

The European Traffic Observatory Measurement Infrastructure (ETOMIC): A testbed for universal active and passive measurements

D. Morató, E. Magaña, M. Izal, J. Aracil, F. Naranjo, F. Astiz and U. Alonso
Universidad Pública de Navarra, 31006 Pamplona, Spain
{eduardo.magana,daniel.morato,mikel.izal,javier.aracil,ulisses.alonso}@unavarra.es
I. Csabai, P. Haga, G. Simon, J. Steger and G. Vattay

Collegium Budapest, H-1014 Budapest, Szentháromság u. 2.
{csabai,haga,simon,stegeer,vattay}@colbud.hu

Abstract

The European Traffic Observatory is a European Union VI Framework Program sponsored effort, within the Integrated Project EVERGROW, that aims at providing a paneuropean traffic measurement infrastructure with high-precision, GPS-synchronized monitoring nodes. This paper describes the system and node architectures, together with the management system. On the other hand, we also present the testing platform that is currently being used for testing ETOMIC nodes before actual deployment. [†]

Keywords: Traffic measurement platforms, active and passive measurement experiments

1 Introduction and system architecture

Internet traffic is growing at an extraordinary pace and, as a consequence, traffic control and forecasting are becoming fundamental issues for network operators. The European Union VI Framework Program is sponsoring a major research effort that aims at building a solid understanding of the Internet evolution until 2025, from a complex systems perspective. The EVERGROW ¹ Integrated Project provides a multifaceted approach to forecasting the Internet evolution. The different subprojects that make up Evergrow cover a wide range of topics, from distributed systems and message passing to measuring Internet traffic in real time. Precisely, one of the subprojects is specifically focused on realizing a paneuropean measurement infrastructure, consisting of 50 measurement nodes which will be deployed at selected European locations. Such measurement infrastruc-

ture is called *European Traffic Observatory Measurement Infrastructure* (ETOMIC).

ETOMIC is targeted to provide the scientific community with a measurement platform that is i) fully open and reconfigurable and ii) extremely accurate (nanoseconds) and GPS-synchronized. First, ETOMIC has been designed to allow researchers to perform *any kind* of measurement experiments. To do so, researchers are provided with an interface from which software upload to the measurement nodes is possible. A choice of measurement scripts are also provided, that cover the most popular measurement techniques (packet pairs, etc). On the other hand, the researcher may also provide his/her own code for the experiments, that will be automatically compiled by the ETOMIC management system. Then, node reservation can be performed through the web-based user interface. Finally, the ETOMIC management kernel takes care of the software upload and experiment execution, in a fully automated fashion. On the other hand, ETOMIC is a high-precision infrastructure, due to the fact that Endace DAG cards [1][2] are incorporated to the nodes. Such cards are specifically designed to transmit packet trains with strict timing, in the range of nanoseconds. Actually, the DAG card transmit code has been specifically modified according to the ETOMIC requirements. Furthermore, a GPS is also incorporated to the measurement nodes, so that the whole measurement infrastructure is synchronized to the same reference clock. For example, a possible application is one-way delay measurements between nodes.

The system components are depicted in figure 1. A central management system is in charge of system control, comprising not only the scheduling and execution of measurements experiments, but also system monitoring (survivability) and configuration. The central management system is composed of a workstation and a traffic repository to which measurements can be downloaded for subsequent

[†]This paper is accompanied by a demo of ETOMIC to be delivered at the conference

¹<http://www.evergrow.org>

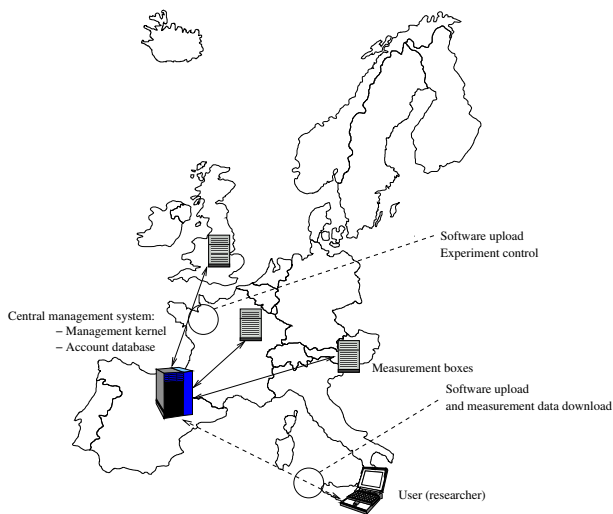


Figure 1. ETOMIC

processing. The software running in the management system is called the *management kernel*. On the other hand, researchers interact with the management system through the *user interface and account database*. ETOMIC provides an account to any registered researcher, so that he/she can upload software and download measurement results. Finally, the interface also serves for the system manager to enter configuration data about the measurement nodes and researcher accounts.

The paper is structured as follows. The remaining of the section is devoted to the state of the art. Section 2 presents several examples of active and passive measurements that can be performed with ETOMIC. Section 3 presents the testing environment for ETOMIC. Section 4 presents the node architecture and management kernel and section 5 is devoted to the user interface. Finally, section 6 presents the conclusions that can be drawn from this paper.

1.1 State of the art

ETOMIC is not the only measurement infrastructure that is currently available in the state of the art, which includes the Surveyor, Felix, Skitter [3], IPMA and AMP [4]. On the other hand, NIMI (National Internet Measurement Infrastructure) [5] is a large-scale measurement infrastructure with diverse administered sites. ETOMIC differs from the previous measurement infrastructures since it combines in the same platform i) high-precision and GPS-synchronization and ii) capabilities to run all kinds of measurement experiments. ETOMIC incorporates dedicated measurement hardware and dedicated machines that are administered by a single management site, in contrast to NIMI. The aim is to create a high-precision measurement platform that can cope with the requirements of high reso-

lution measurements in high-speed environments. Note that the timing requirements for Gbps speeds are very strict both for passive and active measurements. For such scenarios, a dedicated hardware approach is often adopted [6]. On the other hand, high-precision and GPS-synchronization is also offered by the RIPE NCC infrastructure [7]. However, the measurement software available in the boxes is not uploaded by the end user on a per-measurement basis.

2 ETOMIC applications

ETOMIC is an *open infrastructure*, i. e. the final goal is to provide a service to the network measurement community. To this end, ETOMIC does not incorporate specific measurement software. On the contrary, nearly any measurement program can be run in ETOMIC nodes, even though off-the-shelf software is also provided for rapid prototyping. In this section we explore possible applications of ETOMIC, in order to assess the value and usefulness of our measurement platform.

A plethora of network monitoring software has been implemented to date (see for instance [8]), showing an increasing interest in network monitoring. On the other hand, such monitoring software is very diverse, from delay and bandwidth estimation tools [9] to Internet tomography tools [10]. Due to this diversity, a universal network measurement platform such as ETOMIC proves very useful to tackle a wide range of open topics in Internet measurement.

Furthermore, not only ETOMIC is open but it also provides high-precision, which is becoming a requisite due to the increasing bandwidth of Internet core links. For example, packet probing methods [11] can significantly benefit from high resolution timestamping, since larger available bandwidths can be measured.

On the other hand, note that ETOMIC is also a GPS-synchronized platform. This is specially interesting to provide a time reference service for a wide variety of applications. For example, Internet coordinate systems [12] attempt to assign a coordinate vector to hosts in the Internet, such the distance approximates the network distance between them. Such coordinates are useful for content delivery networks, peer-to-peer applications, etc. To implement a coordinate system, some solutions advocate for the use of landmarks, which consist of a set of reference hosts by which other hosts can establish their coordinates. The ETOMIC platform may serve to this purpose since the node locations are known beforehand and the links between nodes can be characterized accurately (delay, bandwidth) at a one-way level.

Finally, ETOMIC may also serve as a reference distributed clock for testing network time protocol [13] enhancements. As a conclusion, we note that ETOMIC is an open platform that allows for a wide variety of applications

for the measurement community.

3 Example network testbed for ETOMIC

ETOMIC is currently being developed by the EVERGROW consortium, with a scheduled development time of 18 months. The design approach is rapid prototyping and testing. The prototype is developed and tested over a laboratory network configuration, that mimics the behavior of an Internet link. Then, changes are introduced to the prototype according to the experience gained with the link emulator.

3.1 Emulator components

The link emulator consists of the following components:

1. Router 3660: A Cisco router with several interfaces. Those used for this emulator are an ATM OC-3c interface and two Fast Ethernet interfaces.
2. Router 7204: Another Cisco router with two Fast Ethernet interfaces, an ATM OC-3c one and a POS interface not used in the emulator
3. An ATM link that connects both routers. The link is configured for CBR service with 100 Kbps bandwidth
4. Four ETOMIC measurement nodes that are attached to the routers' Fast Ethernet networks

The link emulator components are depicted in figure 2.




Figure 2. Testbed block diagram

3.2 Physical layout

The emulator is currently located at the research laboratories of Public University of Navarra, Spain. The aim of the emulator was to have an Internet link, with varying load and delay conditions. By doing so, we were able to extensively test the ETOMIC agents on the link, under controlled congestion conditions.

The physical layout of the network emulator is presented in figure 3.




Figure 3. Testbed layout

3.3 A test experiment

An example of test experiment is presented next. The experiment objective was to test the ETOMIC delay measurement capabilities. To this end, the link was artificially loaded with a step function, so that the load was suddenly increased one minute after the experiment start. On the other hand, a measurement agent was sending unidirectional UDP packets on the link, whose delay was measured at the other end. The expected result is a sudden delay increase that should correspond with the time instant the load is increased.

Accordingly, the following measurement agents were coded:

1. Agent 1: The Agent was in charge of sending a low speed flow to Agent 2. The packets in this flow contained a timestamp telling when they were sent
2. Agent 2: It received the packets sent by Agent 1
3. Agent 3: The Agent sent a high speed flow to Agent 4
4. Agent 4: It received the packets from Agent 3

Agent 1 was uploaded to node ever3 in figure 2, whereas agents 2, 3 and 4 were uploaded to node ever1, ever4 and ever2 respectively. The flow from Agent 3 to 4 produces congestion in the ATM link and so a high delay in the path from nodes ever3 and ever4 to nodes ever1 and ever2. On the other hand, Agents 1 and 2 produced a unidirectional UDP flow. Agent 1 inserted a timestamp per packet and Agent 2 recorded the unidirectional delay. Note that ETOMIC nodes are GPS-synchronized, thus allowing for one-way delay measurements.

The experiment results are presented in figure 4. As expected, a delay increase happens by the time the load is increased (50 seconds after the experiment start). The experiment allowed not only to test the basic ETOMIC functionalities (node booking, software upload and results download) but also the ETOMIC synchronization capabilities and delay measurement resolution.




Figure 4. Experiment results

4 Node architecture and management kernel

In this section, the hardware and software architecture of the measurement node is described, together with the ETOMIC measurement kernel.

4.1 Node architecture

The measurement nodes are based on standard PC hardware with Debian GNU/Linux 3.0 (Woody) with enhanced kernel capabilities for low level network access without root privileges. For the network monitoring interface an Endace DAG 3.6 GE is used, which is specifically designed for active and passive monitoring. Such cards do not use interrupts to signal packet arrivals to the kernel, and, thus, packets can be captured at gigabit speeds. Shared memory is used as a means to relay packets to the analysis program running in user space, in such a way that interrupts are avoided. Furthermore, the GPS reference signal is used to timestamp the incoming packets with high resolution.

The GPS reference is based on a Garmin GPS 35 HVS. It is a low-cost water-resistant GPS receiver that is used to synchronize the measurement nodes. Specifically, the GPS provides a PPS signal (pulse per second) directly to the DAG card. The resulting accuracy is 100 nanoseconds in the packet timestamps and interframe generation intervals. An RS232-RS422 converter has been designed in order to bridge the mismatch between GPS receiver and PC ports. Furthermore, this converter transforms the RS232 signal from the GPS to a RS422 signal that allows larger distances between the GPS and the PC. This is a very important issue because the GPS receiver has to be placed outdoors with direct sky visibility, mainly on the roof. The converter, DAG card and GPS interconnections are displayed in figure 5.




Figure 5. Monitoring node

4.2 Management kernel

The management kernel constitutes the core of the central management system. It is in charge of scheduling the experiments and keeping the corresponding results in the traffic repository. The software is implemented for GNU Linux operating system.

The researcher is expected to use the (web-based) interface in order to book several nodes during a certain time interval and in order to upload the measurement client and server. Such bookings are stored in the central database, that will be described in the next section. Note that the database contains all the relevant information and software about the measurement experiment. Thus, the management kernel is continuously checking the database for new measurement requests. Then, SSL (Secure Sockets Layer, RFC 2246) is used by the kernel to securely command the measurement nodes.

Once a new experiment has been defined and the deadline for execution approaches, the management kernel performs the following tasks:

- Software upload and measurement node configuration
- Experiment execution
- Results download

Such tasks are internal to the system and transparent to the researchers, that will only use the web interface to upload software to the central management system and retrieve measurements.

The *software upload and measurement node configuration* task is performed before the experiment starting time. In this task, the measurement programs are uploaded to the measurement nodes. Accompanying data files may also be uploaded that contain traces or measurement parameters. This allows to upload many different parameters for one single experiment. The data transfer has been enhanced with transfer resume functionalities in order to be able to manage very large data files. Once the files are uploaded, the measurement nodes are programmed with the starting and ending time for each of the executables that are going to be run for the experiment. This time scheduling is performed using

the *atd* service in Linux. Several subtasks are performed for this task, such as quota reservation for the experiment, file upload and time programming for each of the measurement nodes.

During the *experiment execution* task, the management kernel is on standby until the end of the experiment. It only has to ensure that no other experiment is using the same measurement nodes in the same time interval. The aim is to completely isolate the measurement nodes so that the high-precision measurement hardware can be exclusively devoted to a single experiment.

Once the experiment is finished, the management kernel will *download* the resulting data files from each measurement node. The experiment files will be deleted at that time. Note that during the download phase the measurement node cannot be assigned to another measurement experiment, since the measurement and the download would interfere with each other. An estimation of the download time for an experiment is performed by the management kernel and stored in the database as part of the reserved time interval for such experiment. However, since download time is highly dependent on the network connectivity, transfer resume capabilities are necessary. Thus, the download phase may be interrupted by the management kernel and resumed in the future, depending on the number of pending measurement requests.

Besides the experiment management tasks described above, maintenance tasks are also performed by the management kernel. Such tasks have low priority and they take place when no other task is scheduled for a given measurement node. An example of a maintenance task is to configure a newly incorporated measurement node to the network. To do so, the public key and basic software configuration is uploaded to the node. Finally, other maintenance tasks are related to keepalive functions such as connectivity checking and bandwidth estimation.

5 Database and user interface

ETOMIC provides an interface for researchers and administrators which fulfills the different requirements they may have. The former require capabilities to define new experiments, to reserve the measurement nodes and to download results from the management system. The latter require functionalities to manage users, measurement nodes, software and experiments.

It turns out that the database and management kernel (section 4.2) are strongly interdependent. Note that the database provides all the necessary software and data to keep track of measurement experiments. For example, the measurement nodes data (IP address for instance) is stored in the database. In addition, the database is used to store the

experiment results, the experiment status and the measurement node status.

The interface is based on the ubiquitous Web service. Using HTML4, Javascript, PHP4 and Apache Web server (using a Secure Sockets Layer) we were able to provide interface capabilities which are close to those offered by end-user applications. The aim is to facilitate the definition of the experiments as much as possible. This targeted simplicity and ease of use posed more stringent challenges on the researchers interface. On the other hand, the ETOMIC administrators interface was easier to define. Snapshot from the interface are shown in figure 6.

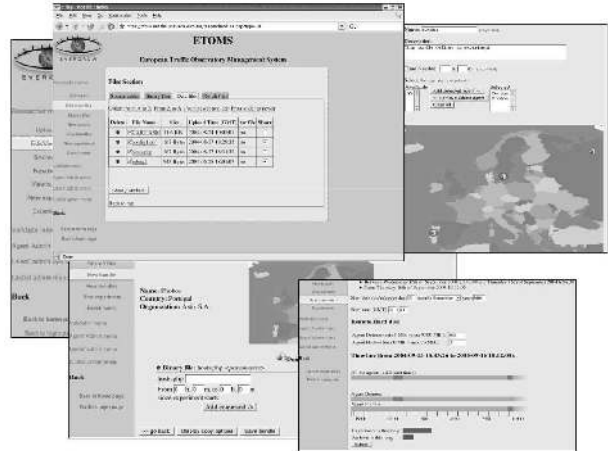


Figure 6. User interface

ETOMIC will mostly be used by researchers willing to define measurement experiments. This kind of user is offered the following functionalities through the interface:

- **Adding a new program:** In order to run simple experiments, several scripts for sending and receiving IP packets with the DAG card high-precision timestamping are provided. However, not only *off-the-shelf* programs are offered but also an API for creating new programs for whatever experiment the researcher may have in mind. Hence, the interface offers capabilities to upload new programs, as source code. The programs are automatically built and made available as soon as the compilation process is over. Shell scripts can also be uploaded in such a way the researcher can easily create *batch-works* based on already existing programs.
- **Uploading data files:** User programs may require configuration parameters or input data files. Such files can be uploaded at any time.
- **Creating an *experiment-Bundle*:** The definition of a new experiment comprises several measurement nodes

and several programs running on each node, with different starting and ending times. This is called an "Experiment Bundle". First, the monitoring nodes that make up the experiment must be selected. Then, the programs that run on each node are chosen. A choice of programs is offered, from the scripts mentioned before to the traditional unix network tools (ping, traceroute, etc), besides the source code written by the user. As a last step, for each of the programs that make up the experiment bundle a starting and ending time can be provided *relative to the starting time of the experiment*. Thus, full flexibility is provided for defining asynchronous execution of several programs within the same experiment. It is noteworthy that the programs that make up the bundle can be reused for other experiment bundles, since they are permanently stored in the database.

- **Booking ETOMIC time:** Once the bundle definition is performed, the researcher should define the starting time for the experiment. To do so, the interface shows the availability intervals per node, together with a best-fit suggestion. However, it is up to the user to select any available time interval, depending on the specific experiment requirements (day time versus night time, for example).
- **Obtaining results:** The management kernel is in charge of downloading the data files created or modified during the experiment. Once the download is over, the researcher will have access to such files through the interface. It is expected that he/she will be willing to download them to his local computer, in order to process the results offline. Hence, functionalities to download a variable number of files, in compressed format, are offered.

The steps followed to create and run an experiment are represented in figure 7.

6 Conclusions

In this paper, the European Union VI Framework program sponsored European Traffic Observatory has been described. The traffic observatory will provide a high-precision fully opened measurement platform for the benefit of traffic measurement researchers. A testbed has been developed for extensive testing of the traffic observatory, that serves to identify the different network conditions and impairments that will affect ETOMIC, before it is deployed at the paneuropean level.

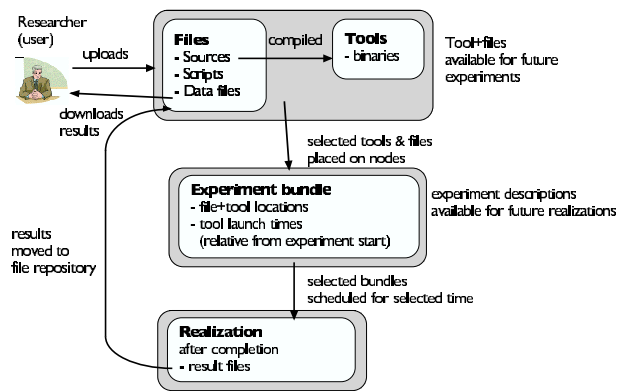


Figure 7. Life-cycle of an experiment definition

References

- [1] Endace Measurement Systems, "http://www.endace.com."
- [2] DAG Project at the Waikato University, "http://dag.cs.waikato.ac.nz/."
- [3] Skitter, "http://www.caida.org/tools/measurement/skitter/."
- [4] Active Measurement Project, "http://amp.nlanr.net/."
- [5] V. Paxson, J. Mahdavi, A. Adams, and M. Mathis, "An architecture for large-scale internet measurement," *IEEE Communications Magazine*, vol. 36, no. 8, pp. 48–54, 1998.
- [6] J. Aracil, D. Morató, and M. Izal, "Analysis of Internet Services in IP over ATM networks," *IEEE Communications Magazine*, vol. 37, no. 12, pp. 92–97, December 1999.
- [7] RIPE Network Coordination Center, "http://amp.ripe.net/."
- [8] N. monitoring tools, "http://www.slac.stanford.edu/xorg/nmtf/nmtf-tools.html."
- [9] C. Dovrolis, R. Prasad, M. Murray, and K. Claffy, "Bandwidth estimation: metrics, measurement techniques, and tools," *IEEE Network*, vol. 17, no. 6, 2004.
- [10] D. M. K. Claffy, T. Monk, "Internet tomography," *Nature*, January 1999.
- [11] A. Pasztor and D. Veitch, "On the Scope of End-to-End Probing Methods," *IEEE Communications Letters*, vol. 6, no. 11, pp. 509–511, 2002.

- [12] E. Ng and H. Zhang, "Predicting Internet network distance with coordinates-based approaches," in *Proceedings of Infocom, 2002*, 2002.
- [13] D. Mills, "A brief history of NTP time: confessions of an Internet timekeeper," *ACM Computer Communications Review*, vol. 33, no. 2, pp. 9–22, April 2003.