Community Research

Specific Targeted REsearch Project

PRISM

D3.2.1: State of the art on monitoring applications

Project acronym: PRISM
Project full title: Privacy-Aware secure Monitoring
Contract No.: 215350
Project Document Number: IST-2007-215350-WP3.2-D3.2.1-R1
Project Document Date: 30/06/2008
Workpackage Contributing to the Project Document: WP3.2
Deliverable Type and Security: Public

Author(s): Felix Strohmeier, Peter Dorfinger (Salzburg Research)
Esa Hyytiä, Ivan Gojmerac (ftw.)
Brian Trammell (Hitachi)
Andrea D. Pietro (CNIT)
Georgios Lioudakis, Fotis Gogoulos, Anna Antonakopoulou,
Dimitra Kaklamani, Iakovos Venieris (ICCS)
Sathya Rao (Telscom)
Enrico Stinco (Nettare s.r.l.)

Abstract:

This deliverable reviews and analyses representing monitoring applications in the domains of performance monitoring, anomaly and intrusion detection, traffic classification as well as lawful interception. As relevance to PRISM, the main focus has been put on their impact to privacy. It also analyses, which type of applications requires what input data to perform their specific tasks. Furthermore an analysis of monitoring application in operational environments has been included.

Keyword list: PRISM, IST-2007-215350, Monitoring Applications

History

Version	Date	Description, Author(s), Reviser(s)
1.0	30/06/08	Version 1 for delivery in Month 4, Felix Strohmeier et al.

Table of Contents

Abbreviations

AAPI	Anonymisation Application Programming Interface
ACK	Acknowledgement Number
AES	Advanced Encryption Standard
AIDS/AIPS	Anomaly and Intrusion Detection / Prevention System
AMD	Advanced Micro Devices, Inc.
AS	Autonomous System
ASCII	American Standard Code for Information Interchange
ASIC	Application Specific Integrated Circuit
BIOS	Basic Input/Output System
CPU	Central Processing Unit
DBA	Database Administrator
DBMS	Database Management System
DES	Data Encryption Standard
DFA	Deterministic Finite Automaton
DHCP	Dynamic Host Configuration Protocol
DNS	Domain Name System
DOS	Disk Operating System
DRAM	Dynamic Random Access Memory
EPAL	Enterprise Privacy Authorization Language
ETSI	European Telecommunications Standards Institute
EU	European Union
FCAPS	Fault, Configuration, Accounting, Performance, Security
FP7	Seventh Framework Programme
FPGA	Field-Programmable Gate Array
FTP	File Transfer Protocol
GGSN	Gateway GPRS Support Node
GPRS	General Packet Radio Service
GPS	Global Positioning System
GTP	GPRS Tunnelling Protocol
GUI	Graphical User Interface
HP	Hewlett Packard
HTTP	Hypertext Transfer Protocol
ICMP	Internet Control Message Protocol
ID	Identifier
IDS	Intrusion Detection System
IETF	Internet Engineering Task Force
IP	Internet Protocol
IPSEC	IP Security
IPFIX	Internet Protocol Flow Information eXport
ISP	Internet Service Provider
LDAP	Lightweight Directory Access Protocol
LLC	Logical Link Control
MAC	Media Access Control
MIB	Management Information Base
MPLS	Multi-Protocol Label Switching
MRTG	Multi-Router Traffic grapher
MD5	Message-Digest algorithm 5
NAT	Network Address Translation
NGSCB	Next-Generation Secure Computing Base
NIDS	Network Intrusion Detection System

OAEP	Optimal Asymmetric Encryption Padding
OEM	Original Equipment Manufacturers
OID	Object Identifier
OPES	Order Preserving Encryption Schema
OS	Operating System
OT	Oblivious Transfer Protocol
P2P	Peer-to-peer
PC	Personal Computer
PDP	Policy Decision Point
PEP	Policy Enforcement Point
POP3	Post Office Protocol, Version 3
PRISM	PRivacy-aware Secure Monitoring
PSAMP	Packet Sampling
QoS	Quality of Service
RADIUS	Remote Authentication Dial In User Service
RBAC	Role-based access control
RMON	Remote Monitoring
RRD	Round Robin Database
RSA	A crypto-algorithm by Ron Rivest, Adi Shamir, and Leonard Adleman
RTP	Real-Time Protocol
RTSP	Real-Time Streaming Protocol
RTT	Round Trip Time
SEM	Secure Execution Mode
SEQ	Sequence Number
SGSN	Serving GPRS Support Node
SIP	Session Initiation Protocol
SLA	Service Level Agreement
SMB	Server Message Block
SMC	Secure Multi party Computation
SMI	Structure of Management Information
SMTP	Simple Mail Transfer Protocol
SNMP	Simple Network Management Protocol
SOM	Self-Organising Map
SQL	Structured Query Language
SRAM	Static Random Access Memory
SSH	Secure Shell
TC	Trusted Computing
TCP	Transmission Control Protocol
TCPA	Trusted Computing Platform Alliance
TMP	Trusted Platform Module
TSA algorithm	Time Slot Assignment algorithm
TSTAT	TCP Statistic and Analysis Tool
UDP	User Datagram Protocol
UDR	User Data Record
URL	Uniform Resource Locator
UTC	Coordinated Universal Time
VLAN	Virtual Local Area Network
VoIP	Voice over Internet Protocol
WP	Workpackage
XACML	eXtensible Access Control Markup Language
XML	Extensible Markup Language

List of Tables

List of Figures

1 Introduction

Monitoring is all times an essential part for network operation. In the early days of the Internet with a limited amount of mainly friendly users, plain best-effort service and the absolute equal treatment for all packets, network monitoring was only of minor importance. Within the last decades, with increasing bandwidth demand, QoS differentiation for voice, video and data traffic, network attacks, discussion on net neutrality, etc. different monitoring applications were created, each dedicated to fulfil a special task. As monitoring data collection (and retention) cannot be used for the dedicated tasks only, but also for marketing, surveillance, tracing, etc., discussion about privacy issues and network monitoring became more important within the last years not only on the technical but also on the political level.
The goal of this deliverable is to analyse the state of the art on monitoring applications with relevance for the PRISM project. It mainly focuses on monitoring applications based on passive network monitoring on a single link, due to the design of the PRISM system. One of the main results is to show which application requires what information from the network or from external data sources.
Although the increasing demands on IPv6, most of the existing measurement applications are only available in IPv4 versions by now. This deliverable does not explicitly differentiate between the use of IPv4 and IPv6 protocols.
In order to follow a structured approach, we introduced four monitoring application domains within this deliverable:

a) Performance monitoring, performed by service providers and network operators, but also private organisations in order to ensure and evaluate the network service.
b) Intrusion and anomaly detection and prevention performed by professional end-users or network operators in order to secure their networks from external intruders.
c) Traffic classification, performed by network operators, in order to provide different service levels for different applications.
d) Lawful interception, performed by law enforcement agencies in order to preserve national security and combat serious criminal activities.

A considerable amount of measurement tools already exist [SCH06], many of them build on passive link monitoring and packet capturing. In this deliverable the goal is not to collect and compare all of those tools but limit us to selected representative monitoring tools and applications, for each of the above mentioned four monitoring application domains.
In order to classify the collected data regarding their privacy issues, we introduce five data categories:

1. *Packet data*: any field from layer 2, 3, 4 data or payload of packets seen on the wire
2. *Measurement data*: based on the results of the measurement but not on bits from inside the packet data
3. *Task Meta-data*: any setting used for performing the measurement task and information about the measurement setup
4. *Derived data*: based on statistical analysis of the *measurement data* and/or *packet data*
5. *External data*: required by monitoring applications to perform analysis and evaluation of the *measurement data* and *packet data*.

Data can appear in different data formats, which can in principle be applied to all of the above data categories. However, not all possible combinations are useful in practice. While format conversions can be applied, information can be lost, depending on the conversion.
We differentiate between three basic data formats:

a) *Plaintext*: data in ASCII or binary form.
b) *Encrypted*: data encrypted by a cipher
c) *Anonymised*: data where private information about users is reduced, by deletion or a one-way function

Note, that the above definition does not imply any technical format, like pcap or erf.

The deliverable is structured as follows. Chapter 2 provides an overview on selected monitoring application domains, each followed by a matrix to show, which monitoring applications rely on which monitoring information. Chapter 3 describes how operational monitoring is currently implemented, and therefore delivers the starting point for the selection of useful monitoring applications. Chapter 4 afterwards describes the state of the art of selected monitoring applications, structured into the above introduced application domains, including a selection of applications and tools available for each application domain. The conclusions drawn in Chapter 5 provide input for the requirements analysis and architecture specification of the PRISM project.

2 Overview of Monitoring Applications

This chapter should give a brief overview on monitoring applications where the focus is on the question what privacy relevant information is needed to perform analysis on the monitoring results. Nowadays monitoring is used by networks providers, where often user privacy is not taken into consideration (e.g. analysis based on full packet payload). The four identified monitoring application domains will be described in more detail in the following subchapters. For each monitoring application domain an overview table identifying (privacy sensible) required data for selected monitoring results as well as a rough estimation of the computation demand is provided. The computation demand is stated as low (l), medium (m) or high (h). The requirement of the data for the monitoring result can be optional (o) or required (x).

2.1 Performance Monitoring

Performance Monitoring covers the monitoring applications to continuously monitor the status of the network, to enable and improve network operation and planning. Can be done either online (in real-time) or offline based on result aggregations.[1]

- *Link utilisation*: To identify bottlenecks to detect network limitations to support network planning or routing optimisation.
- *Protocol shares, Protocol load shares*: To get information about actual shares of transport layer protocols, based on number of packets or bytes.
- *TCP/UDP port number shares*: To get rough estimation about running applications. For more detailed analyses, traffic classification algorithms (e.g. based on deep packet inspection) are necessary.
- *Link utilisation per source/destination IP address*: What customer gets what amount of bandwidth? To identify misbehaving users (or virus/worm infected machines).
- *IP traffic matrix*: To identify the source-destination pairs of the monitored traffic.
- *Packet size distribution*: To know what packet sizes are on the network.
- *AS traffic matrix*: To identify possible new peering partners, if large amounts of the bandwidth go or come from the same source or target AS. Can be differentiated into source, destination or transit traffic.
- *RTT estimation*: To identify end-to-end performance problems, based on TCP SEQ/ACK analysis. Bidirectional traffic monitoring is required for such estimations.

[1] The selected applications can be done on the basis of passive monitoring at a single link, while operational performance monitoring usually includes also other information sources (e.g. router MIBs requested via SNMP), as described in the chapter 3.

Table 1: Matrix to display required data for performance monitoring applications

Data required / Monitoring purpose	computational demand	pktsize	protocol field	src/dst port	src/dst ip	seq./ack.nr.	tcp options	timestamp	link capacity	IP->AS Mapping
UDP/TCP port number shares	l	x	x	x				o		
AS traffic matrix	m				x			o		x
Link utilization	l	x						o	x	
Link utilization per src/dst	l	x			x			o	x	
IP traffic matrix	l				x			o		
Packetsize distribution	l	x						o		
Protocol load shares	l	x	x					o		
Protocol shares	l		x					o		
RTT estimation	m					x	o	o		

Fehler! Verweisquelle konnte nicht gefunden werden. shows the required data for different performances monitoring purposes. The timestamp is always optional. It is not required when the monitoring is performed live. Contrary, when performing a post analysis where absolute time values are necessary the timestamp is required. Sometimes, a relative timestamp might be sufficient. The computation demand of performance monitoring applications is low up to medium.

2.2 Anomaly and Intrusion Detection and Prevention – (A)IDS/(A)IPS

In the presence of increasing threats to network security, Anomaly and Intrusion Detection/Prevention Systems have nowadays become indispensable tools for network operators in order to detect ongoing malicious activities and to prevent them from causing harm either to their own infrastructure or to their customers' networks and systems. A fundamental enabler for (A)IDS/(A)IPSes are of course traffic measurements which allow for the inspection of data before it reaches its targets. In this sense, we can subdivide the data required for such systems into two groups:

1. Those directly obtainable from the traffic traces without stateful processing of packets (e.g., network layer header, transport layer header, and application payload signatures),
2. Those which can only be obtained during the measurement process (like e.g., the timestamp) or by post-processing and though outside sources (like e.g., flow data and other metadata).

Accordingly, the different A(IDS)/A(IPS) can employ all of these data in different combinations in order to fulfil their task, as shown in **Fehler! Verweisquelle konnte nicht gefunden werden.**. For more detailed information about the currently most widely used (A)IDS/(A)IPSes, please refer to Section 4.2.

Table 2: Matrix to display required data for (A)IDS/(A)IPS systems

Data required (typically!)	computational demand	network layer header	transport layer header	application payload	timestamp	flow data	other metadata
Anomaly and Intrusion detection based on:							
Network/transport layer header (ip/tcp/udp/icmp)	l	x	x		o	o	o
Application layer (tcp/udp payload):							
- deep pkt inspection, signatures	m	o	o	x	o	o	o
- machine learning techniques, neural networks	h	o	o	o	o	o	o
- statistical methods, Bayesian analysis	h	o	o	o	o	o	o

2.3 Traffic Classification

Network traffic classification is a fundamental component to several activities carried out by the operators such as network security monitoring, accounting, and for forecasts for long-term provisioning. Typically, the traffic classification is performed on flow level basis, i.e., the core task is to identify the type of application responsible for a given UDP or TCP flow. To this end, the traffic classification algorithm is provided with the headers of corresponding packets and in ideal case also the full payload, which in particular leads to privacy concerns. Due to numerous new applications encrypting the transmitted data, the traffic classification is not a trivial task even if the payload is available. Several different approaches for the traffic classification have been proposed, e.g., machine learning techniques (SOM, neural networks), Bayesian analysis and other statistical approaches, and also signature based methods looking for particular patterns in the payload.

Table 3: Matrix to display required data for traffic classification applications

Data required (typically!)	computational demand	data link header	network layer header	transport layer header	application payload	timestamp
Traffic classification based on:						
Data link layer header (e.g., ethernet)	l	x				o
Network/transport layer header (ip/tcp/udp/icmp)	l		x	x		o
Application layer (tcp/udp payload):						
- deep pkt inspection, signatures	m	o	o	o	o	o
- machine learning techniques, neural networks	h	o	o	o	o	o
- statistical methods, Bayesian analysis	h	o	o	o	o	o

2.4 Lawful Interception

Lawful Interception refers to the functionalities that were once referred to as "wire tapping". It concerns the procedures followed in order for some law enforcement officials to be granted with access to communications-related data by the corresponding network or service providers. While formerly the interfaces for connectivity between the providers and the law enforcement entities were being assembled on a case-by-case basis, during the last few years, the Lawful Interception means have been subject to standardisation. Especially in the European area, the European Telecommunications Standard Institute (ETSI)[2] is in charge of standardisation, while a regulatory framework has been developed in EU level (e.g., [EUR96], [EUR06]); therefore, this document relies on the ETSI standards and the EU legislation.
Two different Lawful Interception modes are identified: interception of the Content of Communication and interception of the so-called Interception Related Information[3]. In addition, strongly related to Lawful Interception, is the collection of data for the purpose of enforcing the European provisions regarding the obligatory retention of electronic communications' data.
Summarising, Lawful Interception may serve the following purposes:

- Interception of the Content of Communication.
- Collection of Interception Related Information.
- Collection of data for the enforcement of the data retention regulatory provisions.

The following Table summarises the data that are required for the execution of the three Lawful Interception scenarios.

Table 4: Matrix to display required data for Lawful Interception applications

Data required / **Monitoring purpose**	Identity	Content of Communication	Interception Related Information	Network or service address	User ID(s) allocated	User IP address	Data & time of log-in and log-off	Service identifier	Identities of peer users	Equipment identifier
Interception of Content of Communication	x	x	o							
Interception of Interception Related Information	x		x							
Data retention of communication data	x			x	x	x	x	x	x	x

Intercept Related Information shall contain in the general case:

- The identities used by or associated with the target identity, that is, the user[4] subject to interception.
- The identities that have attempted communications with the target identity, successful or not.
- The details of services used and their associated parameters.
- Information relating to status.

[2] http://www.etsi.org
[3] The terms "Content of Communication" and "Interception Related Information" are part of the ETSI Lawful Interception terminology.
[4] The term "user" refers here to the definition provided in [EUR06]: "any legal entity or natural person using a publicly available electronic communications service, for private or business purposes, without necessarily having subscribed to that service".

- Time stamps.

However, it should be stressed that the exact types and nature of the data comprising the Interception Related Information differs on a case-by-case basis, based on the type of the service in question. For instance, in the case of an e-mail send event, the following data constitute the Interception Related Information [ETSI102233]:

- Server IP
- Client IP
- Server Port
- Client Port
- E-mail Protocol ID
- E-mail Sender
- E-mail Recipient List
- Total Recipient Count
- Server Octets Sent
- Client Octets Sent
- Message ID
- Status

Regarding the enforcement of data retention, it should be noted here that there are additional data types that are requested, such as the name and the address of the user, as well as her/his location. However, these data types are not directly collected through the network monitoring procedure but require additional "back-office" processing for their generation.

3 State of the Art in Operational Monitoring

Operational monitoring represents an important issue for service providers: it allows revealing eventual failures of the network, to verify the compliance of users and ISPs with the Service Level Agreements and to provide accounting and billing functionalities. For this reason, monitoring requires to retrieve information concerning several different variables of a network (routing tables, NAT status, etc.). However, this project is mainly concerned with the design of a packet capturing system, and, therefore, with those variables which are measurable on the basis of traffic traces.
Several monitoring applications are available, but, even if proprietary solutions are adopted in some cases, most of them use standard protocols in order to gather information from the network devices.
In the following we analyse the kind of data which can be conveyed by such protocols and briefly discuss the privacy concerns involved. In particular, we will examine the possibility of disclosing sensitive information by exporting trace related data to a third party management application through the standard protocols.
Subsequently, we will discuss some relevant examples for different classes of monitoring applications, in order to provide an overview of how operational monitoring is performed by operators and network managers within different tiers and scenarios.

3.1 Protocols for gathering monitoring-related information

3.1.1 SNMP

One of the most well known protocols for monitoring information gathering is SNMP [STA99], which is often supported even by lower class devices. Several network monitoring tools, such as HP OpenView and OpenNMS make use of information gathered through the SNMP protocol.
SNMP is based on a large variety of MIBs (Management Information Base) which define in a standard way the information which is retrievable from each network device.
Each SNMP MIB is associated to an Object Identifier (OID) which is used to refer to the specific information base during queries and responses. OIDs are arranged in a structure of management information (SMI) tree defined by the SNMP standard.
Standard MIBs (defined in the MIB2 RFC) allow to export the amount of packets/bytes in ingress/egress on each interface of the router and to convey cumulative information concerning layer 4 protocols.
Each SNMP MIB has a particular access mode (read-only, read/write, write-only) defining the actions allowed to a user. Each user is associated to a given community, identified by a unique string.
Furthermore, it is possible to configure (SNMP trap) a network device in order to send a report to the network manager if a given condition is verified (e.g. interface or link fault).
In addition it is possible to access a list of the different IP addresses which have been “seen” by the device, but no correlation between a couple of addresses is provided; in an ordinary scenario, therefore, it is not possible to reveal a transaction between two given hosts. As a consequence, only minor privacy issues appear.
On the contrary RMON, an extension of SNMP, is appositely conceived to export data from probes, thus conveying consistently more detailed information:

- the cumulative packet/byte count for a given host
- the cumulative packet/byte count for a given sender/receiver pair
- the cumulative packet/byte for a group of hosts

In addition, RMON supports the definition of filters in order to actually capture packets conforming to a given pattern. Thus sensitive information can be disclosed, since transactions

between a pair of hosts can be revealed. In addition, a further concern for privacy arises from the possibility of exporting entire packets, including the payload, which can possibly convey private information.
Finally it should be mentioned, the RMON functions are normally putting a high load on the devices if enabled, because they are not embedded into the standard ultra-fast packet forwarding path. Therefore these monitoring functions are mostly used only reactively in the operational networks.

3.1.2 Cisco NetFlow

The Cisco NetFlow protocol conveys more detailed information about traffic flows than SNMP. For each traffic flow seen by the router, NetFlow exports:

- IPv4 source address
- IPv4 destination address
- Source transport port
- Destination transport port
- IP protocol identifier
- IP Type of Service
- Router or switch interface

Cumulative packet and octet counts and flow start and end timestamps are also exported for each flow. Though on its face only the source and destination address would appear to be personally identifiable, all fields in this data model can potentially disclose personally identifiable information: for example, it is possible to determine that a given user started an HTTP session with a given server at a given point in time.
Cisco's Flexible NetFlow, which is based on NetFlow, defines different kinds of traffic flows, each of them associated to a different set of per-flow information (e.g. multicast flows, security flows, peering flows). In particular, security flows are of interest here in that they support the export of payload sections in order to allow deep packet inspection. This represents potential disclosure of further personal information.
NetFlow is implemented by Cisco routers, and a variety of open source tools use it as a de facto standard flow interchange format.

3.1.3 IPFIX and PSAMP

The IETF has recently standardised the IP Flow Information export (IPFIX) protocol, which is based on Version 9 of Cisco NetFlow. IPFIX can be seen as a generalisation of NetFlow V9. It allows flexible and extensible definition of flow types and information elements within those flows, and defines a rich set of standard information elements, including:

- "Five-tuple" (IPv4, IPv6) and standard counters
- Packet treatment: e.g., routed next hop and AS
- Detailed counters: e.g., sum of squares, flag counters
- Timestamps down to nanosecond resolution
- Any ICMP, TCP, UDP header field
- Layer 2, VLAN, MPLS, and other sub-IP information

Theoretically, IPFIX can be used to export any set of observable properties for a flow. This expands the opportunities for monitoring applications – indeed, IPFIX is designed without a single application or vendor implementation in mind – but also increases the complexity of protecting the exported data, as the number of potential record types is basically unlimited.
PSAMP extends IPFIX to add export of sampled packets and meta-information about the sampling techniques used. It can be used instead of or in conjunction with IPFIX flow export. Since it enables the export of arbitrary sections of raw packet headers or payload, it has a privacy impact equivalent to full packet capture.

As IPFIX and PSAMP are emerging protocols, they will see implementations by the vendor community in the coming year or two. Several open source monitoring tools already implement IPFIX, and there are two interoperable open source libraries, as well.
Note that flow monitoring is generally deployed within larger networks, as it is more resource intensive than SNMP. Flow monitoring either requires more powerful routers or dedicated devices to observe network traffic at a set of observation points, owing to the greater inspection per packet and the state requirements per flow.

3.1.4 sFlow

sFlow[5] is a multi-vendor technology (very scalable and with a low cost in term of memory and CPU) that provides data for network management and control. The main difference, with respect to the other solutions that we have described, is that sFlow makes an extensive use of packet sampling techniques, thus greatly reducing the overall volume of data flows that have to be exported. Such a feature enables tens of thousands of interfaces and links of speeds up to 10 Gbps to be monitored from a single location without impacting the performance of core Internet routers and switches, and without adding significant network load.
For example, such data could be used for:

- Detecting and diagnosing network problems;
- Real-time congestion management;
- Classification of various applications;
- Usage accounting for billing and charge-back;
- Identifying unauthorised network activity and tracing the sources of DoS attacks;
- Route profiling and peering optimisation;
- Trending and capacity planning.

The sFlow Agent is a software process that runs as part of the network management software within a device. It combines interface counters and flow samples into sFlow datagrams that are sent across the network to an sFlow Collector. Packet sampling is typically performed by the switching/routing ASICs.
The differences among sFlow and other widely used solutions, in terms of implemented functionalities, are shown in the Table 5 below.

Table 5: Comparison of sFlow with other techs

	SNMP/ RMON II	**NetFlow**	**IPFIX/ PSAMP**	**sFlow**
Data Exported	counters	flows	flows. biflows, sampled pkts	sampled hdrs (pseudoflows)
Reporting delay	instantaneous	flow expiry	flow expiry, instantaneous (pkts only)	instantaneous
Packet Capture			sampled/full	sampled
Interface Counters	partial		optional	yes
Protocols Measured				
ICMP/UDP/TCP	yes	yes	yes	yes
IPv4	yes	yes	yes	yes
IPv6	yes		yes	yes
Ethernet/802.3	yes	partial (V7 only)	yes	yes

[5] http://www.sflow.org

Specific Information Exported				
Input/Output Interface		yes	yes	yes
Input/Output VLAN			yes	yes
Host and Network		yes	yes	yes
Address		yes	yes	yes
Routed Next Hop				
Configuration				
via SNMP	yes		read-only	yes
proprietary		CLI	CLI/netconf	yes

3.1.5 Proprietary Protocols

In addition to the standard protocols mentioned above, proprietary solutions for data retrieval can be used for network management. For example, access devices for tier 3 networks (e.g. wireless base stations) run proprietary agents keeping track of traffic count for each connected user. In this case, however, privacy related issues need a specific evaluation for each scenario.

3.2 An overview of present and future monitoring applications

3.2.1 Functional monitoring applications

Applications belonging to this class are adopted in almost all operational scenarios, even in small enterprise or operator networks; many well known applications can be classified within this class: OpenNMS[6], Zabbix[7], Cacti[8], Zenoss[9], Nagios[10] are just examples from a long list of both open source and proprietary applications.

Basically they report the status of the links attached to each switch or router of the network, and, therefore, convey cumulative information about the traffic crossing a link in each direction. Such information can be aggregated at different time scales and is often stored in a round-robin database (RRD). Network managers can therefore easily obtain a time series, at the desired time scale, showing the variation of the traffic load on each link.

In addition, such applications can retrieve further information about the internal state of each network device (CPU load, buffer occupation, etc.) and about the higher level services (http, ftp, DNS) which are made available by each host. In some cases, a topology discovery service is provided as well.

Besides, some of these applications can be configured in order to raise alarms when a given condition is verified; this feature is often based on the underlying mechanism of the SNMP traps.

Most of the open-source applications basically gather data through the SNMP protocol. In this case, as already pointed out, little privacy concerns arise.

However, applications such as HP Openview[11] or IBM Tivoli[12] are designed to accommodate different sources of management related information, such as RMON, NetFlow and even XML messages. In some cases, even data obtained by active measurements (for example Cisco PING or Traceroute) can be processed by this kind of tools, which can also be extended by developing proper plug-ins in order to support other information sources.

In these cases, no a priori remarks about privacy can be made, since it is not possible to limit the amount of information that such applications can convey.

We will thereafter discuss some examples of monitoring applications falling within this class.

[6] http://www.opennms.org
[7] http://www.zabbix.com
[8] http://www.cacti.net
[9] http://www.zenoss.com
[10] http://www.nagios.org
[11] http://www.managementsoftware.hp.com
[12] http://www.ibm.com/itsolutions/servicemanagement

3.2.2 Selected applications

3.2.2.1 OpenNMS

OpenNMS is an enterprise grade network management platform developed under the open source model license. The goal of OpenNMS is to provide a truly distributed, scalable platform covering all the aspects of the FCAPS (Fault, Configuration, Accounting, Performance, Security, as required by the ISO model and framework for network management [CIS07]), and to make this platform available to both open source and commercial applications.
Since OpenNMS is written mainly in Java, it can theoretically run on any operating system: Linux, Solaris, Mac OS X, Windows; thanks to a graphical interface that offers simple and intuitive navigation.
OpenNMS is mainly based on information gathered through the SNMP protocol. In particular the Open.Linkd daemon collects data from the MIB associated to data link level, while Open.collectd collects performance data. When OpenNMS runs, it carries out a discovery procedure, which consists in pinging all the hosts whose address falls within a given range, in order to reveal the active interfaces in the network.
It presents the results of such a procedure through a web based graphical interface, showing a list of nodes which presents a general summary of the node Categories of the network together with their availability percentage over the past 24 hours; Figure 1 shows a typical screenshot of the OpenNMS GUI.

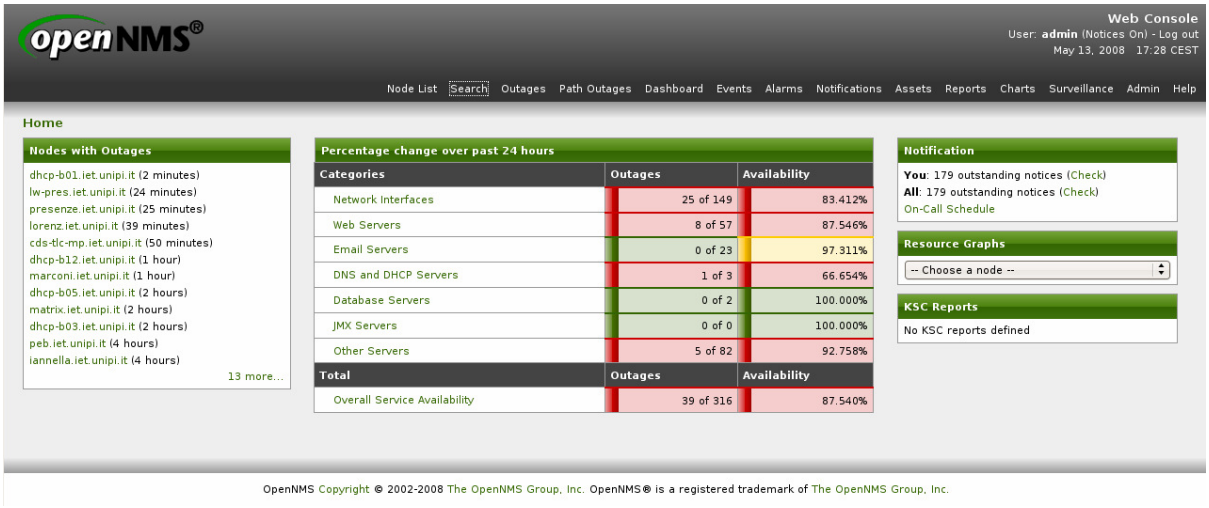

Figure 1: OpenNMS GUI

In particular OpenNMS probes each node in order to reveal the presence of the following protocols and applications: Citrix, DHCP, DNS, Domino IIOP, FTP, HTTPS, HTTP, ICMP, LDAP, Microsoft Exchange, Notes HTTP, POP3, SMB, SMTP, SNMP, and TCP.
The polling process, after revealing the presence of a given service on a network host, verifies its presence periodically.
Furthermore, it is possible to plot the response time for all the nodes in the network, together with the performance information provided by SNMP agents; for every node a complete report of every performance parameter for each interface is shown, together with a description of recent outages and other events.
In particular, for each host running an SNMP agent it is possible to obtain a graph showing the variation of the following performance indicators:

- Bytes In/Out;
- TCP Open connections;
- Current TCP connections;

- TCP Errors and Failures;
- TCP segments;
- Percent Discards;
- Percent Errors In/Out;
- Discards In/Out; Errors In/Out;
- Unicast Packets In/Out;
- In/Out Traffic Utilisation;
- ICMP Response Time;
- POP3 Response Time;
- SMTP Response Time;
- SSH Response Time.

3.2.2.2 Multi Router Traffic Grapher (MRTG)

MRTG[13] is a widely adopted free application for Operational Monitoring of every type of network interface.

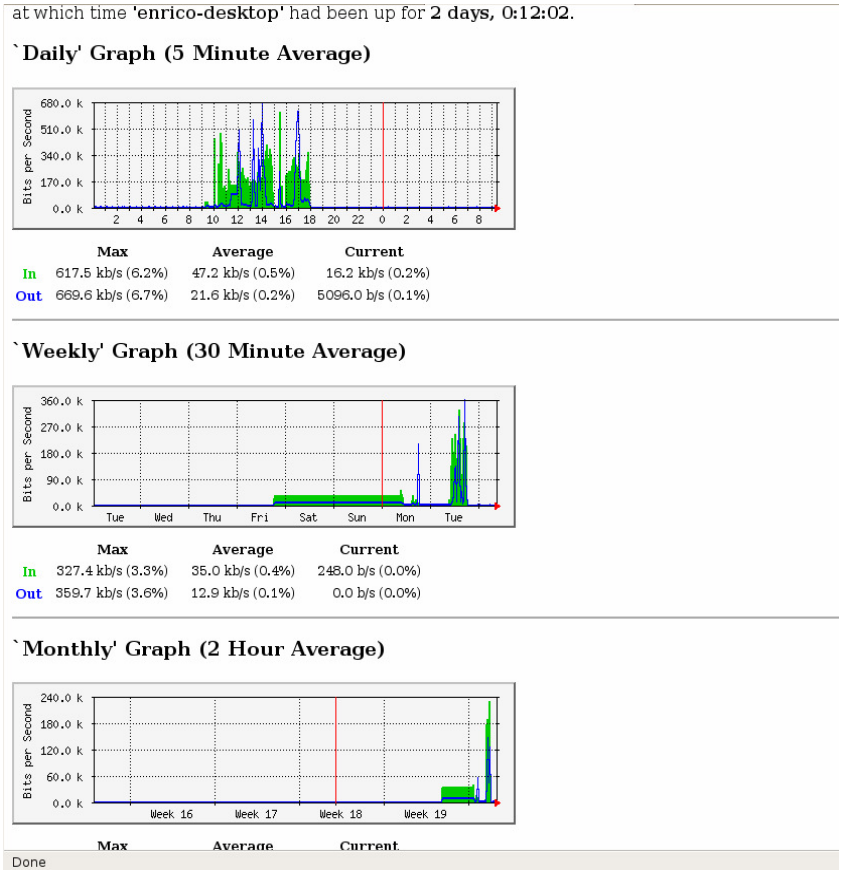

Figure 2: Typical screenshot from the MRTG GUI showing the total amount of in/out traffic for a given interface at different time scales.

This application is SNMP-based and it monitors the network devices implementing SNMP agents, providing cumulative graphs which show the amount of traffic which has passed through each monitored interface.
MRTG is written in PERL and can run over Linux, UNIX, Windows, Mac OS, and NetWare. As a consequence of MRTG being SNMP-based and thus capable of gathering only

[13] http://www.mrtg.com, http://www.mrtg.org

cumulative data for each interface, it isn't possible to reveal a transaction between two given hosts. In particular MRTG shows, for a given network interface four graphs that represent the volume of input traffic and output traffic at different time scales. The first is a "Daily graph", then "Weekly", "Monthly" and "Yearly" graph. Figure 2 shows a typical screenshot from the MRTG GUI.
Moreover it is possible to plot graphs related to Server CPU load, Free Memory, Disk partition usage. Besides, it is possible to configure MRTG to send warning e-mails if target values rise above a certain threshold.

3.2.2.3 HP OpenView

HP OpenView is a commercial software suite which is used for network monitoring. It can be extended to integrate different measurement applications, but its original architecture is composed by the following modules:

- HP Network Node Manager i-series Software
- HP Network Node Manager Smart Plug-ins
- HP Network Automation Software
- HP Process Automation Software
- HP Route Analytics Management System
- HP Live Network

However, in its basic version, OpenView does not encompass all of these blocks; it is based on the HP OpenView Extensible SNMP Agent that allows the monitoring of basic network devices and critical systems and applications. Additional blocks provide other functionalities: for example the HP Route Analytics Management package provides topology discovery capabilities and allows the management system to gather information from routers implementing the most common routing protocols. However, several kinds of third party measurement related applications can be integrated into the OpenView framework; development tools are made available in order to accommodate every source of data.

3.2.3 Multi domain performance analysis

Monitoring applications are generally concerned with the management of a network falling within a particular domain. However, when a communication crosses several administrative domains, the QoS and performance monitoring tasks have to be performed by applications which retrieve information from sources scattered among different administrative units.
Complex QoS/SLA analysis in inter-domain environment is aimed to study, validate and optimise (offline) the network resource usage in order to fulfil the QoS/SLAs required by applications customers. Multi-domain monitoring applications must be able to track where the causes of a performance fault resides. As a consequence, they need to access traffic measurements taken by different agents operating in different domains. For this reason, authentication and access restriction issues are especially sensitive for the correct deployment of such applications.
Another approach to inter-domain QoS analysis is based on active probing and is often adopted when an enterprise network relies on an Internet Service Provider in order to connect different sites, and the provided quality of the service has to be verified.
Such a task is generally performed by the border routers of each site and usually relies on active measurements; probe traffic flows for each type of supported service (VoIP, http, ftp) are generated by one of the border routers and received by a software agent running on the border router of another site. The quality of service perceived by every application can thus be verified without any cooperation of the service provider.
An example of this kind of application is Cisco IOS IP Service Level Agreement[14], which, in turn, makes available the data obtained by active measurement through an SNMP MIB.

[14] http://www.cisco.com/go/ipsla

Passive monitoring is not yet widely used for SLA monitoring, and therefore is not required to be in the main focus of the PRISM project.

3.2.4 Selected applications

3.2.4.1 Perfsonar

An interesting example of multi-domain operational monitoring is provided by the tool Perfsonar[15]. This measurement platform has been designed to be used within the European Research Network GÉANT2[16], which is composed of several national research networks, each one belonging to a different domain.
The architecture of Perfsonar is composed by several components, but we briefly describe here only those which are of interest to our research:
The actual packet capturing is performed by agents called measurement points, which collect data including delay, loss, jitter, flow-based measurements, active stress-type achievable bandwidth measurements, active probe-type available bandwidth measurements and whole packet traces.
Such information is stored by archive service agents and made available to subscribers after authentication by a proper authentication service, which enforces role-based authentication and different levels of trust depending also on which domain the subscriber belongs to.
Before being published, measurement data can be processed by a proper transformation service agent, who can perform several kinds of operations: compression, aggregation, correlation.
Such architecture is conceived in order to locate a fault affecting an inter-domain service (e.g. videoconferencing among different European universities).
Evidently, privacy related issues are involved in this kind of application, and data protection is a constraint that has been taken into account in the system design.

3.2.4.2 INTERMON

The functional components of the INTERMON[17] toolkit are based on a common data base relating topological, measurement and modelling information for different kind of parameters (end-to-end QoS, inter-domain performance metrics, traffic).
The main focus of the INTERMON toolkit is the integration of tools covering different aspects of QoS analysis in large scale Internet environments such as inter-domain topology discovery, QoS and traffic measurement, traffic modelling, QoS prediction, load scenario simulation, pattern and traffic matrix analysis.
In particular, INTERMON include many tools such as:

- Inter-domain route monitoring and quality analysis – InterRoute tool (used to discover the inter-domain path of an end-to-end connection by querying a common inter-domain routing repository and BGP-4 protocol data);
- CM Toolset for proactive end-to-end QoS monitoring and active tracing of connection topology on intra/inter-domain level.
- Traffic measurement tools using IETF IPFIX traffic flow export format (in order to study the impact of particular flows);
- Border router monitoring tools collecting MIB information (Tilab SNMP poller);
- Software for spatio-temporal QoS data mining and QoS pattern analysis for the area of network planning;
- Anomaly detection tools;

[15] http://www.perfsonar.net
[16] http://www.geant.net
[17] http://www.ist-intermon.org

- Inter-domain simulation tools for network planning and management. This set of network simulators can be used to simulate the impact of certain changes to networks monitored by INTERMON. Simulation scenarios are based on measurement data (e.g. input traffic and inter-domain topology) provided by other tools of INTERMON architecture.
- In particular, INTERMON simulation toolkit is composed of:
  - Hybrid packet based simulation (NS2) with integrated analytical models;
  - Rate and time continuous fluid flow simulation based on differential equations (RTC-FSIM simulator);
  - Efficient time series data simulation for inter-domain environment;
- Inter-domain analytical tool.

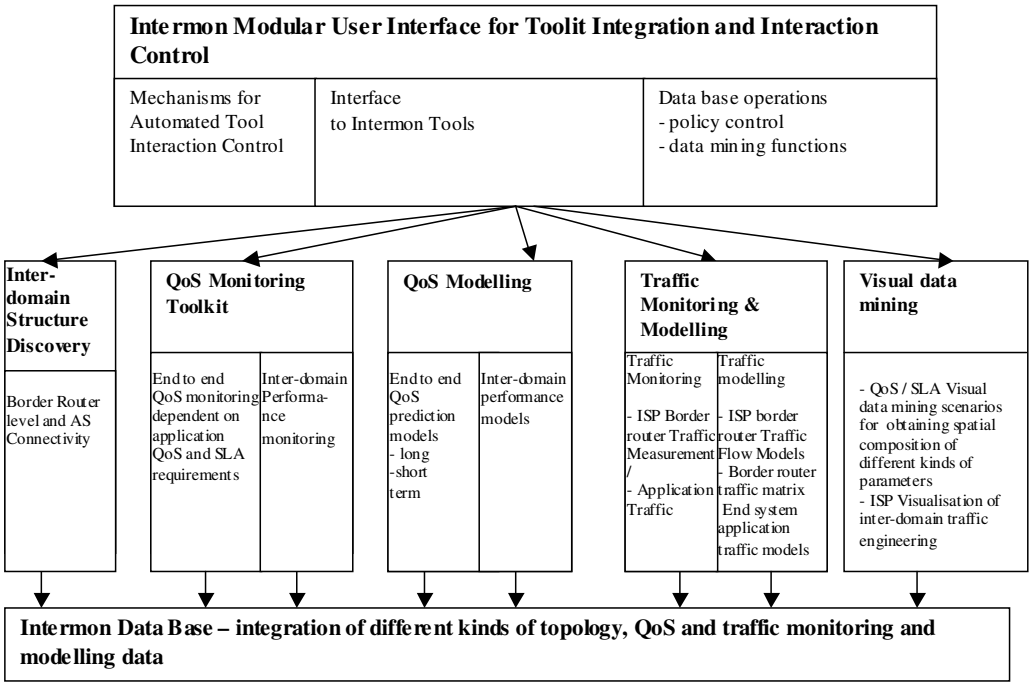

Figure 3: Organisation of the INTERMON toolkit

Moreover, the design of the INTERMON data base provides information for:

- Inter-domain QoS and SLA verification, especially to obtain spatial decomposition of the perceived end-to-end inter-domain performance.
- Inter-domain traffic engineering, based on border router traffic flow measurements and modelling considering different flow granularities.

Thus, using the INTERMON toolkit it is possible:

- To monitor the QoS/SLA automatically and for different aggregation intervals (i.e. INTERMON automated measurement reports per hour, day, week, month, year)
- To provide insights of the impact of different factors on the QoS perceived by an application in different time scales;
- To identify the border router(s) which are responsible for the degradation of end-to-end QoS of applications;
- To predict the end-to-end and inter-domain QoS for different inter-domain routes, by integrating measurement and modelling.

The organisation of the INTERMON toolkit is shown in Figure 3.

3.2.5 Accounting and billing applications

Generally, small operators offer SLAs based on a flat rate pricing model. On the other hand, bigger providers account for service based on actual measurements of the data exchanged through their network.
In order to perform such task, trace of the amount of traffic exchanged for each origin-destination flow has to be kept. NetFlow and IPFIX agents offer an ideal tool to retrieve such information, since they automatically generate and manage a data structure for each traffic flow seen by a given router, and register both the amount of data exchange and the duration of the transaction.
Also RMON agents can constitute a source of accounting related information, in that they can keep track of the per pair data exchanges.
In general, such agents are run by the edge and access router of a network, while they are not deployed in the core, where further processing requirements could affect the performance.
Billing applications generally gather this kind of data and analyse them at different levels of aggregation (by organisation, by location).
Furthermore, billing applications generally need to associate an IP address with the corresponding customer or organisation: this involves correlating data issued from traffic analysis with information retrieved from other sources, such as DHCP and RADIUS servers. Such a process definitely involves a considerable amount of sensitive information.

3.2.6 Selected applications

3.2.6.1 Evident enterprise

Evident Enterprise is a widely deployed commercial solution for billing.[18] Such a tool, after gathering transaction related data from various sources, exports them using a unique format, called "Usage Data Record" (UDR), representing a single "measurement event" or a single usage transaction. Such a format registers detailed information about each data exchange, including:

- source address,
- destination address,
- protocol metered,
- time period,
- amount of bytes transferred,
- source location,
- destination location,
- owner of source address,
- owner of destination address.

Since not all of this information is included in the traffic packets, its retrieval may require interaction with other network services, such as DHCP and directories such as Active Directory and LDAP.
UDR information is then normalised, correlated, aggregated and used for the actual pricing.
From a look at the UDR structure, it is evident that a great deal of potentially personal information is involved in the billing operation. Data protection issues are therefore very relevant for this purpose.

[18] http://www.evidentsoftware.com/products/enterprise.aspx

3.2.6.2 XACCTusage

The XACCT platform[19] is a foundation technology that provides a single point of interface between the operation and business support systems of network service providers and the physical network.
Several big enterprises adopt XACCT for the management of their networks: among them are Bell Canada, British Telecom and Siemens.
XACCT Technologies Inc. introduced support for a wide range of open industry standard-based network elements and it accommodates different sources of traffic data, such as RMON probes. The XACCTusage platform enables Service Providers and enterprise customers to use this information for a variety of business support applications such as network/application/user profiling, Quality of Service (QoS) monitoring, and billing.
XACCTusage collects in real time and aggregates network usage and traffic data from network elements such as routers, switches, firewalls, servers and gateways and synthesises that data into the formats required by the operations and business support systems of network service providers.
In particular, this platform it is possible to support:

- Billing record creation and account provisioning
- Service-specific pricing models
- Network resource planning
- QoS Metering
- Self-care through user-account provisioning

The architecture of the XACCT system is composed of basically three building blocks, which are introduced in the Table 6 below:

Table 6: XACCT system building blocks

XACCTusage	Core system: multi-layer, service, and vendor data collection and enhancement
XIS - XACCT Interface Server	Interface for flexible bi-directional data transport and formatting.
ISMs - Information Source Modules	Network element-specific data collection and provisioning modules

In particular, XACCTusage is in charge of the data collection process: it captures the Internet protocol transaction information produced and logged by the individual network elements and, by processing them in real time it transforms such raw data into meaningful business information. It collects traffic information from a variety of device-specific, as well as general purpose, software agents called information source modules, or ISMs.
These software modules provide an interface for access to data collected from different network elements such as routers, switches, firewalls, authentication servers, and lightweight directory access protocol (LDAP) servers, domain name servers (DNS), web servers, email servers, video servers, voice over IP gateways and hundreds of other network elements. Indeed, the ISMs provide usage information from all layers of the network, from the physical layer to the application layer.
The XACCT Interface Server (XIS) extends the functionality of the XACCTusage platform, acting as a gateway between the platform and the business applications, thus enabling direct integration with external applications. The XIS can be configured in order to export data in any format. Additionally, XIS can also be configured to offer interface functions to the client applications.

[19] http://www.xacct.com

3.2.7 Traffic classification

Many network operators often configure their routers in order to classify traffic related to particular network services, such as VoIP and peer to peer file sharing; such traffic, after being classified, can be subject to particular shaping or queuing policies.
Since, in most cases, such services dynamically allocate port numbers, the classification is often based on a stateful inspection of the packet payload, in order to detect particular signatures. In many aspects, this task is very similar to performing intrusion detection and presents the same privacy related issues, which will be thoroughly discussed in the next sections.
An example of this kind of classification application is Cisco Network-Based Application Recognition NBAR [CIS05], which is available on many Cisco router models and can classify, among others, traffic flows generated by Gnutella and Citrix.
Another interesting system belonging to this class, and especially conceived for peer-to-peer traffic classification, is provided by XACCT, as described above in Section 3.2.6.2 [GIV03]
This solution is primarily based on XACCT PacketSight, that it is a carrier-grade software-based application-level network probe that typically runs on Sun servers.
If applied for peer-to-peer communication protocols, PacketSight can, for example, decode the KaZaA protocol and identify all the communications and generate appropriate events. This is a method to generate a log of all P2P file searches, requests, and actual transfers. PacketSight is highly modular. In order to analyse and process new protocols, it uses a modular approach that currently supports over 750 protocols/applications. Each protocol/application has its independent State-Based Decoder (SBD). When protocols are added or evolve over time (new version), XACCT analyses the changes or specification of new protocols and develop appropriate SBDs to adhere to the latest industry practices.
The other major piece of the technology employed by the XAACT DATMS (Digital Asset Transmission Monitoring System) is the XACCT "Network to Business" (N2B) Platform.
This platform is a carrier-grade, scalable platform that performs arbitrary data collection, data storage and gathering in a distributed fashion. By using the XACCT N2B Platform, it is possible to dynamically perform real-time lookups related to information contained in the log of events generated in real-time by PacketSight.
Another important capability of the XACCT N2B Platform is the ability to perform arbitrarily complex filtering and aggregation on the events produced by PacketSight.

3.2.8 Traffic monitoring system for mobile 3G networks

An example for a traffic monitoring system for mobile 3G networks is the METAWIN system[20] originally developed by the Telecommunications Research Center Vienna (ftw.). It provides extensive traffic monitoring capabilities tailored for the today's mobile GPRS/UMTS networks. The traffic monitoring system relies on non intrusive packet capturing methods, i.e., passive network monitoring. Based on traffic traces, network performance can be evaluated, network troubleshooting can be carried out, and models for user- and control plane traffic can be derived.
The METAWIN traffic monitoring system can be attached to different parts of the mobile network including all 3G specific interfaces such as Gi, Gb, Iups and Gn links at the SGSNs and GGSNs. Thus, in addition to monitoring the end-user traffic, the system allows the operator to monitor signalling traffic in the control plane, and to carry out cross-layer analysis, e.g., to correlate abnormal events the signalling/user plane. Due to the nature of the mobile 3G networks, also the traffic monitoring system is distributed to the given area. The time

[20] The first components of the METAWIN system were deployed in 2004 and it has been evolving since then (similarly as the whole field of 3G mobile networks). Currently the system has already been deployed in large scale covering the most aspects of the today's 3G networks. The original system was developed by the ftw., while the current development is carried out in co-operation with Kapsch CarrierCom.

synchronisation is based on GPS signals, which allows, e.g., accurate one-way delay measurements between different components of the 3G network.
The privacy aspects have been carefully taken into account in the design of the METAWIN system. For example, the packet payloads are not recorded in the packet traces. Similarly, for network troubleshooting purposes, the system is capable of parsing the GTP layer on the Gn interface allowing tracking the establishment and release of each PDP context, and consequently to uniquely identify the mobile subscriber sending or receiving each packet. In order to protect the user's privacy, the individual identifiers are chosen as arbitrary strings, decoupled from the real identity (anonymisation).

3.2.9 Network-based IDS/IPS

Over the last decades, end-user equipment (i.e., the Windows PCs and laptops) was usually equipped only with antivirus software which was able to detect malware (i.e., viruses, worms, trojans, spyware, etc.), followed by a broad default activation of inbound traffic firewalls on the end-systems with the introduction of Service Pack 2 for Windows XP only in 2004, complementing the base of the already installed NAT devices and other firewall software[21]. While the combination of antivirus and firewall software, if maintained and operated properly, usually indeed offers a reasonable level of protection against Internet threats, it does not completely cover all possible vulnerabilities of an end-user system. A good example for such malware entry points are newly discovered browser bugs and the corresponding exploits, which potentially open the system to malicious intruders in the time interval until they are patched. Therefore, in recent years a multitude of malware protection software vendors have integrated Intrusion Prevention Systems (IDSes) with their antivirus and firewall products; these new software packages are usually named Internet Security suites, and the IDS functionality therein basically offers an on-the-fly scanning of traffic for known malware/attack patterns. An independent comparison of Anti-Virus software suites is maintained by AV comparatives[22].
In order to alleviate the danger of security software misconfiguration or a lack of maintenance, the latest trend is to move the per-customer (i.e., per Internet access) traffic scanning to the network operator's domain, and in this way to assure the highest achievable level of protection for the customers. A recent example of such a system deployment is the A1 Internet Security by the mobile network operator Mobilkom Austria[23], which offers network-based Internet protection in combination with a client Internet protection software[24] which is to be installed on customers' Windows computers. In order to maximize the effectiveness of the system in the presence of increasingly encrypted traffic in the Internet, the A1 Internet Security system, e.g., also scans HTTPS connections, and in this way enables the same level of protection even for encrypted session.

[21] http://www.microsoft.com/switzerland/windows/de/xp/sp2/default.mspx
[22] http://www.av-comparatives.org/
[23] http://www.a1.net/privat/internetsecurity
[24] http://www.ikarus.at/history/a1.html

4 State of the Art on Monitoring Applications

In this chapter, the monitoring applications are described, structured into four main application domains. For each application domain, its role and context is described with some indication on usage scenarios. An extra section for each application domain is dedicated to raised privacy concerns and possible countermeasures. For each application domain one or more tools or applications relevant to the PRISM project have been selected and described.

4.1 Performance Monitoring

4.1.1 Role and context

In order to ensure a certain service level or minimum performance, monitoring is a crucial task of network operation. Furthermore, performance monitoring provides important information for network planning and maintenance. Thus the main users of performance monitoring are network operators, ISPs, professional end-users (to monitor the delivered service performance), but also network planners.
Service level agreement (SLA) monitoring is related to a certain service requester, i.e. customer or peering partner and thus necessarily includes data that directly points to a certain user. This possibly involves information about individual users.
Performance monitoring for network operation aims to provide valuable information about the networks status. This information helps with finding bottlenecks and unused resources. Another important aspect is the identification of troublesome or misbehaving users, e.g. sources of viruses and spam emails. Further details are covered by Section 4.2 about anomaly and intrusion detection.
Finally, performance monitoring is required to retrieve basic input for network planning. In this context, the findings are not only used to identify bottlenecks or unused resources, but are further analysed in order to provide suggestions for network optimisations. This can be either upgrading of the existing network or the installation of new routes.

4.1.2 Functionalities

Performance monitoring is performed with various levels of data detail, depending on the focus. In the following we break it down into three main levels of required packet data details:

- No packet details required (only general network usage statistics)
- Single packet details required (evaluation of single packets, independently)
- Multiple packet details required (evaluation over multiple packets, e.g. by aggregation into flows)

The first, most basic functionality of performance monitoring is to observe transmission speeds and volumes. These values are the basis for the analysis of the network status and give rough information, if the network is used to full capacity or still has some reserves left. Aggregated values of all simultaneous flows are sufficient for this kind of evaluation and thus, hardly any privacy critical data is necessary. Such kind of information is usually captured by requesting MIBs from the router using SNMP, but can also be provided by the PRISM system.
In a more advanced manner, the observed values can be split up into their shares by e.g. source or destination IP/IP range or layer 3 and 4 protocols. As an example, the knowledge about source or destination IP addresses is useful to identify the IP ranges with major contribution to the traffic and optimise the routing accordingly. This requires fundamentally more information than in the first stage and some of it is possibly person-related (such as IP addresses).
At last, parameters such as per-flow performance, which cannot be extracted from a single packet, are calculated or estimated from a number of packets. Such values are of great

importance for troubleshooting but can also reveal potential optimisation potentials. This stage depends on a whole flow and evaluates a lot of information, which might include (parts of) the payload. Thus, this evaluation is very critical with respect to privacy issues.

4.1.3 Privacy Concerns

The privacy concerns for performance monitoring are not only related to end-user privacy, but also about business privacy of the network operator, who has to take care to keep a good position in competition (e.g. the knowledge about 90% loaded links should not be available to competitors or end-users). Therefore providers are very careful about which data is made available for which parties, and even don't inform about which data is monitored on a regular basis, and which is done only in case of troubleshooting.
Depending on the focus of the analysis, the protocol headers and/or the payload are investigated. As significant amounts of this data can be directly or at least indirectly related to individual users, suitable anonymisation techniques have to be applied. The value of performance monitoring data very much depends on the input data, but is usually very flexible, which means, that even highly anonymised (e.g. all IP-Addresses are overwritten with '0') data can be used to gather enough performance metrics for network evaluation (e.g. link utilisation), but will not be suitable to get other metrics like a traffic matrix.

4.1.4 Selected Applications

4.1.4.1 Tstat

Tstat – TCP STatistic and Analysis Tool[25]**:** By the use of RRDtool, Tstat is similar to the well known Multi Router Traffic Grapher (MRTG), which is used for network monitoring and measuring based mainly on data collected by SNMP. In contrast to this, Tstat extracts data from packet traces, by extending the functionality of "tcptrace". Tstat is able to produce all kinds of statistics, ranging from separation of traffic into different VLAN IDs (layer 2) up to classification of traffic into applications (above layer 4). In the PRISM context, Tstat is usable to measure and evaluate the performance and utilisation of the network. Depending on the desired level of detail, Tstat requires at least protocol headers up to layer 4 protocols, but some evaluations use (parts of) the payload as well, and can work on both unidirectional and bidirectional traffic. Tstat is an open source tool and can be downloaded from http://tstat.polito.it

4.1.4.2 PasTmon

PasTmon – Passive Application Response Time Monitor[26]: PasTmon is a TCP/IP passive network application response time analysis tool. PasTmon is based on the libpcap an can be used to measure the response time of web and application servers (e.g. mail, http). It is designed for system administrators to measure the performance of their application servers and operates on the server end of the network. Historical data of the measurements are stored in a backend database. PasTmon is an Open Source Tool.
In the context of PRISM PasTmon can be used as performance measurement application for example to measure server SLA conformance. For detailed results the payload or parts of the payload of the packets is needed, but even with anonymous data fundamental server performance monitoring can be performed.

[25] http://tstat.tlc.polito.it
[26] http://pastmon.sourceforge.net

4.2 Anomaly and Intrusion Detection

4.2.1 Role and Context

The central task of Intrusion Detection Systems (IDSes) is, as the name suggests, detecting attacks and unwanted manipulations of computer networks and systems. In addition, Intrusion Prevention Systems (IPSes) have the role of not only detecting such attacks, but also preventing those attacks form causing harm. Complementing the traditional set of IDS and IPS systems, which are based on pattern or signature recognition, Anomaly-based Intrusion Detection Systems (AIDS) are based on heuristic methods which monitor system behaviour and issue alerts (or in the case of AIPS, actions) whenever they detect important deviations from regular behaviour patterns.

4.2.2 Functionalities

Intrusion prevention and detection, as well as anomaly detection and prevention can be performed at various levels in different system architectures. The most common classification of these systems in IP networks is into *network based* and *host based* (A)IDS/(A)IPS systems.
As far as network based systems are concerned, they are usually placed at the edge of company/corporate or university networks, and often their functionalities are integrated into Network Address Translation (NAT) and firewall equipment. Host based (A)IDS/(A)IPS are software components running within the end devices, i.e. clients and serves, monitoring the traffic, data patterns and signatures, and possibly even behavioural patterns of the host machines.

4.2.3 Privacy Concerns

Some (A)IDS/(A)IPS which are used for the recognition and prevention of security violations and attacks tend to record a lot of data which is either related to individuals, or which can at least indirectly be associated to them [BSI02]. In this sense, the following (A)IDS/(A)IPS records are critical:

- Unauthorised data access attempts,
- Unauthorised application access attempts,
- IP addresses or names of hosts/domains which have launched attacks.

On the other hand, one of the most interesting pieces of information these systems potentially offer is the source of the intrusion or security violation, as this data might be crucial for preventing or reacting to such events appropriately.
Therefore, based on the application scenario, one can imagine different types of data to be stored and made available for evaluations. For example, if an intrusion detection system is placed at the border of a private network, and if it is configured in such a way that it prevents incoming security violations and attacks, then it is most likely safe to store detailed data such as IP addresses, as they do not directly point to individual users, but at most to external domain names through reverse DNS lookups.
On the other hand, if (A)IDS/(A)IPS systems also monitor the company/corporate or university network itself, then a multitude of different legal and ethical aspects must be taken into consideration. In such cases, based on the concrete legal situation and the contracts of the individuals with their respective institutions, the acquired data must be carefully analysed and stored, and possibly also far reaching anonymisation techniques have to be applied.

4.2.4 Selected Applications

4.2.4.1 Snort

Snort[27]: As the principal example (A)IDS/(A)IPS we have chosen Snort as an open source network intrusion prevention and detection system which utilises a rule-driven language. According to its website, Snort combines the benefits of signature, protocol, and anomaly based inspection methods. Based on its GPL licensing and a very active and large user community, Snort is reportedly the most widely deployed intrusion detection and prevention technology worldwide and has become the *de facto* standard in this field.
Snort can be configured to run in several basic modes:

- *Sniffer mode.* In this mode, Snort simply reads the packets of the network and writes them in a continuous stream on the console.
- *Packet logger mode.* In packet logger mode, Snort writes the observed traffic to a hard disk.
- *Network Intrusion Detection System (NIDS) mode.* This is the most complex and configurable mode of Snort operation, which allows for the analysis of network traffic for matches against a user-defined rule set and performs several actions based upon what it sees.
- *Inline mode.* In this mode, Snort obtains packets from *iptables* instead of *libpcap* and then causes *iptables* to drop or pass packets based on Snort rules that use inline-specific rule types.

Whereas the functionalities of the sniffer mode are straightforward and self-explicable, it is worth noting that the packet logger mode already offers a wider spectrum of configuration options, like e.g. writing the captured data to disk either in ASCII or in binary format, the latter of which is particularly important in the context of high-speed interface capturing. The NIDS mode of Snort certainly represents the most interesting and widely applied usage scenario of this software, supported by a vast international community of security researchers and open-source enthusiasts. From a practical point of view, especially in environments requiring strong network-based security, the intrusion *prevention* capability of the inline mode which relies on *iptables* is of particular importance, as it enables dropping dangerous packets in real-time, and thus ensures that malicious patterns never reach the end systems in the first place. For more detailed information about Snort, please refer to the Snort User Manual[28].

4.2.4.2 Topaz

TOPAZ (Open source Intrusion Detection System for IPv4 and IPv6) [GEI05]: TOPAZ is one of the first IDSes for IPv6 networks and is available as *open source*. It was developed by Telefonica in the framework of Euro6IX project in FP5. It is used at Telefonica Investigacion Y Desarrollo (TID) laboratories as one of the main tools used to secure its IPv6 network. TOPAZ offers both a text and graphical language to configure new pattern attacks, including specific future IPv6-based attacks. Another interesting characteristic about TOPAZ is that it may detect both IPv4 and IPv6 intrusions, making it ideal for dual stack networks, or networks that are migrating to IPv6.
The TOPAZ system follows the classical guidelines for a Network IDS (NIDS), i.e., it detects attacks against a system by monitoring network traffic and looking for well-known attack patterns matches or anomalies in the protocol interchanges. In order to achieve this, several sensors are installed all over the network, with the objective of providing timely and accurate information about the observed traffic to a central management console.

[27] http://www.snort.org
[28] http://www.snort.org/docs/snort_htmanuals/htmanual_282/

4.2.4.3 Bro Intrusion Detection

Bro Intrusion Detection [PAX99][29]: Bro is an open-source Network Intrusion Detection System (NIDS) running on UNIX systems that passively monitors network traffic and looks for suspicious activity. According to the official descriptions of its functionalities, Bro detects intrusions by first parsing network traffic to extract its application-level semantics and then executing event-oriented analysers that compare the activity with patterns deemed troublesome. This analysis includes detection of both specific attacks (defined by signatures or in terms of events) and unusual activities (e.g., certain hosts connecting to certain services, or patterns of failed connection attempts).
Bro uses a specialised policy language that allows a site to tailor Bro's operation, both as site policies evolve and as new attacks are discovered. If Bro detects something of interest, it can be instructed to either generate a log entry, alert the operator in real-time, execute an operating system command (e.g., to terminate a connection or block a malicious host on-the-fly). In addition, Bro's detailed log files can be particularly useful for forensics.
Bro especially targets high-speed (Gbps), high-volume intrusion detection. By leveraging packet-filtering techniques, Bro is said to achieve the necessary performance while running on commercially available PC hardware, and thus it is supposed to serve as a cost-effective means of monitoring a site's Internet connection. As Bro has been primarily developed as a research platform for intrusion detection and traffic analysis, it is not intended to represent an "out of the box" solution. Bro is designed for use by experts who require the ability to extend an intrusion detection system with new functionality as needed, and thereby track the evolving attacker techniques as well as inevitable changes to the protected network environment and security policy requirements. For more details about Bro, please refer to the Bro User Manual[30].

4.3 Traffic Classification

4.3.1 Role and Context

Traffic classification typically involves analysing traffic flows (a sequence of packets with matching well-known IP five-tuple). The motivation for an ISP to classify traffic can be, e.g., to filter away malicious traffic, or just simply to gain a better understanding of the used application for network planning purposes. In the process of traffic classification each packet or flow is assigned to one of the chosen categories such as voice-over-IP (VoIP), web traffic, etc. For example, a typical scenario could be that an operator would like to know how popular some P2P-protocol is and if it is becoming more or less popular, i.e., what is the trend. Note that typically P2P type of applications have a finite life time in the sense that at some point another (perhaps similar) application is introduced which "kills" the earlier one in terms of overall data volumes.

4.3.2 Functionalities

The traffic classification can be carried out at various levels. For example, the coarsest view on traffic could be to estimate the proportions of different types of IP traffic, that is, e.g., between TCP, UDP and ICMP. At the next level one can study traffic aggregates based on the well-known port numbers per protocol type (e.g., DNS/UDP or HTTP/TCP). Assuming that the traffic volumes in the link are large enough then one can safely say that neither of these discloses any private information.
However, quite often the traffic classification is performed at the flow level granularity. In particular, it is worth noting that the five-tuple used to identify each flow by definition also

[29] http://www.bro-ids.org/
[30] http://www.bro-ids.org/wiki/index.php/User_Manual

identifies both the sender and the receiver of each flow by their IP addresses, and thus the privacy concerns must be taken into account. Once a flow is extracted the next task is to identify the corresponding traffic class. To this end, most applications rely on packet payload inspection, .e.g., by looking for particular signatures, as listed in [KAR04], and on the Snort webpage.

4.3.3 Privacy concerns

The privacy concerns become an issue as soon as the traffic classification is carried out per IP address pair basis. In particular, it is worth noting that it is not always necessary to go into such details, but instead it may be sufficient to know, e.g., the proportion of P2P traffic at certain link. Given that the number of concurrent users in the link is large enough this information hardly reveals any private data to the public.
However, if excluding the aforementioned simple schemes such as traffic classification based on the well-known ports, the basic step is to classify each flow identified by the five-tuple using the deep packet inspection. Thus, the crucial point with regards to the privacy is which entity carries out this task, and if the IP addresses and/or the payload are obfuscated by anonymisation and/or encryption. In particular, from the privacy point of view it would be ideal if the traffic classification can be carried out based on the encrypted data (cipher text) without encrypting the possible sensitive parts during the process (cf. homomorphic encryption).

4.3.4 Selected Applications

4.3.4.1 Appmon

Appmon [ANT06]: One example application of this category is the Appmon developed in the Lobster project[31]. The Appmon scans a live link (or a pcap traffic trace file) and tries to identify the flows belonging to a certain set of the most popular applications. This includes several peer-to-peer applications (e.g., BitTorrent and eDonkey), web traffic, RTSP, FTP, SMTP, and DNS, just to name a few. The identification relies (mainly) on the deep packet inspection, i.e., the packet payload must be available. This puts forth the privacy concerns at the same time. As an output the Appmon gives the uplink and downlink traffic rates per application. Additionally the most active hosts in terms of download and upload traffic are listed. As IP addresses are sensitive information they are anonymised in the public view. However, in the “system operator” view, which protected by a password, the real IP addresses are shown for network troubleshooting purposes.

4.3.4.2 Tstat

Tstat: Another example application for traffic classification is the TSTAT, which has been already described in Section 4.2.4. In contrast to Appmon, the TSTAT is also capable of recognising Skype traffic (as it is at the moment) [BON07][BON08].

Common to all these approaches is that they first identify the flows in the traffic, and then try to classify these flows to different application classes, in general, by payload inspection. However, there are also several proposed solutions which aim to achieve capability to perform traffic classification without the payload inspection. Typically, one tries to apply different machine learning techniques and statistical analysis. See, e.g., BLINC [KAR05] and InFeCT [TEU08].

[31] http://www.ist-lobster.org/

4.4 Lawful Interception

4.4.1 Role and context

Lawful Interception is the legally authorised process by which a network operator or service provider (hereafter, the provider) grants some law enforcement officials with access to communication data (such as, telephone or VoIP calls, e-mail messages, etc.) of private individuals or organisations. Lawful Interception is becoming crucial to preserve national security, to combat terrorism or other serious criminal activities, as well as to investigate these kinds of social mishaps. In the typical case, some Law Enforcement Agency orders to a provider the delivery of the content of communication and/or communication data. The provider should be capable to intercept these data using special equipment and without enabling the interception subject to become aware of the interception, and make available this information to the requesting Law Enforcement Agency.

Lawful Interception applications constitute a very special type of monitoring applications. The field of Lawful Interception constitutes probably the only so strictly legislated area of network monitoring; the providers are being asked to meet legal and regulatory requirements for the interception of voice as well as data communications in IP networks. National and international laws impose rules to the providers to be able to mediate between the network and the law enforcement entities. Providers need to be able to quickly and efficiently identify a target, isolate its traffic and get it to law enforcement entities in a standard, reliable and legally compliant manner. On the other hand, the law enforcement entities should be able to perform detailed analysis of the given information, in order to build the story of the target's activities and interactions.

4.4.2 Functionalities

The functionalities that should characterise a Lawful Interception system are mostly derived by the legislation and are specified to a great extent by international standards. For example, the American National Standards Institute (ANSI) has published the ANS J-STD-025 standard [ANS03], a joint standard developed by the telecommunications industry in the USA devised to meet requirements according to the well known Communications Assistance for Law Enforcement Act (CALEA) law [USC94].

Especially in the European area, a fundamental role is being played by the European Telecommunications Standard Institute (ETSI), which has been providing the communications community with standards related to Lawful Interception for several years and now has a special Technical Committee for this purpose. The ETSI, following the foundational European Council Resolution of January 1995 [EUR96], has published standards covering the whole spectrum of Lawful Interception aspects, from the specification of a general architecture to the requirements for data retention.

The following Figure 4 illustrates the general, high level Lawful Interception architecture, as has been proposed by the ETSI ([ETSI201158], [ETSI102232], [ETSI102233], [ETSI102234], [ETSI101331], [ETSI102656], [ETSI101943], [ETSI101944]). The architecture concerns the interception of both the Content of Communication (CC) and the Interception Related Information (IRI), that is, signalling information, source and destination of the communication, etc.

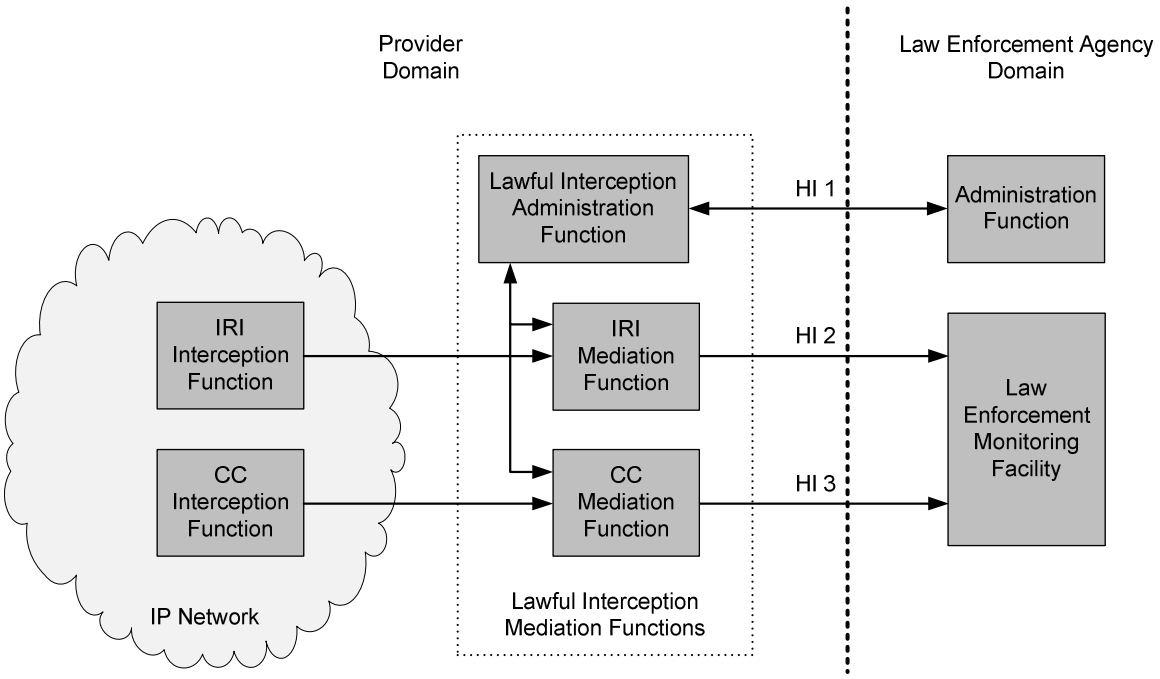

Figure 4: ETSI General Lawful Interception Architecture

From the Figure 4 above, the following high level, discrete but interrelated functionalities are identified:

Interception Related Information Interception Function: The purpose of the IRI Interception Function is to generate IRI information associated with sessions, calls, connections and any other information involving interception targets identified by Law Enforcement Agency sessions.

Content of Communication Interception Function: This is the function that causes the Content of Communication to be duplicated and passed to the Mediation Functions and –finally– to the Law Enforcement Agency.

Interception Related Information Mediation Function: This function is firstly in charge of receiving IRI information related to active intercepts from the IRI Interception Function. Secondly, it correlates and format this IRI in real-time, in order to be delivered to the Law Enforcement Agency, through the corresponding Handover Interface.

Content of Communication Mediation Function: This function is similar to the IRI Mediation Function, targeting the Content of Communication.

Lawful Interception Administration Function: This function administers the requests for Lawful Interception that the provider receives from the Law Enforcement Agency. It ensures that an interception request for IRI or CC or both is provisioned for collection from the network and subsequent delivery to the Law Enforcement Agency.

4.4.3 Privacy Concerns

Lawful Interception constitutes natively kind of privacy violation; in fact, not only some data subject's data are intercepted, but the data subject must by default not be aware of the interception. That is, albeit legitimate, Lawful Interception contradicts to some extent to the personal data protection legislation, since it infringes fundamental provisions. In essence, it constitutes an exemption of the law.

Therefore, the notion of "privacy concerns" in the context of lawful interception has a slightly different meaning and refers to the potential abuse of the legal means that are put in place for legitimate purposes.

The major privacy issue concerns the potential of unauthorised activation of the interception mechanisms. This has been the case in Greece in March 2006, when it has been found out that

a number of cell phones (including these of the Greek Prime Minister and other members of the Greek government) have been monitored for several months, as a result of unauthorised access to the Lawful Interception facilities of a mobile operator [PRE07].
The second privacy issue concerns the limitation of the interception only to very specific types of traffic. Normally, the Lawful Interception means are able to intercept all traffic, including the Content of Communication and the Interception Related Information. In order for the privacy rights of the interception subject to be protected, the interception should be strictly limited to the very specific traffic types that are requested by the Law Enforcement Agency.
Another significant privacy issue regarding Lawful Interception is the protection of the intercepted data, with respect to privacy, during their transmission and consequent storage, as well as their further processing. In that respect, the provider shall not monitor or permanently record the results of interception. All the communications of intercepted data to the Law Enforcement Agency must be secured; strong end-to-end encryption is required.

4.4.4 Selected Applications

In contrast to the other types of network monitoring areas, in the case of Lawful Interception it is not really correct to speak about "applications", but rather about architectures. From the analysis above, it comes out that in order to lawfully intercept IP communications, different modules spanning across the network and performing data interception and mediation operations are necessary. To the best of our knowledge, no open products exist; all the existing solutions (such as the two briefly presented below) are commercial and their common characteristic is that they claim to be compliant with the standards, especially the ETSI ones. A major issue regarding Lawful Interception products is security, which constitutes a very critical requirement. In fact, Lawful Interception systems are very sensitive due to the results caused by unauthorised access.

4.4.4.1 Aqsacom solution

Aqsacom[32] provides a solution for Lawful Interception that is compliant with the ETSI standards. The Aqsacom Lawful Interception System (ALIS) consists of two functional modules: ALIS-M is the management platform which manages the interception sessions, instructs the network elements to start / end interception, while constantly monitoring the status of the network elements to alert the network operator of a fault condition; ALIS-D is the collection platform that receives the interception content and data from the network elements, formats this information according to the standards, then sends the results to the Law Enforcement Agency. ALIS provides mediation capabilities for all types of traffic over IP, data streams (e.g., web traffic), e-mail, VoIP, etc. The security of the connections between the different elements constituting the system is guaranteed by the usage of trusted paths with support for open standards like IPSEC and SSL, while smart tokens and biometric technologies are used to assure secure access to system operational functions and interception data.

4.4.4.2 Siemens solutions

Siemens[33] provides two professional solutions regarding Lawful Interception. The product denoted as Monitoring Center constitutes the basic monitoring and mediation solution. It has been designed in order to be fully compliant with the ETSI standards and to permit integration within all telecommunications networks which use any type of modern standardised equipment compatible with ETSI recommendation. On the other hand, the product referred to as Intelligence Platform serves for integrating and analysing different types of information

[32] http://www.aqsacomna.com/.
[33] http://www.siemens.com/.

originated from different information sources. The overall objective is to perform intelligent analysis of the collected data and draw conclusions that are not obvious from the data collection phase.

5 Conclusions

In this document, we examined the state of the art of network monitoring applications; this was an essential exercise in the definition of the problem space to be addressed by the PRISM project. We focused on four application areas: performance monitoring, network security applications including anomaly detection and intrusion detection and prevention, traffic classification and quality of service, and lawful interception. Each of these application areas requires and generates different types of data, using different protocols to transport and store them; each of these has its own potential impacts on the privacy of the networks' end users.

The applications we examined in this area are a very heterogeneous group. Most application areas have freely available as well as commercial implementations available. Each has different requirements and economics that influence their design and deployment. The application areas identified in this deliverable will be further analysed and refined into their constituent operations in order to define the problem space during the requirements specification of the PRISM architecture.

References

[ANS03] American National Standards Institute. Lawfully Authorized Electronic Surveillance. J-STD-025-A, April 2003.

[ANT06] D. Antoniades, M. Polychronakis, S. Antonatos, E. P. Markatos, S. Ubik, and A. Øslebø. Appmon: An Application for Accurate per Application Network Traffic Characterization. BroadBand Europe 2006. Geneva, Switzerland. http://www.ist-lobster.org/publications/papers/antoniades-appmon.pdf

[BON07] D. Bonfiglio, M. Mellia, M. Meo, D. Rossi, and P. Tofanelli, Revealing Skype traffic: when randomness plays with you, SIGCOMM Comput. Commun. Rev., vol. 37, no. 4, pp. 37-48, 2007.

[BON08] D. Bonfiglio, M. Mellia, M. Meo, N. Ritacca, and D. Rossi. Tracking down Skype traffic. In Proc. Infocom'08, April 2008.

[BSI02] Bundesamt für Sicherheit in der Informationstechnik (BSI). Einführung von Intrusion-Detection-Systemen, Rechtliche Aspekte. Version 1.0, 31. October 2002.

[CIS05] Cisco Systems, Inc. Capacity and Performance Management: Best Practices White Paper. Document ID: 20769. October 2005. http://www.cisco.com/application/pdf/paws/20769/performwp.pdf

[CIS07] Cisco Systems, Inc. Network Management System: Best Practices White Paper. Document ID: 15114. July 2007. http://www.cisco.com/application/pdf/paws/15114/NMS_bestpractice.pdf

[ETSI101331] European Telecommunication Standards Institute. Lawful Interception; Requirements of Law Enforcement Agencies. Technical Specification 101.331, v1.2.1, June 2006.

[ETSI101943] European Telecommunication Standards Institute. Lawful Interception; Concepts of Interception in a Generic Network Architecture. Technical Report 101.943, v1.1.1, July 2001.

[ETSI101944] European Telecommunication Standards Institute. Lawful Interception; Issues on IP Interception. Technical Report 101.944, v1.1.2, Dec. 2001.

[ETSI102232] European Telecommunication Standards Institute. Lawful Interception; Handover Specification for IP Delivery. Technical Specification 102.232, v1.5.1, Oct. 2006.

[ETSI102233] European Telecommunication Standards Institute. Lawful interception; Service-specific details for e-mail services. Technical Specification 102.233, v1.1.1, Feb. 2004.

[ETSI102234] European Telecommunication Standards Institute. Lawful Interception; Service-specific details for internet access services. Technical Specification 102.234, v1.1.1, Feb. 2004.

[ETSI102656] European Telecommunication Standards Institute. Lawful Interception; Retained Data; Requirements of Law Enforcement Agencies for handling Retained Data. Technical Specification 102.656, v1.1.2, Dec. 2007.

[ETSI201158] European Telecommunication Standards Institute. Lawful Interception; Requirements for Network Functions. ETSI Standard 201.158 v1.2.1, April 2004.

[EUR06] European Parliament and Council. Directive 2006/24/EC of the European Parliament and of the Council of 15 March 2006 on the retention of data generated or processed in connection with the provision of publicly available electronic communications services or of public communications networks and amending Directive 2002/58/EC. OJEC, No. L 105, pp. 54-63, Apr. 2006.

[EUR96] The Council of the European Union. Council Resolution of 17 January 1995 on the Lawful Interception of Telecommunications. OJEC 329, Nov. 1996.

[GEI05] R. Geib, E. Azañón-Teruel, S. Donaire-Arroyo, A. Ferrándiz-Cancio, C. Ralli-Ucendo, and F. R. Bueno. Service Deployment Experience in Pre-Commercial IPv6 Networks. European Journal for the Informatics Professional UPGRADE Vol. VI, issue No. 2, April 2005.

[GIV03] T. Givoly, J. Plishker. Response to RFI Regarding P2P File Sharing on University/College Campus, Revision 1.0. May 2003.

[KAR04] T. Karagiannis. Application-specific payload bit strings. Nov. 2004. http://www.cs.ucr.edu/~tkarag/papers/strings.txt.

[KAR05] T. Karagiannis, K. Papagiannaki, and M. Faloutsos. BLINC: multilevel traffic classification in the dark. SIGCOMM CCR Volume 35, Issue 4 pages 229-240. October 2005.

[PAX99] V. Paxson. Bro: A system for detecting network intruders in realtime. Computer Networks 31(23-24) pages 2435-2463. December 1999.

[PRE07] V. Prevelakis; D. Spinellis. The Athens Affair. IEEE Spectrum, Volume 44, Issue 7, pp. 26 – 33, July 2007.

[SCH06] C. Schmoll, J. Quittek, A. Bulanza, S. Zander, M. Kundt, E. Boschi, J. Sliwinski. D11 – State of Interoperability. MOME Project Deliverable. January 2006.

[STA99] W. Stallings. SNMP ,SNMPv2, SNMPv3, and RMON 1 and 2 (3rd Edition). Addison-Wesley Professional. January 1999.

[TEU08] P. Teufl, U. Payer, M. Amling, M. Godec, S. Ruff, G. Scheikl, and G. Walzl. InFeCT - network traffc classification. In Proc. ICN 2008, pages 439-444, April 2008.

[USC94] United States Congress. Communications Assistance for Law Enforcement Act (CALEA), Pub. L. No. 103-414, 108 Stat. 4279, codified at 47 USC 1001-1010. October 1994.