10-12-2018

Deliverable D5.7
Continuous Service Improvement Best Practices

Deliverable D5.7

Contractual Date:	30-11-2018
Actual Date:	10-12-2018
Grant Agreement No.:	731122
Work Package/Activity:	5/SA2
Task Item:	Task 4
Nature of Deliverable:	R (Report)
Dissemination Level:	PU (Public)
Lead Partner:	AMRES
Document ID:	GN4-2-18-4103E2
Authors:	B. Marović (AMRES/UoB); M. Adomeit (AMRES); I. Golub (PSNC)

The research leading to these results has received funding from the European Union's Horizon 2020 research and innovation programme under Grant Agreement No. 731122 (GN4-2).

Abstract
This document integrates the experiences from the optimisation and improvements of Trust & Identity and Multi-Domain Services Service Activity (SA2) services and describes the general approach to Continual Service Improvement (CSI) that has been adopted, tested and refined in practice by the Production Optimisation Task (Task 4). It presents the approaches and tools developed, and provides recommendations for CSI in further work.

Table of Contents

Table of Figures

Table of Tables

Executive Summary

Service improvements are sets of closely related and coordinated service changes that are implemented in order to improve the effectiveness and efficiency of IT services and related processes, realign them with the evolving needs and environment, and improve their quality and service levels. Continual Service Improvement (CSI) is a systematic approach that uses specific techniques and methods that support these goals and are based on lessons learned from past experiences.

Service Activity 2 Trust & Identity and Multi-Domain Services (SA2) Task 4 Production Optimisation has identified, analysed and elaborated potential or needed improvement opportunities in SA2. Some initial, common service improvement aspects were summarised in *Deliverable D5.1 Analysis of Service Elements and Optimisation Opportunities*, including harmonisation of service operational data; web visibility of services; measurement, monitoring and reporting; and harmonisation of processes.

The SA2 Task 4 team presented its findings to the individual teams responsible for implementing the service improvements and the actual steps taken are elaborated in this deliverable. In addition to the service improvement opportunities presented in deliverable D5.1, practical improvements in the domain of measurement, monitoring and reporting were further pursued with perfSONAR and eduGAIN teams.

Deliverable D5.7 Continuous Service Improvement Best Practices integrates the experiences from the optimisation and improvements efforts described in *D5.6 Service Improvement Report* and its predecessor *Milestone M5.3 Service Improvement Interim Report* with those from the subsequent work in July–October 2018. The general CSI approach outlined in *M5.1 Initial Service Improvement Guidelines* was tested, adjusted and refined in practice, and the main purpose of this deliverable is to present a snapshot of the approaches and tools developed during GN4-2 and the experience-based adjustments, and to provide recommendations for CSI in further work.

It covers the following topics: overall CSI process (Section 2); recommendations on how to identify improvement opportunities (Section 3); best practices related to the various aspects of the data that is used by the CSI – its collection, processing, aggregation, visualisation, analysis, privacy and security (Section 4); approaches that should help to increase services' readiness for change, including microservices architecture (Section 5); and best practices related to improvement management (Section 6).

Further process details are provided in the appendices: Appendix A describes the iterative data collection process, while Appendix B describes the data aggregation, visualisation and analysis

process. Appendix C presents a use case, based on addressing GDPR requirements, and Appendix D contains the improvement description template developed by SA2 Task 4.

1 Introduction

Trust and identity services, multi-domain services and perfSONAR are developed and operated in a federated environment. Such a community approach, in which GÉANT project partners develop and operate services and at the same time use them, has placed service teams in a unique position to drive services to high levels of quality. These teams are therefore well-established centres of excellence, and their daily interaction with the services and their users has resulted in a perspective that regards the improvement of services as an accepted part of the daily routine. Together with the teams responsible for managing services in production, a separate team was created within the Trust & Identity and Multi-Domain Services Service Activity (SA2) to focus on cross-service improvement.

The main objective of the Production Optimisation Task (Task 4) in SA2 is to support and promote the continuous improvement of the GÉANT production services and infrastructure for which SA2 is responsible through analysing operational processes and services architecture, applying industry standards and best practices, and overseeing the implementation of improvements. Task 4 has been working with the service and development teams to advance particular initiatives, such as service analytics, and to supervise the aspects of services in production related to these initiatives, at the same time producing various analyses, such as the recommendations on perfSONAR, and developing the corresponding templates and process approaches. Although its main goal is to propose recommendations for services that are in production or in transition towards production within SA2, the scope of its work is not limited to this Activity. SA2 Task 4 collaborates closely with Networking Activities (NAs), other SAs, and/or Joint Research Activities (JRAs) and their Tasks towards the continuous improvement of processes, tools and services, as appropriate.

The aim of this deliverable is not to detail individual improvements that have been undertaken (which was done in *Deliverable D5.1 Analysis of Service Elements and Optimisation Opportunities* [D5.1], *Milestone M5.3 Service Improvement Interim Report* [M5.3] and *Deliverable D5.6 Service Improvement Report* [D5.6]), but to capture the GN4-2 experiences and synthesise the related recommendations and practices. In this way, it represents a closure of the methodology outlined in the Task's first Milestone in GN4-2, *M5.1 Initial Service Improvement Guidelines* [M5.1], documenting SA2 Task 4's approach to collaborating with GÉANT services and related processes.

The document presents recommendations and best practices for applying the Continual Service Improvement approach (CSI) to any organisational context. However, examples from GÉANT practice in CSI and service delivery are given to indicate how this approach is currently followed. The subsections also discuss how to recognise what could be improved and how to conduct the improvement process. Each best practice presented consists of an overall description of the issue, details of the key concerns and related reasonings, and recommendations interspersed with

illustrative examples from the GN4-2 improvements undertaken by SA2, Task 4 and GÉANT in general.

The document addresses the CSI process in Section 2, and in Section 3 provides recommendations on how to identify improvement opportunities. Section 4 presents the best practices related to the various aspects of the data that is used by the CSI – its collection, processing, aggregation, visualisation, analysis, privacy and security. Section 5 describes the approaches that should help to increase services' readiness for change and Section 6 presents the best practices related to improvement management.

2 Traceable and Inclusive CSI Process

The overall service improvement process should be structured and aligned with proven practices and widely recognised and adopted frameworks and vocabularies. For its own work, the GN4-2 SA2 Task 4 team has adopted the Information Technology Infrastructure Library (ITIL) Continual Service Improvement (CSI) seven-step approach [ITIL CSI], illustrated in Figure 2.1. It is also very useful to overlay the individual steps with the more general plan-do-check-act (PDCA) model (also shown in Figure 2.1) and to appraise these and the approaches applied elsewhere or recommended by the existing frameworks.

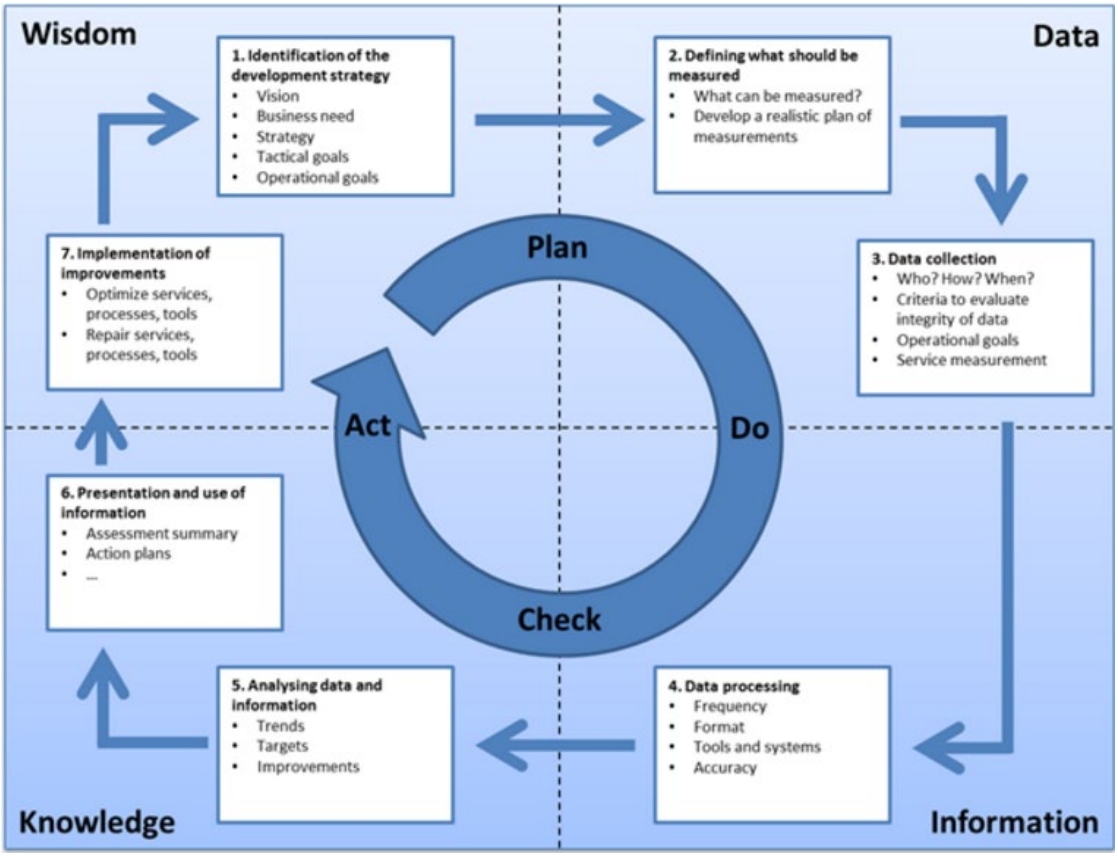

Figure 2.1: Seven-step improvement process, as defined in ITIL CSI

The improvement process should be traceable, i.e. have a recorded audit trail that can be referred to subsequently, and be organised in several iterations, each with defined goals and priorities. Quick wins help with getting buy-in, as even small benefits lend strength for further work. Progression

through the steps is supported by testing, proofs of concept or pilots when necessary for decisions to be made. Besides executing several CSI cycles sequentially, parallel improvements should also be considered when there are sufficient resources. For example, several pilots may run in parallel in order to select the best option.

Data collection and data processing improvements (discussed in Section 4.1 and 4.2) are also better done in small and incremental steps, while the analysis and presentation (also Section 4.2) can be supported by interactive and adaptable tools that may be shared by several services. Appendix A provides a generalised example of the iterative data collection and refinement process applied in several SA2 Task 4 interventions.

The implementation of an improvement may become a project of its own, with expected results, identified risks and obstacles, costs, time estimates, allocated resources, appointed manager, assigned responsibilities and (when needed) allotted budget and formal approval.

After the service has been upgraded and stabilised, the improvement should be reviewed in order to establish and record the conclusions, identify other places of application and suggest subsequent improvements. It is also very important to give credit where due and to mark and praise achievements, even if this can be done only within the inner group.

While the operational processes are typically sufficiently guaranteed if they are regulated, firmly assigned, monitored and kept within the agreed parameters, the implementation of improvements is usually disruptive, so it will benefit from inclusiveness. An inclusive attitude enhances support, adoption and ownership, reduces omissions and risks, and helps to ensure that the transition to a new state is effective, purposeful and as smooth as possible.

2.1 Summary of Recommendations

- Use a structured, proven, widely recognised and adopted service improvement process, e.g. ITIL CSI [ITIL CSI].
- Ensure the process is traceable, iterative, incremental and inclusive.
- Use (shared) tools to support and enhance the process.
- After each service improvement, review what happened, record conclusions, identify other places of application, lessons learned and process improvements.
- Acknowledge contributions and achievements.

3 Identification of Improvement Opportunities

Some service improvement opportunities are known or obvious. Other are hidden or difficult to notice. In order to select the optimal course of action, it is necessary to identify and prioritise the most important and relevant ones. Sources that are most often used to identify improvement opportunities and relevant problems are:

- **Tacit knowledge** – Opportunities are often well known by the service personnel – practitioners, maintainers and managers who are acquainted with the problems or needs and are willing to talk about them. The same applies to expert users, who can describe the issues they have encountered with the services and how things could be done in a more effective, efficient and useful way. Those who actually use or operate the service often have the best knowledge about what should be improved, and the changes they request are sometimes easy to implement and very visible.

- **Available records** – Sometimes the improvement opportunities can be extracted from the existing issue tracking systems and records of service incidents, problems, known errors and even risk registers. The information provided by those tools should be analysed in order to gain knowledge, identify recurring events and interpret improvement opportunities.

- **Unstructured information** – Suggestion boxes, forums, mailing lists and feedback surveys with open-ended questions offered to end users, customers or service personnel are another source of improvement possibilities. They should be analysed and the ideas they contain captured and processed by structuring them into support desk requests, improvement proposals, feedback to the partner relations and service teams, etc.

- **Exploratory research** – Observations, interviews, focus group discussions and brainstorming may be focused on a certain stakeholder group or aspect of the service. In order to obtain useful results, it is important to include those who could competently contribute and complement each other, and for the research coordinator to direct the exploration in a productive way, but without imposing their views. Partners, customers, expert users, experts in neighbouring domains and those working in business relationship management are likely to contribute novel views.

- **Reports and analyses** – Data from monitoring, logs, surveys with closed-ended questions, etc. can result in usage statistics, performance and trends reports, and predictions from the information about how the service and resources are used, where the problems are, who the users are, what they are focused on and what could be expected in the future.

- **Senior management** – Top-level and cross-service management may provide crucial input on strategic service changes and concerns. Portfolio management may indicate that the service is about to be focused on (e.g. made a business/PR priority towards customers or other stakeholders, scaled up, extended, enhanced or worked on further), merged or retired; demand management can provide information for longer-term planning; changes in the management of service design and transition may increase the requirements in the related areas.

- **External references** – Benchmarking against standards, competing solutions, audits, examples from other services (discussed in Section 3.2), industry practices, guidelines and best practices may also indicate what could be enhanced.

Obtaining insightful and pertinent feedback can be a delicate matter. Direct input from service practitioners and managers, business developers, support staff and partners can save time-consuming sifting of data or wild-goose chases. The information collected should be jointly interpreted with the service personnel. If the approach to the topic is too inquisitorial, the interlocutor may become defensive and withdraw, so it is more effective to discuss the problems and opportunities in a positive, supportive manner and gradually help the person articulate their proposal. Also, a discussion about problems should not turn into an argument about responsibilities or blame for past mistakes. The timeframe of the interaction also matters; in order to ensure quality – and repeated, future – input, people's concerns and proposals should be addressed in a timely manner with meaningful answers, thoughtful consideration, and frequent and regular service and software updates that reward the contributors by incorporating their ideas and resolving the issues.

If the idea is not fully articulated or the problem's root cause is unknown, an in-depth investigation may be required.

The resulting improvement candidate is recorded in the Continual Service Improvement Register (CSIR) [ITIL CSI]. Once the potential opportunities worth pursuing are identified, they need to be quantified and prioritised. This decision making can be done on the basis of the data available within the services and existing monitoring and analytical systems, or through targeted measurements and surveys. Only after sufficient details have been captured, numbers analysed and outlier and borderline cases taken into consideration, does it become possible to make an informed decision.

3.1 Illustrative Examples

All of these methods of collating the service improvement opportunities are used in the SA2 service portfolio. Gathering and addressing the feedback from users is handled using various channels. For both eduroam and eduGAIN, there are several mailing lists from which the service teams obtain ideas on further improvements. Ultimately, a well-defined governance model is in place: federations using eduroam and eduGAIN services are represented by their steering group members, and they jointly decide on the improvements. Wider research community outreach and requirements gathering are conducted through collaboration with groups such as AARC, REFEDS, FIM4R, RA21, etc.; GÉANT partner relations surveys also provide valuable structured feedback. perfSONAR also has a well-defined governance model that enables strong collaboration with and input gathering from

users and the development community, through mailing lists, surveys, training and workshops. These inputs are always used to guide further development and roadmaps for any new releases.

Service-related measurements and monitoring were implemented by the trust and identity (T&I) services from the beginning of their operations and are a valuable source of data about services adoption and performance. Some of that data, prepared for the stakeholders to consume, is available in the eduroam and eduGAIN technical sites [eduGAIN-T, eduroam-M]. The same holds for the perfSONAR service [perfSONAR], which implemented Lookup Service so that the public installations of the perfSONAR nodes could be followed up. Some of the CSI initiatives conducted by Task 4 were designed to make additional use of this information by improving the data analytics and presentation.

Standards are a further significant source for identifying improvements to SA2 services. The eduroam development team is participating in IETF and the Wi-Fi Alliance standardisation bodies, while eduGAIN is based on SAML2int and community-based standards are evolved in bodies such as REFEDS. OpenID Connect Federation implementation is also embedded in the standardisation process of the Open ID Foundation. As well as best practice for standards, these bodies also provide exposure to industry and enable know-how from outside the community to influence services. More generally, the ITIL and TM Forum frameworks are constantly referred to, evaluated and selectively used in order to improve how the service lifecycle is managed. In particular, SA2 benefited from the TM Forum work on the definition of certain business processes for SA2 services, for example, the eTOM process flows Order-to-Payment (i.e. Order-to-Closure) and Problem-to-Solution. Bearing in mind the heterogeneity of the federated service development and offering, such an approach provides the possibility for the comparison, alignment and consolidation of business processes of different services.

Finally, the GÉANT Product Lifecycle Management (PLM) team, Activity leaders and the project management team track the overall cross-service portfolio through regular meetings and periodic reviews. This provides an opportunity to identify service improvements looking across service families as well.

3.2 Application of Analogies and Precedents

Although analogies and precedents may sometimes be overstretched, they are an extremely powerful tool for exploration and adoption of complex concepts and planning of actions. They serve to reach beyond the familiar and particular to the new and more general in order to improve understanding and highlight similarities, differences and options. The differentiating aspect helps in establishing useful abstractions and coherent narratives that are applicable across individual cases. Comparison with and between analogous examples helps in finding creative, more general or simplified approaches to complex problems. This applies to both the identification of opportunities and the implementation of solutions.

Precedents are directly applicable analogies from previous experiences in similar circumstances. They can be considered as examples or guides, as they usually provide convincing arguments, model solutions or illustrative results. Most guidelines, frameworks, toolkits and reusable components are built by elaborating or addressing what has been done or produced before.

3.2.1 Illustrative Examples

Comparison of services was one of the reasons for assigning the CSI to a separate GN4-2 SA2 Task that is not linked to any particular service; its separation also from the development and operations teams further enables it to uphold an independent view. The application of analogies, examples and design patterns has been at the core of Task 4's approach to the CSI. They were identified during the gathering and analysis of the data on service operation, web visibility, measurements and processes for deliverables D5.1, D5.6 and Milestone M5.3, as well as while addressing General Data Protection Regulation (GDPR) requirements, and were used to produce templates, conduct work in predefined stages and provide cross-service examples.

Precedents from data-related iterative collection efforts conducted during various interventions of this task are summarised in Appendix A and Appendix B, and a use case, based on addressing GDPR requirements, is presented in Appendix C. Descriptions of the individual process flows in SA2 services were collected through surveys and interviews, translated into diagrams, catalogued, mutually compared and matched against the process patterns provided by the TM Forum Business Process Framework (BPF) [BPF] in order to identify the similarities and differences and create generalised flows. This approach was also used for improving the data analysis of services in SA2, where practices, solutions and outputs from the analysis of data from the perfSONAR Lookup Service provided a precedent for eduGAIN and other services. Similarly, the implementation of SA2 service data aggregation and processing tools was inspired by the principles of microservices architecture, thus promoting the application of this architectural approach in GÉANT.

3.3 Summary of Recommendations

- Be exhaustive in identifying improvement opportunities, exploring sources that are hidden or difficult to notice as well as those that are known or obvious.
- Be respectful, supportive and positive in obtaining input and feedback, address concerns and suggestions in a timely manner, and provide updates on progress and outcomes.
- Discuss and interpret the information obtained with service personnel.
- Record potential improvements in the Continual Service Improvement Register (CSIR) [ITIL CSI], quantify and prioritise.

4 Data for the CSI

Any improvement opportunity – its direction, significance, applicability and the impacts that it has – is driven and directed by data. Data about the service, from the service, about the outputs and payloads that it creates, the value for the users, data about the users – any data can be relevant. Even the ITIL CSI seven-step improvement process [ITIL CSI] is largely focused on data that is needed to identify, implement and track improvements. However, since only some of the potential data sources might be useful, this section addresses specific data-related aspects, such as data collection (in Section 4.1) and data processing, visualisation and analysis (in Section 4.2), which both provide practical examples, from throughout the SA2 Task 4 work, of dealing with data. The extent of data collection undertaken needs to be carefully judged taking into account its practical and justifiable purpose, which is also regulated by the GDPR; Section 4.3 outlines the most important principles in that regard.

4.1 Data Collection and Processing

As shown in Figure 2.1, obtaining the input data is the third step of the ITIL CSI seven-step process, which is conducted after defining what should be done (Step 1. Identification of the development strategy) and what data is needed (Step 2. Defining what should be measured). It provides the necessary inputs for enhancing quality and optimising delivery and it is also needed in formulating and tracking the progress of improvement initiatives. This data may already be present within the service operations tools, technical management systems and service-related databases. The relevant data is typically related to service faults, performance, security, usage, configuration, process or technical components. When data about the performance of the service and its elements is missing, it is not possible to judge its current state, progress or the effects of an improvement; instead of being based on numbers, the conclusions can only be inferred from anecdotes.

Ideally, the relevant data that needs to be collected should be specified during the service design and its acquisition integrated into the service operations. The appropriateness of data gathering should be judged on benefits against costs, and on whether it will ultimately provide the required evidence about the infrastructure, problems, needs and how, where and by whom the service is used. Possible sources of data include service logs, monitoring and supporting tools. Long-term data collection and processing in order to track and predict trends are sensitive to changes in the availability, content, structure and naming of the records. The same applies to the related specifications and reports containing the resulting analyses. Also, subsequent changes in data collection may cause resistance, not only because of the additional work involved and discomfort with change, but also because of a fear of the conclusions that could be drawn. Such changes should therefore be introduced carefully, after thoroughly exploring and evaluating the data that is already

available. Although the software tools can help in data realignment, it is easier and more effective to establish and enforce conventions early.

Service users, development and operations personnel and other stakeholders are also valuable sources of data as they can indicate their interests and the significance of individual service elements and suggest enhancements and new features. If such data is needed on a regular basis and from a large number of individuals, it can be collected through permanent and dedicated mechanisms such as surveys, general ticketing systems and support forums. Surveys are particularly useful to capture information in a structured way. Such data capture mechanisms should be accompanied by supporting explanatory guidelines and tested with one or two groups before wider distribution. This helps in identifying misunderstandings and ambiguous or insufficiently specific questions, provides examples and stimulates others to join the effort. Data gathering in several iterations should start with easy and typically technical questions. This should result in an accessible path from the technical level towards the targeted service, organisational, procedural or legal levels, and help service managers in stepping out of the comfort zone.

4.1.1 Illustrative Examples

Data collection is present in each of the four themes identified in [D5.1] and all Task 4 deliverables and milestones deal extensively with its various aspects. The Task gathered information on currently performed measurements and monitoring of services in production in SA2 and collected the additional data that was needed in order to implement various service improvements. Two sets of such data that were a particular focus for the Task come from the perfSONAR Lookup Service and eduGAIN.

The perfSONAR data about the public nodes is continuously recorded in the database of its Lookup Service. All Lookup Service instances are queried, exported to a JSON format and saved on an external server every day. The individual snapshots are then streamlined, aggregated and put into the data processing and enhancement pipeline described in Section 4.2. This snapshot approach is suitable for all cases where there is a database that maintains the actual state without preserving the history. On the other hand, the suite of eduGAIN tools offered at [eduGAIN-T] preserves the history of collected records and performed tests, so the analysis and visualisation of trends can be performed by directly using the data provided through the tools' APIs.

The above recommendations on data collection from service personnel were applied in the SA2 Task 4 work on the service template. The work on systematising data on service definition and specification was initiated during the GN4-1 project with a simple service definition and pointers to the service's technical resources. During GN4-2, the service template has been significantly extended and improved to include more information. In the opposite direction, the collection of data about service measurement and reporting extends from simple technology and processes-related monitoring and measurements towards service metrics, key performance indicators (KPIs) and critical success factors (CSFs).

Another illustrative example of iterative data collection is the comparison of the findings on the web visibility of distinct services by using the common information grid. The findings and initial recommendations on eduroam, eduGAIN, FaaS, eduPKI and perfSONAR prepared by Task 4 members were offered to service managers for review, which resulted in a number of service-specific practical

suggestions. Cross-checking of these materials helped in identifying a few shared concerns, formulating several general recommendations and creating an outline for a product or service narrative that could be used in further development of web content and promotional materials.

The most complex iterative data gathering and analysis were carried out in addressing EU General Data Protection Regulation (GDPR) requirements for SA2 production services. A detailed account, which stands as a use case, is provided in Appendix C. In summary, the subject was approached by raising general awareness of the issues and requirements, and forming the dedicated GÉANT GDPR team, which worked with SA2 Task 4 and individual services. Understanding of the subject was first enhanced by clarification of concepts and publishing supporting materials on the GÉANT wiki. The service managers filled in the data inventories, which were used as input for data-mapping tables. These tables provided the key elements to be considered in producing service privacy notices, the writing of which was also supported by comprehensive examples of privacy notices. All these steps were assisted by the GDPR team and included gradual refinement of the materials as a result of the continuous feedback. The wiki templates and pages completed for services in production are integrated within a summary dashboard, which serves for navigation, progress check, and as a starting point for GDPR audits. In addition, the documents produced were incorporated into the service documentation in the corresponding service wiki pages.

4.2 Data Aggregation, Visualisation and Analysis

In order to be useful, usable and effectively used, the data that is collected as described in the previous section should be processed within smart analytical platforms. Their general architecture follows the structure of the service knowledge management system (SKSM), of which they are part, and which has four layers: data, information integration, knowledge processing and presentation. The already available or purposely collected data can be imported into a data warehouse by implementing the extract-transform-load (ETL) process.

Ideally, data processing implementation and analytical work should be jointly undertaken by data analysts and visualisation specialists and service practitioners who include both the technical domain experts and members of service operational teams. The analysis should start with information that is already regularly collected. This data should be processed in order to maximise usability and consistency; most interpretation and enhancement can be based on the data's statistical analysis and visualisation. These results should then be used to draw the first conclusions and explore the usability and quality limits of the available data. However, for a more precise analysis and to answer the arising and remaining questions, services and data gathering mechanisms typically need to be enhanced in terms of data content or resolution.

4.2.1 Illustrative Examples

In pursuance of both immediate and long-term developments and benefits in this area, SA2 and its Task 4 implemented a practical demonstration of the significance and usability of the data from and about services. The approaches, techniques and tools that were jointly developed and used for the data that is regularly collected by the perfSONAR Lookup Service and eduGAIN can be applied to

many other services. The eduroam service has already implemented the data analytics and visualisation as part of its monitoring and support tools at [eduroam-M].

In the case of eduGAIN, several data sources are available through eduGAIN databases and tools [eduGAIN-T], and it is probably easier to extract and transform data from various sources independently than to move them to one database or hide behind a single façade or API. The Splunk tool [SPLUNK], which is already used in GÉANT, was considered by SA2 as a possible solution for data processing and analysis. There are alternatives, such as Elastic Stack (formerly the "ELK stack") [ELK]. Both platforms are usually seen as log management solutions, but are actually complex analytical platforms that could be used to load and use database exports and purposely collected and curated data from services. In a consolidated setting, the shared data warehouse and the analytical tool could serve a number of services and processes. Task 4 made a proof of concept of data analysis for eduGAIN and perfSONAR services using Splunk, which proved to work well in terms of data onboarding, federated user management, dashboard features, report generation and their expedient export and embedding within external websites and portals. Based on those results, the core operational team in SA2 is looking into the possibility of establishing such an analytical platform to manage, aggregate and visualise data. Splunk supports the features needed in a federated setting, such as access control via the federated eduGAIN authentication, role-based authorisation and scheduled creation of charts and maps. In development of the F-Ticks solution for eduGAIN, JRA3 Trust and Identity Development also looked at using ELK as an analytical platform.

The approach developed can be applied to other services by tapping into data and logs they already regularly collect and by relying on the findings from the work with perfSONAR and eduGAIN:

- Data analytics tools enable effective investigation of service data through efficient creation and customisation of dashboards, charts and maps. Although the visualisation of values and distributions can be achieved with pie and bar charts, the use of data filters and trend charts can greatly help in understanding the explored phenomena. Of even greater help is the placement of the information into a more natural context by using heat maps, maps with clickable points and timelines with annotated events.
- Periodic snapshots allow trend analysis, even if the original sources were not designed to record the historical data, and preserve the data when processing and visualisation tools are changed.
- Even incomplete, imprecise and partially inaccurate data can be enhanced or filtered. This is particularly the case with data that can be linked and cross-checked using other services. For example, perfSONAR node location data was validated and significantly enhanced using Google Maps, OpenStreetMap and GeoIP [GeoIP] databases and services.
- Information presentation should be tailored to the audience. Exploratory data visualisation that analytical tools directly provide to expert users is extremely useful. However, other target groups need tailored, prepared and less flexible visualisations and wider audiences may need only selected and predefined charts.

The individual steps for data aggregation, analysis and visualisation, grouped into the PDCA stages, are described in the model given in Appendix B.

4.3 Privacy and Security in Data Collection and Analysis

Privacy, security and the extent of the collected service data should be incorporated into service design and operational processes. Data gathering initiatives and arrangements should be aligned with service policies, particularly the user-facing privacy policy. The merits of collecting and analysing any personal data should be weighed against the associated risks, while the purpose of data collection should be clearly stated at the places where the information is collected from service users or organisations. It is also important to declare the limits of visibility and downstream sharing of the data, limiting the potential uses to the actual and legitimate purposes such as further development of the service, platform or infrastructure. This may also require preparing and publishing acceptable use policy (AUP) documents for specific data sets and services. Published privacy and acceptable use policies enhance data quality by helping the users to understand the goals and benefits of data collection and making them less reluctant to provide the information requested.

Following the exposure minimisation principle, if any personal data is present at the source and is not used in the measurements or analysis, it should not be pulled from the source repositories at all, which ensures that it cannot be misused or exposed through the measurement and analytical pipeline.

However, control over proper usage and propagation of data must be exercised beyond personal data protection and extended to usage and propagation of all potentially sensitive data. Protection of the information within GÉANT services is not only about ensuring privacy protection, proper data use and data integrity. It is also about data security, which should be ensured even when identifiable personal information is not present. Concerns about data security apply even to open data, not to mention production data about the locations, usage, availability and other security-related characteristics of the resources managed or used by GÉANT services. It is quite easy to imagine the possible undesired consequences of indiscriminate access to the addresses and technical attributes of the infrastructure managed by this large international collaboration.

4.3.1 Illustrative Examples

SA2 trust and identity services are good examples of where these methodologies were followed and special care was taken to apply privacy by design principles since the services were first conceived.

4.4 Summary of Recommendations

- Ensure data requirements are clearly defined and specified – ideally at service design stage – and the data collected through the most appropriate, effective, efficient mechanisms, e.g. as an integral part of service operations where possible.
- Use smart analytical tools to process, aggregate, visualise and analyse the data.
- Follow the data collection steps in Appendix A and the aggregation, visualisation and analysis steps in Appendix B.

- Ensure privacy and security policies and regulations are adhered to when collecting and processing data.

5 Design for Change

Implementation of any kind of improvement necessarily introduces changes to service. Making a service highly changeable can be a difficult principle to adopt and justify because this takes a far greater effort than just to make the service fit for the current purpose and use, which is often the focus of the initial immediate plan. Prototypes and pilots often gradually evolve into production solutions, with the assumption that the initial postulates will persist. This actually means that systematic consolidation for production should start as soon as the probability of longevity and/or expectations of higher quality and greater durability become apparent. The key concern is to remove accidental and unmanaged redundancy and arrangements that could impede later changes. The source code, interface definitions, documentation, data persistence specifications and configuration templates often share the same considerations, such as target domain categories, which could be moved to a single point of responsibility from which all dependent elements could be derived. This helps in preserving consistency, both horizontally and across layers, which is a crucial feature of change-ready design.

For these reasons, services must be designed in a way that allows for ongoing development, integration, testing, delivery and deployment, and must incorporate facilities that support such change cycles. The same applies to the collection and processing of user feedback (Section 3) and service data (Sections 4.1 and 4.2), which require the service designs to plan for integration with smart analytical platforms by at least exposing the relevant data and metrics at the key points within services.

Design for change requires a dedicated effort to factor out and insulate the variable elements. This can be hard to justify when there are shifting expectations that are a challenge for the service to meet, and when it is difficult to tell where to start with (re)factoring the design. The only way to achieve maintainability is to start with a substantially flexible design and to keep on improving it in an agile and disciplined manner. This means that every change that increases entropy should be matched with an effort aimed to reduce it. For example, instead of making a direct replacement in the service design, it is better to isolate the point of inflexibility or uncertainty that is the reason for the change and also to improve the practical steps in a way that will allow easier and faster modification the next time. This requires a marginally larger – but justifiable – effort the first time an element is changed, but gradually increases the overall maintainability and adaptability. And since subsequent similar changes will become easier and less laborious, and the overall design is cleaner, more resources and time will become available to keep it this way.

5.1 Illustrative Examples

The GÉANT community and project are constantly advancing in the way they deliver services to meet users' increasing and evolving expectations. GÉANT services and the people involved in their delivery align with the standards and practices adopted by the industry. The need for data definitions, conventions, their curation, maintenance of interfaces between components and use of supporting tools are also drivers for change. A constant effort is made to adapt individual components to operate in various domains (i.e. services) and environments, toolkits, new technologies, and configuration and integration facilities.

5.2 Microservices Architecture

As many internal and external drivers can trigger the changes in services, they must be designed to withstand modifications, which stimulates the growing interest in operational agility, DevOps and microservices architecture, as they promise more efficient and flexible service design, management and operation, as well as greater technological adaptability. Disentangling the components makes changes much easier and makes it possible to work in parallel while preserving the ability to integrate the changes easily and often. Therefore, adoption of microservices architecture in the GÉANT context could provide a number of benefits. They include:

- Flexible service design, management and operation.
- Greater technological adaptability.
- Catering for controllable but adaptive evolution, addressing user needs and scalability.
- Facilitation of agility and coordination in federated environments.
- Affirmation of common service functions, tools and components.
- Easy update, rearrangement and replacement of components.
- Promotion of documentation and accessible interfaces between components.
- Support to decentralised service governance, service level agreement (SLA) enforcement and change management.
- Alignment of the GÉANT ecosystem with the industry.

5.2.1 Illustrative Examples

Service Architecture Elements and Microservices, an internal report/guidance document on service architecture elements [SAEM], was prepared by SA2 Task 4 in order to provide an overview of and arguments for the microservices architecture and thus enhance the understanding and coordinated adoption across GÉANT of this emerging trend. The same subject is covered by the TMF member white paper *Transforming BSS/OSS systems to Microservices Architecture* [T2MSA], but the Task 4 report attempts to provide a more general and accessible argument that is not TMF specific. It also aligns ITIL and TMF Frameworx terminologies, and groups applications, software platforms and tools into ITIL-related categories, at the same time offering a subset of TMF's Application Framework (TAM) taxonomic classification in order to enumerate the potential key elements of the SA2 service

architecture. *Service Architecture and Microservices* identifies and details the following elements of the microservices architecture, which are also recommended for the application of further improvements in SA2 services:

- Decomposition of services into simpler but self-contained microservices. These services are discoverable, easy to link and can be arranged in complex configurations.
- Separation of the infrastructure and functional service groups into domains encompassing their own tools, components and operations. Each domain sets a scope and administrative boundary within which functions work together to fulfil a complex business function. The domains also form natural boundaries for the definition of individual security policies, anomaly detection and security enforcement, and associated high-level service concepts.
- Complex inter-component/service data and process interactions are achieved through decentralised or, where suitable, centralised orchestration. Integrations are supported by association rules on loose coupling and interfaces, consistent interfaces and data sets, tools and control mechanisms.
- Individual services are loosely coupled and provide well-documented lightweight interfaces between components. While REST and JSON interfaces are the most popular, portable remote procedure call (RPC) protocols can be used when performance is a concern. API definition frameworks and languages such as Swagger and RAML enforce documenting and publishing of semantics.
- Decentralised service governance is achieved with virtualisation and containers, automated configuration management, wiring and discovery of service instances. Ability to deploy multiple instances increases availability and supports vertical scaling.
- Monitoring and reporting capabilities subscribe to business processes boundary events at the provided interfaces or API gateways or exit points of microservices, providing loosely coupled smart analytical platforms and versatile reporting, analysis and alerting capabilities.
- Mechanisms related to integration and coordination with other independently developed services, and shared concerns such as monitoring, instance discovery and tracking, are made available in a common and platform-neutral way.

The underlying drivers for the microservices architecture are already resulting in the ongoing organic adoption of its elements by some GÉANT developments and it was embraced in SA2 Task 2 Trust and Identity Operations. Task 4 does not advocate a massive migration of all services to the new paradigm but the evolution of services and the gradual adoption of the microservices architecture in new developments.

5.3 Summary of Recommendations

- Follow the principle of designing for change, e.g. by maintaining consistency, identifying shared dependencies, implementing a flexible design and developing it in an agile, disciplined manner.
- Adopt microservices architecture where possible, as described in *Service Architecture Elements and Microservices* [SAEM].

6 Management of Improvements

Service and improvement leadership is the key aspect of successful service management and governance. The significance of service ownership was highlighted by GN4-1 SA4 Production Application Services and Infrastructure, Task 3 Production Optimisation and Continuity; it has been strengthened through the refinement of the Product Lifecycle Management (PLM) process and continues to evolve due to the dynamic, multi-faceted and federated nature of GÉANT. It is also challenged by the highly technical and developmental nature of SA2 services, which often leads to the distribution of high-level service management and developmental, technical and operational management.

Improvements, like all other changes, should be judged on their usefulness and benefits versus one-time and recurring costs and effort, but should also take into account potential user dissatisfaction, resistance and problems during the implementation. Their success depends upon the understanding and involvement of those who approve, carry out, promote and support the improvement implementation. Establishing a consensus could be challenging in a heterogeneous setting with a number of stakeholders, development groups, users and goals. In multi-faceted initiatives, of particular importance is the ability to combine the service and strategic perspective in order to select the most suitable improvements, and make them happen by removing or overcoming all impeding barriers. And when the efforts to build consensus or secure the mandate are successfully achieved, it is important to assign improvement leadership as someone has to refine the scope, allocate resources, select the pilot sites and actually conduct the improvement.

The service improvement plan must address the perspectives and interests of all targeted groups, and therefore an extended stakeholder analysis may be required. Ideally, all stakeholders should be beneficiaries in some way. Where this is not possible, those who do not benefit should at least be persuaded from opposing or at least blocking the initiative. There are several ways to do this:

- Find a champion who would benefit the most or is willing to preach the benefits to others.
- Start pilots with enthusiasts or those open to the intervention. After the first positive results, it becomes easier to involve those who are reluctant.
- Spread the effort across participants or over time by conducting an iterative process.
- Evidence, arguments and examples emphasise the expected benefits or prestige.
- Foster pride and satisfaction in participation and achievements and in the “magic” of the result. Tangible or symbolic rewards should also be included.
- Prepare for incidents or crises in participation.

Service managers may lead clearly delineated single-service improvements or nominate those who are responsible for the affected service functions or processes. However, personnel who are focused

on operations and service stability rarely initiate or manage large cross-service improvements. Service managers should liaise with those responsible for the improvements in order to negotiate the most effective paths and effectively manage the conflict between change and stability.

Leadership and improvement ownership are critical whenever those involved need to be stimulated to endorse the change. Identification, development and implementation of improvements are generally difficult as they require a departure from established habits and practices. By the same token, the retirement of a solution produced by those who should be involved in its replacement can also be very challenging. It is, therefore, crucial to find gains for almost everyone and win the support of those who are about to suffer any negative impact.

The complexity of improvements should also be addressed at the leadership level. The continuous nature of services underlines the value of a participative leadership style, which is also more effective in a distributed and multi-organisational setting. This leadership style involves the team members in decision making and relies on delegation of extensive and complex tasks. It is based on values and behaviours rather than organisational authority and its highest priority is to encourage, support and enable the team members in their work and development.

It is necessary to eliminate any blame-culture tendencies, accept the possibility of failure and make it clear that mistakes along the way will not be punished, in order to encourage risk taking and thinking outside the box. Anything should be open to critique, which should never be personal – everyone is doing the best they can, given what they know and are able to do at a given point in time. The main challenge of the participative leadership style is to balance the development of rapport and initiative with the use of authority in order to achieve the optimal results and accountability. On the other hand, it instigates performance, satisfaction and team cohesion, while the respectful and inclusive treatment of team members leads to more satisfying interactions with service users and partners and positive perception of the service and others. The participative leadership style can be supported by training on the subject and promotion of suitable and appropriately prepared individuals to managerial positions. The ongoing support of team workshops, gatherings and retreats within GÉANT should be continued as such events foster the development of rapport.

6.1 Management of CSI Outputs

Individual improvement opportunities (identified as described in Section 3) should be recorded in the CSI Register (CSIR) [ITIL CSI], which is designed to capture basic information about them and prevent duplication. The CSIR is a database or structured document used to store and manage possible improvements throughout their lifecycle. It contains important information for the overall service lifecycle and should be regarded as part of the service knowledge management system (SKMS). The listed improvement opportunities are classified as small, medium or large undertakings and also according to whether they can be achieved quickly or in the medium or longer term. Their descriptions should also indicate the benefits that will be achieved if they are implemented. This information is used to prioritise the listed opportunities. The items with lower priorities should be periodically reviewed in order to potentially increase their priority or to cancel them and thus prevent their further consideration.

Where possible, CSI should rely on outputs it produced earlier in the making of new ones, even when they are created for other services. Since the CSI outputs are also the documents or data sources that are in due course handed over to services and improvement implementation teams, it is wise to pass them on for post-CSI stewardship in terms of the related actions. Of course, responsibility for the implemented improvements can later be transferred to a third team.

The individual CSI outputs and their uses in other lifecycle phases, as defined in ITIL, are listed in *M5.1 Initial Service Improvement Guidelines* [M5.1]. Some CSI outputs are logically primarily produced by the service CSI; these, and their downstream stewards and users, are shown in Table 6.1:

CSI output	Passed to
Financial information regarding improvement initiatives for input to budgets	Service Strategy
Input to change evaluation and change advisory board meetings	Service Transition
Results of customer and user satisfaction surveys	Service Operation, but all phases use them to analyse the feedback
Data required for metrics, KPIs and CSFs	Service Operation, while all phases should use the derived reports to compare the achievements against the goals; Service Design, Service Transition and Service Operation should make sure that the data is collected
Automated service reports and dashboards	Service Operation, while all phases should use the selected reports to compare achievements against metrics, KPIs and CSFs

Table 6.1: Outputs produced by the service CSI, and their downstream stewards and users

Other CSI outputs could be produced by either the service CSI or shared CSI function:

CSI output	Passed to
RFCs for implementing improvements	Service Transition; should be reviewed by change evaluation or change advisory board while considering the perspective of all phases, with the participation of their representatives if needed
Service reports	Potentially to all phases, depending on the report's subject and scope; this also determines its post-CSI steward
Feedback on strategies and policies	Service Strategy
Input to business cases and the service portfolio	Service Strategy
Feedback on service design packages	Service Design

CSI output	Passed to
Input to design requirements	Service Design
Input to testing requirements	Service Transition

Table 6.2: Outputs produce by either the service CSI or shared CSI function, and their downstream stewards and users

These CSI outputs could initially come from the shared CSI, since they are the inputs to the early phases of the service lifecycle, when the service CSI may not yet be in place. Their later iterations primarily come from the service CSI. The corresponding outputs from CSI of similar services can potentially be useful also. For example, service reports from other assessments may provide useful examples or indications of possible gaps. Requests for Change (RFCs) are an exception to this, as they should never come directly from other services' CSI.

The management practices and policies, such as periodic service reviews, are outside of the CSI scope. They may also be used to review the CSIR or detailed descriptions of significant improvements proposed in Appendix D. Within the current version of the ITIL framework, service reviews are an activity/practice associated with Service Level Management (SLM). The service reviews (for example, reviews with the PLM team or project management, or internal reviews within IT) certainly need to take into account the opportunities identified by CSI as their inputs, particularly when the authorisation for the proposed action exceeds the range of a single service or service activity.

Also, CSI provides guidance on how to structure periodic service reviews in terms of participants, inputs and outputs, as well as on the structure of related process assessments and service benchmarking that can be used as checklists and templates for service reviews.

6.1.1 Illustrative Examples

In SA2, Task 4 established the shared CSIR, which also includes the work from the previous project phase as the earlier recorded CSIR items were transferred. In addition, some products and services, such as perfSONAR, are adopting their own service-specific CSI Registers in order to structure and track individual improvements, including those that only have internal significance. Beyond the CSIR, Appendix D provides the improvement description template developed by Task 4 that could be used to further develop and detail individual initiatives.

6.2 Summary of Recommendations

- Assign an owner to each service.
- Incorporate effective leadership into the improvement process to arbitrate on cost / benefit matters, manage stakeholders, combine perspectives, and manage and address potential user dissatisfaction, resistance and implementation problems, especially with regard to cross-service improvements.

- Conduct a stakeholder analysis to identify perspectives and interests.
- To help ensure effective management and traceability of the improvement process, maintain a Continual Service Improvement Register (CSIR) [ITIL CSI] and use a structured improvement description template such as that shown in Appendix D.
- Reuse CSI outputs where possible, and ensure they are passed on to and fit for purpose in other lifecycle phases.

7 Conclusions

This *Continuous Service Improvement Best Practices* report reflects the experiences resulting from the actual work on improvements identification, design, and implementation of the optimisation opportunities that were pursued or supported by SA2 Task 4. It synthesises the recommendations and best practices that were applied or developed, and the useful observations made while working together with SA2 service teams and that are based on real-life situations. They cover the identification and promotion of improvement opportunities; use of precedents and existing models and available or obtainable information; and dealing with the people, organisation and processes involved in services. The relevant examples from GÉANT services and the work within SA2 given as illustrations should facilitate the adoption of collected best practices. The key messages to take away are:

- Data is the key enabler – available operational and service data, or additional data gathered for the purpose, measurements and metrics can and should be used to prioritise the interventions and quantify the merits and outcomes of improvements.
- To ensure and increase willingness for change, service teams must be convinced that the improvement intervention that disrupts already running services will bring tangible and long-term benefits.
- An inclusive and open attitude, collaborations and coalitions are key to success.
- Improvement is an iterative process – iterations are useful and welcome in many phases of the seven-step improvement process, which is also iterative itself.
- The presentation and promotion of improvements and sharing of accomplishments encourage and facilitate the involvement of stakeholders.
- Services can benefit from and are ready for advanced analytics and interactive visualisation.
- Commonalities and cross-service improvements associated with shared technical functions and ownership need top-level support and commitment in order to be introduced into several services.

As teams have already adopted and been implementing SA2 Task 4's improvement proposals, the Task is optimistic that the Continual Service Improvement best practices outlined here will be accepted, applied and useful in addressing the forthcoming improvement opportunities across GÉANT activities during its future project phases, and that some of the improvements carried out or initiated will be developed or continued in the future.

Appendix A Iterative Data Collection PDCA

The iterative data gathering process outlined in this Appendix is based on examples and recommendations provided in Section 4.1. It consists of one overall plan-do-check-act (PDCA) cycle and two or three embedded data entry PDCA iterations, where one iteration can be horizontally extended across all target groups, or all iterations can be vertically tested on one group first and then expanded to other groups.

A.1 Overall PDCA

The recommended approach consists of the following steps, grouped into the corresponding PDCA stages:

Plan

1. **Clarify data gathering objectives and provide explanations of terms and concepts**. Having a shared vocabulary and understanding is crucial.

Do

2. **Carry out initial data gathering, interpretation and validation**, usually at the technical level and closely involving the practitioners. See A.2.
3. **Optionally, carry out intermediate data gathering, interpretation and validation**, which includes rearranging, reinterpreting and supplementing the previously supplied data. Staging and structuring the effort helps in achieving a subsequent objective. See A.2.
4. **Carry out final data gathering, interpretation and validation** – this is sometimes just a revision of the previously supplied, possibly rearranged and aggregated data. This iteration can be assisted by access to the inputs provided by other target groups or preliminary generalisations, and usually includes a reinterpretation of the earlier inputs at the service, organisation, process, governance or legal level. See A.2.

Check (or Study)

5. **Establish a knowledge database, and produce the examples or guidelines** for the work to be conducted by services.
6. **Arrange implementation advice** or support.

Act (and Adjust)

7. Develop and **implement individual or shared data gathering solutions** and platforms.

8. **Establish permanent arrangements** as supporting services or separate service functions.
9. Establish and **enforce data-related policies**.

A.2 Data Entry, Interpretation and Validation PDCA Iteration

Every data gathering, interpretation and validation iteration listed in the "Do" stage of A.1 is a small PDCA cycle of its own, comprising the following steps:

Plan

1. **Refine and clarify inputs from reference materials** or domain experts.
2. **Produce data-entry template, introductory notes, instructions** on using the template, and terminology **explanations**.

Do

3. **Data entry** by the selected target group, with support from the coordinating team and domain experts.

Check (or Study)

4. **Review the data-entry process, data obtained**, data quality and the target group's **interpretation and understanding** of the assignment and individual data items.
5. **Refine the inputs** obtained **from the target group** as they may need to be streamlined, explained, contextualised or made more accessible in another way.

Act (and Adjust)

6. **Improve the template, instructions, or explanations** – this may include adaptation of these artefacts from other data gathering, interpretation and validation iterations.
7. **If needed, review input refinements or add** the **data** that was not originally provided, with the help of domain experts.
8. **Extend the data collection and interpretation** horizontally to other target groups, or continue it vertically, by performing subsequent data gathering iterations.

Appendix B Data Aggregation, Visualisation and Analysis PDCA

The approach to data analysis outlined in this Appendix is based on examples and recommendations provided in Sections 4.1 and 4.2. The individual steps are grouped into the PDCA stages:

Plan

1. **Learn about the goals of the exercise** in terms of the kind of conclusions or decisions that will be based on the data.
2. **Identify the relevant data** that is already routinely collected by the service or operations.
3. **Learn how to collect the data** and make the necessary arrangements.
4. **Learn about the formats used, available data fields and their semantics**. This includes locating and studying the available related user or developer documentation, exploring sample data and consulting with the service team.

Do

5. Once you have a general overview of the data, **make sure that the relevant data is regularly recorded**. Although some services may accumulate historical data, others may only maintain the current state or keep just a short history needed for day-to-day operations. The recommended approach is to record database or API output snapshots containing the relevant data. Periodic snapshots allow trend analysis even if the original sources were not designed to record the historical data and secure the data even when the source systems or processing and visualisation tools are changed.
6. Once the data format and semantics are sufficiently understood, **import one snapshot** or current database/output **into the analytics and visualisation (business intelligence) tool**. This serves to validate the understanding and assumptions, assess the data quality and usability, and learn about the limitations and problems. The potential basic problems may include data or precision losses, format problems, and misinterpretation of some specific values. Data quality, semantics and interpretation problems are usually more serious, and the analytical platform at this step serves to identify and assess them.
7. **If required** for processing the snapshot, **create a custom processing tool** in order to reformat, join, trim, filter, validate, remap, add missing data, etc. Besides mere conversion, this step provides an opportunity to **perform first-hand data consolidation** and enhancement. This processing tool should produce files that are ready for import into the visualisation platform, but also **produce daily logs or reports on data quality, errors or**

suspect values. These outputs could also be visualised in order to identify the most frequent issues and track trends.

Check (or Study)

8. **Review the processing logs, distributions and patterns of the original data** in order to understand it further, and, if needed, enhance the processing tool. Automated data quality reports produced by the processing pipeline allow first-hand assessment and data quality tracking that may result from some interventions or improvement initiatives.
9. Once the first snapshot is imported, it is quite easy to **import other collected snapshots**. Depending on the data and reporting timeframes, **automate the import process** in order to perform it on a periodic basis. **Check again** for potential data or format problems in the subsequent series, as these may be caused by changes in the imported data.
10. **Produce the data visualisations** and make them available to the authorised members of the service team. The data analysis tool should provide powerful data processing functions such as grouping, selecting and drilling down through the data. It can then be used to obtain comprehensive metrics displayed with interactive charts and maps. It is easiest to start with distribution visualisations of the most interesting fields, for any available point in time. It is useful to **group the produced visualisations into thematic dashboards**; for example, one dedicated to the visualisation of value distributions within the individual or combined snapshots, one for tracking of trends, and one dedicated to the analysis of issues. The trends analysis is easy to add once a number of snapshots have been imported and the corresponding temporal information is available, as the most relevant views used for individual points in time hint at what trend-tracking charts should be created. Some of the values and distribution statistics produced by the processing tool could be compared with the numbers obtained through the analytical platform, thus ensuring that no data was lost during the import.
11. **If appropriate, make the simplified form of selected charts available to a wider audience**. These charts may not be interactive and could be periodically recreated in order not to put a load on the analytical platform.
12. Upon the implementation of the analytical solution, **produce a dedicated report summarising what was achieved and how**. It should detail the findings, possible options and actions. It should suggest **what could be improved** in the service, but also in data generation and collection. This document may contain some sensitive details or discuss options that could be interpreted as commitments, and should therefore be distributed on a need-to-know basis. The implications of the suggestions should be discussed, as they may lead to changes in the data to be collected in the next improvement cycle and, more importantly, lead to major service changes. The document and debate should outline what needs to be done or additionally supplied in order to improve the service and enhance or extend the analysis of the existing data.

Act (and Adjust)

13. The service team should **review the proposed changes, the current plans and available resources and decide on the further course of action**, incorporating the agreed proposals into the **service roadmap**. The best changes are those that would both enhance the service performance and improve the data collection and analysis. When a thorough analysis of

service data is performed for the first time, it is likely that there will be plenty of opportunities to improve the data and it is wise to start by focusing on the subsets of data that could be enhanced with minimal effort and user involvement.

14. **Assess whether the overall approach and some improvements could be relevant to other services**, including the application of some of the developed data transformation tools. Even the visualisations, dashboards and reports provided for one service may provide inspiration to another.
15. **Promote further the adopted approach to data analytics**, as its elements could be of interest to a wider group of services and to infrastructure, product and project management.

Appendix C Use Case: GDPR

In some cases, an improvement is triggered and accelerated by practical needs or strict deadlines, as was the case with addressing EU General Data Protection Regulation (GDPR) requirements. The action roadmap and workflow used in dealing with it were applied across SA2, SA3 and JRA2 services, and provide a use case of staged and complex data gathering that was conducted with several participant groups, which resulted in modifications and realignment of internal processes within several services. Addressing GDPR requirements is an example of a widely relevant improvement area where coordinated work was required in order to adjust the existing services and processes. Furthermore, the developed workflow can also be used to align services with any requirement or regulation.

This case also illustrates how an emerging challenge can become a motivator to identify, mobilise and engage early adopters. After their feedback and involvement were used to validate and refine the approach, others were able to join more easily, gaining access to fully developed procedural steps, artefact templates, and practical examples produced by the GDPR team and service predecessors, among which the first were the eduroam, eduGAIN and eduPKI teams.

Addressing the subject was initiated by raising general awareness of the issues and requirements at project meetings and by forming the dedicated GÉANT GDPR team, which worked with SA2 Task 4 and individual services. The coordination across services ensured a consistent approach, and avoided duplication of work. After the initial awareness raising, understanding of the subject was enhanced by clarifying concepts, basic terms and the main areas of change resulting from the requirements. This information was captured and published on the GÉANT wiki, which is the accessible platform the community is familiar with.

The services developed their GDPR privacy notices by first filling in the data inventories. These were then used as input for populating data-mapping templates, which in turn provided the key elements that were considered and used in writing the service privacy notices. The services were able to build on the outputs of previous phases, had access to prepared supporting materials, and were assisted throughout by the GDPR team, who also provided a comprehensive example of a privacy notice, on which the privacy notices for individual services were based. The SA2 Task 4 team helped to create the privacy notices for those service managers who requested additional support. The materials produced in each of these steps were gradually refined based on feedback obtained.

The GDPR-related work for all the services addressed was tracked on a shared summary dashboard on the GÉANT wiki (as the area is restricted access, no link has been provided here). This page serves as an internal checklist and navigational tool, and also as a starting point for documentation in the event of GDPR audits. Each service has a dedicated row in the dashboard table, where the presence of a link or value indicates that the output corresponding to the step was produced, or a decision

made. All produced wiki templates and pages are made available from the dashboard. It also links to comprehensive GDPR information about GÉANT's role, legal basis, data processing, privacy notices, contracts with data processors, records of processing activities, consent management, data protection impact assessment (DPIA) pre-evaluation and incident management. All these GDPR-related steps and resources also help the service teams to complete the DPIA pre-evaluation checklist and thus determine whether they need to conduct the data protection impact assessment.

The GDPR team is now ideally placed to provide advice and consultancy on GDPR-related tasks, thus maintaining the momentum and sustaining the proper handling of the subject within the emerging services even after the completion of the campaign to align the production services with the GDPR requirements and related practices. It also supports other activities that are impacted by this subject. For example, the GDPR team provided consultancy to SA2 Task 1 for addressing GDPR "privacy by design" and "privacy by default" principles within Secure Code Training (SCT) and secure code audits.

Data inventories used in addressing GDPR requirements are additionally valuable during the migration to microservices architecture (see Section 5.2 of this document). The information model that was captured in the process is also very helpful in the integration of service components, orchestration and automation of workflows, which are all additionally supported by a broader formalisation and harmonisation of business processes.

C.1 PDCA in Practice

The work described above can be used to illustrate how the generalised and possibly slightly abstract iterative data-entry plan-do-check-act (PDCA) model described in Appendix A reflects the practical steps of dealing with GDPR. Furthermore, the more elaborate steps of the Plan and Act stages in Appendix B are, after a few minimal modifications in wording, also applicable here, the same way the work on GDPR is related to data aggregation and analysis.

Plan

1. **Clarify data gathering objectives and provide explanations of terms and concepts**.
    - Raise general awareness of the issues and requirements, e.g. in project meetings.
    - Form the dedicated GÉANT GDPR team.
    - Define GDPR requirements for SA2 production services, namely, to produce privacy notices and align internal procedures, data-related tools and practices.
    - Clarify concepts and publish supporting materials on the GÉANT wiki.
    - Draft general information regarding the GDPR and glossary documents.
    - Participate in awareness-raising events on GDPR.
    - Identify early adopters.

The corresponding Plan steps from Appendix B, the data aggregation, visualisation and analysis PDCA, can also be recognised in the conducted work (here changed to a, b, etc. to differentiate more clearly from the steps of the data collection PDCA):

a. Learn about the goals.
b. Identify the relevant data that already exists and is routinely collected.
c. Learn how to collect the data.
d. Learn about the available data and their semantics.

Do

2. **Carry out initial data gathering, interpretation and validation**.
   - Produce the data inventories by performing the PDCA outlined in A.2. One service often incorporates several data sets. The service manager fills in the table with the predominantly technical and familiar details about all service-related data sets (data collections) that are collected or processed by their service, each with a list of the data items contained.
   - Review the table. This is done by a member of the GDPR team.
   - If necessary, enhance or clarify the data inventory with the service manager.
   - Decide whether an item is (a part of) personal data. This done by the GÉANT Data Protection Officer.
3. **Carry out intermediate data gathering, interpretation and validation**.
   - Produce data mappings by performing the PDCA outlined in A.2. This was undertaken by the service managers, assisted by the GDPR team, using the data inventories as input for the data-mapping tables.
4. **Carry out final data gathering, interpretation and validation**.
   - Draft and approve privacy notices by performing the PDCA outlined in A.2. The data-mapping tables provided the key elements to be considered while producing service privacy notices, the writing of which was also supported by comprehensive examples of privacy notices.

Check (or Study)

5. **Establish a knowledge database, and produce examples or guidelines**.
   - Initiate GDPR dashboard development by establishing transitional wiki pages with instructions, explanatory documents, other educational materials, templates, samples, abbreviations list and team members.
6. **Arrange implementation advice**.
   - Make advice and support from GDPR team, SA2 Task 4 and early adopters continually available as needed by new services.

Act (and Adjust)

7. **Implement individual or shared data gathering solutions**.
   - Consolidate templates; add tracking matrix and progress colour coding and tagging conventions on the GDPR dashboard.
8. **Establish permanent arrangements**.
   - Finalise general information on GDPR, glossary and FAQ documents.

- Continue work on data inventories, data mappings and privacy notices for new services and data processor agreements for services in production.
- Liaise with other GÉANT and external activities.

9. **Enforce data-related policies**.
   - Production services incorporate the documents produced during the GDPR exercise into the service documentation in the corresponding service wiki pages, enact and follow the policies described in their privacy notices.
   - Every new service must produce a data inventory, data-mapping table and privacy notice during its transition into production.

The corresponding Act steps from Appendix B can also be recognised in the conducted work (as before, changed to the equivalent alphabetical value to differentiate more clearly from the steps of the data collection PDCA):

m. Review the proposed changes, the current plans and available resources and decide on the further course of action.
n. Assess whether the overall approach and some improvements could be relevant to other services.
o. Promote further the adopted approach.

Additional information is available in Section 5.1 of *Deliverable D5.6 Service Improvement Report* [D5.6].

Appendix D Improvement Description Template

This appendix describes the improvement description template developed by SA2 Task 4 that can be used to further develop and detail individual initiatives.

<Improvement Name>

A short summary of the intervention subject, results, and possible outcome and impact (in terms of business significance). In the executive summary style.

D.1 At Initiation

This may range from a brief memo about the new idea to a comprehensive project brief, but should at least contain the following elements:

Problem/Subject

- The state before the improvement and where the problem lies.
- Approximate size/complexity.
- Deadline or approximate timescale.
- An indication of priority.

Objectives & Measures

Description of the metrics related to the improvement and, where present, their target values.

Immediate (and specific, measurable, achievable, relevant, time-bound (SMART)) objectives, particularly:

- Reduction of:
  - Risks.
  - Rework.
  - Defects.
  - Delivery errors.
  - Cost.
- Increase in:
  - Customer satisfaction.

- Productivity.
- Decision-making ability.
- Accuracy.
- Competitive advantage.

Implementation Proposal

High-level plan, resources, man effort, skills, job descriptions, etc.

D.2 At the Beginning of Implementation

The nature and extent of changes to the information in the template at the beginning of implementation will depend on the scale of the improvement "project", but should at least include updating the initial description and additionally describing:

- Intervention content (what is to be done).
- Improvement process and milestones and timeline.
- Anticipated issues, obstacles, risks and possible responses.
- Expected outputs and results.

D.3 After Completion

Action Description

- Intervention content, i.e. what has been done.
- Improvement process (dates, plan, phases, steps).
- Issues, obstacles, risks and related adjustments: anticipated, addressed and remaining problems or obstacles and measures to overcome them.
- Elaboration of concrete outputs (with a list of outputs, if possible).

Impact Achieved

- Elaboration of achieved impacts such as on: capabilities, processes, solutions, use of resources, documentation, repeatability, and key achievements (qualitative and, if possible, quantitative)
- Change in observed metrics.
- Buy-in and sustainability of the change or improvement. This aspect is often problematic!

Further Work

- What is next?
- Wider applicability.
- Lessons learned.

References

[BPF]	https://www.tmforum.org/business-process-framework/
[D5.1]	*Deliverable D5.1 Analysis of Service Elements and Optimisation Opportunities*. Confidential. Available to GN4-2 project participants at: https://intranet.geant.org/gn4/2/Activities/SA2/Deliverables%20Documents/Analysis%20of%20Service%20Elements%20and%20Optimisation%20Opportunities/D5-1_Analysis-of-Service-Elements-and-Optimisation-Opportunities.pdf
[D5.6]	*Deliverable D5.6 Service Improvement Report*. Confidential. Available to GN4-2 project participants at: https://intranet.geant.org/gn4/2/Activities/SA2/Deliverables%20Documents/Service%20Improvement%20Report/D5.6_Service-Improvement-Report.pdf
[eduGAIN-T]	https://technical.edugain.org/
[eduGAIN-PR]	https://edugain.org
[eduPKI]	https://www.edupki.org/home/
[eduroam-M]	https://monitor.eduroam.org
[ELK]	https://www.elastic.co/elk-stack
[GeoIP]	http://www.geoip.co.uk/
[ITIL CSI]	*ITIL Continual Service Improvement: 2011* (Best Management Practices), Vernon Lloyd, The Stationery Office; July 29, 2011, ISBN-10: 011331308X, ISBN-13: 978-0113313082
[M5.1]	*Milestone M5.1 Initial Service Improvement Guidelines*. Confidential. Available to GN4-2 project participants at: https://intranet.geant.org/gn4/2/Activities/SA2/Milestones%20Documents/Initial%20Service%20Improvement%20Guidelines/M5-1_Initial-Service-Improvement-Guidelines.pdf
[M5.3]	*Milestone M5.3 Service Improvement Interim Report*. Confidential. Available to GN4-2 project participants at: https://intranet.geant.org/gn4/2/Activities/SA2/Milestones%20Documents/Forms/AllItems.aspx
[perfSONAR]	https://www.perfsonar.net/
[SAEM]	*Service Architecture Elements and Microservices* https://intranet.geant.org/gn4/2/Activities/SA2/Shared%20Documents/Service%20Architecture%20Elements%20and%20Microservices.docx
[SPLUNK]	https://www.splunk.com/
[T2MSA]	*Transforming BSS/OSS systems to Microservices Architecture*, https://www.tmforum.org/resources/whitepapers/transforming-bss-oss-systems-to-microservices-architecture/

Glossary

AARC	Authentication and Authorisation for Research and Collaboration
API	Application Programming Interface
AUP	Acceptable Use Policy
BPF	Business Process Framework
BSS	Business Support System
CSF	Critical Success Factor
CSI	Continual Service Improvement
CSIR	Continual Service Improvement Register
DPIA	Data Protection Impact Assessment
ETL	Extract-Transform-Load
eTOM	Business Process Framework (TMF)
Faas	Federation as a Service
FIM4R	Federated Identity Management for Research
GDPR	General Data Protection Regulation
GN4-1	GÉANT Network 4 Phase 1, a project part-funded by the EC's Horizon 2020 programme under the Specific Grant Agreement No. 691567
GN4-2	GÉANT Network 4 Phase 2, a project part-funded by the EC's Horizon 2020 programme under the Specific Grant Agreement No. 731122
IETF	Internet Engineering Task Force
ITIL	Information Technology Infrastructure Library
JRA	Joint Research Activity
JRA2	GN4-2 Joint Research Activity 2 Network Services Development
JRA3	GN4-2 Joint Research Activity 3, Trust and Identity Development
JSON	JavaScript Object Notation
KPI	Key Performance Indicator
NA	Networking Activity
OSS	Operations Support System
PDCA	Plan-Do-Check-Act
PLM	Product Lifecycle Management
RA21	Resource Access for the 21st Century
RAML	RESTful API Modelling Language
REFEDS	Research and Education FEDerations
REST	Representational State Transfer
RFC	Request for Change
RPC	Remote Procedure Call
SA	Service Activity
SA2	GN4-2 Service Activity 2 Trust & Identity and Multi-Domain Services
SA3	GN4-2 Service Activity 3 Network and Services Assurance

SA4	Service Activity 4 of the GN4-1 project
SAML	Security Assertion Markup Language
SCT	Secure Code Training
SKMS	Service Knowledge Management System
SLA	Service Level Agreement
SLM	Service Level Management
SMART	Specific, Measurable, Achievable, Relevant, Time-bound
T&I	Trust and Identity
TAM	The Application Framework (TMF)
Task 1	SA2 Task 1 Service Transition and Software Management
Task 2	SA2 Task 2 Trust and Identity Operations
Task 4	SA2 Task 4 Production Optimisation
TMF	TeleManagement Forum (or TM Forum)
UoB	University of Belgrade