g. [i.14]), from statistical analysis of network traffic [i.15], or even from "unsuspicious" header fields, such as the IP ID alone [i.16].

Several mature approaches have been proposed for addressing these issues, mostly concerning the anonymization of the traffic. The AAPI [i.17] offers a generic and flexible anonymization framework that provides extended functionality, covering multiple aspects of anonymization needs. Recently a Java wrapper has been made available to facilitate the integration of the AAPI primitives in semantic-oriented environments, which employ the OWL APIs. Although frameworks such as the ones described in [i.17] and [i.18] are aimed to be quite generic, a significant drawback is that they base on quite "static" anonymization policies specification; in all cases, "someone" must define in an explicit manner the policies that will regulate the execution of the underlying anonymization APIs. Additionally, although they are quite practical and effective for applications referred to as offline, such as related to internet research based on packet traces, they are not applicable to applications' domains that are referred to as online, such as intrusion detection.

On the other hand, only a limited amount of work has been specifically targeted to design privacy-preserving operations on data gathered from monitoring and measurement systems, and specifically privacy-preserving operations related to access control and elaboration and data extraction.

Regarding data extraction, a large amount of work has been carried out in the field of Privacy Preserving Data Mining (PPDM) [i.19], [i.20], [i.21], [i.22], but most of this work has been discussed in general terms and has not been specifically addressed to the area of network monitoring. Some proposed solutions may provide a good starting points for being adapted to the specific TMA domain. Guidelines to accomplish computation of statistics on data such that the client will only learn the desired statistics, but not the values of the data nor partial computations, have been provided in [i.23] and an application of these techniques to large-scale network measurements has been recently proposed in [i.24]. Techniques devised to modify original data values in order to have a new version of the database, which can be safely released to the public, such as additive perturbation, replacement with meaningless symbols, aggregation to a coarser granularity, or even sampling, have been proposed [i.61], [i.62]. The crucial point is then to accurately reconstruct the aggregate distributions and to easily perform data mining, as data perturbation usually results in a degradation of the database performance.

Regarding access control, the privacy-preserving semantic interoperation and access control of heterogeneous databases (PACT) toolkit [i.25] addresses the problem of syntactic and semantic heterogeneity of access control policies among different systems. PACT uses encrypted ontologies, encrypted ontology-mapping tables and conversion functions, encrypted role hierarchies and encrypted queries. The relevance of this work is its capability to still provide acceptable performance. Other approaches being grounded on a semantic basis have been specified in [i.26] and [i.27], as well as in a work aiming at the enforcement of legislation-aware access control (e.g. [i.28]).

An important contribution is given by those XML extensions specifically devised for access control and privacy protection. These include the OASIS eXtensible Access Control Markup Language (XACML) [i.29], a general-purpose access control standard written in XML, and its Privacy Policy Profile specified in XACML v2.0 [i.30].

The FP7 projects PRISM [i.31] and MOMENT [i.32] are presently dealing with unified, privacy-aware access to network data, proposing two different, albeit complementary, approaches. They will be detailed in the following clauses.

5.1 MOMENT Project

Under the umbrella of the MOMENT project, an ontology comprising all aspects of the IP measurement domain has been developed: it includes a Data ontology as well as a Metadata ontology, an Upper ontology and an Anonymization ontology. The Data ontology describes the hierarchy of the different kinds of measurements and their relations to the physical components of the network; the Metadata ontology describes all the information about how the measurements are stored, how and when they were measured, etc. All concepts common to those ontologies, such as time, units and context, are described in an Upper Ontology and finally, the Anonymization ontology describes dependencies between the possible anonymization strategies that data have to undergo, prior to being released to the user requesting them, and the role and purpose of such a user within the community of people interested in network measurements.

This design allows for information to be placed at different abstraction levels, including the definition of specific class of measurements that are derived from generic ones and introducing the concept of meta-information (so that users can also request a view of what the system knows and what they are allowed to ask).

Several iterations were needed to achieve a generic and powerful enough model, which is able to accommodate for all the schemas contained inside all data repositories that are connected to the MOMENT mediator. First iterations on the ontology were designed based on a strict hierarchy of network measurements categories, onto which many data sources failed in smoothly mapping their data. This approach has changed in later revisions of the semantic model, see [i.33], and the final Data ontology gives much more importance to the details of the information the measurement carries within itself, rather than trying to assign the Measurement class to one of a set of predefined categories under a fixed hierarchy. Specifically information carried by the measurement is modelled through the hasMeasurementData property and the instances of MeasurementData subclasses. There is a subclass of MeasurementData for every possible measurement value. Other high-level concepts such as Route, Capacity, etc., which cannot be determined with single values, are represented with the Metric class.

Moreover, since its definition, the MOMENT project has also been concerned about obfuscation of certain fields of the data, which are passed to the end-user, in order to enforce a layer of anonymization for protection of the data originator. The PolicyObject is the cornerstone of the Anonymization ontology. It can be viewed, in OWL terms, as an N-ary relation that associates together a number of UserRoles and a number of UsagePurposes, applied to a number of PrivacyScopes. The PolicyObject specifies a well-defined AnonymizationStrategy and an associated AcceptableUsePolicy. The AnonymizationStrategy is composed of a group of AnonymizationTargets and an AnonymizationBackend to support and implement that strategy, i.e. the specific external tool that will ultimately be invoked to do the real anonymization job.

Two innovative ideas are applied to the Anonymization ontology: the Network Data Age and the Acceptable Use Policy. The first technique is employed to capture the common concept that, when a Measurement was generated a long ago, it usually becomes less sensitive, so that a looser anonymization scheme could be enforced. The Acceptable Use Policy simply represents an informative document, although structured, about what the provider expects from the user regarding the usage of the data that the provider itself is willing to release. This approach can be regarded as a kind of End User Legal Agreement, but in the field of network measurement.

5.2 PRISM Project

The PRISM project proposes a two-tier architecture which mediates between the source of information (i.e. the communication channel) and the entities that consume data originating from network monitoring and enforces an access control model specifically designed for this context. For the specification of the access control policies, it relies on a semantic model that enables the dynamic, real-time evaluation of the provisions, depending on the particular characteristics of each request for information.

The PRISM ontology has been developed in order to become the semantic implementation of a privacy-aware access control and authorization model, specifically devised for the protection of network monitoring data. The PRISM access control mechanism constitutes a two stage approach, where the underlying policies are described by means of semantically defined X.509 Attribute Certificates [i.34], as far as their static part is concerned, while the dynamic "privacy context" is evaluated in real-time, by means of direct ontological reasoning. For a detailed description of the PRISM access control and authorization approach, the reader is referred to [i.35].

Essentially, the PRISM ontology defines access control rules as associations of data types, monitoring purposes and operational roles. In that respect, the instances of the Rules class implement the access control provisions, by connecting instances of the PersonalData, Purposes and Roles classes, respectively. Each of the latter three classes is characterized by a number of OWL object properties, defining different directed graphs over the same sets of instances.

The instances defined as members of these three classes have been selected after the elaboration of data models, as well as the actual needs identified in the context of the project and described as requirements. In that respect, the PersonalData class contains types coming from the network (IPv4 and IPv6) and transport layers and application-specific data, as well as data types of the IPFIX protocol [i.36], which is extensively used by the PRISM project. The Purposes class includes a variety of monitoring purposes coming from heterogeneous domains, while the Roles class reflects a simplified, yet realistic, operational structure within a network operator.

Each rule specifies three aspects of data management: read access rights, write access rights and data retention. When the rule is subject to conditional provisions, the appliesUnderCondition property links the rule with some instance of the Conditions class, which specifies spatial, temporal and history-based conditions for the enforcement of the rule in question. Moreover, each rule defines possible complementary actions that should be executed along with the rule's enforcement, frequently referred to in the literature as "privacy obligations" [i.37]. Finally, each rule is characterized by certain meta-rules, reflecting concepts such as whether the rule is inherited by the descendants of the personal data, purpose and role types.

The PRISM ontology introduces also the concept of exclusive combinations of data. In this context, the semantic definitions of different data types may be members of relations that are defined as ExclusiveCombinations instances and impose restrictions on the disclosure of some data types depending on prior disclosure of other types.

For the specification of anonymization norms, the PRISM ontology incorporates the DataTransformations class, the instances of which specify the processing actions that lead from a set of data types to another. Each transformation is linked to an instance of the Components class containing the "semantic signatures" of PRISM processing components, being either proprietary Java modules or wrapped functionalities offered by the AAPI [i.17]. It should be noted here that what this class' instances specify is the capabilities offered by the underlying PRISM systems. The actual transformations that take place comprising the anonymization strategy, are determined in real-time based on the ontological reasoning and take the form of a workflow specification.

6 Specific Information Models of Existing Infrastructures

This is a specific clause to show the data schemas of some important, large-scale IP traffic measurements database that exist as of today and continuously probe the Internet. In this clause we provide a description of data representation schema of the following measurement infrastructures that operate on the Internet:

- MOME Meta data repository
- LOBSTER
- DIOR
- RIPE
- ETOMIC at UPNA
- ETOMIC at ELTE
- DIMES
- MINER

They will provide a basis to understand the key concepts in the domain of IP traffic measurements.

6.1 Metadata database (MOME)

The design of the MOME database [i.38] has been made to enable the addition of almost any meta-information available, broken down into 6 main measurement data types (PacketTrace, FlowTrace, Routing, HTTP, QoS, application-level trace) plus a very generic data type of a "web repository", linking to any web page providing measurement data. This led to a design where many of the database fields are simple strings using the "text" or "varchar" data type. Providing such data in a service-oriented manner would be useful, in order to be able to access the data in a machine-readable way. This would require putting additional context to identify the single data fields and further specialization of them. As example the MOME database does not strictly define a "Collector Location". This can be GPS coordinates, the address, the institution, the room, etc. The advantage is, that the user can decide to provide what meta-information he has available (or is allowed to provide to respect privacy issues) to his measurements. The drawback is obvious, that this kind of meta-data is rather worthless for a machine, as this information is only interpretable by human readers.

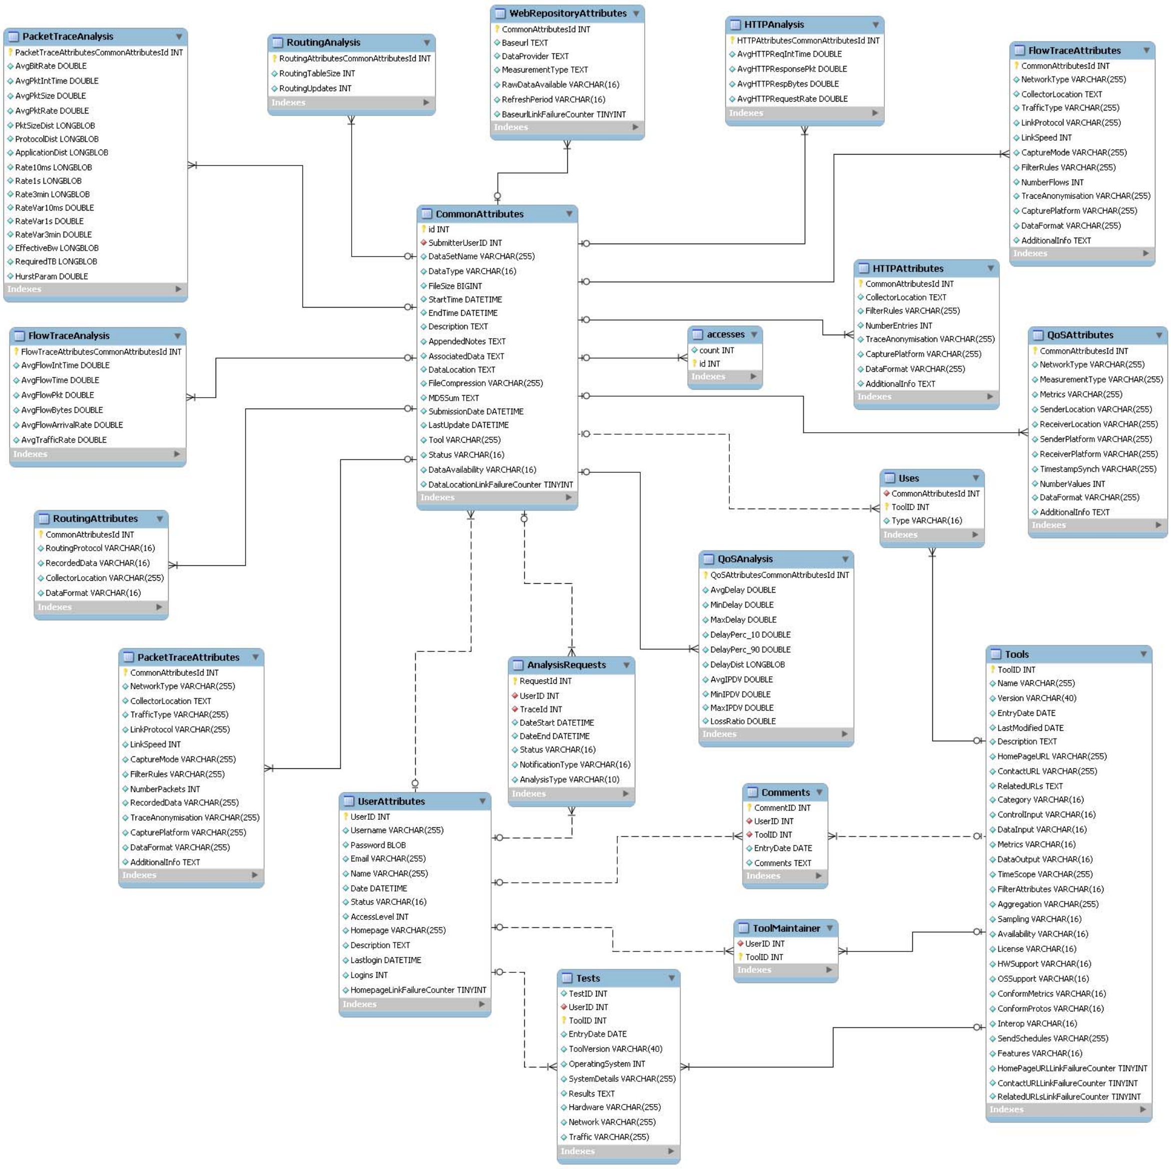

Figure 1: Schema used in MOME metadata repository

6.2 Aggregated passive trace database (LOBSTER)

The LOBSTER project [i.39] has built an advanced pilot European Internet traffic monitoring infrastructure based on passive network monitoring sensors. LOBSTER has also developed novel performance and security monitoring applications, which have been enabled by the availability of the passive network monitoring infrastructure, and has realized the appropriate data anonymization tools for prohibiting unauthorized access or tampering of the original traffic data.

The passive monitoring applications running on the sensors have been developed on top of MAPI (Monitoring Application Programming Interface) [i.40], [i.41], an expressive API in C for building network monitoring applications, which has been developed in the context of the SCAMPI and LOBSTER projects. The LOBSTER sensors monitor the network traffic using different measurement applications, such as traffic categorization, packet loss measurement, and intrusion detection. Depending on the type of measurement or traffic processing, the gathered data or the computed results are stored in files, a database, or Round-Robin Database (RRD) archives, while usually they can be viewed through a Web interface.

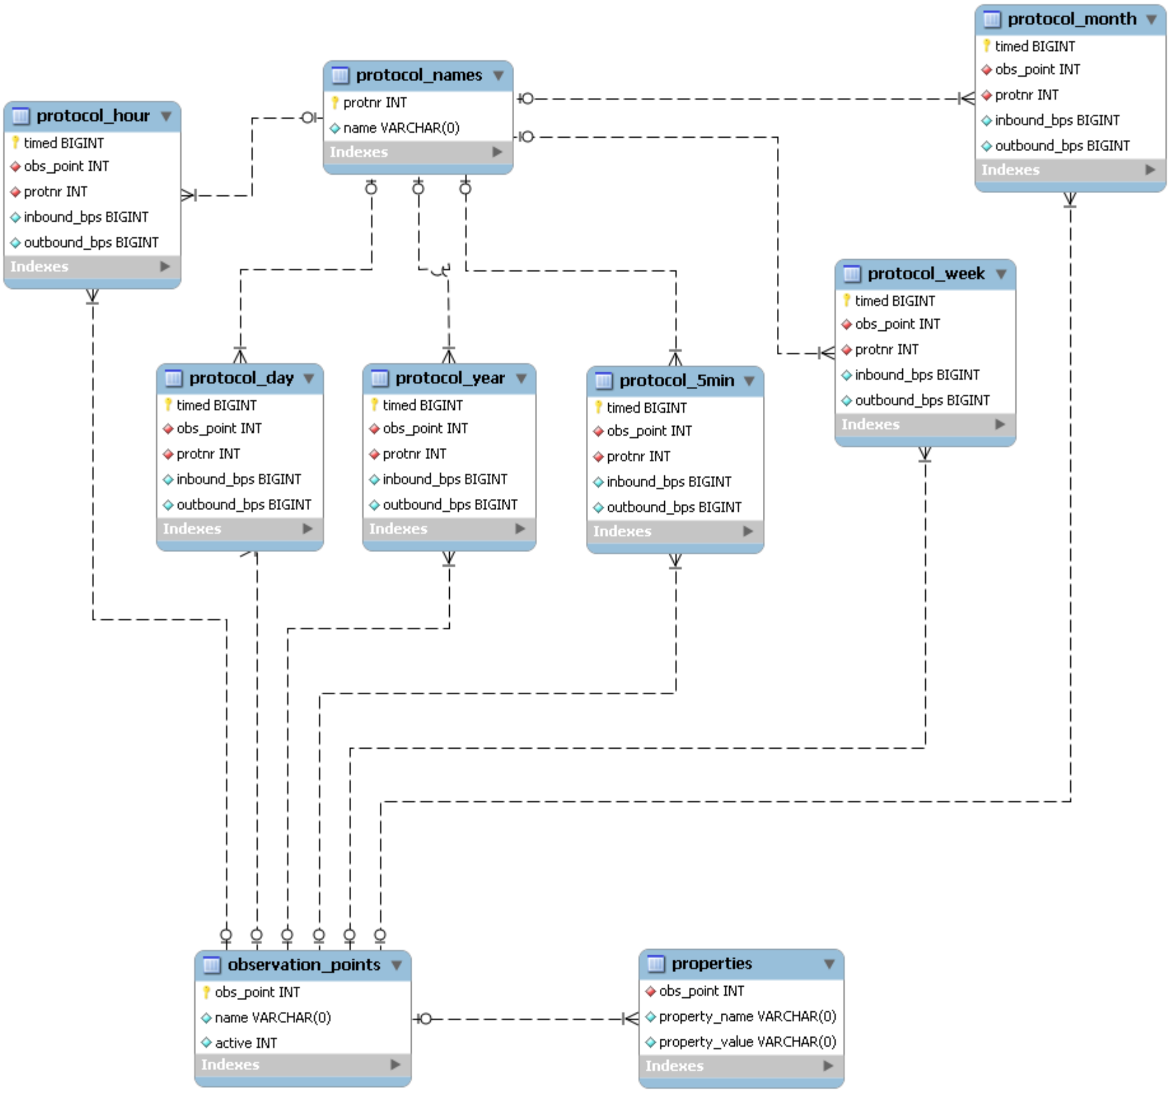

Figure 2: Schema used to store passive measurement data in LOBSTER

6.3 Aggregated passive trace database (DIOR)

The Spanish NREN serves more than 260 institutions, mainly universities and research centres, and comprises 18 Points of Presence across the country [i.42]. First, the *Flow-Tools* software package was used for data collection at the repository. Then, a number of statistics were obtained by the processing subsystem, which included total bandwidth consumption per university, peak-hour bandwidth requirements and most active IP addresses and port numbers. Finally, the Monitoring System provides a graphical interface, whereby such processed information can be accessed via web and properly visualized (this is the third stage).

Figure 3 shows the database structure, which consists of a single table with the typical MRTG fields.

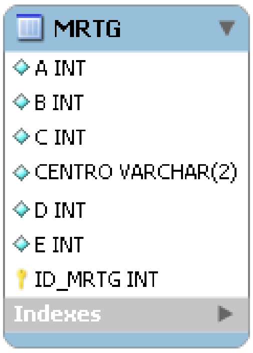

MRTG
A INT
B INT
C INT
CENTRO VARCHAR(2)
D INT
E INT
ID_MRTG INT
Indexes

Figure 3: Schema used to store aggregated passive measurement data

6.4 BGP routing information database (RIPE)

Although RIPE (Réseaux IP Européens) [i.43]seems to be focused on IPv6, its database offers a wide repository of BPG routing announcements and withdraws (also for IPv4). Such data can be used to understand IP traffic troubles and learn to set traffic engineering guidelines. In [i.44] one can find an extensive set of documents that describe RIPE´s data exploitation. As for BGP, two tables are of interest, the ones shown in Figure 4:

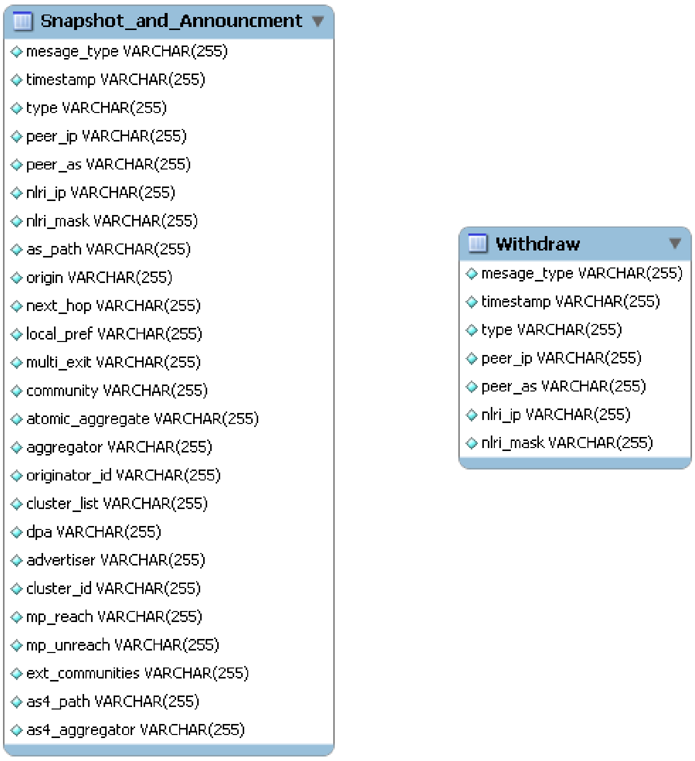

Snapshot_and_Announcment
mesage_type VARCHAR(255)
timestamp VARCHAR(255)
type VARCHAR(255)
peer_ip VARCHAR(255)
peer_as VARCHAR(255)
nlri_ip VARCHAR(255)
nlri_mask VARCHAR(255)
as_path VARCHAR(255)
origin VARCHAR(255)
next_hop VARCHAR(255)
local_pref VARCHAR(255)
multi_exit VARCHAR(255)
community VARCHAR(255)
atomic_aggregate VARCHAR(255)
aggregator VARCHAR(255)
originator_id VARCHAR(255)
cluster_list VARCHAR(255)
dpa VARCHAR(255)
advertiser VARCHAR(255)
cluster_id VARCHAR(255)
mp_reach VARCHAR(255)
mp_unreach VARCHAR(255)
ext_communities VARCHAR(255)
as4_path VARCHAR(255)
as4_aggregator VARCHAR(255)

Withdraw
mesage_type VARCHAR(255)
timestamp VARCHAR(255)
type VARCHAR(255)
peer_ip VARCHAR(255)
peer_as VARCHAR(255)
nlri_ip VARCHAR(255)
nlri_mask VARCHAR(255)

Figure 4: Schema used to store BGP routing information

6.5 Periodic active measurement database (ETOMIC)

ETOMIC [i.45] has been designed to allow researchers to perform any kind of extremely accurate measurement experiments. To do so, researchers are provided with a set of monitoring nodes with Endace DAG cards and a web-based graphical user interface from which distributed experiment definition must be done. ETOMIC takes advantage of the time periods when no researcher is using the monitoring nodes in order to collect some pre-defined periodic measurements. At the date of the present document the available measurements in the repository contain one-way-delay measurements, NTP and GPS measurements and several different types of traceroute measurements among all the nodes. The measurement data is stored in a SQL database, while a web interface is opened for public users to reach these data sets. Using HTML forms any researcher can specify the range of measurements to download, which are provided in a simple plain text format.

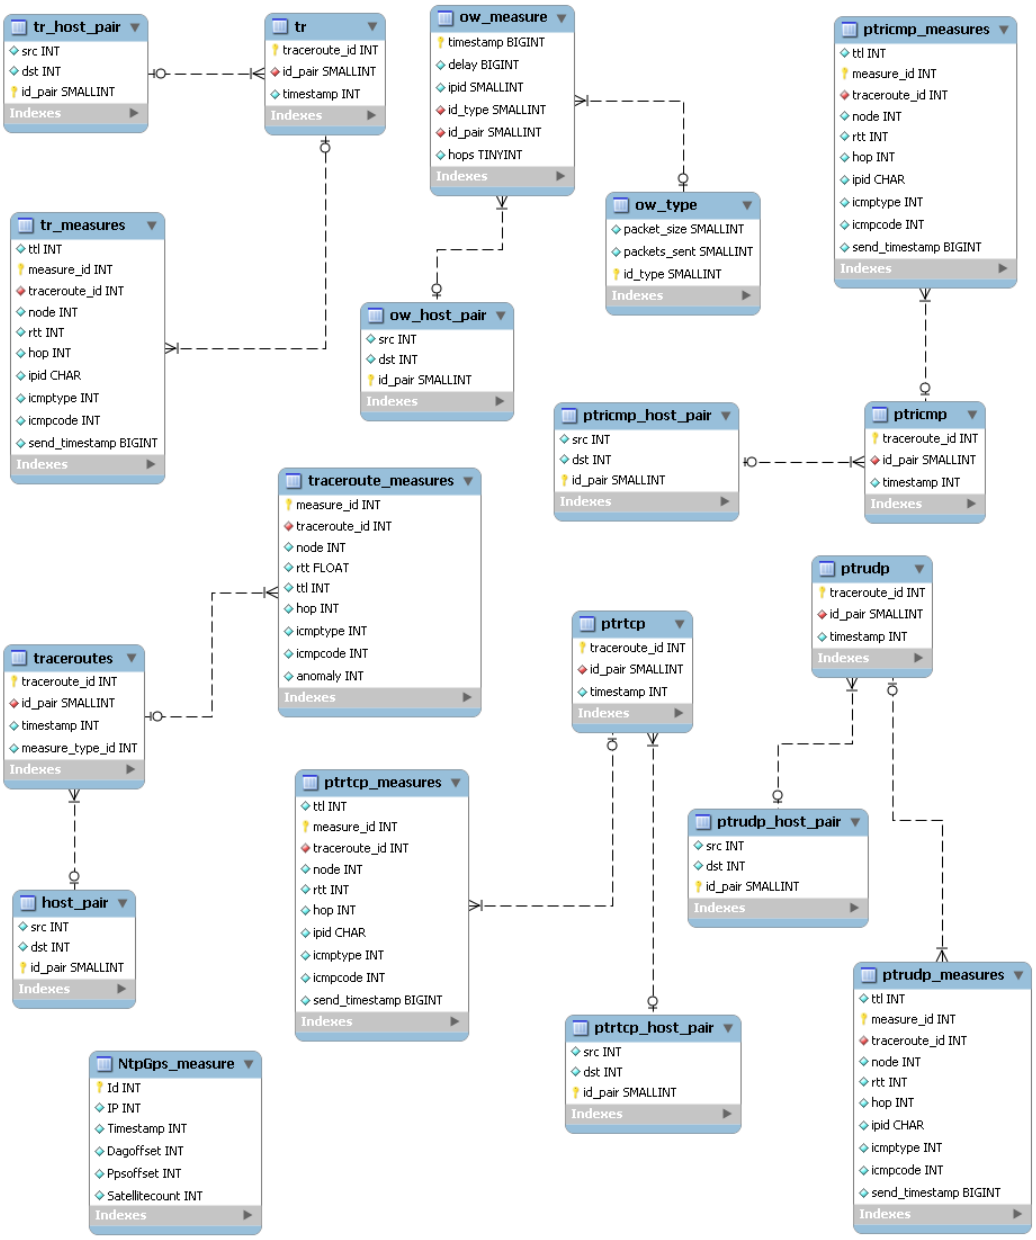

Figure 5: Schema used to store ETOMIC@UPNA pre-defined active measurement data

The ETOMIC infrastructure is very general, and external researchers can create their own measurement experiments, by receiving an account and using the provided interface. This way, external research activities end-up using the interface to run experiments, but collect measurements in a database they host, which is based on a different information model and schema.

The so called ETOMIC-UPNA schema is the database model that hosts the results of the pre-defined periodic measurements, which are executed when no external experiment is running. These results are offered in a database hosted by Universidad Pública de Navarra, and are automatically scheduled.

Any researcher running external experiments must do it by hand, using the web interface to program the experiment, download the results, and store them in his own database. For this reason the ELTE interface also offers a Web Service for downloading obtained results.

The ETOMIC nodes are in common (thus ensuring a common infrastructure) but schemas for organizing and accessing information, and experiments themselves are different, even when the different measurement campaigns measure the same metrics. Thus information collected from ETOMIC nodes is often scattered among different databases.

For instance in [i.46] the collected data from the ETOMIC nodes is organized in an SQL database in a virtual observatory fashion. Most of the measurements are executed in an inter-ETOMIC fashion, but there can also be found datasets gathered in the context of ETOMIC-DIMES and ETOMIC-PlanetLab measurements. Most of the regular ETOMIC measurements were launched in the spring of 2005. As a major instance, there were approximately 1 200 distinct queueing delay tomography measurements performed since April 2005. There are various types of active measurements stored here, the most significant ones are:

- regular experiments to trace temporal changes in the inter-ETOMIC network topology – traceroute logs and extracted topologies;
- regular one way delay time-series and their statistics;
- queuing delay tomography to draw congestion maps of the internal network;
- evaluated end-to-end queuing delay distributions;
- ping measurements and their statistics;
- joint experiments in cooperation with DIMES and PlanetLab – one-way delay, traceroute and tomography;
- GPS logs.

It is also possible to extract end-to-end loss probability information from these stored datasets.

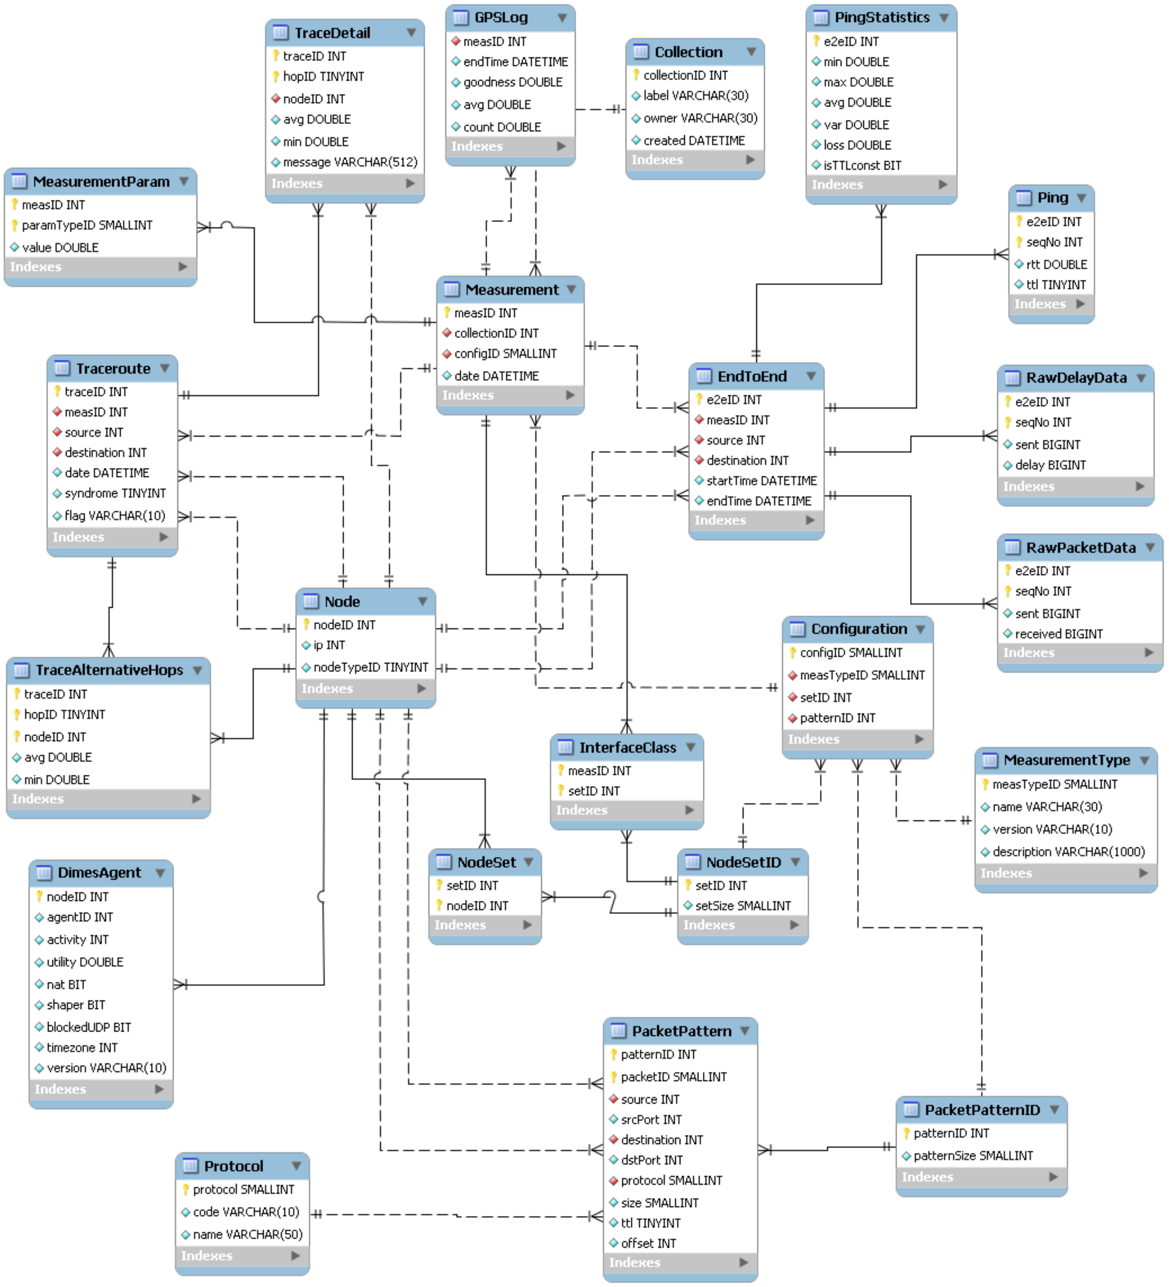

Figure 6: Schema used to store ETOMIC@ELTE active measurement data (A)

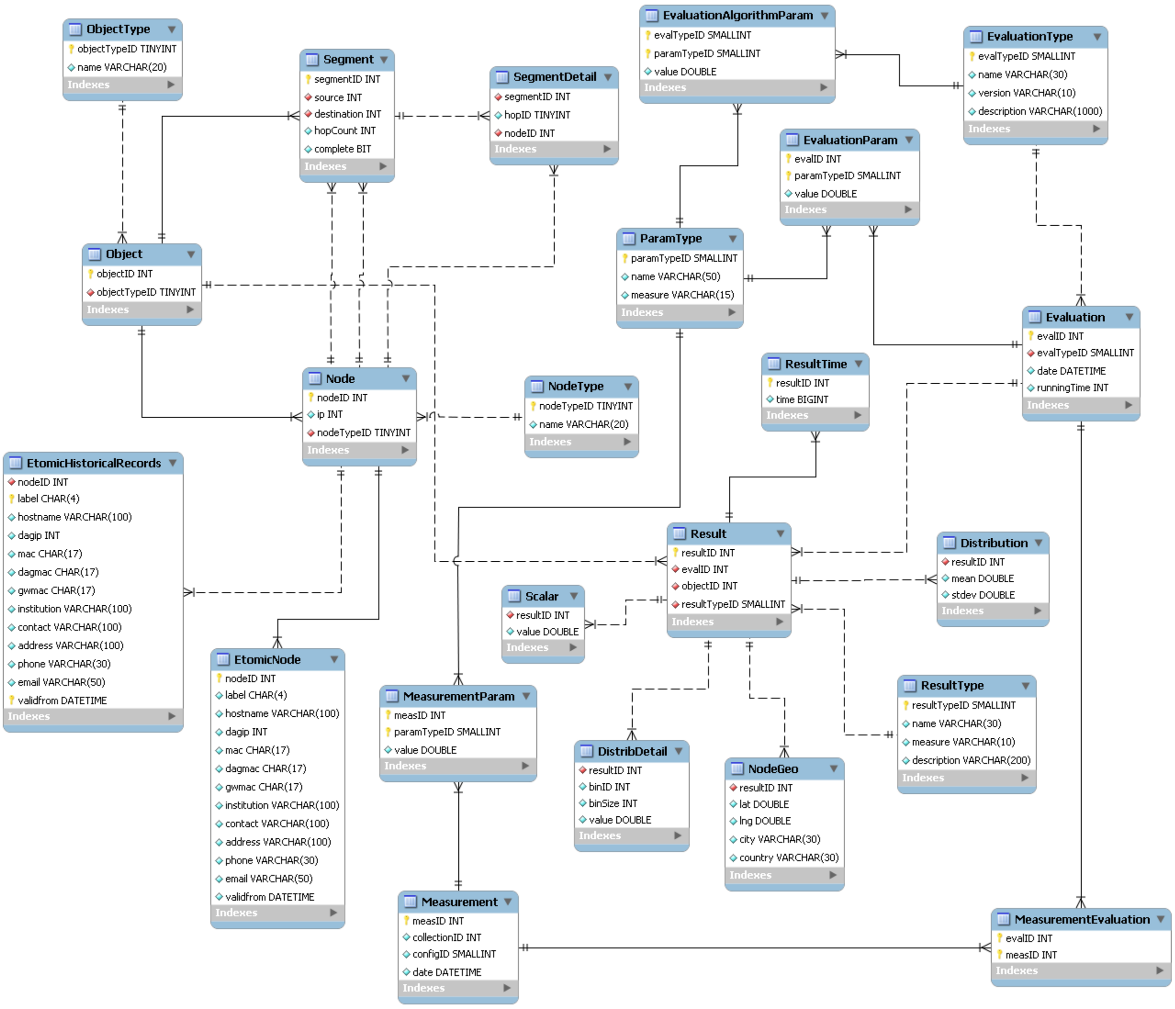

Figure 7: Schema used to store ETOMIC@ELTE active measurement data (B)

6.6 Large scale topology database (DIMES)

DIMES (Distributed Internet Measurements and Experimentation System) [i.47]is a large-scale distributed active measurements effort that measures and tracks the evolution of the Internet from hundreds of different view-points, in an attempt to overcome the "law of diminishing returns". DIMES traceroute and ping measurements are targeted at a set of over 5 million IP addresses, which are spread over all the allocated IP prefixes. On a weekly basis over a thousand measuring clients access the central server and perform measurements to various destinations. These measuring clients are spread on over 250 different ASes around the globe.

The DIMES infrastructure consists of thousands of measurement clients, called Agents. Each of these agents perform Traceroute and Ping measurements using ICMP or UDP packets. The agents return the measurement results to the server, every time they have completed running a script. The results are then stored in the central server to enable later processing for the network research community.

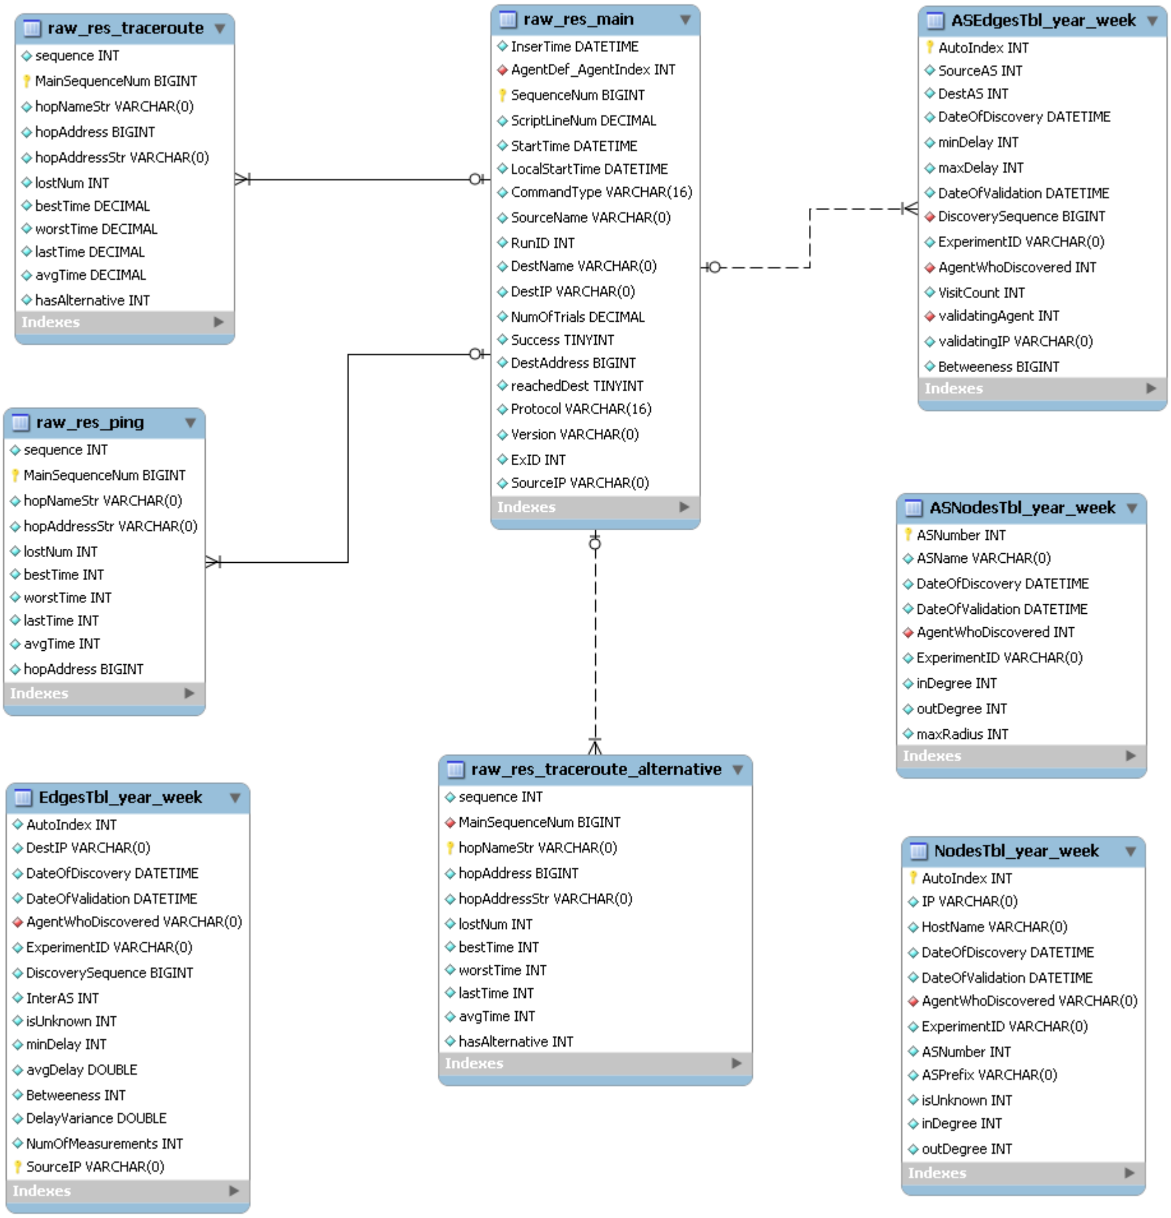

Figure 8: Schema used to store DIMES measurement data

6.7 MINER

MINER [i.48] differs from the infrastructures described above in several ways. First, there is (currently) no deployment in the public Internet with an existing infrastructure consisting of servers, databases, measurement nodes etc. Second, MINER is not designed for some specific set of measurements to be made but is instead built upon the concept of enabling the integration and orchestrated use of any available kind of monitoring/measurement tools (and even tools that do not measure at all but e.g. configure a component of the system under test). MINER is thus completely agnostic to what kind of measurements it actually conducts.

In brief, the value of MINER is that provides a fully programmable measurement infrastructure that supports users in the process of carrying out extensive distributed measurement studies. MINER enables the specification of measurement activities, so called MINER scenarios, that define which tools have to be executed on which measurement nodes at what given time periods and which results have to produced. MINER then does the heavy-lifting of executing the scenario.

MINER has been designed with extensibility from the outset and there is a strict separation between the core functionality and the measurement tools that are plugged into the system. It is a major requirement that the integration of a tool must not require any changes to the core system.

As a direct consequence, the MINER core which also contains the database can not make any assumptions about the kind of results it will have to store. This explains why the result tables of the DB schema (blue area in Figure 9) relate to basic types like integer, long, double, string, text, and binary data and there are no tables named IPAddress, delay, loss or the like.

In support of documentation and the ability to rerun scenario executions MINER additionally stores the complete scenario specification (green) as well as information about the state of the infrastructure itself (red area in Figure 9).

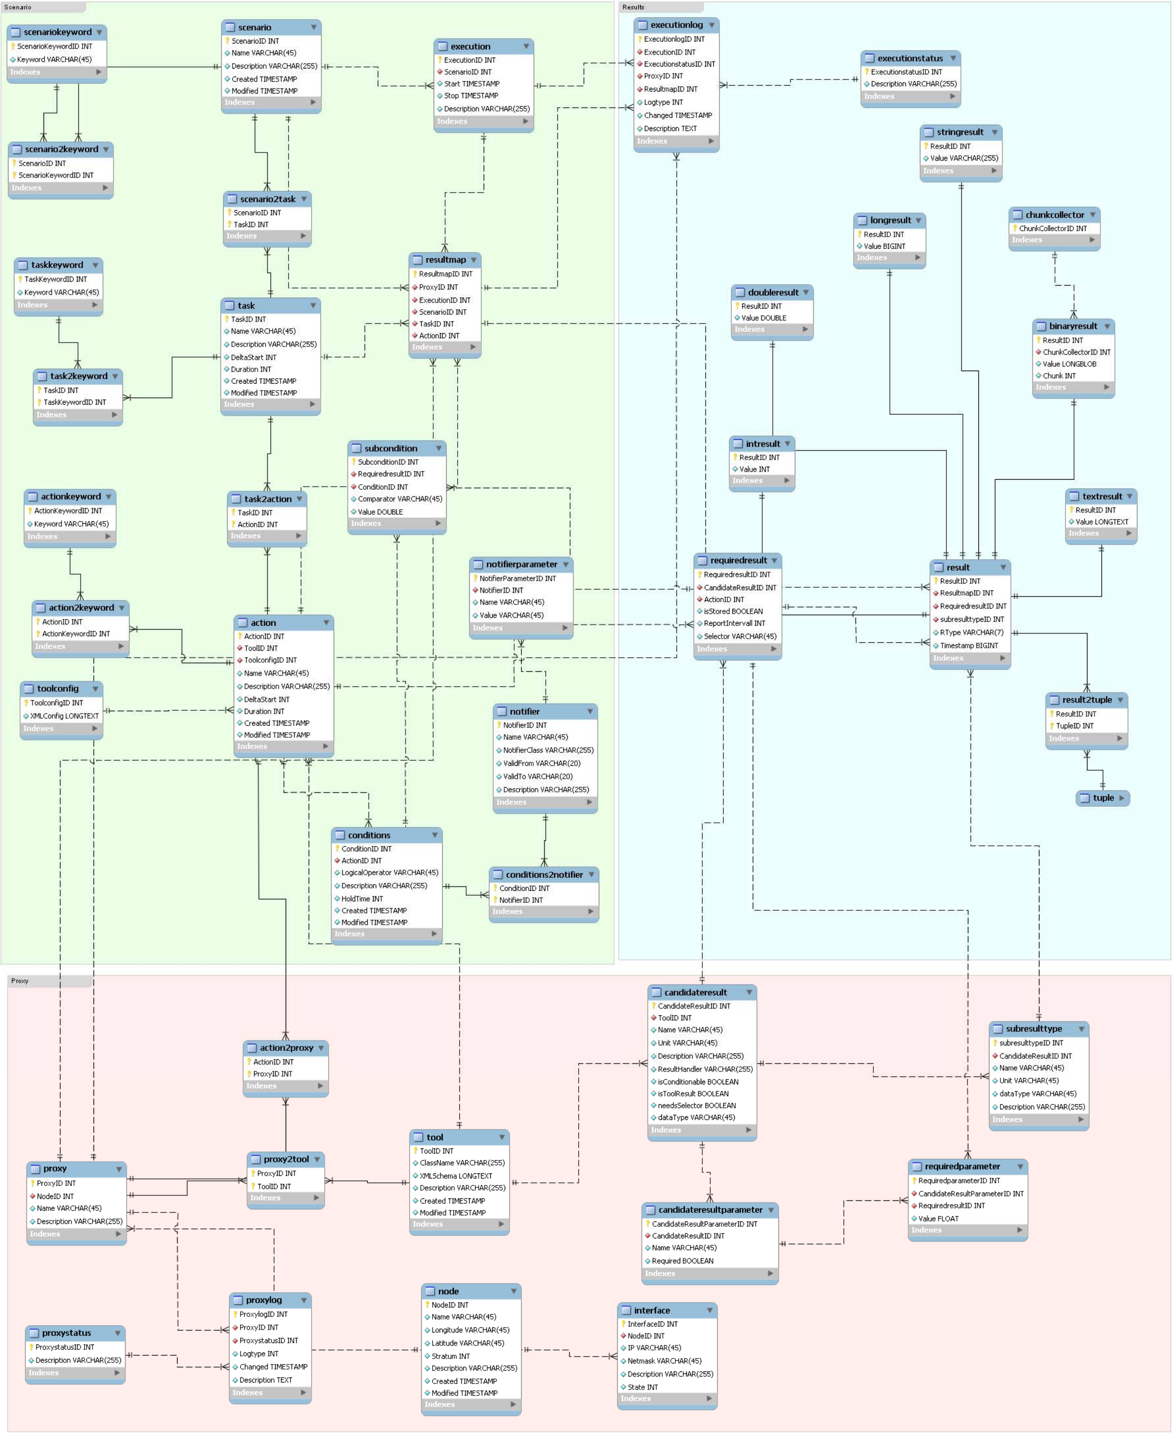

Figure 9: Schema used to store MINER measurement data

7 Models for measuring Quality of Experience in Network Services

Quality of Experience (QoE) is a subjective measure of a user's experiences with a service [i.49]. It is related to but differs from Quality of Service (QoS), which attempts to objectively measure the provided service, or its score against a formalized contract with the end-user. Conversely, QoE is a purely subjective measure from the user's perspective of the overall value of the service provided, and it is the only measure that counts for customers of a service. This is why the concept of QoE is also known as Perceived Quality of Service (PQoS), in the sense of the QoS as it is finally perceived by the end-user. The crucial issue is that, although subjective, as an important measure of the end-to-end performance at the service level from the user's perspective, the QoE is also a valuable metric for the optimal design of the service.

ETSI defined in [i.50] two different aspects of quality (see Figure 10), consisting of the user or consumer vision and the vision of the service provider. That report identifies two main aspects for the user QoE: the quality requirements defined by the consumer and the perceived quality. From the standpoint of the service provider the offered quality of service and the quality actually delivered are defined. Both the offered and distributed quality depend on criteria related to network and independent of it. The report's proposal includes the need to establish goals and measures related to network quality defined and delivered by the provider. It suggests the use of matrices to capture user requirements

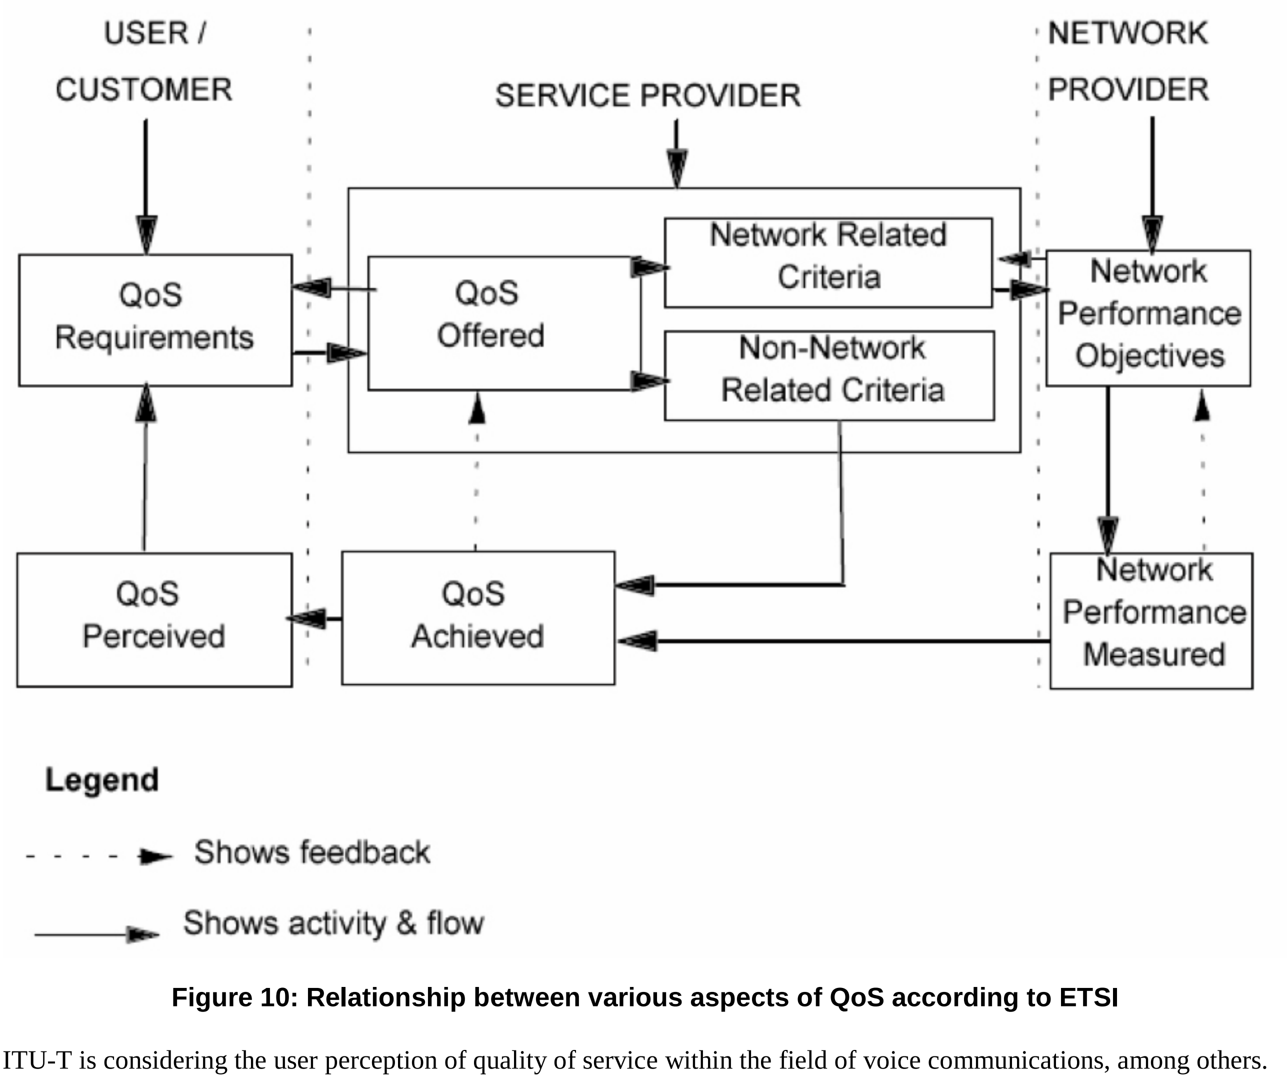

Figure 10: Relationship between various aspects of QoS according to ETSI

ITU-T is considering the user perception of quality of service within the field of voice communications, among others. While the first standards were developed focusing on circuit switching, it is now widely applied in the voice over IP services.

The E.800 recommendation [i.51] states that "the essential aspect of the global evaluation of a service is the opinion of the users of the service. The result of this evaluation expresses the users' degrees of satisfaction". This recommendation provides a framework for describing the quality of service concept, for relating quality of service and network performance and introduces a set of performance measures, since it is desirable that the provider has a detailed knowledge about the quality of the offered service. The document says that the user's degree of satisfaction with the service provided depends on the perception of four performance metrics of the service: support, operability, serviceability and security, which all are, to different degree, dependent on network characteristics. All performance concepts may be related to instant of time (instantaneous, etc.) or expressed as a mean value over a time interval. Measures are connected to events (failure, restoration, etc.), states (fault, up state, down state, outage, etc.) or activities (e.g. maintenance), with their time durations.

There are several recommendations that have tried to express these concepts for use within frameworks for quality. The ITU X.641 [i.52] defines a framework for QoS that presents and defines a set of key terms and concepts:

- Quality of Service characteristic. Quality of service of a system, service or resource that can be identified and quantified. There are generic characteristics, specialized ones, and derivatives of those.

- Quality of service information. Any information relating to QoS. This information is divided, according to their nature, in *data* and *requirements*. The data are consistent with information on the actual behaviour of the system in terms of quality. The requirements are objectives that must be met in relation to quality.

- Quality of service measurement. Observed value of a quality of service characteristic.

- QoSAttribute. An attribute of a managed object that relates to the quality of service.

- QoSManagement. Any set of activities by a system or communication service to support the monitoring, control and management of service quality.

- QoSPolicy. A set of rules that determines the characteristics of service quality and management functions to be used.

- QoSManagementFunction. A function specifically targeted to meet the QoS requirements of a user or application, provided by one or several QoS mechanisms.

- QoSMechanism. A specific mechanism that can make use of elements of protocols, QoS parameters or a QoS context, possibly in combination with other mechanisms in order to support the establishment, monitoring, maintenance, quality control or consultation of a service.

Measuring the QoE, or perceived QoS, still remains one of the major goals which current research efforts are focusing on. There exist both specific approaches for a certain restricted category of services (i.e. in the context of video streaming over IP) and generic frameworks that try to capture the semantics of QoS and how services are perceived by the user, in order to design the useful matching functions between measures of service parameters and their corresponding levels of enjoyed QoE from the user.

Subjective assessments of video quality are done using a panel of humans rating a series of videos according to their personal opinion, while objective assessments use algorithms and formulas to assess the quality automatically in a repeatable way. Subjective video quality estimations are standardized by ITU in the recommendation ITU-R Recommendation BT.500-11 [i.53]. There are several variants within this standard: Single Stimulus (SS), Double Stimulus Impairment Scale (DSIS), Double Stimulus Continuous Quality Scale (DSCQS), Single Stimulus Continuous Quality Evaluation (SSCQE), Simultaneous Double Stimulus for Continuous Evaluation (SDSCE), Stimulus Comparison Adjectival Categorical Judgment (SCACJ). Subjective evaluation methods usually call the outcome of the evaluation process under the term MOS (Mean Opinion Score). The main problem with subjective evaluations is their cost in terms of both time and manpower, which makes them hard to repeat often and impossible to be part of an automatic process.

Thus, to perform real-time video quality evaluation, a hybrid class of techniques called Pseudo-Subjective Quality Assessment has been proposed [i.54]. This class of techniques allows to automatically learn the values obtained from a test (be it subjective or objective), and generalize from them, inside a software module. The PSQA technique combines the advantages of both approaches (subjective and objective) by evaluating a series of pre-distorted videos (where some of the parameters affecting audio/video quality are set to specific values) that are used as database to represent a significant sample in the space of all possible combinations of input parameters (affecting video quality) and the output given by the evaluation of the received video. This empirical function is used to train a neural network, in order to approximate the result of the evaluation outside the sampled points in the space of parameters.

On the other hand, frameworks for modelling the QoE semantic domain in generic telematic services are quite scarce, and this is a fast developing topic in both the research community and the standardization bodies. In [i.56] a work is presented that aims to contribute to the measurement of the perceived quality of a telematic service by using a framework based on semantic representation of the quality of service itself, and the relationships between the technical quality layer with its objective parameters, and the layer of the quality perceived by the user. It defines a semantic schema capable of representing the concepts of the quality of service and the perceived quality scope, using ontologies and the knowledge from the quality of service standards and recommendations. Furthermore, in [i.57] methods are described to capture and evaluate perceived quality information, in order to define an automatic system for the calculation of the quality of experience of the users, using the objective characteristics measured by the technical QoS layer of the telematic service.

Annex A (informative): Authors & contributors

The following people have contributed to the present document:

Rapporteur:

- Mr. Giuseppe, Tropea, CNIT

Other contributors:

- Prof. Jorge, Lopez de Vergara, UAM
- Dr. Georgios, Lioudakis, NTUA
- Dr. Daniel, Morato, UPNA
- Prof. Javier, Aracil, UAM
- Mr. Alfredo Salvador, UAM
- Mr. Felix, Strohmeier, SRFG
- Mr. Anastasios, Zafeiropoulos, GRNET
- Dr. Athanassios, Liakopoulos, GRNET

History

Document history		
V1.1.1	May 2010	Publication