30-06-2019

Deliverable D6.1
Network Management and Monitoring Portfolio

Deliverable D6.1

Contractual Date:	30-06-2019
Actual Date:	30-06-2019
Grant Agreement No.:	856726
Work Package	WP6
Task Item:	Task 1
Nature of Deliverable:	R (Report)
Dissemination Level:	CO (Confidential)
Lead Partner:	UoB/AMRES
Document ID:	GN4-3-19-1495F9
Authors:	Pavle Vuletić (UoB/AMRES), Ivana Golub (PSNC), Tim Chown (JISC), Szymon Trocha (PSNC), Antoine Delvaux (PSNC), Kurt Baumann (SWITCH), Lukasz Lopatowski (PSNC), Marinos Dimolianis (NTUA/GRNET)

© GÉANT Association on behalf of the GN4-3 project.
The research leading to these results has received funding from the European Union's Horizon 2020 research and innovation programme under Grant Agreement No. 856726 (GN4-3).

Abstract

At the beginning of the GÉANT Network 4, Phase 3 (GN4-3) project, the Monitoring and Management Task (Task 3) of the Network Technologies and Services Development Work Package (WP6) takes care of several network monitoring and management services in production, and continues the relevant developments from the GN4-2 project. This document presents the network management and monitoring service portfolio at the project start, and details work and new initiatives in this project phase.

Table of Contents

Table of Figures

Table of Tables

Executive Summary

As of the beginning of the GN 4 Phase 3 (GN4-3) project, the *Monitoring and Management* task (Task 3) of the *Network Technologies and Services Development* Work Package (WP6) is responsible for several production network monitoring and management services, i.e, perfSONAR, perfSONAR consultancy and expertise, the Performance Measurement Platform (PMP) and Network Management as a Service (NMaaS), and is continuing developments towards production from the previous GN4-2 project generation, i.e., NetMon, WiFiMon.

This document presents the network management and monitoring service portfolio at the start of the project, ongoing work and new initiatives in this project phase. It also presents a service status baseline for WP6 Task 3, which, in the later stages of the project, can be used to monitor the progress of the service development and the path the Task will follow.

1 Introduction

In GN4-3, WP6 Task 3 works on the further development and evolution of network monitoring and management tools and services. A focus of this Task is on monitoring new network services that are being developed by the GÉANT partnership and helping to improve network management automation for campus and NREN end user institutions, although in principle monitoring any network services is of interest to the Task.

Task 3 inherited several in-production services from the previous project, as well as the development of some tools and services which are still in the earlier phases of their lifecycle. The services which entered their production phase before the GN4-3 project started are perfSONAR, perfSONAR Consultancy and Expertise, the Performance Measurement Platform (PMP) and Network Management as a Service (NMaaS). Another two tools that were piloted and tested successfully in real network environments during GN4-2 are WiFiMon and NetMon.

Three perfSONAR-related services are supported by the WP6 Task 3 team:

- The **perfSONAR International Project** - perfSONAR software development and user support within the international collaboration with ESnet, Internet2, Indiana University and the University of Michigan.
- **perfSONAR Consultancy and Expertise** - supporting the GÉANT community in deploying perfSONAR in their environments, providing advice, training and support on designing and deploying a perfSONAR-based measurement architecture.
- The **Performance Measurement Platform (PMP)** - providing an open, trusted monitoring and measurement information infrastructure (this is currently based on perfSONAR but may in the future adopt additional tools).

The GÉANT project partners providing resources for the development, maintenance, support, training and dissemination of perfSONAR-related services are AMRES, CARNET, DFN-FAU, GÉANT Association, GRENA and PSNC. PSNC also provides Service Managers to perfSONAR and the PMP, who are responsible for managing and supervising development, operations and support.

In parallel with the operation and development of the above services and tools, Task 3 is exploring the potential for new service developments proposed by the GÉANT community, such as Campus Network Management as a Service (CNaaS). CNaaS facilitates campus network management outsourcing services provided by NRENs. Task 3 is also evaluating the potential to enhance network flow monitoring and flow analysis tools, the use of virtualised environments for network element and service monitoring, and the use of programmable network element features for network service monitoring.

This document presents a snapshot of the status of the Task 3 monitoring and management services at the start of the GN4-3 project. The document focuses on those services and tools which are either in production or close to production (WiFiMon and NetMon). Other potential services that are at an early stage of the service lifecycle, where their value and demand for their development is being assessed, are not described in this document, as it is still not certain that they will progress to the later phases of the lifecycle (design, transition and operation). The outcomes of the strategic work on campus network management and flow monitoring activities will be summarised in separate milestone documents towards the end of the first year of the GN4-3 project.

The document consists of two main parts: Section 2 which describes the current status of the services in production and Section 3 which summarises the current status of NetMon and WiFiMon development.

2 Services in Production

2.1 perfSONAR International Project

perfSONAR is an open-source, modular and flexible architecture for active network performance monitoring that provides a view of network performance across multiple domains, allowing Network Operations Centre (NOC) and Performance Enhancement Response Team (PERT) engineers to analyse and diagnose network behaviour across an entire end-to-end path. The tools provided in the perfSONAR suite perform active measurements of throughput, packet loss, delays and jitter, and record network route and path changes. There is also provision for application level measurements like HTTP, DNS and disk to disk transfers.

The global perfSONAR team develops, maintains, distributes and provides support for the full perfSONAR tool suite that is installed and used on numerous Research and Education (R&E) networks around the world to perform active measurements and monitor network performance.

Public-facing information about perfSONAR development and support provided by the GN4-3 project can be found online [Wiki].

A number of information sources are available for the perfSONAR project:

- Main website [Website].
- Installation and usage documentation [Docs].
- Support request can be submitted through the user mailing list using perfsonar-user@internet2.edu or webmail [Webmail].
- Developers' resources are available on Git [Git].

2.1.1 Users, Uptake and Usage

Active network measurements are useful to network engineers, PERT engineers, system administrators, research communities, researchers and students. perfSONAR users include:

- Organisations (e.g. universities, GÉANT NRENs and GÉANT itself) that want to provide active network measurement capabilities to their users or to any collaborating organisation's users (enabling multi-domain measurement capabilities).
- Organisations that want to perform active measurements within their own domain or any other perfSONAR-enabled domain.

- Research communities that undertake data-intensive science, need to move large volumes of data between sites and want to have telemetry on network characteristics.
- Individual users who want to monitor end-to-end performance or performance on particular links of interest.
- Network researchers interested in developing or monitoring and assessing the performance of new high-speed networks, technologies and protocols.

perfSONAR users are located worldwide and form the global perfSONAR community. The current usage map is available on the perfSONAR website [Usage]. Figure 2.1 shows the 2130 nodes with 10164 services implemented worldwide as of April 2019 (the number of services depends on the perfSONAR version and on which services the users have activated).

Figure 2.1: perfSONAR installations worldwide

There are more than 400 perfSONAR deployments in European countries, including Armenia, Austria, Belarus, Belgium, Bulgaria, Croatia, Cyprus, Czechia, Denmark, Estonia, Finland, Germany, Greece, Hungary, Ireland, Italy, Latvia, Lithuania, Luxembourg, Moldova, Montenegro, Netherlands, Norway, Poland, Portugal, Romania, Serbia, Slovakia, Slovenia, Sweden, Switzerland, the United Kingdom and Ukraine.

2.1.2 Key Performance Indicators

The key performance indicator (KPI) for perfSONAR measures the number of major perfSONAR releases per year. Table 6.1 shows the current KPI value.

KPI	Target	Baseline	Measured
Number of major perfSONAR releases per year	1	1	1

Table 2.1: perfSONAR KPI measured for the period 01 May 2018 – 30 April 2019

The last major perfSONAR release, 4.1, was issued in August 2018. Recently, two minor releases were issued: perfSONAR 4.1.5 in December 2018 and 4.1.6 in February 2019.

2.1.3 Activities and Issues

The work of the global perfSONAR team is currently focused on preparing the next major release of perfSONAR (version 4.2) and upgrading the build and release infrastructure in order to increase the quality and reliability of releases. The GÉANT team contributes to building, testing and fixing all Debian/Ubuntu packages, improving the TWAMP tools, introducing a new traceroute visualisation tool and testing and debugging perfSONAR 4.2. The team is also responsible for upgrading the whole Debian build and distribution infrastructure to the latest Debian and Ubuntu standards. perfSONAR minor release versions 4.1.5 and 4.1.6 focused on minor updates and bug fixes. The roadmap for the perfSONAR project can be found on the perfSONAR website [Release_Notes].

The GÉANT team is also providing user support via the regular channels. Queries are mostly related to upgrades of existing installations and feedback about usage of new perfSONAR releases or, in the case of new users, regarding the perfSONAR installation and setup.

In the research and development area, work is being done by the team in prototyping new perfSONAR integrations with existing visualisation tools, for example with Grafana, Kibana and some common JavaScript libraries.

The perfSONAR software suite has recently been presented as a part of the WP6 work at the following events:

- Project Management Convention, 12-13 February 2019, Cologne, Germany
- Service and Technology Forum, 26-27 February 2019, Bucharest, Romania
- SIG-MSP, 05 March 2019, Nicosia, Cyprus
- Jisc Networkshop conference, 09 April 2019 Nottingham, UK

perfSONAR will also be presented at the TNC conference in Tallinn, Estonia, in June 2019, both within a broader 90 minutes session on Tuesday 18 June about WP6 activities in the GN4-3 project, and also as a Lightning Talk ‘Five Things You Didn’t Know About perfSONAR’.

2.2 perfSONAR Consultancy and Expertise

perfSONAR Consultancy and Expertise aims to provide support and disseminate knowledge about perfSONAR usage for the GÉANT community. It offers four different activity types:

- Helping to ensure that designed measurement architectures and infrastructures based on perfSONAR fit the performance monitoring and measurement needs of the requesting party.
- Providing specific training and running workshops on perfSONAR deployment, usage and best practices.

- Providing extra support to assist GÉANT and NRENs deploy and operate perfSONAR, as requested.
- Maintaining and operating a set of perfSONAR services useful to the global perfSONAR community and GÉANT-area perfSONAR users in particular.

Contact details for perfSONAR Consultancy and Expertise are available on the GÉANT wiki [Wiki2].

2.2.1 Users, Uptake and Usage

Target users of this service are teams and individuals from the GÉANT community. Since active network measurements and network performance monitoring require specific and advanced knowledge, users are mostly from Network Operating Centres (NOCs) and/or the Performance Emergency Response Teams (PERTs) of NRENs, the NRENs' constituencies or cross-domain projects that they might participate in. However, the service's availability is not limited to a specific user group.

2.2.2 Key Performance Indicators

The key performance indicator for perfSONAR Consultancy and Expertise is the number of requests for consultancy that the team fulfils during the reported period.

KPI	Target	Baseline	Measured
Number of requests for consultancy fulfilled	3	3	3

Table 2.2: perfSONAR KPIs measured for the period 01 May 2018 – 30 April 2019

Since the perfSONAR Consultancy and Expertise team overlaps with the development team, the resource to serve requests is relatively limited.

2.2.3 Activities and Issues

The perfSONAR Consultancy and Expertise service is currently providing support to:

- The PRACE project to deploy perfSONAR on the MD-VPN service provided by GÉANT to measure the VPN performance.
- The LoLa project to help measure latency and jitter between LoLa project sites.
- The CNaaS group to provide application level monitoring in a multiple VLAN setup from a single perfSONAR instance.

The Consultancy and Expertise service has recently been presented as a part of the WP6 work at the following events:

- Project Management Convention, 12-13 February 2019, Cologne, Germany
- Service and Technology Forum, 26-27 February 2019, Bucharest, Romania
- SIG-MSP, 05 March 2019, Nicosia, Cyprus

- Jisc Networkshop conference, 09 April 2019 Nottingham, UK

The team has also started to prepare the first European perfSONAR user workshop, which will bring together perfSONAR user and developer communities to explore use cases for the toolkit, share best practices, and discuss future features and the perfSONAR development roadmap. The workshop is being hosted by Jisc in London on 5-6 June 2019.

The Consultancy and Expertise service will also be presented at the TNC conference in Tallinn, Estonia, in June 2019, within a broader 90 minutes session on Tuesday 18 June about WP6 activities in the GN4-3 project.

2.3 Performance Measurement Platform - PMP

The Performance Measurement Platform (PMP) has developed from the successful perfSONAR Small Nodes project initiated during GÉANT 4-1 in order to contribute the following aspects of GÉANT strategy in the domain of network and performance monitoring:

- Maintain GÉANT's position as a provider of infrastructure services ahead of the commodity market.
- Focus on collaborative efforts to innovate the service portfolio, creating more advanced services.
- Organise the whole ecosystem of service delivery from NRENs to other NRENs.
- Harness the services of the NRENs to create a one-stop-shop for international organisations.

The Performance Measurement Platform is set up as an open, trusted monitoring and measurement information infrastructure, provided to network engineers, NOC operators, research communities, network researchers and NREN participants to monitor, explore, practice and learn how network performance monitoring can contribute to better and more efficient usage and understanding of the existing multi-domain infrastructure.

The Performance Measurement Platform includes 37 distributed measurement points with pre-installed perfSONAR. It thus extends the perfSONAR footprint in Europe and beyond (Ghana, Nigeria, Senegal) and enables, from an NREN's perspective, better visibility of the performance of the established GÉANT network infrastructure. PMP also changes the perspective of how perfSONAR measurement points are designed and deployed - from stand-alone servers to cheaper, readily available and readily configurable small boxes.

The small nodes perform regular measurements towards a few perfSONAR Measurement Points (MPs) located in the core of the GÉANT network and operated by the GÉANT network operations centre. The central components that manage the platform elements and gather, store and represent the performance data, are operated and maintained by WP6 Task 3. The project team has also implemented automation (using Ansible) to address maintenance routines and replace sequences of manual tasks. Small node users can shape the predefined setup and configure additional measurements to their needs and get more familiar with the platform.

2.3.1 Users, Uptake and Usage

PMP has three user groups:

- NRENs (or associated organisations and universities within the NREN) that are hosting a node and thus participate in a PMP mesh, setting up measurements of interest and using the PMP measurement results for their own benefit.
- End users or communities that are not hosting a node but use the measurement results that are publicly available via the PMP dashboard.
- Individuals or organisations that are implementing the PMP model and the software to install their own public or private measurement mesh.

Nodes provided by the project are distributed primarily to NRENs, but also to individual organisations willing to host the node and to institutions generally cooperating with the GN4-3 project. Some of the nodes are used for research and development as well as testing purposes by the PMP service team. Independently of their hosting organisation all nodes are managed by the PMP team in WP6 Task 3.

The third group procures and manages their nodes themselves. This also includes organisations that have started with a single PMP node provided by the project and then extended the footprint of perfSONAR nodes in their domains based on this initial deployment. Example countries are Belarus (deployment for EaPConnect project), Poland and the UK.

2.3.2 Key Performance Indicators

Critical success factors that can be used to measure the success of the infrastructure are:

- Availability of the central infrastructure.
- The number of GÉANT project partner countries that host a small node in their network.
- User satisfaction.

Availability of the central infrastructure will be measured based on the availability of the MaDDash (graphical user interface) and measurement archives. Availability measurements were implemented at the beginning of the project, therefore baseline figures are not available. The measured availability between 1 Jan 2019 and 30 April 2019 was 98.8%, while the target availability is 99%. The availability was temporarily lower than the target due to a GÉANT VM infrastructure outage in March 2019. In addition, PMP experienced short connectivity issues between the monitoring station and the central infrastructure. This did not impact the MaDDash user interface, only the availability monitoring.

Having as many as possible GÉANT NRENs taking part in this project is an important goal, therefore distributing the PMP nodes to GN4-3 project partners (NRENs) is the priority. However, it can be expected that due to limited resources some NRENs will currently not be interested in taking part. Also, for some countries more than one node is distributed. Additional nodes go to organisations particularly interested in hosting perfSONAR and using the infrastructure provided by GÉANT and the NRENs. Therefore, the total number of nodes distributed will be higher than the number of GN4-3 project partners countries (38). At the time of writing a total of 37 nodes have been distributed amongst PMP service participants, 22 of which went directly to GN4-3 project partners (NRENs). 15

nodes are hosted amongst other interested organisations, institutions or used as testing instances by the PMP service team. The KPI will focus on the geographical service coverage and count just the number of countries seen as GN4-3 project partners. The target participation by the end of the project is 80% of the NREN project partner countries. The target value is not 100% as it can be expected that not all countries are willing or able to host a node.

User satisfaction will be measured via a survey that will be performed once a year among the participating organisations. Target satisfaction is at least 85% of users who rate their experience with PMP service as good. Measuring how users describe the service allows the team to assess if it meets the users' needs and provides high customer satisfaction. The first user satisfaction survey is planned towards the end of the first year, therefore current data is not available at this moment.

KPI	Target	Baseline	Measured
Availability of the central infrastructure	99%	N/A[1]	98.4%
The number of GÉANT project partner countries	80%	48%	63%
Overall user satisfaction	85%	N/A	N/A

Table 2.3: PMP KPIs measured for the period 01 May 2018 – 30 April 2019

2.3.3 Activities and Issues

From the beginning of the project, the team has worked on the distribution of additional small nodes. This is usually organised during GÉANT-related meetings and conferences together with dissemination activities. In 2019 four new NRENs and one organisation received their pre-configured nodes and connected them to the measurement mesh: Cyprus, Georgia, Moldova, Spain (2 nodes).

The team also continued with the transition of existing small nodes to the new Performance Measurement Platform. This involved upgrading the PMP node to a new operating system, installing the latest version of perfSONAR and migrating (configuration update) to the new central infrastructure. Parts of the activities (upgrade of operating system, installation of perfSONAR and restoring its configuration from backup) must be done by the hosting organisation. The WP6 Task team provides support to the hosting organisation during this process.

Migrated nodes are updated to support the latest perfSONAR version and OS - currently pS 4.1.6 running on CentOS7.

In order to properly monitor the usage, availability and performance of central infrastructure and nodes, a dedicated Grafana-based monitoring solution was implemented. This not only allows the WP6 task team to monitor the current parameters but also provide historical data.

[1] The central infrastructure for the previous measurement mesh was provided using different platforms.

The Performance Measurement Platform was presented as a part of the WP6 work at the following events:

- Project Management Convention, 12-13 February 2019, Cologne, Germany
- Service and Technology Forum, 26-27 February 2019, Bucharest, Romania
- SIG-MSP, 05 March 2019, Nicosia, Cyprus
- Jisc's Networkshop 47 conference, 09 April 2019 Nottingham, UK

PMP will also be presented at the TNC conference in Tallinn, Estonia, in June 2019, including a training session on Sunday 16 June 2019, and a presentation 'Building Performance Measurement Platform - Lessons learned from operations' on Wednesday 19 June, as well as live demonstration sessions at the GÉANT booth.

2.4 NMaaS

Network Management as a Service (NMaaS) provides a solution for network management by making network monitoring and management tools available through a virtualised, cloud-based platform which is accessed via the NMaaS portal. Using NMaaS, a user can create a containerised environment with pre-installed network management and monitoring tools, solely for their own need, securely separated from other users. In this way, the user can focus on the management of the selected infrastructure they are responsible for, instead of taking care of the platform itself.

Establishing a Network Management System (NMS) requires complex, robust, reliable and usually costly infrastructure, with skilled and knowledgeable experts who need to take care of the platform and associated tools. With NMaaS, the user does not have ownership of specific costs and risks associated with the management of the NMS itself.

The NMaaS service includes three aspects:

- Providing, managing and maintaining the infrastructure of the portal, platform and selected tools.
- Supporting users in using the platform and the selected tools, and in monitoring their networks via NMaaS.
- Supporting users in contributing their software to the NMaaS platform. The NMaaS team decides which tool will be offered via the platform and supports contributors in providing their software via the platform.

The GÉANT project offers NMaaS in two forms:

- Developing and maintaining the NMaaS software for organisations (usually NRENs) who wish to download and install an NMaaS instance on their own network for their users, and providing support for those organisations.
- Providing a central NMaaS installation managed and operated by the GÉANT Operations team and at the second level supported by the WP6 Task 3 NMaaS team.

Since December 2018, NMaaS has been a production service in the GÉANT project service portfolio.

Although the Kubernetes cluster hosting the NMaaS central instance at GÉANT can be scaled up to serve an increased load, its resources are finite. Where an NREN or project foresees that it will need to serve multiple customers on its own or run multiple monitoring applications, and have full control over their instance, it is advised to have a dedicated instance of NMaaS deployed on that NREN's premises.

2.4.1 Users, Uptake and Usage

NMaaS is offered primarily to the GÉANT community. Its user groups are the following:

- Organisations that:
  - Do not want to own an NMS infrastructure themselves.
  - Want to outsource network management.
- Organisations and/or individuals:
  - Looking for quality network management software.
  - Who want to share their software with the community.

This project starts with NMaaS already installed by RENATER and Uninett to be used and managed within their NREN. Moreover, the GÉANT organisation deployed NMaaS on their infrastructure for the use of the GÉANT project, NRENs (including organisations approved by NRENs) and other GÉANT projects and services.

On the central GÉANT NMaaS installation, one customer domain is already installed to allow PMP to monitor central components. CyNet and DFN-FAU are pending deployment, awaiting the installation of a VPN connection to be able to use the GÉANT NMaaS instance securely. Other NRENs and institutions that expressed interest in NMaaS have been contacted and discussions are currently being held to establish how NMaaS can address their particular network management and monitoring requirements and use cases.

2.4.2 Key Performance Indicators

The NMaaS software stack is planned to be continuously developed in order to address recognised improvement opportunities and enhancement requests from NMaaS users. Further developments will also ease the installation process and operation of the software to improve user experience. This will result in new functionalities being added, either as requested by new users for addressing requirements or as found necessary by the NMaaS team to promote service uptake.

New tools will be integrated with NMaaS as the project progresses to enhance the offer for the existing and new users. This effort is considered a crucial factor to achieve greater service adoption. The tools will be selected based on direct user feedback and popularity among NRENs.

The availability of both the NMaaS platform and deployed instances of tools will be monitored and measured by a dedicated monitoring tool capable of querying the Kubernetes cluster API and retrieving information about the state of installed services.

The NMaaS team proactively promotes the service among NRENs and conducts meetings with interested institutions to provide detailed information on NMaaS and use cases it addresses.

KPI	Target	Baseline	Measured
Number of releases in production per year[2]	1	1	0
Number of tools added to the portfolio per year in response to customer demand[3]	2	6	0
NMaaS platform availability	99%	100%	100%
NMaaS service uptake per year (deployed on GÉANT or NREN instances)	2	2[4]	0

Table 2.4: NMaaS KPIs measured for the period 01 May 2018 – 30 April 2019

2.4.3 Activities and Issues

Since the beginning of the project, when the NMaaS service went into production, the team has been working on addressing the feedback received from PMP users.This involved, for example, enhancing the view of the configuration wizard for Prometheus and Grafana with additional custom parameters to be passed during tool deployment. As part of the continuous improvement activities and development of richer test cases, the quality of the software is being maintained at a high level.

The team has also completed various research and development tasks focused on the enhancement of NMaaS software components, the tools offered in the NMaaS portfolio and the NMaaS deployment process in Kubernetes.

One of the major development efforts was the enhancement of the configuration process for deployed tool instances with the initial data being provided by the user, using the configuration wizard in the portal. The option for updating the configuration of an already running tool was also implemented.

Significant effort was devoted to end-user-focused enhancements of the portal web user interface. Some new features were added to the portal, for example:

- A tool management panel - used by administrators and tool managers to add new or update information about existing tools offered by NMaaS.

[2] NMaaS had one major release when the service was put into production. This is the baseline value at the beginning of the project. The measured value shows the number of major releases in the first five months of the GN4-3 project.
[3] NMaaS had six tools added to the portfolio when the service was put into production. This is the baseline value at the beginning of the project. The measured value shows the number of tools added to the portfolio in the first five months of the GN4-3 project.
[4] Installed by RENATER and Uninett on their own infrastructure.

- A language management panel - used by administrators to update the content of email notifications and texts displayed across the portal in various supported languages.
- An NMaaS general information page - a dedicated page providing information about the currently deployed version of NMaaS, the changelog and a contact form allowing users to communicate with the NMaaS team.

Other subpages of the portal were also improved, for example, the application deployment view, the application ratings and comments components.

In parallel, efforts were made to define and implement a robust VPN solution in the central GÉANT instance to enable secure user access to the NMaaS tools running in the cluster, and to secure the data exchange between the tools and the managed user equipment. Moreover, the GÉANT Kubernetes cluster is being redesigned and redeployed on an enhanced IT infrastructure in an automated manner with dedicated scripts based on Kubespray [Kubespray] to simplify the cluster lifecycle management and persistent storage management based on Ceph [Ceph].

3 Services at Earlier Lifecycle Stages

within addition to running and developing the production services further, the WP6 T3 team is continuing the work from previous GÉANT projects related to the NetMon and WiFiMon tools.

3.1 NetMon

NetMon is a system for monitoring the performance of any network service (e.g., MPLS based L2 and L3 VPNs, but also new services based on SDN or service function chaining). Its unique features include monitoring each service instance (VPN, a network set up for a single user) separately and the capability to spatially localise the performance degradation or fault within each service instance. NetMon is a system which was created by reusing and integrating several well-known and reliable open source components (e.g., Grafana, InfluxDB, Ansible, RabbitMQ and OWAMP code) into a unique system for network service monitoring, and which only required a relatively small amount of software development from scratch.

NetMon supports Service Quality and Service Performance Management processes in the eTOM Assurance process area, especially the following functions:

- Pre-production service continuity tests with standards-based automated invocation and feedback results. These tests are invoked after the network service is initially provisioned to test whether the network service provides the required connectivity and automates this process.
- Continuous monitoring of key standards-based network service performance parameters, such as delay, jitter, loss or availability for each service instance once the network service is in full production.
- Dashboards for quick insight into the network service health and temporal graphs of all key performance metrics for in-depth analysis.
- The ability to detect network service performance degradation end-to-end and localise the faults and the segments where the performance is below the required level.
- Sending of alarm notifications to human operators and other OSS/BSS systems triggered by performance degradation and a parameter above or below a threshold.
- The ability to detect network element configuration errors that can occur in any virtual network.
- Periodic service level agreement (SLA) report calculation.
- Automated standards-based monitoring invocation, which enables integration with other OSS/BSS components.

The design of the monitoring system is scalable and the number of monitoring components does not grow with the number of network service instances due to smart network virtualisation use on the monitoring agents which are used simultaneously for multiple service instances.

3.1.1 Users

NetMon users are:

- Network service providers who offer monitored network services.
- Network service users.
- Other OSS/BSS components.

The following sections describe the three user groups and individual use cases.

Network Service Providers

The service assurance process area and service quality management are key elements and inseparable parts of network service operations. Network service providers must monitor the network services they provide to be able to verify the health of the products they are offering and the underlying infrastructure. This also supports the expansion of the service portfolio and increases customer satisfaction. Since the NetMon system can monitor each customer and its service instance separately, it can be used to detect configuration errors in some of the service instances. As the NetMon system is vendor and service independent, it enables the use of a single system for different types of network services. Furthermore, the system is able to track SLA and contractual obligations in order to create SLA reports. The capability to detect faults can significantly decrease fault-to-resolution process flow times.

GÉANT Community as a Service Provider

The GÉANT community is a unique network service provider environment in which multiple providers provide network services together (e.g., the MDVPN service which connects customer endpoints in different NRENs via the GÉANT network). The NetMon system recognises the specifics of multi-domain operations and allows modes of operation that support both the export of data out of the single administrative domain, or a domain's opt-out of the data export. Further, especially if monitoring zones are full administrative domains, the NetMon fault localisation feature allows a fast and easy way to detect a domain with performance problems, and so shortens the time to resolve the issues.

Network Service Users

Network service users with strict requirements about performance parameters (like latency, latency variation or packet loss) need to be able to verify that the requirements are met as defined in the SLA. The NetMon system allows read-only access to the user, who can use the dashboard to continuously monitor the overall health of their service or specific parameters.

Other OSS/BSS Components as NetMon Users

In a well-organised, automated, network service management software portfolio, users of the NetMon system can also be one or more components of some other operations support system/business support system (OSS/BSS). Typical examples are in process flows such as:

- **Initiating service testing and monitoring**

  In this process flow (see Figure 3.1), the service provisioning component requests either a pre-production test or a service to be monitored. It is using the standard TMF Service Test API [TMF653] for a specific service instance. NetMon is providing the service provisioning component with references to the service instance and the Service Test Specification (step 1)

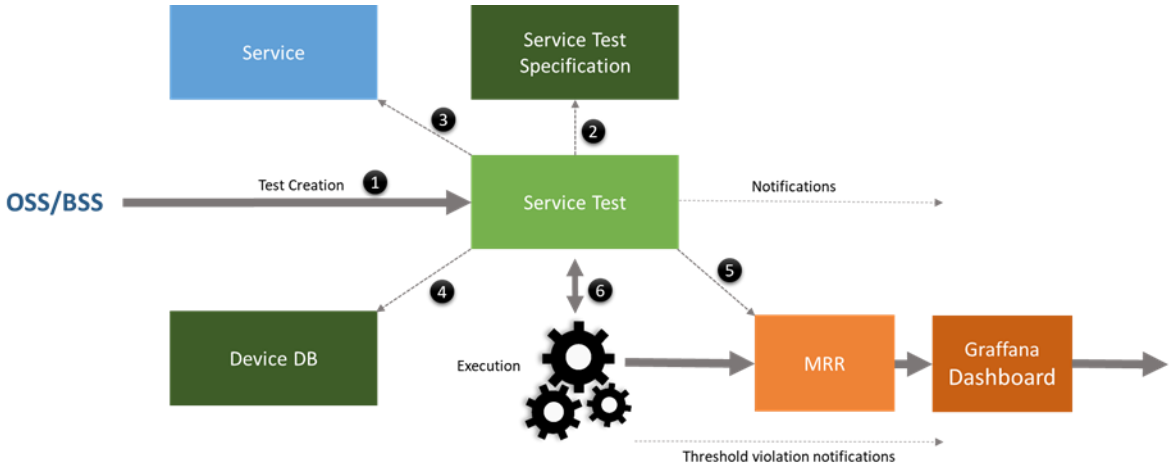

Figure 3.1: NetMon operation workflow

  The NetMon component then uses the reference to the Service Test Specification and queries the local or external Test-API compliant service (step 2) to get a set of predefined settings needed to configure a specific type of test (e.g., capture method, basic metrics and thresholds).

  Based on the data collected by that point, the NetMon component gathers the required details of the specific service instance from the Service and Resource Inventory using the Service Inventory API (step 3). The details include service termination points and other technical details needed for setting up the service (VLANs, IP addresses and so on).

  In addition, SLA data and threshold values can be gathered from the inventory, or dedicated SLA component. With this information and the information stored in the internal database (step 4), NetMon is able to automatically set up all components of the system in a way which allows service instance monitoring (steps 5 and 6).

- **Alert in response to a KPI-threshold violation**

  In this process flow the NetMon system tracks all performance indicators and compares them with predefined threshold values. If any of the values crosses the threshold an alert is created and sent. Standard alerts create a trouble ticket using the TMF Trouble Ticket API [TMF621].

However, NetMon uses a more lightweight approach of sending formatted emails to the trouble-ticketing system and/or Slack messages, as this does not require changing the existing trouble-ticketing system. Other criteria for sending an alert to reduce the number of alerts can also be configured, such as crossing a given threshold value in a specified number of consecutive tests.

3.1.2 Development Status/Lifecycle Phase

NetMon was developed in GN4-2. In the last year of the GN4-2 project, NetMon was deployed in:

- The GTS lab where it monitored CCS circuits automatically upon circuit creation.
- The GTS testbed environment monitoring L2, L3 and IP services.
- The GÉANT network where it monitored MPLS-based L2 and L3 VPNs.

These installations were a part of the NetMon pilot phase where NetMon operated reliably over long time periods. During the previous project phase, the NetMon service was developed in parallel with the NetMon software. However, it was not pushed into production due to the recognised potential opportunities to combine NetMon with other existing services like perfSONAR. The related work in this area is presented in the next section

The future NetMon roadmap depends on the number of network services used and the timeline of planned new network services, as well as the demand for network service monitoring within the GÉANT network and wider NREN community.

3.2 perfSONAR/NetMon Analysis Results

perfSONAR and NetMon have some architectural and technical similarities, such as their overall architecture (both tools have a measurement controller, agents and result repositories as key components) and their use of the same OWAMP protocol and active monitoring approach. Therefore, the GN4-3 WP6 T3 team decided to analyse the similarities and differences between the two tools in order to assess the gaps and define the potential effort required to modify perfSONAR in a way that would enable it to become a service-aware network monitoring tool.

The analysis of the two tools showed the following key differences:

- While perfSONAR was created for monitoring the performance of general (multi-domain) network links, NetMon was designed for monitoring and determining the performance verification (SLA verification) of network services, especially those established as virtualised overlay networks.

- The NetMon controller can work as a standalone tool but can also be integrated into the wider OSS/BSS architecture through TMF-compliant interfaces (to get the order for a new measurement and/or network element and configuration data required to automatically initiate measurements). This makes it possible to fully automate the process of monitoring newly established network services. While perfSONAR's pScheduler has a well-documented

API, and the automation of measurements can be done through pSconfig (the perfSONAR orchestration solution) or by direct calls to the pScheduler API, it lacks the richness of the TMF language to communicate to external OSS/BSS systems.

- perfSONAR, as a production service, has thousands of deployments worldwide, and thus a strong, existing user base, while NetMon is not yet as advanced in its product lifecycle and therefore is not (yet) widely deployed.

- NetMon uses a modified version of open source OWAMP code for its fault localisation capability, while perfSONAR currently uses a standard implementation of OWAMP that as yet does not support that feature.

- NetMon can operate in multiple network namespaces simultaneously, thus supporting virtual network monitoring use cases. Virtual networks can use overlapping IP addresses. perfSONAR can define and initiate measurements in multiple network namespaces via the context construct. At the moment, to operate in multiple namespaces, perfSONAR needs to have multiple OWAMP daemons listening on all interfaces in all the namespaces. These need to be set up once; a pending feature request exists to include this in the perfSONAR workflow.

- perfSONAR has a long-term stable development and support process with well-defined new release schedules, while NetMon has suffered from discontinuities in the development team composition between the GÉANT projects.

This analysis shows that if perfSONAR is supposed to be used for virtualised multi-tenant network services, several changes would have to be made to it. The key issue is perfSONAR's lack of support for out-of-the-box per-virtual network measurement, data storage and results display. This can be overcome using one of the following approaches:

- Customisation and amendment of perfSONAR to mimic the NetMon operation, which includes:

    1. Storage of the virtual network namespace ID in the backend:
        1. Configuration of an archiver transform to store the virtual network namespace ID.
    2. Storage of the results in an alternative backend to Esmond [Esmond]:
        - Configuration of storage to ElasticSearch (already available in perfSONAR) [Elastic].
        - Configuration of storage to use RabbitMQ (already available in perfSONAR) [Rabbit].
    3. Use Grafana for measurement visualisation:
        1. Re-use NetMon Grafana dashboard.
        2. Adapt NetMon dashboard Grafana queries to retrieve perfSONAR measurements objects/data.
        3. Retrieve and display stored OWAMP jitter measurements.
- Changes to the perfSONAR code and addition of new features, which includes:

    - Automatic and fully-featured OWAMP measurements in Linux namespaces [570].
    - Display measurements for a given time range natively in perfSONAR graphs [65].

- Display Jitter natively in perfSONAR.
  1. Retrieve and display it in graphs
  2. Add a MaDDash check and visualisation plugins for Jitter.
- Display measurements for overlapping addresses corresponding to multiple Linux network namespaces:
  1. Natively store the Linux network namespaces ID in the backend.
  2. Use the stored Linux network namespaces ID to differentiate addresses: use it in MaDDash and use it in graphs.

The first approach could produce results faster, but still requires a considerable amount of manual work and testing in order to make all the components operate reliably, an effort probably not smaller than to maintain NetMon itself. The second approach provides natural integration in perfSONAR. If neither of the two options is accepted, perfSONAR and NetMon can continue as separate network monitoring solutions.

3.2.1 Transition Plan/Roadmap

The analysis of perfSONAR and NetMon will continue, to establish whether the integration of features is practical and whether NetMon will continue as a standalone product. The roadmap will evolve to reflect emerging demand for network services. In case of lower demand for virtualised network services, the integration will be the likely option.

3.3 WiFiMon

WiFiMon has been under development within the GÉANT community since the GN4-1 project as a system for monitoring the performance of wireless networks.

The wider scope of WiFiMon is to:

- Deploy HW probes for deterministic performance measurement.
- Establish performance benchmarking by time series analysis.
- Make performance predictions for strategic purposes.
- Build WiFiMon as a service for the R&E community, with plans to examine its application within the industry, e.g., with ISPs.

WiFiMon is primarily targeted as a monitoring solution for R&E and eduroam enabled sites, but the same principles of operation can be taken advantage of in any wireless network based on IEEE 802.1X.

WiFiMon uses the unique approach of crowdsourced gathering of performance metrics - currently bandwidth and latency - from both opportunistic and deterministic measurements. Such a unique approach captures the quality of experience of each user, but also gives an insight into the overall performance of the WiFi network being measured. Measurements are triggered by users (mobile clients) when they visit WiFiMon-enabled websites and/or WiFiMon-enabled mobile apps [WCP_App]. WiFiMon performance measurement is comparatively non-invasive, testing through small downloads

from a specific server, which means that the user is unlikely to notice any performance degradation of their network connectivity while the measurements are performed.

3.3.1 Architecture Building Blocks

Mobile Client: the mobile user (client - laptop or using the app) will be authenticated at the home organisation (IdP) and authorised at the end-user's location (SP, which may be their home organisation or a visited site) on an IEEE 802.1X-based WiFi network campus.

Data Sources: the data sources (see Figure 3.2) shows (A) collected raw data gathered by the network performance tools (Boomerang [Boomerang], NetTest [NetTest] and HTML5), using JavaScript code embedded on chosen web sources that is running tests without client intervention and performance data retrieved when the HW probes visit these web sources, and (B) exporting the raw data from data sources (collectors) like Syslog, Server Log, RADIUS Accounting, DHCP logs, and L2/L3 address binding. A wireless controller and access point identifier may provide the minimum amount of information in a smaller implementation, at the most basic level for performance verification and monitoring on a wireless campus network. Therefore, the data sources will vary depending on the vendor-specific tools integrated on the campus.

Elastic Cloud (Storing, Analysing and Searching Data): the raw data is automatically collected and fed into the WiFiMon database, which uses Elasticsearch, a full-text, distributed NoSQL database (see Figure 3.2). Elasticsearch uses documents rather than schema or tables (that SQL databases use), thus allowing the accessing of data at very high speeds and making it appropriate for WiFiMon. In addition, the Elastic cloud uses technologies such as logstash [Logstash] and filebeat [Filebeat] to pipeline the raw data from the data sources simultaneously, transforming the data and sending it to Elasticsearch. For scalability and security issues, the raw data from all access points is sent to one database per site/campus.

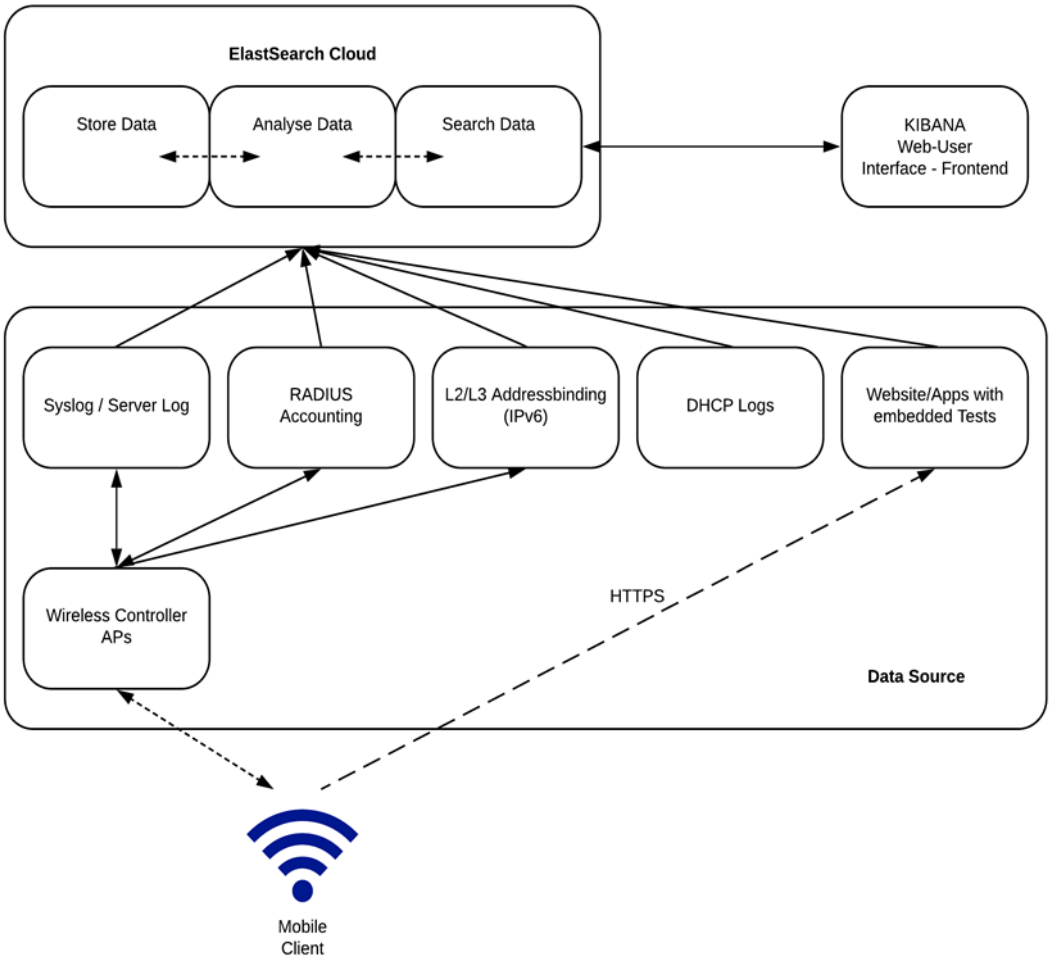

Figure 3.2: WiFiMon Architecture Building Blocks

Visualisation by web-UI Kibana, the network administrator front-end: raw and reference data from the Elastic Cloud accessible through a network admin Web-UI allows data querying (see Figure 3.2). It is the looking glass of wireless performance on the campus, which allows customised status checks of the wireless network. This architecture block is used to collect the data and to allow real-time visualisation options, such as of: collected data for a specific time period, collected data for a specific access point, min-max-mean values of download/upload/latency measurements, anomalies of measurements on time, location etc., or of the top/bottom five performing locations.

Additional functionalities for this architecture block include the ability to discover a specific access point, the ability to select from a list of access points and group them together per location, the ability to type in a name of a user based on username, the MAC address of a device, or the IP address of a device, and to track the performance of a device over time on a graph overlaid with the location information, the ability to graph the results over time in a number of different forms: scatter plot, line graph etc., and the ability to generate/create a report/graph of the performance at a location in the past versus the current performance.

3.3.2 Development Status/Lifecycle Phase

WiFiMon was piloted during the GN4-2 project and tested successfully at several institutions, conferences and workshops (e.g. TNC15 [TNC15], DCU (Dublin City University WiFi infrastructure), HEAnet Conference 2015 [HEAnet15], APAN40 [APAN40], NORDUnet Technical Workshop 2015 [Nordu15], BFH) where it proved its operational stability and capability to gather accurate WiFi performance results.

In GN4-3, the WiFiMon team has started to add new features to the system, which primarily include deterministic measurements on the dedicated small form factor hardware probes (based on single-board computers) and the implementation of the hybrid WiFiMon operation. Besides quality of experience measurements, such an approach will enable a thorough understanding of the relation between the quality of the crowdsourced measurements and the deterministic measurements gathered from fixed hardware probes, establishing an alert system, which can detect real-time performance problems based on the long-term performance history and accurate prediction of performance behaviour.

In parallel to the development activities, the WiFIMon group is working on the WiFiMon service definition which entails the following decisions:

- Deciding whether the WiFiMon service will be offered in a modular format, which means including crowdsourced, opportunistic (by JavaScript Technology), and/or deterministic measurements (by HW probes) as a hybrid concept.
- Deciding the scope of WiFIMon service which could include:
  - Providing the code for crowdsourced measurements for download to any interested NREN and end institution connected to the NREN, along with separate consultancy on how to use the system.
  - Providing the code for the hardware probes and assistance in deploying the probes.
  - Providing hardware probes with the WiFIMon system for events (such as TNC).

3.3.3 WiFiMon Roadmap

WiFiMon hardware probes will be deployed at the TNC 2019 venue with the help of the eduroam WP5 team, and the TNC2019 organisation and technical committee team members. TNC2019 will also host a WiFiMon side meeting where the solution, as well as the measurement results, will be presented, and feedback will be gathered from conference attendees.

Future work includes additional trial deployments and finalising the milestone document M6.6 which will set out the work plan for the WiFiMon service starting from the second project year (M13, January 2020).

4 Conclusions

This document presented an overview of the current status of the four production services perfSONAR, perfSONAR consultancy and expertise, the Performance Measurement Platform (PMP) and Network Management as a Service (NMaaS), and of WiFiMon and NetMon, two tools that were in pilot during GN4-2 and tested successfully in real network environments.

For the production services, the document summarised the status of the services at the beginning of the project, user adoption, KPIs and the ongoing activities on service operations and improvement. It is meant to be used as a baseline of the service status at the beginning of the GN4-3 project and for further evaluation of the service development.

For the services close to production, the document presented the current status of their development and an estimation of the timeline and steps needed for the tools to reach the production phase.

References

[570]	https://github.com/perfsonar/pscheduler/issues/570
[65]	https://github.com/perfsonar/graphs/issues/65
[APAN40]	https://apan.net/meetings/apan40/
[Boomerang]	https://github.com/akamai/boomerang
[Ceph]	https://ceph.com/
[Docs]	https://docs.perfsonar.net
[Esmond]	https://github.com/perfsonar/esmond
[Elastic]	https://www.elastic.co/
[Filebeat]	https://www.elastic.co/products/beats/filebeat
[Git]	https://github.com/perfsonar/
[HEAnet15]	https://conferences.heanet.ie/2015/
[Kubespray]	https://github.com/kubernetes-sigs/kubespray
[Logstash]	https://www.elastic.co/products/logstash
[NetTest]	https://code.google.com/archive/p/nettest/
[Nordu15]	https://events.nordu.net/display/NTW2015/Welcome
[Rabbit]	https://www.rabbitmq.com/
[Release Notes]	https://www.perfsonar.net/release-notes/
[TMF621]	https://www.tmforum.org/resources/interface/tmf621-trouble-ticket-rest-api-specification-r14-5-0/
[TMF653]	https://www.tmforum.org/resources/specification/tmf653-service-test-management-api-rest-specification-r16-5-1/
[TNC15]	https://tnc15.terena.org/
[Usage]	http://stats.es.net/ServicesDirectory
[WCP_App]	https://wiki.geant.org/pages/viewpage.action?pageId=114921524&preview=/114921524/114921527/WiFiMon_GN4-2-App_Measuring_WiFiPerf_as_Experienced_by_endusers.pdf
[Webmail]	https://lists.internet2.edu/sympa/info/perfsonar-user
[Website]	https://www.perfsonar.net
[Wiki]	https://wiki.geant.org/display/perfSONAR
[Wiki2]	https://wiki.geant.org/display/gn43wp6/perfSONAR+Consultancy+and+Expertise
[YouTube]	https://www.youtube.com/channel/UCjK-P49pAKK9hUrrNbbe0Sg

Glossary

BSS	Business Support System
CNaaS	Campus Network Management as a Service
DNS	Domain Name Service
HTTP	Hypertext Transfer Protocol
HW	Hardware
IdP	Identity Provider
IP	Internet Protocol
ISP	Internet Service Provider
KPI	Key Performance Indicator
MaDDash	Monitoring and Debugging Dashboard
MDVPN	Multi-Domain Virtual Private Network
MP	Measurement Point
MPLS	Multi-Protocol Label Switching
NMaaS	Network Management as a Service
NMS	Network Management System
NOC	Network Operations Centre
OSS	Operations Support System
OWAMP	One-Way Active Monitoring Protocol
perfSONAR	Performance focused Service Oriented Network monitoring Architecture
PERT	Performance Enhancement Response Team
PMP	Performance Measurement Platform
R&E	Research and Education
SDN	Software Defined Network
SLA	Service Level Agreement
TMF	TM Forum
TNC	The Networking Conference
TWAMP	Two-Way Active Measurement Protocol
VLAN	Virtual Local Area Network
VPN	Virtual Private Network