09-07-2014

Deliverable D8.6 (DS4.3.2): Initial Design for Refactoring GN3 Tools based on Business Process Modelling

Deliverable D8.6 (DS4.3.2)

Contractual Date:	31-05-2014
Actual Date:	09-07-2014
Grant Agreement No.:	605243
Work Package/Activity:	8/SA4
Task Item:	T5
Nature of Deliverable:	R
Dissemination Level:	PU
Lead Partner:	AMRES
Document Code:	GN3PLUS14-589-31 GN3PLUS14-589-310

Authors:	I. Golub (CARNet), Lj. Hrboka (CARNet), B. Jakovljević (AMRES), F. Liu (DFN), N. Ninković (AMRES), V. Olifer (Janet), B. Schmidt (CARNet), P. Vuletić (AMRES), M. Wolski (PSNC)

© DANTE on behalf of the GN3plus project.

The research leading to these results has received funding from the European Community's Seventh Framework Programme (FP7 2007–2013) under Grant Agreement No. 605243 (GN3plus).

Abstract

This document presents the initial design of the Service Quality Management (SQM) solution for GÉANT services. Similar performance monitoring tools are analysed in order to find potential reusable components. It also presents a solution to one of the major SQM problems – a scalable strategy for measurement point placement.

Table of Contents

Table of Figures

Table of Tables

Executive Summary

The main aim of this document is to provide the initial design of the Service Quality Management (SQM) solution for GÉANT network services. This system enables performance monitoring for each service instance. It is built on top of the resource performance components like the existing GÉANT tools (i.e. perfSONAR), network element measurement systems (such as Cisco SLA, Juniper RPM, Y.1563 or Y.1731 Ethernet OAM measurements) and other commercial or open-source performance measurement points. The key additional value of SQM compared to resource performance monitoring (RPM) designs is that it allows the correlation of raw performance metrics with service-related data; this enables service assurance through the observation of key performance indicators (KPIs) and Service Level Agreement (SLA) monitoring.

Currently, there are few SQM applications available, especially for services that cross domain borders. There is no proper solution for L3VPN service performance monitoring, which makes monitoring GÉANT MD-VPN key service parameters very challenging. There is an ongoing activity within the IETF's L3VPN workgroup that aims to standardise a methodology for in-service L3VPN monitoring, but it is still in the early stages of development. The same applies to the LMAP working group, which aims to create an architecture for large-scale monitoring of broadband links. On the other hand, the perfSONAR architecture is very similar to the current LMAP work, and making end-to-end active measurements in a multi-domain environment is often the only way to check service performance; these are always more accurate than aggregation of per-domain SLA checks, as will be shown in this document. This means that many of the components needed for a GÉANT SQM application are already present, although these need to be evaluated and customised.

Section 2 defines what SQM is, what the functionalities of the SQM system are and which Operation Support System (OSS) components SQM should be integrated with. A definition of SQM is obtained through the analysis of the standard models (TMF eTOM, TAM, OSS/J, IETF LMAP) and open source and commercial off-the-shelf tools. SQM users are defined, and some GÉANT services that could use the SQM application are identified. KPIs and SLA parameters are explained in this section as well as one of the key inputs to the SQM design: what has to be measured and how?

Section 3 analyses the service measurement components: measurement points (MPs), measurement archives (MAs) and user interfaces (UIs). This section provides an input into the later stages of SQM design through the analysis of the reusability of some existing components, dimensioning the size of databases and so on. For this purpose, both GÉANT and commercial tools have been analysed, to provide valuable input to the design considerations.

Section 4 presents the initial design considerations for the GÉANT SQM application. This aims to be as general as possible (especially those parts which do not depend on the underlying technology), but pays particular attention to the GÉANT MD-VPN service. At the end of Year 1 of the GN3plus project, the MD-VPN service is almost at the production phase, even though it still does not have an appropriate SQM application for service monitoring. The reasons for this could be that such solutions are in the initial phases of standards specification, as well as in the early stage of MD-VPN service development. One of the key problems of SQM design is the placement of measurement points. If there is a need of per-service instance monitoring, there is a natural, but not scalable solution of putting a measurement point at each user point of presence in each service instance. The possibility of using multi-homed measurement points that could at one moment reside in multiple service instances with different (usually logical) interfaces was explored. This approach can enable per-service instance performance measurements and also provide a solution to the scalability problem.

1 Introduction

1.1 Objective

The objective of this publication is to give the initial design for refactoring GN3 tools, extended to GN3plus tools, based on the business process analyses from the first phases of the project. In the first phase, SA4 T3 chose three business process groupings for further analysis: Performance Management, Problem Management and Multi-domain Service Interaction [DS4.3.1]. In the second phase, those business processes were further analysed and the Task selected Performance Management process grouping, especially Service Quality Management (SQM) as the business process to analyse in the last phase and for which the appropriate process supporting component will be designed and developed.

1.2 Why Service Quality Management?

There were several reasons for choosing Performance Management for the design in the last phase.

perfSONAR MDM [pS-MDM] is a resource performance software suite, that has been continuously developed within previous GÉANT projects. Its architecture is very similar to the LMAP architecture [LMAP-arch] now in the process of standardisation in IETF. The set of measurements supported by perfSONAR (including one way delay measurements) are the current state-of-the-art, and can be used for service performance management. In a multi-domain environment, end-to-end active measurements are often the only way to check service performance and are always more accurate than aggregation of per-domain SLA checks, as will be shown in Section 2 of this document.

Despite all these properties which suggest that perfSONAR should be widely used, there are few ps-MDM instances in active use in the community. One reason is the fact that existing performance monitoring tools in GÉANT are service agnostic, and cannot be used as-is for performance monitoring of specific services. This problem can be resolved by building a new SQM layer on top of the existing resource performance management (RPM) set of tools. This layer correlates the raw performance data with the service parameters, but also impacts the strategy of RPM measurement point placement and configuration. There are few similar recent approaches which use raw performance data from perfSONAR PS [pS-PS] for problem resolution [Pythia], or for the measurement of specific services like SDN-created circuits [GlobalNOC].

At the moment, the lack of end-to-end SQM solutions to enable per service instance performance monitoring and service assurance is obvious. Providers typically rely on measurement panels where one or a few service instances are monitored and customer experience in other service instances is estimated from the measurements in the selected instances [lmap-use-cases]. In ongoing activities, the IETF LMAP group aims to define the architectural framework and protocols for large scale service performance measurements to enable per service instance measurements, while the IETF L3VPN group aims to find a methodology to measure performance in L3VPNs which are used also in GN3plus. Both IETF activities are in the early stages of protocol specification and the implementation in devices is a few years off.

An impression may have been given that almost all the technical building blocks for SQM already exist within the GN3plus project, despite the fact that there are as yet no standardised solutions. Some components are close to being ready, others are less close to being ready for use in SQM. The development of a final SQM solution does not seem far away and could give a new boost to perfSONAR deployments, as an innovative SQM platform which can fill a gap in service performance monitoring.

The other two business processes chosen by SA4 T3 in the first phase of the work were Problem Management and Multi-domain Service Interaction. Automating Service Problem Management, especially complex problem detection which is treated by [Pythia] will probably stay in the research domain, since defining the baseline or “normal” network operation, which can then be compared to real network data in order to draw reliable conclusions, is problematic. On the other hand, Multi-domain Service Interaction covers a set of diverse inter-domain information exchange processes which cannot be supported by a single OSS component. For all the reasons listed above, SA4 T3 decided to focus on the design of the SQM supporting system.

This document describes the initial design of the SQM supporting system in the following way: Section 2 defines SQM business process, the scope of SQM and the boundaries and interactions with other business processes, like resource performance management, service inventories and alarm and ticket management systems. It gives a brief overview of similar existing tools, as well as defining the key parameters for service performance validation. Section 3 analyses the main components of the performance management systems and their suitability for SQM. Section 4 gives design considerations for the GÉANT SQM solution.

2 Service Quality Management

This section gives an overview of the Service Quality Management (SQM) process and the set of functionalities required. This overview is used as a landmark for defining the set of functionalities GÉANT SQM supporting solutions should have. In order to give a complete overview, SQM process and related components are presented from various viewpoints: business processes supported, application features and existing commercial off-the-shelf (COTS) solutions and open-source tools.

2.1 SQM process overview

According to the eTOM model, the Service Quality Management process is placed in the Operations Process Area which includes processes that support day-to-day customer and network operations and management. Specifically, SQM is located in the Assurance part which is responsible for enabling the continuous availability of the services and their performance according to SLA. The main sub-processes within SQM processes are:

- Monitoring Service Quality.
  - Monitoring and logging service performance data.
  - Comparison of monitored data against quality levels set for each service and detection of quality threshold violations.
- Analyse Service Quality.
  - Analysis of performance data and determination of root causes of performance problems
  - Improve Service Quality.
  - Restoring service quality by the means of service improvement plans, reassigning or reconfiguring service parameters.
- Report Service Quality performance.
  - Monitoring the state of performance degradation reports.
  - Providing notifications to other processes in the Service Management & Operations.
- Create Service Performance Degradation Report.
  - Creating new performance degradation report.
  - Modifying or canceling existing performance degradation report.
- Track & Manage Service Quality Performance Resolution.

  - Coordinating, scheduling and assigning restoration activities.
- Close Service Performance Degradation Report.
  - Closing of the performance degradation report after issues have been resolved.

2.2 TAM SQM and SPM applications

The Application Framework (TAM) specifies two service-level performance related components: Service Quality Management (SQM) and Service Performance Management (SPM) applications which are positioned in the vertical eTOM Assurance process group and the horizontal SID Service Management Domain. [TAM]

SPM applications are built on Resource Performance data and active service performance test data. Their aim is to provide monitoring, analysis and reporting on the end-to-end service performance in order to ensure that each service is functioning correctly. The SQM applications purpose is to establish a quality model, and to allow operators to monitor, analyse and report the levels of service being offered: they are closely connected to service level management applications.

2.2.1 Existing COTS SQM applications

TMF maintains a list of commercial tools which support Service Performance Management [TMF_products]. Some of the SQM-supporting tools listed there are a part of a bigger OSS software suite, while there are few that can be purchased separately. The authors have analysed available information about the following tools: Clarity Performance Manager [Clarity], NetCracker SQM [Netcracker], and Comarch SQM [Comarch], even though only general descriptions are given. It was difficult to find screenshots or use cases, as the SQM tool typically requires a lot of customisation because of the varieties of network resources and specific services which form the environment for the SQM tool. However, all tools share common properties, like the integration with RPM tools, Alarm management tools, Customer facing tools, Trouble ticket management tools and Inventory Management tools. The main aim of these tools is to translate resource-level performance measurement data into service-level performance indicator data, comparing that to the SLA-defined levels. Therefore, the presentation layer was typically similar to the resource performance tools (various mrtg, rrd or similar diagrams) with the addition of RAG indicators of the SLA violation.

It is important to also mention the Cyan BluePlanet View system [Cyan_BPV] as a unique cloud-based SQM solution. The business model of this solution assumes that various Measurement Agents from the BluePlanet or third-party network elements send performance data to BluePlanet cloud servers. The Planet View system processes this data, allowing the network operator to view the data, KPIs and SLA indicators.

2.2.2 Open source SQM applications

A number of open source OSS applications are available for network SQM, however with different levels of maturity in terms of project progresses, functionalities, and maintenance. SA4 T3 task surveyed five mature OSS SQM applications (OpenNMS, NetXMS, Zenoss, Zabbix, Centreon) that shared the following features:

- A large user base.
- Wide functional coverage regarding SQM.
- Established support processes, either through peer users or commercial consulting service.
- Actively developed and maintained.

The open-source tools are presented in more detail in Appendix A, from page 26. One of the main properties of these applications that make them unsuitable for monitoring multi-domain services is that they do not support the inter-domain feature that supports SQM processing across multiple organisations – they mostly have their focus on enterprise-level usage scenarios. As single-domain usage typically implies full control over the whole domain, these tools mainly rely on passive performance management.

2.2.3 SQM scope and integration with other components

SQM is one of the core OSS components and as such is directly connected to various other OSS processes. Possible inputs and outputs of the application supporting SQM, as defined by TMF OSS/J is shown in Figure 2.1. The set of functionalities as defined in [OSS/J SQM] suggests that the main sub-processes supported are: Monitoring Service Quality, Analyse Service Quality and Report Service Quality performance, while service reporting and tracking and management is not the main focus of these tools.

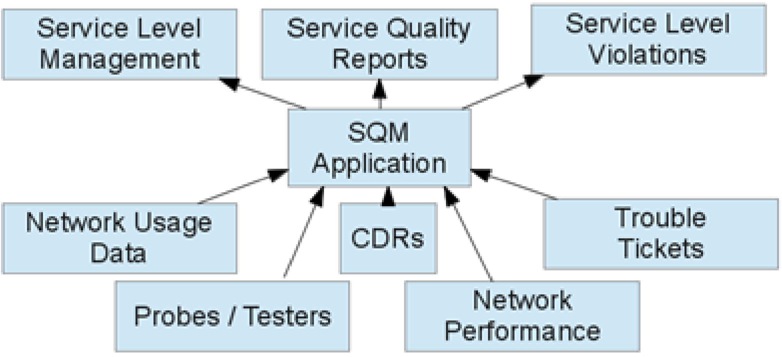

Figure 2.1: SQM application inputs and outputs [OSS/J SQM]

2.3 SQM users

A typical SQM system, as defined by eTOM, provides functionalities to monitor, analyse, improve service quality and report, track and resolve service problems during operation. Those fundamental functionalities of the SQM system can assist network operators to track the performance of network services, detect performance problems and report as well as to locate solutions once service degradation takes place. The target users of a SQM system are typically network operators or network service managers at NOCs and data centres. On the other side, users who have service contracts with strictly defined service parameters might want to be able to check the status of the service and the level of compliance with the contract. In this document, and in the overall design of the SQM system, it is assumed that the main users of the SQM system are service managers and NOCs within a service provider. Also SQM design will specify the possibilities for the extension of the SQM over the last mile (towards the end user).

2.4 What does SQM monitor?

2.4.1 Service Level Agreement – SLA

An SLA generally formalises the business relationship and agreements that should ease the operation and, in some cases the automation, of various business processes between different entities. The SLA content depends on negotiated services between involved parties. It describes the expectations and obligations of each party by providing a resolution to situations when one or both parties fail to fulfil expected guarantees. In general, SLA content is proposed with five basic elements: Service identification, the Business part, the Service part, the Technology part and QoS reporting [M.3342]. In more detail:

- Service identification, as the name implies, should provide a summary description of the service provided.
- The Business part should contain, but is not limited to, terms and conditions, responsibilities, different kinds of procedures (change, termination, upgrade, etc.), penalties and violations of the contract signed between involved parties. The business part addresses legal aspects of the SLA.
- The Service part contains the agreed level of service expected, along with a more detailed description of the service. The Service level should not be detailed as it relies on the technology part of SLA where more thorough description is provided relating to the guaranteed QoS level.
- The Technology part contains details about service level and QoS that two parties agree on. Definitions of the KPI are provided as well as their values that must be guaranteed by the service provider. Depending on the complexity of service, the technology part will vary in terms of detail. Most notably, this part of the SLA specifies technical and technological aspects, which is important during the service provisioning phase in terms of identifying SLA violations.
- The QoS reporting part contains details regarding reporting frequency and the way that QoS measuring is carried out with respect to definitions and metrics defined in the technology part of the SLA. QoS reporting is extremely important when there are conflicts in measurement results taken by both parties. Measurements must be carried out by all parties consistently according to the specifications in the SLA part.

Additionally, [Y.1241] states that a complete SLA may contain details regarding the service level specifications, service performance objectives, monitoring components (with associated measurement details) and financial compensations activated by a violation of the terms of the business part of the SLA.

2.4.2 Key performance indicators – KPI

A service-oriented architecture demands that the service level is guaranteed by quantifying performance. Using the performance metrics, an SLA may be verified according to the requirements presented at the SLA definition stage. Metric values have to be the result of the service provider's offer (depending on the technical capabilities and present technologies) and the user's performance request (depending on the services and applications that will be used). After identifying the user's performance request, KPIs and targeted values are specified that the service provider must guarantee. Furthermore, [Y.1540] and [Y.1541] specify in detail the definition of specific performance metrics and KPIs through service differentiation described using a number of service classes.

Production environments for service oriented architecture usually address three performance metrics: packet latency, latency variation and packet loss rate (PLR). Metric definitions are further addressed in [Y.1540] where, specifically, packet latency and delay variation may be calculated as minimum or mean values during a measurement period. In order to keep track of SLA validation, metric measurements must be carried out in the same meaningful way (specified in QoS reporting part of SLA) by both parties. Exceeding the negotiated metric value constitutes SLA violation which further activates financial compensations and/or penalties expressed in service credits. Initial performance objectives should be defined bearing in mind that KPIs are attainable and reasonable, thus avoiding unnecessary SLA violations.

For different services and applications, there are specific performance metrics that must be guaranteed in order to satisfy service perception. Real-time applications demand guarantees of all three metrics, whereas applications like file transfer and web browsing only need guarantees for PLR. The initial SLA definition requires that services and applications are identified and subsequently, performance metrics and their targeted values determined are attainable and reasonable.

2.4.3 Multi-domain SLA and performance measurement approaches

The bilateral agreement model presents the prevalent way of E2E service negotiation in the Internet today where each provider negotiates service with neighboring providers. Delivering inter-provider QoS is very complex as it depends on the performance guarantees of each provider on the end-to-end path. Consequently, multi-domain SLA takes into account complex interdependencies and responsibilities of each involved party in delivering an end-to-end service. A multi-domain service relies on each provider delivering negotiated performances, where failure of just one may result in a failure of the end-to-end service.

Performance measurements should be carried out in a scalable and non-intrusive manner in order to collect accurate performance information, while being transparent to production services. Scalability must be addressed with end-to-end QoS measurements, especially in multi-domain scenarios in which separate providers cannot be held accountable for failure of performance guarantees outside of its domains. Therefore, two methods need to be used – the end-to-end measurements and the metric composition approach.

The End-to-end measurement approach assumes that measurement points are placed at the topological end points of the service instance measure of key service performance indicators. In this approach the number of measurement instances increases proportionally to the number of service instances and the number of service points of presence (end points). The End-to-end measurement approach is thus inherently less scalable than the metric composition approach described below. It does, however, provide the most accurate performance measurement methodology in multi-domain environments. It assumes that each service instance traversing multiple providers has a separate measurement instance, measuring the actual end-to-end performance as perceived by the specific service instance.

Metric composition combines measured KPI and SLA parameters of service instance components from the domains participating in the delivery of one service instance and provides an *estimation* of the end-to-end service performance. It uses native characteristics of performance metrics, since latency is additive and jitter is approximately additive. The packet loss rate (PLR) is indirectly multiplicative[1] but it may also be considered additive for values less than 10^{-3}. Separate measurements in each domain that the specific service instance is traversing are collected (measurements are made between measurement points between domains), enabling the calculation of end-to-end performance. This is also called “spatial metric composition” as it takes into account the spatial aspect of measurements carried out at the same time in each provider’s domain constituting the end-to-end path. Spatial metric composition in multi-domain scenario is depicted in Figure 2.2.

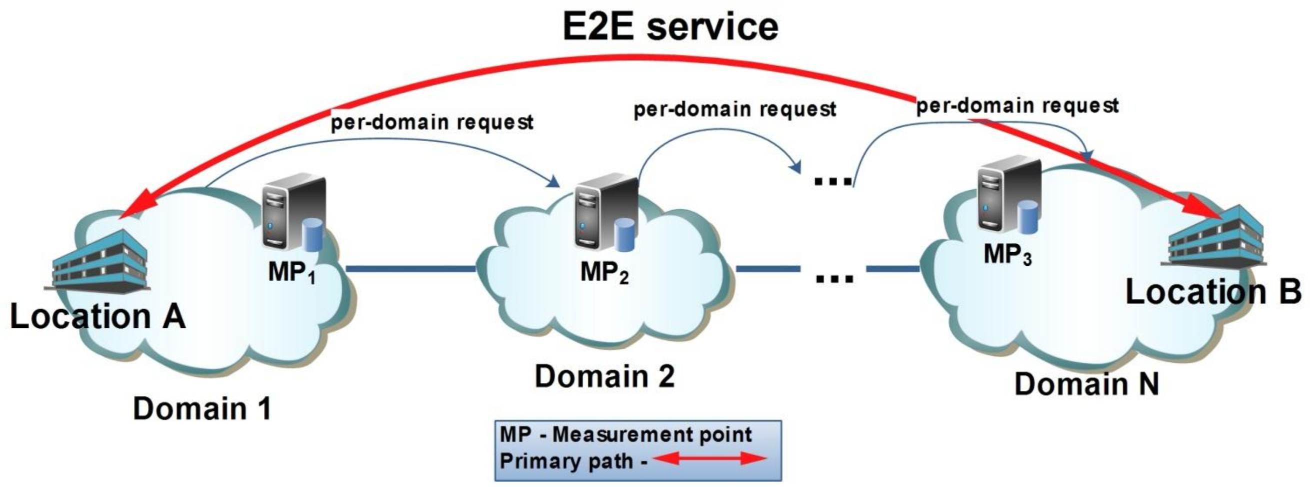

Figure 2.2: Spatial metric composition in a multi-domain scenario.

The metric composition approach alleviates the requirement that each service instance needs a separate end-to-end instance, trading off an increased level of scalability with reduced measurement accuracy. Accuracy and reliability of the SLA verification using this method is lower than in the end-to-end measurement approach because of the difficulty of conducting measurements at the same time in all the participating domains. If measurements are not well coordinated in time in different domains, some temporary short-term performance problems in a service instance would be captured only in a subset of domains which have measurements scheduled during that interval, giving a false overall end-to-end performance estimation. Also, difficulties with the inter-provider link measurement and in stitching all end-to-end measurements reduce the reliability of the composed SLA metrics. Recent analysis [commag] has shown that more accurate measurements are made when performing metric composition using minimal latency and when latency variation values can be measured in separate domains. Additionally, accuracy may be increased if separate measurements in domains are carried out over network paths consistent with paths taken by the end-to-end service. Collecting measurements from separate domains may also encounter difficulty since some providers may be unwilling to provide

[1] PLR is indirectly multiplicative metric indicating that PLR on end-to-end path consisting of *N* domains depends on separate PLRs measured in separate domains according to $\mathrm{PLR} = 1 - \prod_{i=1}^{N}(1 - \mathrm{PLR}_i)$, where PLR_i is packet loss metric measured in domain *i.*

technical details to external entities. Detailed specifications regarding metric composition when considering different metrics may be found in [RFC 5835] and [RFC 6049].

3 SQM system components

This section analyses the main entities of the Service Quality Management architecture: Measurement Agents/Points, Measurement Archives/Collectors and User Interface. The main aim of this section is to analyse whether existing GÉANT-developed components can be reused for a SQM system through the comparison with similar commercial systems, and to provide the input for the specification of these components (such as UI content or dimensioning the size of databases, etc.).

3.1 Measurement points

Measurement points (MPs) are software or/and hardware entities, sitting on network nodes, which produce measurement data characterising network performance. MPs can be of a different nature and able to produce different types of data, contributing to the calculation of different KPIs. The terminology in this area is not completely standardised so synonyms of ‘Measurement Point’ are: measurement agents (IETF LMAP terminology), probes, and measurement hosts.

The more types of MPs an SQM supports the more powerful and flexible it is. Unfortunately, the common situation is that a SQM supports only a few types of MP, or even only one type implemented as a proprietary software agent. Hence, extending the supported MP types is a very important component of improving SQM functionality. In this section a brief overview of the main types of MPs is given.

MPs may be classified according following characteristics.

- Software or hardware based.
  - A software MP is a piece of software working on a host connected to a network. It is relatively easy to deploy as a software MP can work on a host with many other applications. However, a software MP can be used only for end-to-end performance measurements and not for segment-to segment measurements as a host can’t be “put inside” a provider or corporate network.
  - A hardware-based MP is an element of network equipment; it can be embedded into network boxes of any type, for example into routers, switches and multiplexers. An example is an SNMP agent supporting Management Information Base (MIB-II) and capable of producing data about the interface status and throughput. A hardware-based MP can either be part of a network box – doing something other than performance measurement (for example, routing customer traffic) – or it can

be part of a device completely dedicated to network performance measurement. A hardware-based MP can be used both for end-to-end and for segment-to-segment measurements.

- Active or passive MPs.
  - An active MP generates extra traffic to measure some KPI while a passive MP uses existing traffic to calculate a KPI. A SNMP MIB II agent is an example of passive MP while iperf software that generates traffic to measure achievable TCP throughput is an example of an active one. Usually active measurement methods are the only truly available option when the service crosses the boundary of the administrative domain.
- A layer of the protocol stack at which KPIs an MP works.

  According this criterion an MP can work at Layer 1 measuring DWDM KPIs, Layer 1 measuring SDH or OTN KPIs, Layer 2 measuring Ethernet KPIs, Layer 3 measuring IP-related KPIs, or at the Application layer measuring KPIs of a particular application.
- Standards-based or proprietary.

  There are three kinds of standards to which an MP can comply:
  - A KPI-wise standard that defines what KPI or KPIs an MP can evaluate. For example, an MP can measure the IP packet one-way delay metric defined in RFC 2679, or defined in Y.1540.
  - A measurement protocol standard that defines a protocol where a pair of MPs is used to measure a KPI. OWAMP and Y.1731 DMM are examples of a protocol that is used for measurement of the one-way delay KPI at IP and Ethernet layers respectively.
  - An access method standard that defines how an agent can be accessed to obtain measurement data. SNMP is an example of such kind of standard; other methods can be based on ftp, scp or http.

Software-based MPs available in perfSONAR and CMon, and hardware-based MPs available in network equipment are described below.

3.1.1 GÉANT tools perfSONAR and CMon

3.1.1.1 *perfSONAR MDM*

GÉANT perfSONAR MDM, as well as the perfSONAR PS (from ESNet and Internet2) both use the perfSONAR protocol specified by OGF NM-WG to exchange data and has flexibility, extensibility, openness, and decentralisation as its main design goals [3.1-1].

At the IP layer, both support two types of MPs for the measurement of:

- Delay, jitter and IP packets loss.
- Achievable TCP and UDP bandwidth.

The first MP type uses the One Way Delay Measurement Protocol (OWAMP), as defined in RFC 4556. perfSONAR MDM has two different implementation of OWAMP MPs – one as a part of the HADES system and another one as a part of the OWAMP system which is compatible with ps-PS. The second MP for achievable TCP and UDP bandwidth is a wrapper around the Bandwidth Test Controller (BWCTL) Measurement Tool, which is a client/server program developed to simplify running other software measurement tools between hosts. At the moment, BWCTL can run iperf, thrulay, or nuttcp for bandwidth measurement. The perfSONAR architecture allows new MP types to be included relatively easily by writing an appropriate wrapper that supports the perfSONAR protocol while taking into account specific features of the MT.

3.1.1.2 *CMon*

CMon (Circuit Monitoring) is a software system for performance monitoring of Layer 2 Ethernet circuits, developed by the GÉANT project. The CMon architecture is similar to perfSONAR's, as it supports perfSONAR MPs/MTs and allows gathering of measurement data from third-party monitoring proxies or directly from network equipment through AGT (CMon AGenT). At the time of writing, CMon AGT was capable of supporting passive monitoring of Ethernet circuits status (Up or Down) by polling SNMP MIB agents of network interfaces along a circuit path.

3.1.2 Support on network equipment

Modern network equipment can be a powerful source of measurement data. Many router and switch models have embedded agents which can carry out performance measurements and produce useful data for SQM. Quite often these agents are in a dormant state but once activated they become MPs that can give very valuable information without a need to buy and install new equipment. Embedded MPs can be used both for end-to-end and segment-by-segment performance measurements. Table B.1 on page 28 shows some available embedded performance measurement agents of different kinds, from a standard SNMP MIB II agent to additional services of routers/switches OS like Cisco IP SLA.

3.1.3 Dedicated appliances

Dedicated performance measurement boxes are available on the market. They could be hosts with performance measurement software installed or routers or switches with rich monitoring functionality. Both types of devices can be used for end-to-end performance measurements, but routers/switches can also be used for segment-by-segment measurements with the primary purpose as customer-provided demarcation boxes. Some dedicated performance measurement appliances are presented in Table B.2 on page 28.

3.2 Measurement Collectors/archives

Measurement archives are one of the key components of performance management systems. These databases store past measurement data and allow various analyses to be conducted. The current LMAP framework [LMAP-arch] specifies three basic elements: Measurement Agents, Controllers and Collectors. The Collector accepts a Report from an Agent with the Measurement Results from its Measurement Tasks. It then provides the Results to a *repository*. A results repository records all Measurement Results in an equivalent form, for example an SQL database, so that they can easily be accessed by data analysis tools. perfSONAR SQL MA and HADES MA, developed within GÉANT projects, are examples of measurement archives.

3.2.1 perfSONAR MA

The perfSONAR Measurement Archive (MA) service is used to publish historical monitoring data which is stored in an archive. It acts as a wrapper around an existing data archive to provide data to the outside world. The archive can be, for example, a network's Round Robin Database (RRD MA), relational database (SQL MA) or a proprietary database of a Network Management System. Additionally, an MA can publish information produced by MP services. It does not create (generate new raw data) or transform (i.e. aggregate/correlate/filter) any data.

The SQL Measurement Archive (SQL MA) stores link data that is collected by measurement tools. It provides the data from the following measurements:

- IP interface link utilisation.
- IP interface link capacity.
- IP interface input errors.
- IP interface output drops.
- Circuit / lightpath status.
- Achievable throughput (TCP).
- UDP throughput.

Data can be accessed using for example the perfSONAR UI web client (for IP link utilisation). Technical details of the existing perfSONAR MA are given in Appendix C.

3.2.2 Measurement Archive on a cloud – Blue Planet Platform

Cyan Blue Planet is the platform designed for service providers to simplify the development, deployment, and orchestration of scalable network-based services. Blue Planet is comprised of multiple individual components: Blue Planet Platform, Blue Planet Applications, Third-Party Applications, Northbound APIs and Third-Party Element Adapters.

Planet View is Cyan's performance monitoring and SLA assurance application that allows a network operator to provide real-time and historical visibility of the performance of their services to end-customers via a customised web portal. Planet View is a cloud-based application that may be deployed in a SaaS capacity. With this operational architecture, performance monitoring data from network elements is captured by the metrics collector on a Cyan-provided appliance, which then forwards the data to the Planet View application running in the cloud. Planet View then partitions the data by customer/user and makes it accessible via a secure login.

3.3 User interface

The user interfaces of SQM and RPM have many similarities. Tools like NetCracker SQM or Cyan Planet Blue have in addition to the RPM parameter monitoring, RAG indicators for SLA parameter violation and maps of service instances. The perfSONAR user interface allows visualisation of past measurements which are stored in the MAs (HADES MA stores measurement of delay, jitter and packet loss, while SQL MA can store throughput measurements performed in the past). The user can also initiate an on-demand measurement between two MPs. Users can choose the set of measurements they want to monitor from the set of MPs and MAs that are registered to the perfSONAR lookup service (LS) and grouped into so-called “services”. One “service” is defined as the data from one MA or from one MP. The data about all the measurements is available to all perfSONAR users, and there is no possibility of providing different views for different services and restricting the data exposed. The perfSONAR UI can simultaneously display the data of one measurement (e.g. the delay between two endpoints). There is no dashboard which would allow users to permanently monitor the status of their services. However, perfSONAR UI is customisable and the information about the SLA and KPI can be easily added to it as well as to custom dashboards. Such a dashboard representation of measurement results is used for perfSONAR PS [pS_dashboard] and provides a very useful, first-glance indication of the network status.

4 SQM for GÉANT-NREN environment – Design considerations

Previous sections outlined the main features of the SQM system and described some of its key components. Several important conclusions can be drawn from that analysis that impact design decisions of the GÉANT SQM solution:

- SQM architecture relies heavily on the underlying Resource Performance Management (RPM) monitoring system and depends on the type of service. While the service level part of the SQM system can be unique for various services, different services can require different measurement agents depending on the layer of operation (e.g. Ethernet based services might need Y.1731-compliant agents, while L3 services could use OWAMP or similar solutions).
- At the moment there are no standardised protocols for the communication between key SQM components, and the standardisation procedure is still in an early phase.
- Existing perfSONAR architecture is very similar to the developing IETF LMAP architecture. Measurement methods deployed in perfSONAR are the current state-of-the-art, and the active measurement approach in multi-domain environments gives more accurate SLA verification than the approaches with per-domain SLA metric composition.
- Key SQM components like Service/SLA inventory and the engine for the correlation of the raw measurement data and service-specific KPI are missing in the current perfSONAR architecture which is not service-centric. There are commercial SQM solutions available; however such solutions require a customisation effort towards the specific service and SLA parameters.
- Measurement archives for the current and expected number of service instances do not require a significant amount of space that would compel the use of large storages or clouds.
- perfSONAR measurement points do not at the moment support certain specific service performance measurements (e.g. Ethernet OAM measurements from network elements). However, perfSONAR can be easily extended to accept and store measurements from other devices.
- perfSONAR UI can be extended to provide SQM-compliant dashboards and indicators
- Some measurements (such as throughput measurements) available in the present perfSONAR architecture, are not found in typical SLAs. Such measurements are typically not used periodically for continuous performance measurements but rather at the acceptance phase of the new service instance or when some service problems are being debugged or resolved.

These conclusions present a key input into the design of the GÉANT SQM system. The aim of the SQM design is to create a solution for service quality and performance management that will be as general as possible, however with the awareness of the differences that might exist at the resource performance layer. The design also aims to reuse as much as possible of the existing perfSONAR architecture and components which are of an up-to-date design. Since SQM solutions are customised for particular services and MD-VPN is one of the GÉANT services that currently does not have an appropriate service performance management solution (and a standardised solution is unlikely to be made available soon), the GÉANT SQM solution will be designed with the MD-VPN service in mind and MPs will be specified for the case of a multi-domain multi-point MPLS VPN service.

4.1 GÉANT MD-VPN service monitoring – current status

The Multi-domain Virtual Private Network (MD-VPN) service spans over multiple domains of control and administration and includes multiple NOC/NRENs operational teams. In order to be able to monitor the performance at demarcation points as well as the service as a whole, it is required that service quality must be managed under the SQM system, through which performance between the Provider Edge routers can be perceived. NOC/NRENs engineers as targeted users of SQM system can thus gain an overall view of network services in such an inter-domain scenario.

The MD-VPN task currently uses a dedicated VPN instance for service monitoring and the smokeping tool [Smokeping]. Such approach with the dedicated monitoring service instance does not allow per-instance service quality management and cannot accurately capture the quality of experience of all service users, but is the approach nowadays often used by service providers due to the lack of standardised methods for monitoring Layer 3 Virtual Private Networks.

4.2 MPLS L3VPN monitoring – new standardisation efforts

At the moment there are no appropriate standardised solutions for L3VPN monitoring which significantly affects performance monitoring of the MD-VPN service. The IETF Layer 3 Virtual Private Network (L3VPN) working group currently aims to specify L3VPN performance monitoring standards and methodologies, with the drafts prepared by the main equipment manufacturers (e.g. Cisco, Huawei). This fact suggests that the aim is to have L3VPN measurement capabilities (measurement points) built into network devices. The work is in progress, and still in the early stages: the three Internet drafts that the group is actively preparing are in the second or third revision:

- The recently expired draft [draft-zheng] summarises the current performance monitoring mechanisms for MPLS networks, and challenges for L3VPN performance monitoring. To perform the measurement of packet loss, delay and other metrics on a particular VPN flow, the egress Provider Edge (PE) router needs to recognise to which specific ingress VPN Routing and Forwarding (VRF) a packet belongs. But in the case of L3VPN, flow identification is a big challenge. According to the label allocation mechanisms of L3VPN, a private label itself cannot uniquely identify a specific VPN flow and as a result, it is not feasible to perform the loss or delay measurement on this flow. As a conclusion from the draft, performance measurements cannot be performed in L3VPN networks without any extensions or alteration to the current label allocation mechanisms.

- A more recent draft [draft-dong] proposes the framework and mechanisms for the application of L3VPN performance monitoring. In current MPLS technology implementations, for a particular VPN prefix, the directly connected PE routers allocate the same VPN label to all the remote PEs which maintain VPN Routing and Forwarding Tables (VRFs) of that VPN. This concept of work is the reason why performance monitoring cannot be performed on the egress PE, since it is not possible to identify the source VRF of the received VPN packets. To resolve the above mentioned issues it is critical for the egress PE to identify the unique VRF, i.e. to establish the point-to-point connection between the two VRFs. Once the point-to-point connection is built up, current measurement mechanisms may be applied to L3VPN. In this way, the new concept of the "VRF-to-VRF Tunnel" (VT) is introduced. In this concept, each PE router needs to allocate MPLS labels to identify the VRF-to-VRF tunnel between the local VRF and the remote VRFs (labels are called VT labels). For each local VRF, the egress PE router should allocate different VT labels for each remote VRF in PEs belonging to the same VPN. This guarantees that the egress PE could identify the VPN flow received from different ingress VRFs, and the packet loss and delay measurement could be performed between every ingress VRF and the local VRF. When a VPN data packet needs to be sent, the ingress PE router firstly pushes the VPN label of the destination address prefix onto the label stack. Then, the VT label allocated by the egress VRF should be pushed onto the label stack, to identify the Point-to-Point connection between the sending and receiving VRF. At the end of MPLS label stack encapsulation, the outermost LSP label is applied. When the VPN data packet arrives at the egress PE, the outermost tunnel label is popped and then the egress PE could use the VT label to identify the ingress VRF of the packet. After this de-encapsulation, the procedures for the packet loss and delay measurement, as defined in [RFC6374], can be utilised for L3VPN performance monitoring.
- A further draft [draft-l3vpn-pm] introduces and describes the BGP encodings and procedures for exchanging the information elements required to apply performance monitoring in MPLS/BGP VPN. To achieve this, a new sub-address family, called VRF-to-VRF Tunnel (VT) Subsequent Address Family, is introduced.

The description of the status of L3VPN monitoring work given above suggests that currently there are no mechanisms for performance monitoring inside L3VPN service instances. Time will be needed for these drafts to be approved and implemented by different network equipment vendors. Because of this, to be able to measure packet loss, delay, jitter or any other performance metrics inside L3VPNs, some external tools or applications will need to be implemented and used.

4.3 MDVPN SQM – scalability issues

MDVPN is a multipoint multi-domain service. As described in Section 2.4.3, the most accurate approach to measuring the service performance of such services are end-to-end measurements, but in this case, scalability is an important consideration that needs to be addressed. If there are *n* service instances and each service instance has *m* points of presence (although this number can be variable per instance), and if per service instance measurements are required, the total number of measurement points is *mn*, because to each point of presence one measurement point has to be placed. If MPs have large footprint and with the growing number of service instances, it is clear that such solution is not scalable.

A scalable architecture should facilitate the adding of (or removal of) measurement probes for new service instances, without any new MP hardware requirements. In the case of new service instances, an MP should only be reconfigured to support service-specific measurements.

Consequently, a scalable measurement design should provide service-aware architecture and be capable of future expansions only through MP configuration changes related to specific service instances. It is evident that measurement scalability can be achieved using a single MP performing measurements for specific services. However, a number of issues appear as a result of this assumption, mainly on the network and application levels. For a measurement protocol to be used as a scalable solution it should support multiple instances established on a single MP meant to be used for separate services. The use of virtualisation in order to achieve a high level of scalability in a single MP is not an option, as it significantly reduces measurement accuracy to the extent that SLA validation is not possible. At this point, the MP should also implement a certain level of privileges and access to the MA, specifically the part of the MA containing measurement data for a specific measurement instance. A *measurement instance* denotes a single instance of a measurement application conducting measurements for separate service instance with defined user privileges.

An additional problem resides on the network level and the way routing is realised as it requires that separate measurement probes are routed over the network path that a particular service instance is using. MP multi-homing requires a modification in the routing of outgoing packets, since measurement probes generated by separate measurement instance must be placed on the path that a specific service is using. Furthermore, for services like MDVPN, the problems of address overlapping and logical separation have to be addressed. These aspects strongly affect the positioning of the MP when designing measurement solutions for specific services. A measurement solution, bearing in mind scalability, should not require more than a single physical interface. However, due to availability concerns, redundant connections should be provided to achieve the recommended redundancy.

4.3.1 Proposed scalable measurement solution

The proposed high-level design addresses most of the previously mentioned requirements for measurement scalability. For the test scenario the OWAMP application developed by Internet2 was chosen, since this was able to address the majority of the aforementioned issues. OWAMP’s command-line client is an implementation of the OWAMP protocol, as defined by [RFC4656]. AMRES and CARNet used version 4.3 of OWAMP, which is compatible with IPPM performance metric definitions.

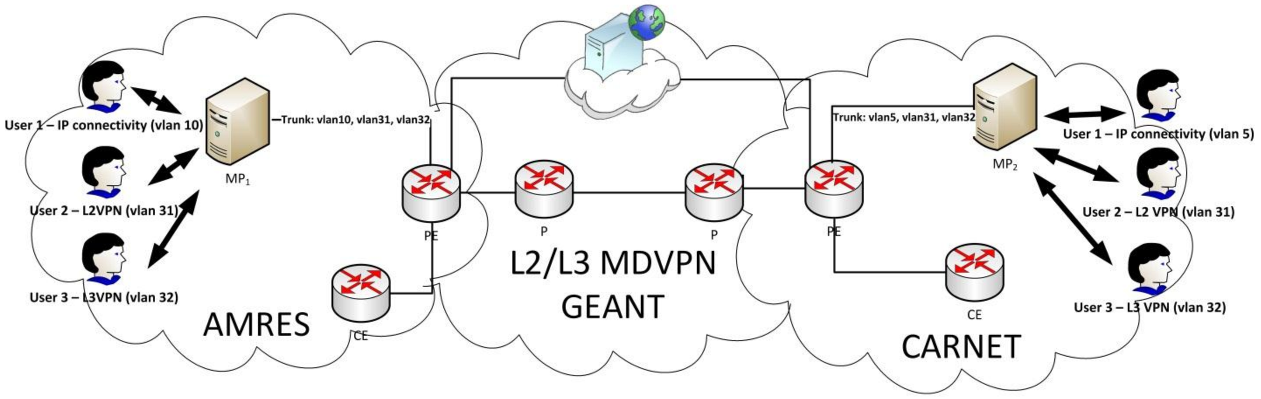
User 1 – IP connectivity (vlan 10)
User 2 – L2VPN (vlan 31)
User 3 – L3VPN (vlan 32)
MP$_1$
Trunk: vlan10, vlan31, vlan32
PE
CE
AMRES
P
P
L2/L3 MDVPN
GEANT
PE
Trunk: vlan5, vlan31, vlan32
MP$_2$
User 1 – IP connectivity (vlan 5)
User 2 – L2 VPN (vlan 31)
User 3 – L3 VPN (vlan 32)
CE
CARNET

Figure 4.1: High-level overview of proposed measurement solution in testing environment.

It was proved that it is possible for OWAMP to create multiple instances by performing replication of OWAMP files and activating separate OWAMP daemons – *owampd*. Created OWAMP instances may be assigned to different users defined on the MP, with each instance having its own separate configuration and not dependent on existing OWAMP instances. The maximum number of instances is practically constrained by the hardware capabilities of the MP. Separate users assigned to manage specific measurements may start measurements using the *owping* command, in practice starting the connection on a pre-configured TCP control port towards an MP that has the *owampd* command activated and configured to listen on that same port. Sockets which separate OWAMP instances are using should be mapped according to the service (service ID, VPN ID, circuit ID, etc.). In this way, a more intuitive configuration and easier management may be achieved.

Following the creation of multiple OWAMP instances and according to user privileges, outgoing routing sends packets through a sub-interface assigned to specific instance. If measurement is performed between MPs that are targeted using public addresses, conventional routing may be applied. However, for services where address duplication and private addressing is possible, such as VPN, conventional routing is not able to determine an appropriate outgoing interface. In order for the proposed model to be applied, the *iptables* command is used, more specifically mangle table[2], where it is possible to classify each packet according to the owner of the application. Packets are marked and using the *iproute2* package it is possible to create multiple routing tables (to which marked packets are directed) for each measurement instance and send packets through an appropriate sub-interface. This resolves the problem concerning address overlapping and private addressing. For the purpose of scalability, sub-interfaces are created and they are assigned to separate VLANs. The MP is connected to the rest of the infrastructure using the trunk. A high-level overview is shown in Figure 4.1.

4.3.2 Measurement solution test results

Continuous delay and packet loss measurement tests were organised between AMRES and CARNet according to the proposed measurement solution. Measurements were made over three sets of service: basic IP

[2] Packet mangling refers to the process of intentionally altering data in IP packet headers before or after the routing process. In this scenario, *iptables* mangle tables are used to classify locally generated packets by marking them before the routing process. Classification can be done in a number of ways, including the definition of packets generated by specific user or user groups.

connectivity between MPs, L2 and L3 MDVPN. On both sides, MPs were positioned on PE routers so that measurements for L2 and L3 VPNs could have higher accuracy.

Results have shown that our model produces very accurate results when basic IP connectivity performance between MPs is measured. Two OWAMP instances were tested and periodic measurements were scheduled. The measured results were constant during the interval in which measurements were made and an additional check was performed using ICMP packets, which supplied the RTT information. The measured delays between AMRES and CARNet using the proposed model is shown in Figure 4.2.

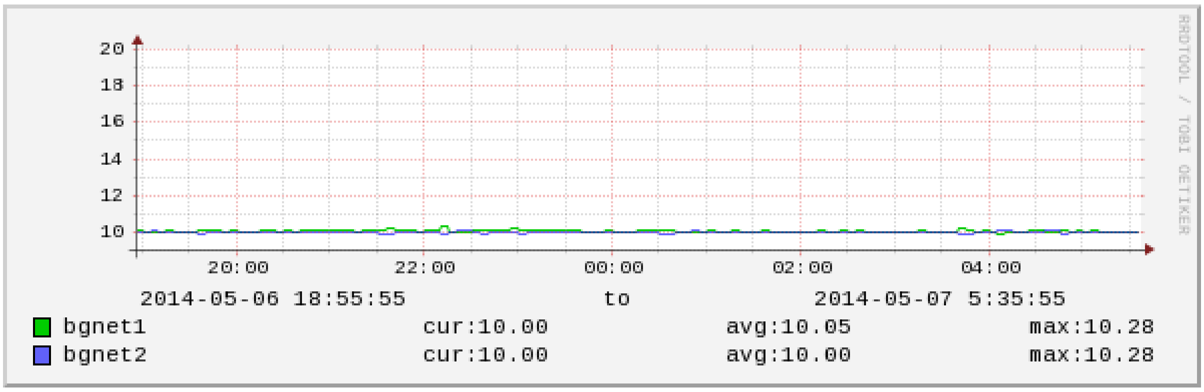

Figure 4.2: Measured delays between AMRES and CARNet

On the other hand, measurements over L2 and L3 MDVPN instances between AMRES and CARNET created certain problems. OWAMP behaviour was not consistent compared to the previous case. Firstly, measurement took a lot longer, which was identified as a problem in the TCP control part of the OWAMP protocol. However, in this phase of the work it was not possible to identify the cause of the delay in generating probe packets over L2 and L3 VPNs. Also, the number of generated probe packets is random although configured as deterministic. OWAMP does not report missing packets as lost, but this problem can be resolved by increasing the OWAMP probe timeout so that probe packets are definitely received during the measurement period. Secondly, there were cases when the MP was able to receive only results from one way and not the other. This was resolved by decreasing the MTU on the MP on the AMRES' side. Measurement results lag still persisted, though.

Despite emerging problems regarding the measurements in L2 and L3 VPN scenarios when deciding to use this solution in production, it should be pointed out that measurement results were consistent with the used transmission path. At this point, the application of our proposed measurement solution regarding L2 and L3 VPN, gave inconclusive results, as the current problems were not adequately countered due to lack of time. A high level of scalability and accurate results for basic IP services provides a strong incentive to resolve these issues.

4.4 SQM overall architecture

Previous analyses of the GÉANT OSS portfolio [GN3 DJ2.1.1] show that some of the components in Figure 2.1 on page 7 already exist within GÉANT (e.g. Probes/Testers, Network performance and Network Usage data are within the scope of the perfSONAR tool), but the rest of the shown portfolio does not exist. Also, there is no GÉANT service inventory that stores the information about service instances and data about SLA or other KPI parameters for those service instances. Such a service inventory is a requirement for SQM process automation, especially with the increased number of service instances, and it provides interfaces with the SQM component of commercial solutions (such as Clarity Performance Manager, mentioned in section 2). perfSONAR has some elements of the service inventory: measurements are grouped into sets based on the MP–MA measurement archive association, but these sets do not have the appropriate service context.

Figure 4.3 (a) shows the set of components that are within the scope of SA4 T3's SQM design.

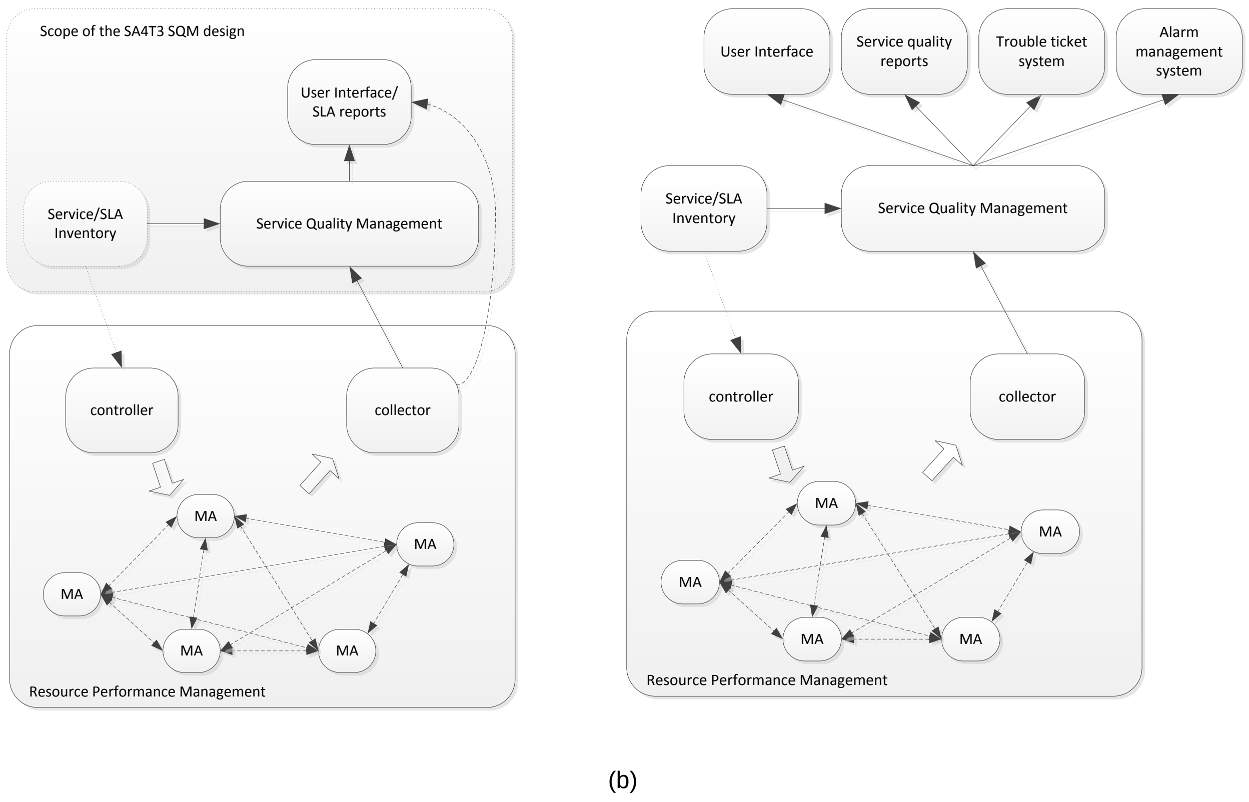

Figure 4.3: SQM application functionality scope: (a) SA4 T3 scope (b) full SQM functionality

SA4 T3 will design and develop a basic SQM solution to create SLA reports and provide a useful user interface for service operators who will be able to track the status of SLAs and KPIs. A simple prototype of service instance/SLA inventory will be created in order to allow for the automated operations of SQM. However, the interfaces with trouble ticketing systems or alarm management systems which typically exist in SQM solutions (but do not exist in GÉANT OSS portfolio at the moment) will not be developed. Figure 2.1 shows Resource Performance Management component using the LMAP architecture terminology [LMAP-arch]. Further activities will be to assess the reusability of components from pS-PS, ps-MDM and CMon for the SQM solution. The components include: how to tag measurements with the particular service instance, archiving measurements with the service instance IDs, and gathering data from archives for a particular service.

Table 4.1, below, summarises the scope of the SQM SA4 T3 design and development.

In scope	Out of scope
• Gathering data from measurement collectors/archives. • Getting service related data SLA and KPI parameters from service inventory. • Correlating measurement data with the SLA and KPI parameters and creating SLA reports • Providing UI for the service operators and/or service users. • SQM architecture, measurement agent placement, scalability issues, choice of measurement agents. • Multi-homed measurement agent design requirements, data/information model and archive specification. • Service and SLA modelling.	• Measurement agent design, measurement methodologies. • Network discovery, topology discovery, measurement agent discovery. • Service performance improvements, corrective activities. • Resource performance measurement architecture. • Creating alarms and tickets upon detected SLA violations. • Defining rules and procedures for specific services, service definition. • Billing, penalty schemes in case of SLA violations.

Table 4.1: Scope of the SA4 T3 SQM design and development

5 Conclusions

This document gives the initial design considerations for the SQM system for GÉANT multi-domain network services. Special emphasis is given to the analysis of the suitability of existing GÉANT tools, particularly the perfSONAR suite for the SQM for GÉANT services. The conclusion is that perfSONAR has the appropriate architecture and supports a set of measurement methods which are necessary for successful SQM for multi-domain services. Some measurement methods, especially for Layer 2 technologies are missing, but can be added to perfSONAR and some measurements available in perfSONAR such as throughput tests, are not relevant to SLA verification. An SQM requires a strict service orientation and comparison of the measurement data against a set of KPIs for the service instances, while perfSONAR is mainly focused on the presentation of raw measurement data.

GÉANT SQM can be built on top of perfSONAR with the addition of a few missing OSS components like Service/SLA inventory, slight customisation of the UI and some changes to the measurement archive databases (data belonging to different service instances should be distinguished). These changes do not require huge development effort for the first usable prototype as there are a lot of ready artefacts such as standard information models and interfaces. One of the more prominent problems of SQM – the scalability of the number of MPs required can be solved using multi-homing measurement agents, and this is probably the only approach that has to be followed for future SQM systems. There are still a few open issues that have to be resolved before the development is started, such as:

- The use of network element-based SLA measurements (e.g. Cisco SLA, Juniper RPM) and gathering data of these measurement for SLA monitoring, the suitability of this approach for MD-VPN, the interoperability on various platforms, etc.
- The potential for using CMon for SQM.
- pS-PS versus pS-MDM: there are ongoing efforts towards the convergence of the two platforms, most probably in a direction closer to the pS-PS version of the tool. Both platforms will be compared, the differences that impact SQM design should be analysed, and the platform that is going to be used as for SQM will be chosen with special attention to the ease of installation and the ease of use of MPs.
- Issues with the reliability and lag of the measurement procedures over multi-homed MPs in MD-VPN (L2 and L3 VPN) that were noticed in the last phases of the scalability problem analysis. Also the possibility of using HADES will be analysed in the multi-homing scenario.

These issues will be resolved before the next Milestone of the SA4 T3 task (due in M15). After that, the development of the GÉANT SQM will begin.

Appendix A Open-source SQM tools

OSS Applicaton	Features
OpenNMS	• Automated and directed network devices and services. • Event and notification Management. • Service assurance from devices to service level. • Performance measurement of networked services.
NetXMS	• Monitoring status of network devices as well as hosts and servers with applications. • Discovery of IP topology and new network devices. • Notification of network events to operators. • NetXMS has business impact analysis tools.
ZABBIX	• Automated discovery of networks. • Distributed monitoring of service quality. • API for two-way integration. • Rule-based problem detection.
Centreon	• SLA metric aggregation. • Configurable frequencies for KPI collection. • Load analysis breakdown by strategy, geography or network topologies. • Hierarchical notification system based on business, network devices dependency. • Ticketing tools interfaces.
Zenoss	• Manages the configuration, health, performance of networks, servers and applications. • An integrated CMDB. • Custom devices like temperature sensors can also be monitored.

Table A.1: Open-source SQM tools

Appendix B Measurement Points

MP/Agent name	Equipment	KPIs measured	Measurement Protocol & Access method	Comments
Cisco IP SLA	Cisco routers and switches; extra feature of IOS	IP layer: - One-way delay and jitter - Two-way delay and jitter - Packet loss Delays for network services: - DNS, FTP, HTTP, DHCP	Proprietary measurement protocols; Access methods – CLI and proprietary MIB traps	Active measurements
Juniper RPM	Juniper routers and switches; extra feature of Junos	IP layer: - One-way delay and jitter - Two-way delay and jitter - Packet loss Delays for network services: - HTTP	Proprietary measurement protocols; Access methods – CLI and proprietary MIB polls and traps	Active measurements
SNMP MIB II	All L3 and L2 equipment	- Throughput (physical and logical interfaces) - Status of physical and logical interfaces	Standard (RFC 1213) Access methods: SNMP polls and traps	Passive measurements, a well-established standard supported on any network box
BFD[3]	Routers and switches	Connectivity status of: - IP nodes - MPLS LSPs - Ethernet nodes	Standard measurement protocol (RFC 5880 – RFC 5885) Access methods:	Active measurements (keep-alive based)

[3] The status of BFD and its suitability for the use for MD-VPN service are described in more detail in Appendix D.

			- CLI - BFD MIB (Internet draft)	
CFM/802.1ag	Routers and switches of most leading vendors; Tested in GN3 on: - Cisco - Juniper - Brocade - Extreme	Connectivity status of VLAN-based Ethernet services	Standard measurement protocol IEEE 802.1ag Access methods: - CLI - CFM MIB	Active measurements, keep-alive based with configurable interval from 3.3 μs
Y.1731 MEP	Routers and switches of most leading vendors; Tested in GN3 on: - Juniper - Brocade	Connectivity status of VLAN-based Ethernet services (compatible with CFM) Delay (one and two ways), jitter and loss of Ethernet frames	Standard measurement protocol Y.1731 Access methods: - CLI; - No standard Y.1731 MIB	

Table B.1: MPs embedded in network elements

Model	Type	KPIs measured
Accedian MetroNID	Carrier Ethernet switch Standalone design with 4 GE ports Small footprint: 150 x 150 x 35 mm	- Ethernet service status by using CFM/802.1ag protocol - Ethernet frames one-way, two-way delay and jitter, frame loss according Y.1731 protocol - RFC 2544-based throughput - Y.1564-based throughput
Accedian NanoNID	SFP with rich Ethernet OAM functionality Could be plugged into any switch or router with SFP slots. Needs a central part (Accedian V-NID Actuator, a small switch) to work as a part of a distributed monitoring system	Same as MetroNID

Table B.2: Accedian measurement agents

Appendix C perfSONAR MA – Technical Characteristics

Service	Storage backend	GÉANT MA database	GÉANT MA measurements ratio
perfSONAR SQL MA	MySQL (default) or PostgreSQL	The current database size is 98 MB. The data store has grown at around 40 MB per year.	The current schedule is 208 measurements per day. There still can be found failed measurements, which results in lack of data and therefore storage usage unused.

Table C.1: perfSONAR MA – basic information

Installation type	Measurement scenario	Required data storage
SQL MA storing data from BWCTL MP scheduled tests between endpoints within GÉANT	50 links in GÉANT covered with tests each endpoint runs test to other neighbouring endpoint, not full mesh bidirectional 4 tests per day 365 days of data archive test: TCP, 30s, 6 intervals	160 MB
HADES MA storing data from HADES scheduled tests between endpoints within GÉANT	40 endpoints in GÉANT 9 packets per minute full mesh bidirectional 365 days of data archive	142 MB

SQL MA storing data from OWAMP MP scheduled tests between GÉANT and other domains	40 endpoints in GÉANT each endpoint runs test to two other endpoints in other domain bidirectional 365 days of aggregated data archive	100 MB

Table C.2: perfSONAR MA – data storage requirements

	Description
Database Management	No data is deleted automatically, all measurements are kept in the DB. This means that deleting or archiving should be done manually or through a custom built script. A script could be written to automatically delete (and archive if needed) all data older than a year, for example.
Performance issues	Querying with the regular perfSONAR API (XML/SOAP) is limited when retrieving a large quantity of data (a lot of different measurements or measurements spanning a very large time frame). Querying through the native MySQL interface is limited by MySQL scalability which should be satisfactory as all data of a certain type is located in a single table. The SQL queries are simple ones and are therefore easy to serve for the SQL engine.
Scalability issues	If scalability issues arises, it is recommended to tackle them through the deployment of multiple MA rather than through any upgrade or upscale to the backend storage of a single MA. MP can be partitioned into groups and having each group store its measurements into a different MA. A service quality analysis tool could then query those multiple DBs in parallel and aggregate the data and results.

Table C.3: perfSONAR MA – operational issues

Appendix D MPLS monitoring – Bidirectional Forwarding Detection

One of the biggest challenges of modern networks is the need to support services that are sensitive to packet loss, transmission errors, delay, and jitter. In order to meet the sometimes ambitious SLAs for these services, networks are designed with backup links, redundant node components, and other resiliency features. Detecting a link failure in a timely manner and resolving it as quickly as possible is of primary importance.

The first step to performance monitoring is verifying the connectivity and checking that its forwarding path is end-to-end operational. In the context of the GÉANT MD-VPN service, MPLS LSP availability and service parameters from an edge router in one NREN to the edge router in another NREN must be measured. Here, the possibility of using Bidirectional Forwarding Detection (BFD) as fast failure detection tool in this type of environment is explored.

The BFD protocol defines a method for detecting liveness in arbitrary paths between systems. These paths may span multiple network hops. It is a simple “hello” mechanism that detects failures in a network, between the forwarding engines of routers. “Hello” packets can be sent at regular intervals; failure is detected when the router fails to receive a reply after a specified interval.

The BFD could be utilised by other network components for which their integral liveness mechanisms are too slow, inappropriate, or non-existent. It is intended as an Operations, Administration, and Maintenance (OAM) mechanism for connectivity checks and connection verification.

BFD packets are carried as the payload of whatever encapsulating protocol is appropriate for the medium and network. BFD can provide failure detection on any kind of path between systems, including direct physical links, virtual circuits, tunnels, MPLS Label Switched Paths (LSPs), multihop routed paths, and unidirectional links.

One desirable application of Bidirectional Forwarding Detection is to detect a data plane failure in the forwarding path of a Multiprotocol Label Switching (MPLS) Label Switched Path (LSP) [RFC 5884]. BFD is used as a periodic OAM feature for LSPs to detect LSP data plane faults.

According to the RFC 5884 document, the MPLS LSP may be associated with any of the following Forwarding Equivalence Class (FEC):

- Resource Reservation Protocol (RSVP) LSP Tunnel IPv4/IPv6 Session.

- Label Distribution Protocol (LDP) IPv4/IPv6 prefix.
- Virtual Private Network (VPN) IPv4/IPv6 prefix.
- Layer 2 VPN.
- Pseudowires based on PWid FEC and Generalised PWid FEC.
- Border Gateway Protocol (BGP) labelled prefixes.

In case of MD-VPN, BGP labelled prefixes are exchanged between PE routers in different NRENs. Unfortunately, network equipment vendors (e.g. Juniper, Cisco) do not currently implement RFC 5884 for all FECs, and they only partially support it. Juniper has implemented BFD for MPLS LSPs that use either LDP or RSVP as the signalling protocol, but not for BGP labelled unicast LSPs, which are important in the case of MD-VPN. Also, Cisco supports BFD over MPLS TE LSPs implementation in Cisco IOS XR Software Release 4.3.1 which is based on RFC 5884, but not BGP labelled unicast LSPs.

Unfortunately, due to the problem of incomplete implementations of RFC 5884 by different network equipment vendors, it is currently not possible to use BFD for testing or implementation of performance monitoring with the GÉANT MD-VPN service.

References

[3.1-1] perfSONAR MDM -perfSONAR PS Comparison. DANTE, 2012, https://indico.in2p3.fr/getFile.py/access?contribId=21&sessionId=2&resId=0&materialId=paper&confId=6900

[3.1-2] perfSONAR Online Overview, slide 36, http://cbt.GÉANT2.net/repository/gn3/perfSONAR_online_overview/player.html

[Clarity] http://www.clarity.com/performance-manager/

[Comarch] http://www.comarch.com/telecommunications/products/service-fulfillment-assurance/service-quality-management/

[commag] R. A. Dourado, L. N. Sampaio, J. A. Suruagy Monteiro, “On the composition of performance metrics in muti-domain networks”, IEEE Communications Magazine, vol.51, no.11, pp.55-62, November 2013

[Cyan_BPV] https://www.cyaninc.com/products/blue-planet-sdn-platform/planet-view

[draft-dong] J. Dong, Li. Z., and B. Parise, "A Framework for L3VPN Performance Monitoring", draft-dong-l3vpn-pm-framework-02 (work in progress), January 2014.

[draft-l3vpn] H. Ni, S. Zhuang, Z. Li, “BGP Extension For L3VPN Performance Monitoring”, draft-ni-l3vpn-pm-bgp-ext-01 (work in progress), February 13, 2014

[draft-zheng] L. Zheng, Z. Li, Aldrin S., and B. Parise, "Performance Monitoring Analysis for L3VPN", draft-zheng-l3vpn-pm-analysis-02 (work in progress), October 2013.

[DS4.3.1] End-to-end management – catalogue of business processes

[GlobalNOC] GlobalNOC OESS Suite, http://globalnoc.iu.edu/sdn/oess.html

[GN3 DJ2.1.1] “Information Schemas and Workflows for Multi-Domain Control and Management Functions”, GN3 deliverable DJ2.1.1, May 2011.

[ITU-T Y.1541] ITU-T Y.1541 “Network performance objectives for IP-based services”

[LMAP-arch] Eardly P., Morton A., Bagnulo M. Burbridge T., Aitken P. Akhter A., “A framework for large-scale measurement platforms” (LMAP), IETF draft, March 2014.

[lmap-use-cases] Lisner M., Eardly P., Burbridge T., Sorensen F. “Large-scale Broadband Measurement Use Cases”, IETF draft, April, 2014.

[M.3342] “Guidelines for the definition of SLA representation templates”, ITU-T Recommendation M.3342, July 2006.

[Netcracker] http://www.netcracker.com/en/products/service_fulfillment_assurance/service_quality_management/

[OSS/J SQM] OSS/J Service Quality Management API – TMF884 – May, 2010

[pS-MDM] http://services.geant.net/PerfSONAR/Pages/Home.aspx

[pS-PS]	http://www.perfsonar.net/
[pS_dashboard]	http://ps-dashboard.es.net
[Pythia]	P. Kanuparthy, D. Lee, W. Matthews, S. Zarifzadeh, C.Dovrolis, “Pythia: detection, localization, and diagnosis of performance problems”, IEEE Communications Magazine, vol.51, no.11, pp.55–-62, November 2013
[RFC4656]	S. Shalunov, B. Teitelbaum, A. Karp, J. Boote, M. Zekauskas, ”A One-way Active Measurement Protocol (OWAMP)”, IETF RFC 4656, September 2006.
[RFC5835]	A. Morton, S. Van den Berghe, “Framework for Metric Composition”, IETF RFC 5835, April 2010.
[RFC6049]	“Spatial Composition of Metrics”, A. Morton, E. Stephan, January 2011.
[RFC6374]	Frost, D. and S. Bryant, "Packet Loss and Delay Measurement for MPLS Networks", RFC 6374, September 2011.
[Smokeping]	http://oss.oetiker.ch/smokeping/
[TAM]	Application Framework (TAM) – tmforum GB929
[TMF_products]	http://www.tmforum.org/ProductsServices/ServicePerformance/2550/Home.html
[Y.1241]	**“**Support of IP-based services using IP transfer capabilities”, ITU-T Recommendation Y.1241, March 2001.
[Y.1540]	**“**Internet protocol data communication service – IP packet transfer and availability performance parameters, ITU-T Recommendation Y.1540”, March 2011.
[Y.1541]	**“**Network performance objectives for IP-based services”, ITU-T Recommendation Y.1541, December 2011.
[Y.1563]	“Ethernet frame transfer and availability performance”, ITU-T Recommendation Y.1563, January 2009.
[Y.1731]	“Performance Monitoring In a Service Provider Network”, ITU-T Recommendation Y.1731, March 2011.

Glossary

BFD	Bidirectional Forwarding Detection protocol
BWCTL	BandWidth ConTroLler
CMon	Circuit Monitoring
COTS	Commercial off-the-shelf
eTOM	enhanced Telecom Operations Map
IETF	Internet Engineering Task Force
IPPM	Internet Protocol Performance Metrics (Working Group)
ITU-T	International Telecommunications Union – Telecommunication Standardisation Sector
KPI	Key Performance Indicator
L3VPN	Layer 3 Virtual Private Network
LDP	Label Distribution Protocol
LMAP	Lightweight MTA (Message Transfer Agent) Authentication Protocol
LSP	Label Switched Path
MA	Measurement Archive
MD-VPN	Multi-Domain Virtual Private Network
MIB	Management Information Base (SNMP MIB-II)
MP	Measurement Point
MPLS	Multiprotocol Label Switching
NOC	Network Operations Centre
NREN	National Research and Education Network
OAM	Operations, Administration and Management
OSS	Operation Support System
OWAMP	One Way Delay Active Measurement Protocol
PE router	Private Edge router
perfSONAR	Performance focused Service Oriented Network monitoring ARchitecture
PLR	Packet Loss Rate
ps-MDM	perfSONAR Multi-Domain Monitoring, part of the GÉANT services portfolio
ps-PS	perfSONAR PS, toolkit developed by I2/ESNET
QoS	Quality of Service
RAG	Red Amber Green
RPM	Resource Performance Management
SDN	Software Defined Networking
SLA	Service Level Agreement

SNMP	Simple Network Management Protocol
SQL-MA	Structured Query Language Measurement Archive
SQM	Service Quality Management
TAM	Telecom Applications Map
TMF	TeleManagement Forum
UI	User Interface
VPN	Virtual Private Network
VRF	VPN Routing and Forwarding Tables
VT	VRF-to-VRF Tunnel