1 perfSONAR tools evaluation[1]

The goal of this PSNC activity was to evaluate perfSONAR NetFlow tools for flow collection solution and assess its applicability to easily subscribe and request different NetFlow sources and store them and/or present to the user.

Section 1.1 provides an overview of perfSONAR architecture. Section 1.2 describes testing of Flow Subscription Measurement Point to request near real-time streams of flow packets and section 1.3 describes testing of Flow Selection and Aggregation Measurement Archive to perform remote flow selection and aggregation requests.

1.1 perfSONAR overview

perfSONAR (Performance focused Service Oriented Network monitoring ARchitecture) [1] is a result of GN2 and GN3 [2] EU-funded projects and aims at providing a framework for performing multidomain measurements. It is a result of international collaboration and is deployed in the European Research Network GEANT and the connected National Research and Education Networks (NRENs) as well as selected international projects e.g. LHCOPN. The name reflects the choice of a Service Oriented Architecture for the system implementation. The architecture of this monitoring framework is depicted in Figure 1.

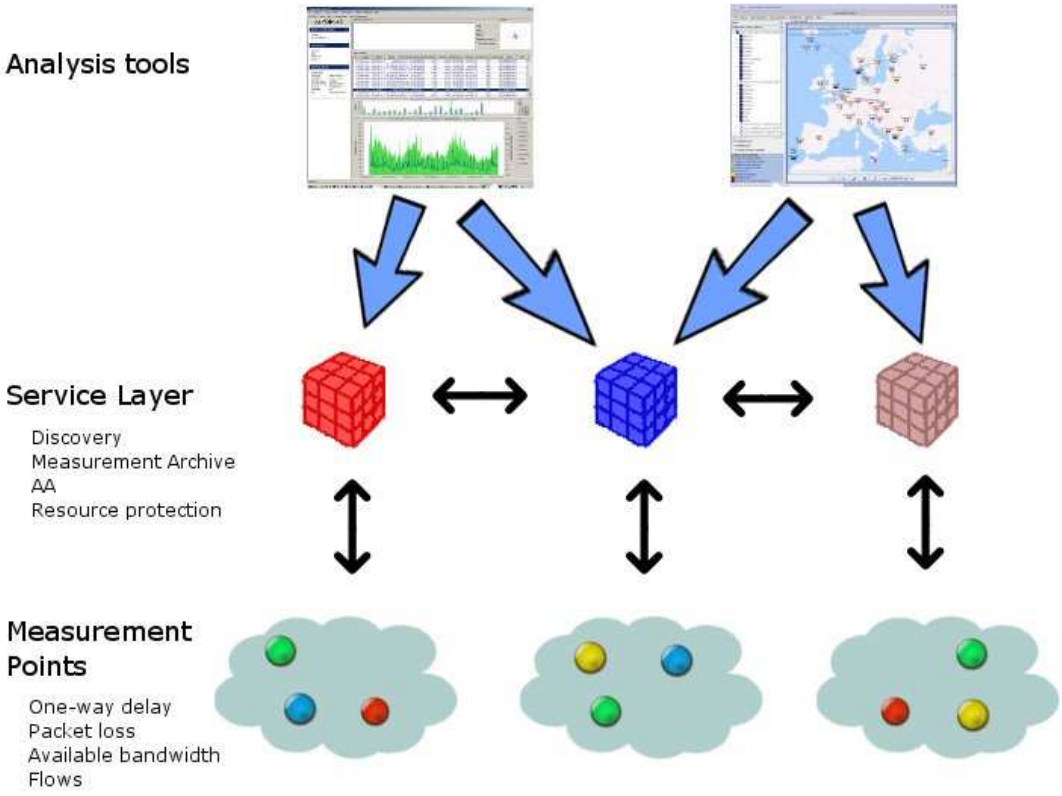

Figure 1. perfSONAR architecture

[1] This work was supported by the EC IST-EMANICS Network of Excellence (#26854).

The Measurement Points are the lowest layer in the system and are responsible for measuring and storing network metrics as well as for providing basic network information. The measurements can be carried out by active or passive monitoring techniques. The Measurement Point layer of a domain consists of different monitoring components or agents deployed within the domain. A monitoring agent provides information on a specific metric (e.g. one-way delay, jitter, loss, available bandwidth, NetFlow data) by accessing the corresponding Measurement Points. Each network domain can, in principle, deploy Measurement Points of its choice. The Service Layer is the middle layer of the system and consists of administrative domains. It allows for the exchange of measurement data and management information between domains. In each domain, a set of entities (services) is responsible for the domain control. Each of them is in charge of a specific functionality, like authentication and authorization or discovery of the other entities providing specific functionalities. In particular, the Measurement Archive Service (MA) is designed as a repository for measurement results. The interaction of the entities inside a domain as well as other domains is using perfSONAR protocol and may not be visible to the end user. Some of the entities contain an interface which can be accessed by the User Interface layer. This layer consists of visualization and analysis tools (user interfaces) which adapt the presentation of performance data to be appropriate for the needs of specific user groups. In addition, they may allow users to perform tests using the lower layers of the framework.

1.2 Flow Subscription Measurement Point

The Flow Subscription MP [3] is a Java application developed by SURFNet which makes it possible to request near real-time streams of flow packets (that is Netflow or Sflow exported by routers), as if they were coming directly from the routers where the information originated. This allows clients of this perfSONAR service to subscribe to flow information from different locations and still use their own favorite flow collector and processing tools.

The MP collects one or more NetFlow streams and clients of this service. Users can specify the router(s) from which they want to receive flow information, and can further tune the amount of information sent by creating a filter. As flow information can be privacy-sensitive, the Flow Subscription MP can anonymize the IP addresses before the information is sent to the client. In addition, when setting up a NetFlow stream between the MP and a client the flow information is sent through an encrypted tunnel to protect the information. Zebedee [4] software is used here to establish encrypted, compressed tunnel for data transfer and the only type of data that goes through perfSONAR layer is a subscription request to set up the tunnel and data stream, keepalives to maintain the data stream and unsubscribe requests to finish transmission. The flow data collected by the Flow Subscription service is replayed, anonymized (when configured) and send through the tunnel based on client subscriptions. On client side the data exits secure tunnel on a specific UDP port as it would be coming from the direct router stream and can be collected by any flow collector tool to store or process the NetFlow data. Parallel client connections can be established to request data from different NetFlow sources.

The MP provides its users with on-demand and real-time access to (a selection of) flow information for a specific amount of time, allowing them to perform their own security analysis, performance monitoring calculations or data collection using the tools of their choice.

1.2.1 Testing

The Flow Subscription MP is available as software package for both Debian GNU/Linux and RedHat distributions. We used .deb package and APT packaging tool in order to install the MP in Debian 4.0. This required adding perfSONAR repositories into the system and installation of Java v5. The installation was very straightforward and installed all required dependencies like nfdump, Tomcat 5.5 and zebedee. After restarting Tomcat application was ready for configuration. Unfortunately after initial configuration it turned out that the packaged version has bug which prevented the software to correctly recognize NetFlow exporters. After contacting with former (but still active) developer we were given by a latest WAR file created from the SVN repository and it worked fine. For the clarity of this description the following steps are using this type of installation. The bug was reported and should be fixed.

For testing purposes we used two routers sending NetFlow v5 to the test server where the service was deployed:

- PIONIER NREN core router 212.191.126.4, using port 9001
- POZMAN MAN core router 150.254.163.165, using port 9002

Initial configuration of the Flow Subscription MP required specifying flow exporters. Configuration of the service is done with graphical interface called WebAdmin. This is a set of Web pages for basic and advanced configuration. It's not required to configure all of the options in order to run the service - some of them are perfSONAR specific (like registration to directory service). The only ones required at the beginning for our purposes were flow exporters data and WebAdmin password change. The rest could be left default. The configuration page was available under http://loco2:8180/ps-mdm-flowsub-mp/ while the web service itself was available under http://loco2:8180/ps-mdm-flowsub-mp/services/FlowsubscriptionMeasurementPointService. Here we found difference between installation guide available in the Web which suggested wrong service URL and exact service URL. This information was corrected based on our test. Figure 2 depicts configuration of the two aforementioned exporters. Last configuration point at service side was to change default security settings. The Flow Subscription MP has been shipped with a filter configured by default and we had to change it in /var/lib/tomcat5.5/webapps/surfnet_java-flowsubscription-mp/WEB-INF/web.xml in order to allow client access.

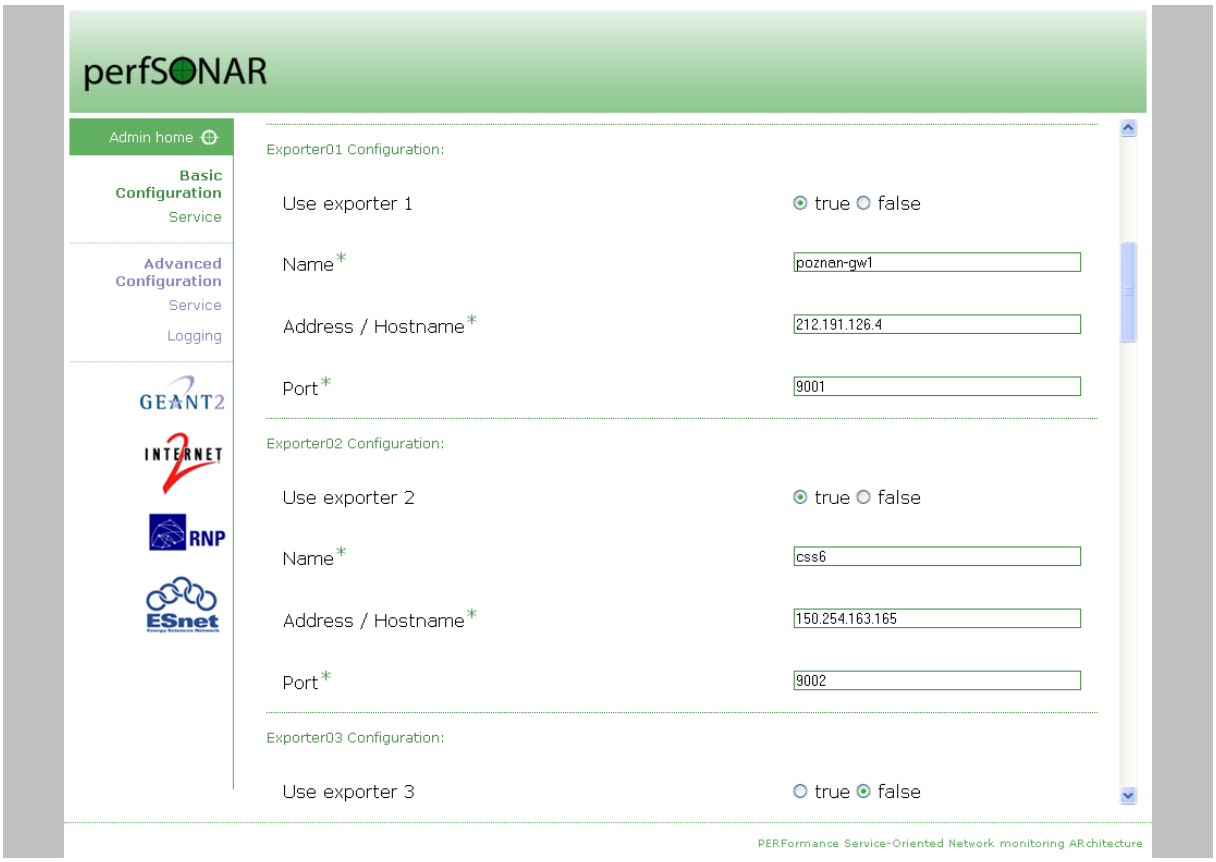

Figure 2 FlowSub MP administration interface

In order to run the client we used one of the EMANICS lab servers where missing zebedee, python and SOAPpy module were installed. Then we downloaded python client directly from perfSONAR repository. The client receives NetFlow data through an encrypted tunnel by subscribing to the MP. The only data that goes through perfSONAR layer are the control messages (web service requests and responses). Before tests started the client had to be configured. This required us Flow Subscription MP URL, IP address of the client, port we want to receive data (7777) and a name of the NetFlow source (as configured at the MP side) we wanted to subscribed NetFlow data from. The client currently support subscription to a single NetFlow stream. By default the service also uses the Crypto-PAn module with 32 characters key to anonymize IP addresses. Figure 3 shows the output of the process running Python client. It starts with subscribing to MP service and then informs us about establishing the zebedee tunnel. Then it periodically reestablishes it in order to receive flow data. When the client is terminated it unsubscribes from the service.

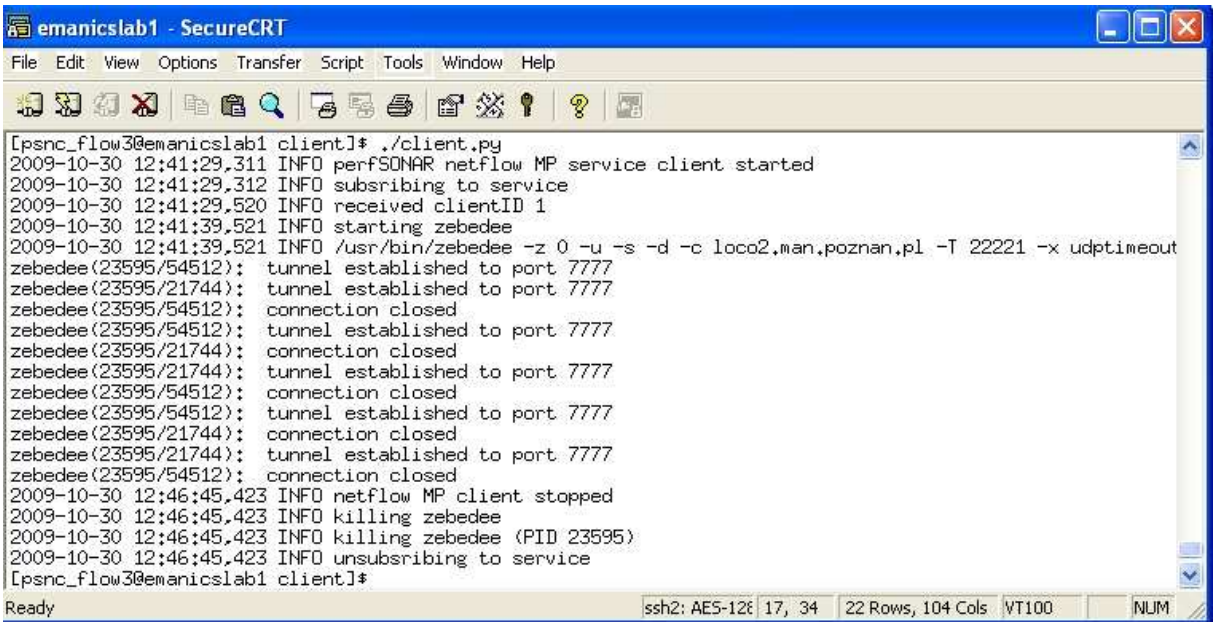

Figure 3 Running subscription client

During our tests we successfully subscribed at the EMANICS lab server to our NetFlow streams from the Flow Subscription MP and received NetFlow packets from aforementioned routers with the use of zebedee tunnel. Figure 4 shows NetFlow data from packets received on port 7777 using the client. Then the NetFlow packets could be easily locally processed by any NetFlow analysis tools.

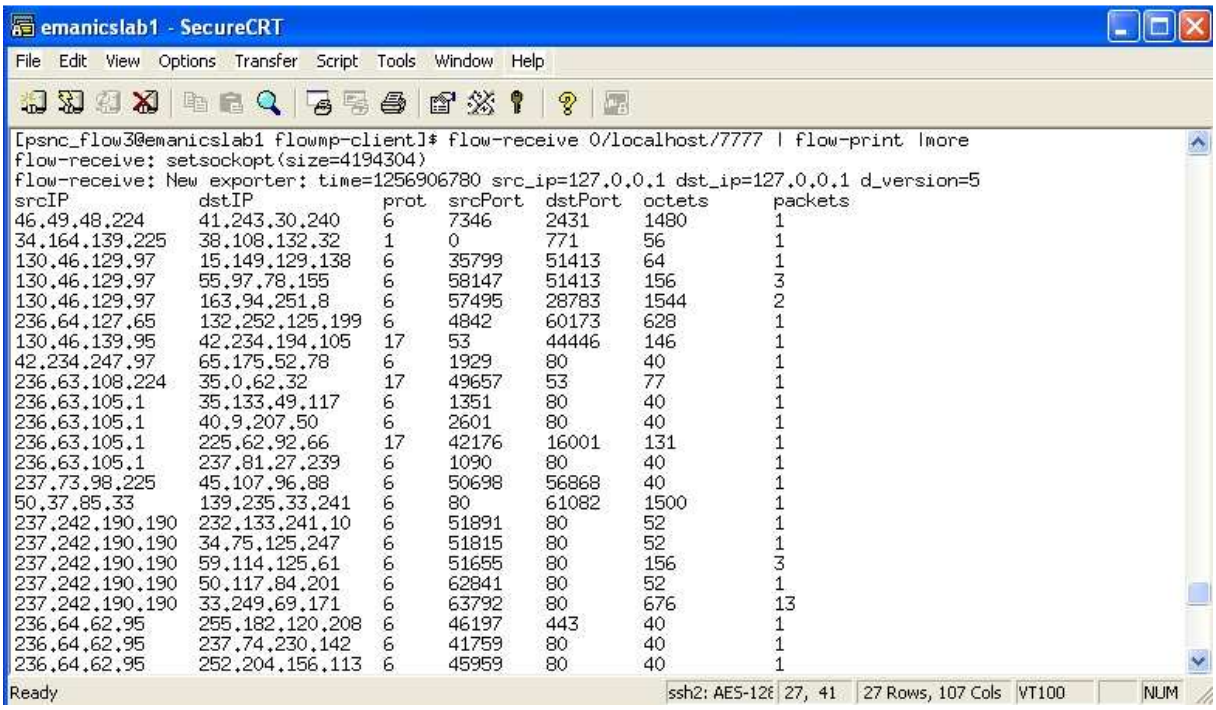

Figure 4 NetFlow data received at the given port

1.3 Flow Selection and Aggregation Measurement Archive

The Flow Selection and Aggregation MA [5], developed in Java by SURFNet, acts as a perfSONAR interface wrapper around the functionality of the flow capturing and processing commonly used tool nfdump [6]. This makes it possible to perform remote flow selection and aggregation requests. The

MA is accompanied by a plugin for perfSONAR visualization tool perfsonarUI [7] that gives access to all available functionality. perfSONAR community uses this dedicated analysis application perfsonarUI - Java Web Start graphical user interface for querying a range of perfSONAR services deployed around the world. The latest version is available here: http://perfsonar.acad.bg/psui_beta/perfsonar.jnlp. With Flow Selection and Aggregation MA service users can perform perfSONAR-based remote nfdump style selection and aggregation queries on stored log files. This enables them to search for flows patterns, security related information, and to debug network related problems. Depending on the storage capacity of the MA, queries can be performed on flow information from weeks or even months ago. The MA supports the following types of data requests:

- flowStatistics return statistical information about the selected flows from a given group of routers
- rawFlows return specific fields from the selected flows. An optional nfdump style filter rule can be used to limit the number of matching flows and/or an aggregation rule can be used to accumulate information
- topFlows return the top N flows given a nfdump style top filter rule

1.3.1 Testing

The Flow Selection and Aggregation MA is available as software package for both Debian GNU/Linux and RedHat distributions. We used .deb package and APT packaging tool in order to install the MP in Debian 4.0. This required adding perfSONAR repositories into the system and installation of Java v5. The installation was very straightforward and installed all required dependencies like Tomcat 5.5. After restarting Tomcat application was ready for configuration.

For testing purposes we used a router sending NetFlow v5 to the test server where the service was deployed:

- PIONIER NREN core router 212.191.126.4, using port 9001

Then for test purposes these NetFlow packets were stored locally by nfcapd service run with the script attached to the installation package and installed in /etc/init.d folder.

Initial configuration of the Flow Selection and Aggregation MA required specifying flow exporters. Configuration of the service is done with graphical interface called WebAdmin. This is a set of Web pages for basic and advanced configuration. It's not required to configure all of the options in order to run the service - some of them are perfSONAR specific (like registration to directory service). The only ones required at the beginning for our purposes were flow exporters data and WebAdmin password change. The rest could be left default. The configuration page was available under http://loco2:8180/ps-mdm-flowsa-ma/ while the web service itself was available under http://loco2:8180/ps-mdm-flowsa-ma/services/FlowsaMeasurementArchiveService. Here we again found difference between installation guide available in the Web which suggested wrong service URL and exact service URL. This information was corrected based on our test. Figure 5 depicts configuration of the exporter we used in our tests. Last configuration point at service side was to change default security settings. The Flow Subscription MP has been shipped with a filter configured by default and we had to change it in /var/lib/tomcat5.5/webapps/surfnet_java-flowsubscription-mp/WEB-INF/web.xml in order to allow client access.

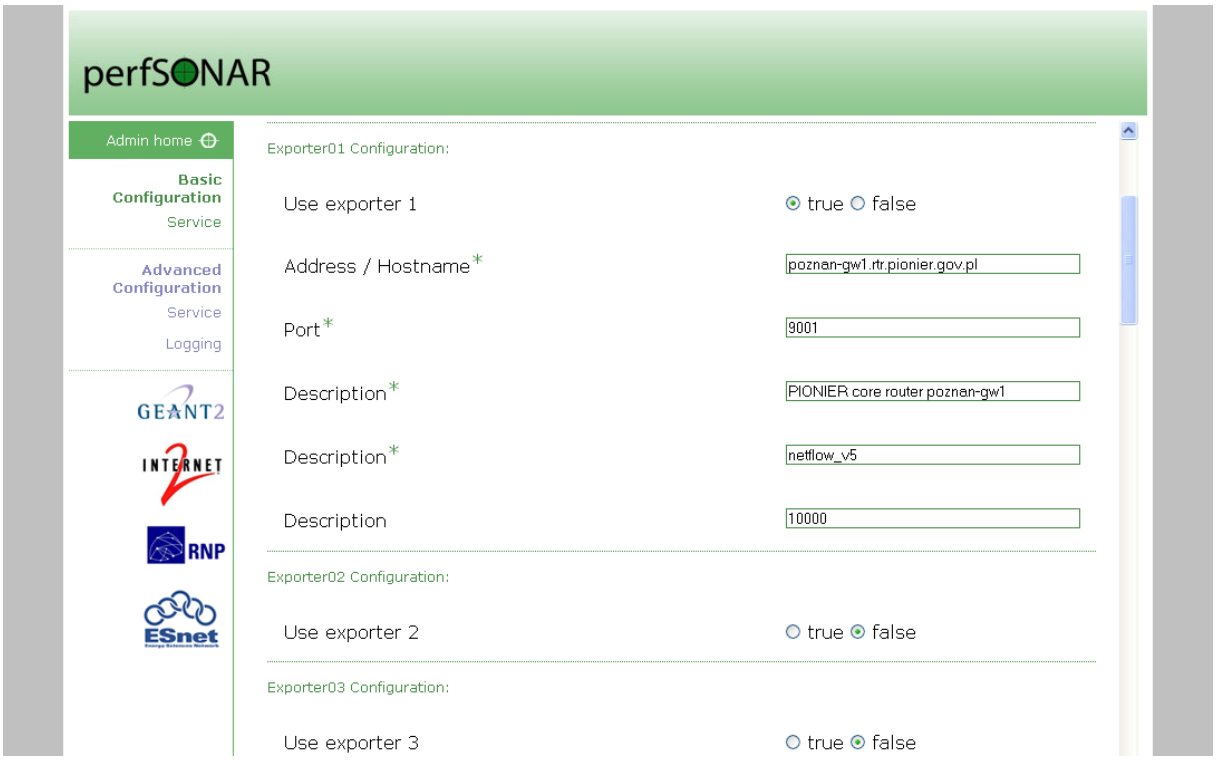

Figure 5 FlowSA MA administration interface

After initial configuration we might have started retrieving data from the service using perfsonarUI. Of course it is possible to query the service with other tools provided they conform to perfSONAR protocol. The user is constructing a query with the client. The client then transmits this query to the service and the web service transforms the query into nfdump commands. These nfdump commands have Netflow data or statistical data as a result and are converted back to NMWG XML (used by perfSONAR protocol), and sent to the client where they are visualized. Before the tests started a new configuration file was created in order to specify Flow Selection and Aggregation MA service URL. The configuration file is read then by the dedicated tab in the GUI. Unfortunately it turned out that the Selection and Aggregation MA tab in perfsonarUI is not working as expected and is missing an important feature. For example "Options" configuration window doesn't currently enable to specify the IP address or DNS name of routers which is crucial to retrieve the data. To avoid this problem router name was entered directly into "Settings" window. Other fields of the "Option" window work well. The other potential problem we noticed is that some of users may want to see an option for anonymization on the service side. These problems were reported to perfsonarUI developers as an RFE.

First we tested different TopN requests as the most useful command. We choose analysis period and then additional parameters like type of statistics (record, IP addresses, ports, AS numbers, interface) and the order by which the statistics is ordered (flows, packets, bytes, etc). Number of statistics was also specified. Although filter and aggregation rule don't apply here these field are still available in the configuration window which may be misleading. Figure 6 depicts the result of an example query using in perfsonarUI using Record type of statistics, aggregating by Packets and displaying 10 results where a few fields are displayed providing information about start time of the flow first seen, duration of the flow, flow details, bytes transferred, number of packets and protocol type.

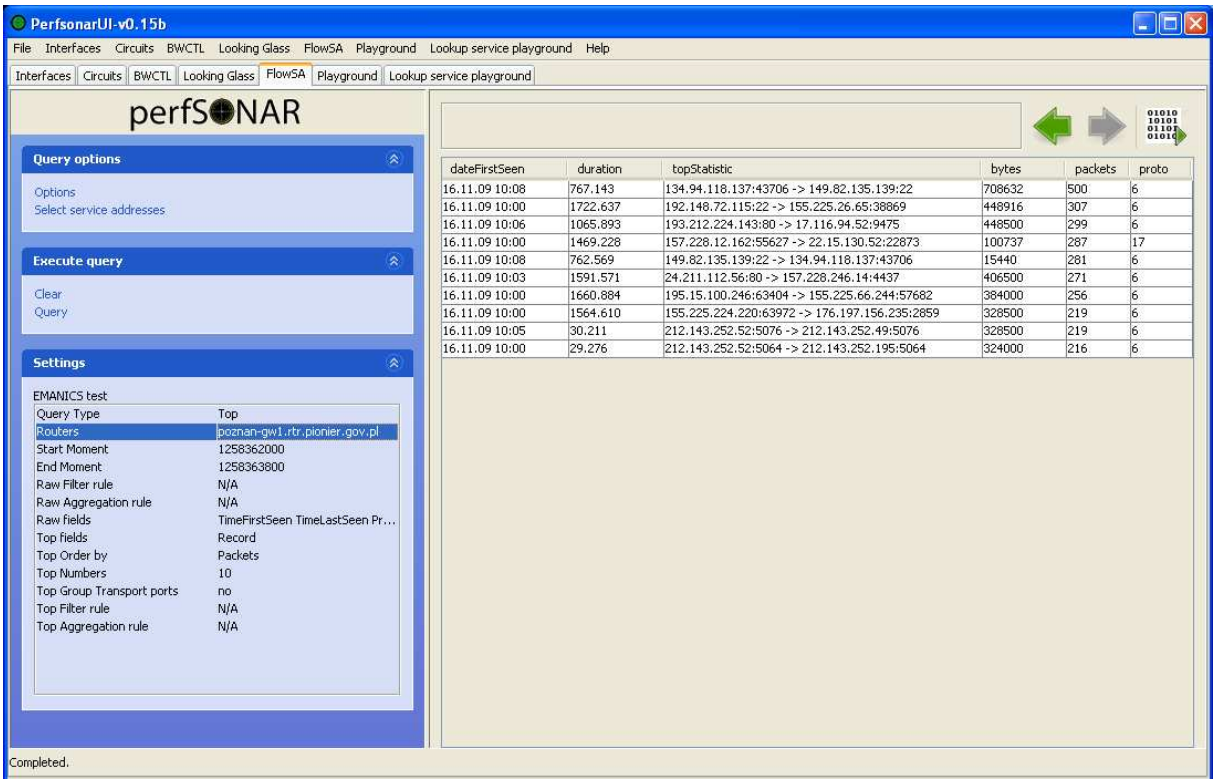

Figure 6 perfsonarUI accessing Flow Selection and Aggregation MA - TopN flows

Then we tested getting statistical information about flows collected by this perfSONAR service within a specified period of time. Therefore no options are configurable apart from time period. Figure 7 depicts the result of an example statistical query to Flow Selection and Aggregation MA using perfsonarUI. The result provides information about total number of flows, packets, bytes and the same information but per protocol, start of the first flow and end of the last flow.

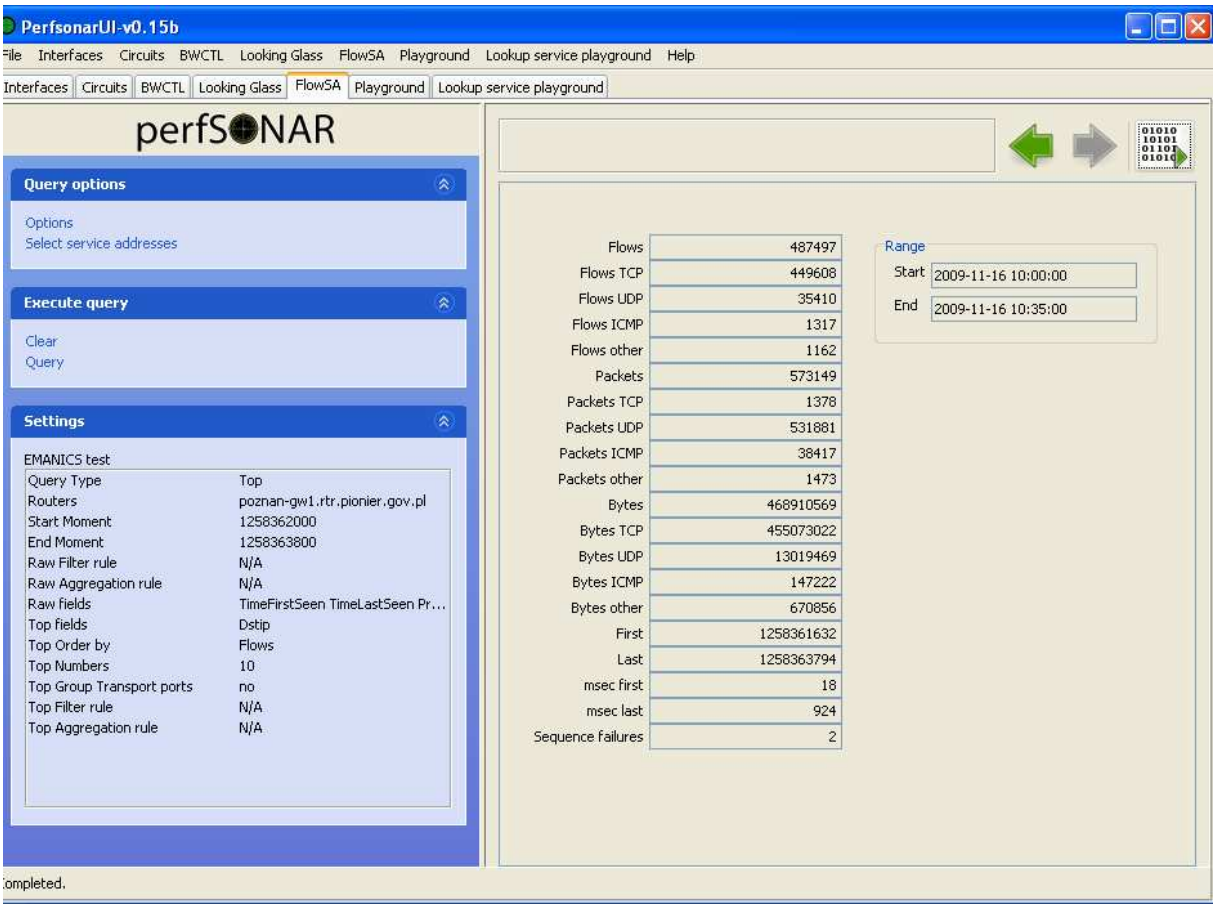

Figure 7 perfsonarUI accessing Flow Selection and Aggregation MA - statistical information about flows

Finally we tested raw flows which could be used for detailed display of each NetFlow record and could be rather rarely used. In the Options window we selected Raw type of query and a filter to limit amount of data (port 8080). We also chose fields to be displayed (sourca and destination IP address, source and destination port number, source and destination as number and number of bytes and packets. Figure 8 depicts the correct result of an raw query to Flow Selection and Aggregation MA using perfsonarUI. If the time range is too wide or no specific filter is applied it results in a lot of raw data and generates an error about too many results retrieved.

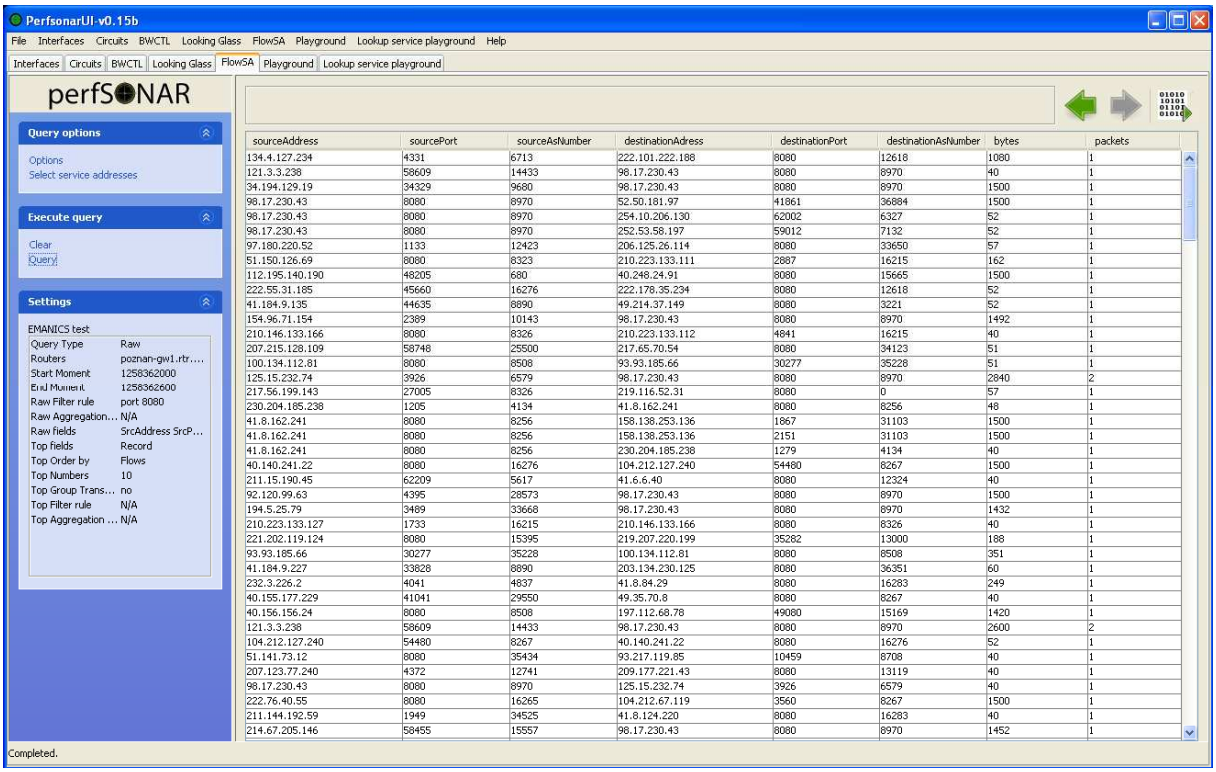

Figure 8 perfsonarUI accessing Flow Selection and Aggregation MA - raw flows

1.4 Conclusions

During our test we tested two perfSONAR tools: Flow Subscription Measurement Point to request near real-time streams of flow packets and Flow Selection and Aggregation Measurement Archive. Both are using perfSONAR architecture as a basis. This European activity seems to be well established in a research community with the ongoing GEANT project which continues working on perfSONAR deployment thus increasing the community of potential users. This should hopefully provide a long term support for the tools. There are also steps in order to standardize the perfSONAR protocol under OGF. During the Flow Subscription MP tests we found a few problems which were reported to the perfSONAR community and we hope will be taken into account during next releases. The configuration part was straightforward although the client is very simple and not sophisticated. We think the service is a useful tool in order to publish NetFlow data in a secure and controlled way especially in the management and monitoring of networks in a multidomain environment or collaboration of distributed partners. Especially for those networks which already use perfSONAR services in order to share other monitoring data. Flow Subscription MP may be used in order to give access to flows streams by simply subscribing to a selected device for a selected period of time only necessary to collect data without the necessity to receive flows all the time. Potential problem could be a zebedee tool which is in a good shape now but seems to be not developed anymore (the last stable and development versions 2.4.1A are from 2005/09/06. The results of Flow Selection and Aggregation Measurement Archive tests were not such optimistic mainly due to analysis tool problems which currently may prevent users from using this set of perfSONAR products. We couldn't properly configure the GUI to access the service. It's obvious that these perfSONAR web services needs bug corrections in order to make it usable. The graphical interface enabling access to the data through perfSONAR interface also needs a few corrections and enhancements necessary to make it easy to use and valuable to the community. We submitted appropriate bugs and we hope that the

feedback from our thorough user tests will help the developers to enhance the product which will become an efficient tool for NetFlow analysis using perfSONAR products.

Bibliography

1. Hanemann, J. Boote, E. Boyd, J. Durand, L. Kudarimoti, R. Lapacz, M. Swany, S. Trocha, and J. Zurawski, Perfsonar: A service oriented architecture for multi-domain network monitoring, Service-Oriented Computing - ICSOC 2005, Springer Verlag, p. 241-254.
2. GÉANT homepage including information about GN2 and GN3 projects, available at http://www.geant.net
3. Flow Subscription Measurement Point homepage, available at https://wiki.man.poznan.pl/perfsonar-mdm/index.php/Flow_Subscription_MP
4. Zebedee: Secure IP tunnel homepage, available at http://www.winton.org.uk/zebedee/
5. Flow Selection and Aggregation Measurement Archive homepage, available at https://wiki.man.poznan.pl/perfsonar-mdm/index.php/Flow_Selection_and_Aggregation_MA
6. The nfdump tools to collect and process netflow data homepage, available at http://nfdump.sourceforge.net/
7. perfsonarUI analysis tool homepage, available at http://iris.acad.bg/perfsonar/perfsonar.jnlp

Szymon Trocha, PSNC
szymon.trocha@psnc.pl

26 November 2009