



SPECIAL

**Scalable Policy-awareE Linked Data arChitecture for
prIvacy, trAnsparency and complIance**

Deliverable 5.5

Pilot implementations and testing plans V3

SPECIAL DELIVERABLE

Name, title and organisation of the scientific representative of the project's coordinator:

Jessica Michel Assoumou t: +33 4 97 15 53 06 f: +33 4 92 38 78 22 e: jessica.michel@ercim.eu

GEIE ERCIM, 2004, route des Lucioles, Sophia Antipolis, 06410 Biot, France

Project website address: <http://www.specialprivacy.eu/>

Project	
Grant Agreement number	731601
Project acronym:	SPECIAL
Project title:	Scalable Policy-awareE Linked Data arChitecture for prIvacy, trAnsparency and compLIance
Funding Scheme:	Research & Innovation Action (RIA)
Date of latest version of DoW against which the assessment will be made:	17/10/2016
Document	
Period covered:	M1 – M35
Deliverable number:	D5.5
Deliverable title	Pilot implementations and testing plans V3
Contractual Date of Delivery:	31.10.2019
Actual Date of Delivery:	30.11.2019 (as agreed with PO)
Editor (s):	Martin Kurze (DTAG)
Author (s):	Benedict Whittamsmith, Rudy Jacob, Martin Kurze
Reviewer (s):	Sabrina Kirrane, Piero Bonatti
Participant(s):	TR , Prox, DTAG/TLABS
Work package no.:	WP5
Work package title:	Use Case Implementation & Evaluation
Work package leader:	TR
Distribution:	public
Version/Revision:	1.1
Draft/Final:	Final
Total number of pages (including cover):	53

Disclaimer

This document contains description of the SPECIAL project work and findings.

The authors of this document have taken any available measure in order for its content to be accurate, consistent and lawful. However, neither the project consortium as a whole nor the individual partners that implicitly or explicitly participated in the creation and publication of this document hold any responsibility for actions that might occur as a result of using its content.

This publication has been produced with the assistance of the European Union. The content of this publication is the sole responsibility of the SPECIAL consortium and can in no way be taken to reflect the views of the European Union.

The European Union is established in accordance with the Treaty on European Union (Maastricht). There are currently 28 Member States of the Union. It is based on the European Communities and the Member States cooperation in the fields of Common Foreign and Security Policy and Justice and Home Affairs. The five main institutions of the European Union are the European Parliament, the Council of Ministers, the European Commission, the Court of Justice and the Court of Auditors (<http://europa.eu/>).

SPECIAL has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 731601.

Table of Contents

1	Summary	6
2	Introduction to pilots and testing plans in SPECIAL	7
2.1	Types and purposes of pilots, business/industry context	7
2.2	Purpose and requirements for tests/testing plans	7
3	Pilot 1 (Proximus)	9
3.1	Summary	9
3.2	Objective and evaluation for Proximus	9
3.3	Data protection considerations	10
3.4	Dissemination within Proximus	10
3.5	Proximus use case	11
3.5.1	Technology choices for the use case iterations	11
3.5.2	Iteration 2 – Full functional pilot (M35)	13
3.6	Use of SPECIAL components	16
3.6.1	Policy Language & W3.org data vocabulary community group	16
3.6.2	Sticky Policies	18
3.6.3	Privacy Dashboard	18
3.6.4	Compliance checker	19
3.6.5	Event log	20
3.6.6	Dynamic Consent	20
3.6.7	Conclusion on SPECIAL components	20
3.7	Results of the user tests	21
3.7.1	Gender & Age	21
3.7.2	Expertise & daily time on the internet	21
3.7.3	Preferred internet browsing device	22
3.7.4	GDPR and your privacy	22
3.7.5	Agreement with popular beliefs surrounding privacy	23
3.7.6	Findings after testing the application	23
3.8	Conclusion	26
4	Pilot 2 (Refinitiv)	27
4.1	Summary	27
4.1	Pilot objectives	27
4.2	Data protection considerations	28
4.3	Pilot components	28
4.4	Evaluation criteria	29
4.5	Policy language, information models, and standardisation evaluation	29
4.6	Business Process	30
4.7	Compliance Checking	33

5	Pilot 3 (Deutsche Telekom)	34
5.1	Remark	34
5.2	Use case description (short recap)	34
5.2.1	Additional use for device (client) based QoS/location data	34
5.2.2	Previous (pre-SPECIAL) situation	34
5.2.3	Target scenario for pilot implementation	36
5.3	Objective of pilot implementation (in short)	38
5.3.1	DT's view on Big Data and AI and their relevance to SPECIAL	38
5.3.2	Objectives for the DT pilot	38
5.3.3	What is NOT in scope	39
5.4	Pilot implementation architecture	40
5.5	Data flow	41
5.6	Data flow specification in GDPR terms for legal assessment	43
5.6.1	Types of user consent/policy	43
5.6.2	Background and role of Motionlogic:	44
5.6.3	Detailed explanations on some of the terms/tasks mentioned above:	45
5.7	Use of SPECIAL components	46
5.7.1	Policy Language	46
5.7.2	Sticky Policies	47
5.7.3	Privacy Dashboard	47
5.7.4	Compliance Checker	47
5.7.5	Event log	47
5.7.6	Dynamic consent	47
5.8	Testing	48
5.9	Results and Evaluation	49
5.9.1	Results and Observations	49
5.9.2	Evaluation	50
6	Conclusion	53

1 Summary

Three industry partners (Proximus, Refinitiv and Deutsche Telekom) describe their pilots and testing plans for components, tools and concepts developed in SPECIAL by the consortium. The pilots and implementation plans are based on previous documents/deliverables (namely D5.1, D5.3 and as well as D5.2, D1.1, D1.2, D1.3, D1.4 and subsequently D2.4, D3.2, D3.3 and D3.6)

This document was designed for “staging”, i.e. it was reworked on a regular basis. The present version “Deliverable 5.5 V3” is the third and final iteration and includes implementation plans, results of the (already implemented) pilots and the test plans/results. Due to the different objectives and types of implementation, the industry partners were assigned one chapter each. These chapters may be read relatively independently of each other. However they all have certain aims and findings in common. These are collected in the overarching chapters 2 and 6.

Since this work is based on “pilots” (not on products), neither a profound summary nor ultimate conclusions can be expected. The reader is kindly asked to take advantage of the current (mature but still preliminary) state of findings.

The present document aims at giving the reader a comprehensive view into the current (final) state of implementation of the pilots, testing plans as well as expected and achieved findings.

2 Introduction to pilots and testing plans in SPECIAL

2.1 Types and purposes of pilots, business/industry context

Three different industry partners, each with their own use case and differing overall objectives designed pilots for implementation. This resulted in different types of pilots, tests and even different types of descriptions in this deliverable:

- **Proximus** is aiming (short- to mid-term) towards a new, innovative commercial product in the recommender-business based on personal data. In this early phase of the product development, the prototype aims at implementing and testing a relatively complex use case with relatively small amount of data subjects.
- **Refinitiv** is interested in applying methods and tools to automatically read and evaluate privacy policies and a company-internal workflow. Therefore the focus is very much on the policies, policy language and policy checking tools. The number of planned users is not so much in focus since the system will mainly be applied internally.
- **Deutsche Telekom** (with its research & development units VTI/T-Labs) is aiming more at applications that use the huge amount of (personal) data that exists in DT and today cannot be monetized due to the current tight interpretation of applicable privacy legislation, e.g. GDPR. Thus, DT needs a PoC (proof of concept) that shows that GDPR compliant collection, processing and even sharing of data is feasible in a legal way, once tools like those developed in SPECIAL are available and mature. Therefore, DT's pilot is relatively simple and straightforward (supplementing an existing process): to prove feasibility, test usability and estimate effort to reach this goal. While the use case is relatively small, the number of potential users is relatively high which will give DT a sound basis for user acceptance.

Due to the inherent differences of the pilots, the respective chapters of this document are under full authorship of the individual industry partners. This also explains the slightly different styles of writing.

2.2 Purpose and requirements for tests/testing plans

All industry partners plan to evaluate all relevant aspects of SPECIAL's results during the course of the project. Given the novelty of the approach and the complexity of implementation of any aspect of it in an "industrial" environment, industry partners decided to stepwise evaluate parts of the results or intermediate results. This not only minimizes the evaluation work, it also allowed (and still allows) the consortium to enhance work results during several iterations.

Currently, the industry partners focus on the most relevant aspect for their business, not neglecting the fact that several other aspects need to be evaluated and are planned to be implemented in the operational business succeeding the project duration. Fundamental requirements will be tested repeatedly (e.g. security, privacy and usability; see below).

Iteration 1 (initial document, deliverable D5.1)

Proximus mainly evaluated the overall architecture and the composition of building blocks, general feasibility needed to be checked, and first experiences with corporate internal processes needed to be collected. Furthermore tests in conjunction with Proximus' recommender engine have been carried out.

Refinitiv has a strong focus on the policy language and checked its' expressivity and compliance in detail.

Deutsche Telekom (DT) used this phase for a complementary approach: internal corporate (IT) processes were checked or revised to implement core aspects of SPECIAL. This lead to an in-depth

evaluation of SPECIAL's modularity. Also scalability was (and is) tested relatively early with a "minimal viable prototype" (put in quotation marks because DT does not expect a product level tool but rather a prototype or PoC – Proof of Concept –that delivers "viable"/expedient results) focussing on a small set of privacy relevant data.

Iteration 2 (previous document, deliverable 5.3)

Proximus extends the user group and thus plans to focus on scalability of the approach. Focus will be on transparency and compliance.

Refinitiv in this phase looks towards generalising some of the lessons of applying the SPECIAL Technology to its use case into wider/arbitrary use cases in the financial data supply chain.

Deutsche Telekom focuses on technical ease of implementation and applicability of DT's corporate design and corporate identity rules of the tools. Detailed alignments with corporate regulation and affected departments/subsidiaries took place and lead to a comprehensive view and product (feature) opportunity. In addition, DT also ran a first evaluation of business relevance.

Iteration 3 (this document, deliverable 5.5)

All industry partners will evaluate the applicability of their respective pilots in their organizations. Also the overall concepts of SPECIAL (linked data, usage policies etc.) will be evaluated utilizing first experiences with the pilots. Finally, a first approximation of the cost/effect ratio will be sketched, i.e. the question of whether the effort of using SPECIAL technology pays off in the given pilots and potentially in other industry use cases will be given a first (preliminary) answer.

Testing and test plans are based on use cases and requirements. Since the use cases and objectives are very specific for each industry partner, purpose and testing requirements for the tests are use-case specific as well.

Nevertheless, two requirements apply to all pilots:

Security and Privacy: While these are "hygiene features" of usual products (security and privacy need to be present and sufficient with respect to applicable laws and state-of-the-art technology), in the case of SPECIAL, particular attention needs to be put on fulfilling all privacy needs. SPECIAL needs to provide methods and tools to implement GDPR compliant applications, while the business purpose of the applications is not "privacy". This results in the need for special attention to security and privacy on the conceptual as well as on the implementation level.

Usability: Even though the number of users of the initial versions of pilot implementations is limited, it is obvious that usability needs extra devotion and effort since users will not make use of their control forces if they are not willing or able to use the respective tools/interfaces. Thus, bad usability will either result in poor functionality of the products (because users will not allow any use of personal data) or it will leave the impression that user's don't care anyway.

3 Pilot 1 (Proximus)

3.1 Summary

A detailed description of the Proximus use case can be found in deliverable D1.5 (2nd version of the requirements definition). The use case concentrates on the following concepts:

- Recommender Engine application for events at the Belgian coast
- Mobile responsive user interface
- Data subjects of Proximus
- Personal data requested from data subjects:
 - Location
 - Television viewing behaviour

The first iteration, an initial end-to-end version, was ready for internal Proximus beta testing in 2018. In 2019, a fully functional second iteration was finalised. The second iteration focused on the transparency and compliance functionality based on deliverables D3.4 & D4.5, as well as on the integration via linked data with the “W3C data privacy vocabulary and controls community group”¹ (deliverable D6.5). The prototype was tested by 12 data subjects resulting in very interesting insights in the understanding of personal data privacy.

Proximus also evaluated the usefulness of all the SPECIAL components (called building blocks) for its use case.

3.2 Objective and evaluation for Proximus

As a general objective, Proximus wants to test the willingness of its customers to share personal data. The current use case is not commercial, precisely to avoid being valued for any immediate commercial return. What is more important, is the question as to whether sharing personal data attracts or scares people. We expect this answer to be dependent on several variables:

- Intuitive interface
- Getting something in return
- Age of the data subject (e.g. Is there a difference between millennials and 50+ people?)
- (Mobile) IT awareness of the data subject
- Privacy awareness of the data subject

In the first iteration (with 5 data subjects), we were looking into technically testing whether the solution works end-to-end, and focused highly on the **intuitive interface**, to “attract” enough data subjects for the second iteration. This was successful as all data subjects agreed to sign up for the second iteration of tests. **The data subject’s age range was from 30 to 70, and all were very much privacy/mobile/IT aware.** They found the interface intuitive, but it must be said that the interface only asked for basic consent on three types of personal data and did not yet have a privacy dashboard.

The “**getting something in return**” was not good, as the quality of the recommendations was rather low, due to a capacity issue within Proximus on the data science side to improve the recommender engine. This was addressed to the Consortium, but the building of the recommender engine is considered not part of SPECIAL and therefore, impossible to reassess the use of Consortium capacity/funds. As an alternative, the Proximus team got the help from a trainee student whom developed the first version of the recommender engine and gave suggestions how to further enhance it. For the second iteration of the tests with the data subjects, the recommender engine itself was not

¹ <https://www.w3.org/community/dpvcg/>

improved, but we were able to use existing, even though more 'general', calculated profile categories delivered by the data intelligence team.

The second iteration focused more on the Privacy Dashboard and the impact of seeing real personal data being used. The size of the test group was 12 persons, having a better age and IT skills distribution. While almost all the participants thought of GDPR as a good thing, they felt helpless and it seemed that they had to learn a new vocabulary first. Words like controller, processor, third party were confusing them more than helping them. The test results further in this chapter will explain why.

3.3 Data protection considerations

The SPECIAL project team made the following assessment as to the identity of the Data Controller and the Data Processor.

- Proximus is the Data Controller.
- Proximus itself uses many external cloud services when it comes to prototyping, as it is simply faster and cheaper. For the SPECIAL project, budget and time constraints led us also in this direction and therefore we use Microsoft Azure as a Data Processor for the Data Science recommendations. (Please note that this will be different when running in production)

On July 19th, 2018, the first iteration of the pilot implementation was presented to the internal Proximus Privacy Council and approval was obtained to formally approach 5 data subjects (internal to Proximus and all part of - or very close to - the SPECIAL project team who showed informal willingness to participate) and to explore the technical solution to retrieve and store the TV viewing records.

Having access to location/TV viewing/browsing history is a very sensitive matter for the Privacy Council, and to get an initial Go, it was decided to withhold location and browsing history for future iterations. This incremental approach was recommended to avoid having to deal with the more complicated governance of the Proximus project methodology with a release calendar beyond our control.

On Aug 23rd, 2019, the project team presented the second iteration of the pilot to the internal Proximus Privacy Council asking for approval for a higher number of data subjects (15) and adding location as a data source. Approval was obtained.

The project must also undergo a Data Privacy Impact Assessment (DPIA) which is a standard procedure for any application that deals with personal data. This is also a SPECIAL project requirement as mentioned in WP8.

The DPIA assessment involves:

- (1) Filling out a Privacy Inventory document
- (2) A risk assessment and double-checking of contracts with all (sub)processors
- (3) Presentation to a Legal Committee

3.4 Dissemination within Proximus

The involvement of the SPECIAL Proximus project team is disseminated within Proximus via:

- (1) Alignment with the Proximus project team working on GDPR compliance for all IT/CRM systems. Feedback from SPECIAL legal advisors lead to updates in the MyProximus APP consent page.
- (2) Alignment with the Enterprise Architecture team (one SPECIAL project team member is part of the Enterprise Architecture team).
- (3) Alignment with the Proximus Privacy Council. The last meeting was held in August 2019, to get approval for performing tests with 15 data subjects, as well as to start the Data Privacy Impact Assessment process.

- (4) Alignment with the Consumer Business Unit with its innovation initiatives.
- (5) Final debrief meeting in 2020 with all interested parties upon conclusion of SPECIAL.

The Proximus use case was also presented in Oct 2019, at the ETIS Data Privacy Task Force² of which Proximus is a member.

3.5 Proximus use case

3.5.1 Technology choices for the use case iterations

The volume and streaming of payload consumer data for the SPECIAL Proximus use case obviously qualifies to be called Big Data at Proximus. In 2017, Proximus decided that all (old and new (Big) data) initiatives will simply be called Data initiatives.

Somewhat different choices (Table 3.1) were made between iteration 1 (fast deployment of an end-to-end application) and iteration 2 (more robust and scalable solution with access management within our control). Node.js was chosen for future scalability to be able to run on a big data system.

Table 3.1 Technology choices for the SPECIAL use case at Proximus

Functional area	Iteration 1 (M20): (End-to-End Proof of Concept with 5 data subjects)	Iteration 2 (M35): (fully functional pilot with 15 data subjects)	POST SPECIAL: Possible Integration within Proximus
Identity & Access Management	Firebase (manual)	Based on JSON webtokens https://jwt.io/ (open, industry standard RFC 7519); Secure but not production ready	MyProximus Identity & Access Management
Frontend UI	(developed by TUB) One basic vertical page (React.js)	(developed by TUB) Compliant interface validated by the Consortium Legal partners (incl T&Cs) (React.js)	MyProximus libraries with input from Iteration 2
Privacy Dashboard (TUB)	Was not implemented	Implemented and contains the data subject's rights (right to access/rectification/erasure)	Integrate in MyProximus which currently only has a basic Privacy Dashboard (no right to access/rectify/erase)
Layered Privacy Statement/Notice	Dynamic Consent suggested by SPECIAL team but still in ideation phase	Dynamic Consent suggested by SPECIAL team is still in ideation phase and not applied in the use case.	Evaluate once mature

² <https://www.etis.org/page/DPTF>

Functional area	Iteration 1 (M20): (End-to-End Proof of Concept with 5 data subjects)	Iteration 2 (M35): (fully functional pilot with 15 data subjects)	POST SPECIAL: Possible Integration within Proximus
Consent Datastore	Firestore Database	MongoDB with https://nodejs.org/en/ and http://expressjs.com/	MyProximus Identity & Access Management
Event Datastore	MS Azure	MS Azure	MS Azure
Profile Datastore	MS Azure (MySQL / MongoDB)	MS Azure (MongoDB)	MyProximus Interest Profile dB
Payload data filtered input stream	MS Azure MySQL for (real) TV viewing records, or manually marked TV viewing list by each data subject	MongoDB with real location data, TV viewing records and derived profile data. Browsing history is synthetic data.	Kafka + datastore (MongoDB)
Machine learning	Python	Python	Python
SMS / email sending	Proximus Enco SMS API	Will be in the APP; No more need for additional SMS/email	In the APP

3.5.2 Iteration 2 – Full functional pilot (M35)

Figure 3.2 below represents the high-level architecture of the solution.

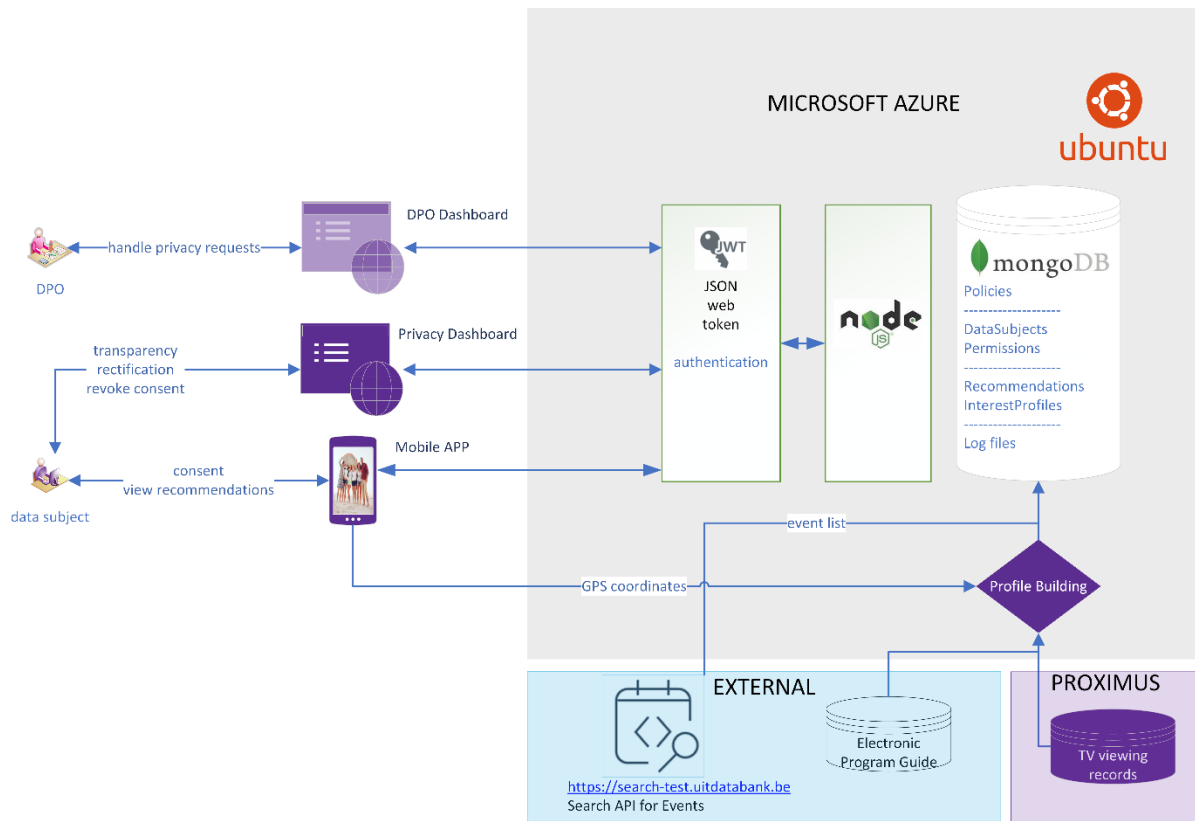


Figure 3.2. High-level architecture of the solution.

Important to note is that only TV viewing records are used directly from Proximus. For the internet browsing, we used synthetic data. For the location records, we used the GPS coordinates of the mobile phone or the location of the web browser using the W3 geolocation framework³.

There were several reasons for not using Telecom radio location records (= the location of the antenna with which the mobile phone communicates):

- (1) Accuracy is not the same as GPS. Triangulation with various antennas would be needed.
- (2) Technical availability: Those records are only briefly available for guaranteeing GSM network quality (legitimate interest) and no link is made with the subscriber. Therefore, impossible to filter per data subject.
- (3) The asymmetric regulatory approach to the use of GPS location data subject to GDPR versus mobile location data subject to e-Privacy rules has not been resolved, adversely affecting the telecoms sector's ability to innovate and compete with providers of apps and mobile operating systems.⁴

³ <https://www.w3.org/TR/geolocation-API/#introduction> framework


⁴ <https://www.etno.eu/news/all-news/8-news/654-eu-digital-leadership-eprivacy.html>

The application will be available at the following link⁵ until Feb 2020 after the final project review. Personal data will have been deleted and only a generic user with dummy data will remain as of Dec 1st, 2019.

It has a Technology Readiness Level 4 which means “technology validated in a lab environment”. Below are three screen shots of the “Events Nearby” application that show the landing page (Figure 3.3), the consent page (Figure 3.4), and the interest profile page (Figure 3.5).

⁵ <http://special-webapp.westeurope.cloudapp.azure.com/>

Events nearby



Capoeira in Knokke

Braziliaans dansgevecht voor kinderen (+3 jaar), jongeren en volwassenen. Combineert dans, gevecht, ritme, muziek, kracht, lenigheid, waarden en normen. ...



Tai Chi

Elke ochtend ziet men miljoenen Chinezen trage harmonieuze bewegingen uitvoeren in de parken. Zij beoefenen een kunst, Tai Ji Quan genaamd. Deze kunst werd gedurende eeuwen geheim gehouden en wordt sedert een vijftigtal jaren in de westerse wereld verspreid. Het ...



Figure 3.3



Choose your consent settings

You can benefit from better event recommendations by telling us what you're interested in. But you can also use our service without giving us any data.

General events

For general event recommendations we don't need personal data. However, it could happen that we inform you about events you're not interested in.

Personalized events

For personalized event recommendations we use your location information, TV and browsing behavior in order to inform you about events that interest you.

Configure personal data sources

Below you can configure which data we can use to personalize your events recommendations.

- Location and places
- Television data
- Browsing history

Figure 3.4



Interest profile

These are keywords that describe your interest.

Based on TV preference scoring

- kids
- sport
- foot
- movies
- series
- romance

Based on visited places

- (work) Brussels
- (home) Aalst

Based on browsing

- Concert
- News
- Music

Figure 3.5

3.6 Use of SPECIAL components

The following section describes where and how the Proximus use case is using components from SPECIAL.

3.6.1 Policy Language & W3.org data vocabulary community group

As mentioned in D2.5 (Policy Language V2), the usage policy language is meant to express both data subjects' consent and data usage policies of data controllers in formal terms, understandable by a computer, to automatically verify that the usage of personal data complies with data subjects' consent. For an industry partner the challenge is to convert the theoretical basis of the policy language into a workable product. This has been accomplished with the help of Consortium partners CeRict and TUB. Figure 3.6 explains the relation that exists for a data subject between permission (=consent), policy and data.




Figure 3.6. Relation between Data Subject and Data/Policy/Permission (in RDF notation).

An example of a possible implementation of such a policy for the Proximus use case (source CeRict and D2.5) is given here:

```

P2: Use personal data for recommendations
{
  hasData: {
    Personal Identifiable Information
      AudiovisualActivity + Location + OnlineActivity
    }
  hasProcessing: Collect + Copy + Transfer
  hasPurpose: PxsEventRecommendation
  hasRecipient: Ours
  hasStorage: {
    hasDuration: StatedPurpose
    durationInDays < 42
    hasLocation: {EU ControllerServers}
  }
  hasLegalBasis: Art6_1_a_Consent
  hasDuty: GetValidConsent
}
  
```

The above example is formulated in an extension of JSON, that can be related to the standard OWL2 formulation described in D2.5 as follows:

- curly brackets correspond to ObjectIntersectionOf
- '+' corresponds to ObjectUnionOf
- has XXX:Y corresponds to ObjectSomeValueFrom(hasXXX,Y)

A more practical implementation of a policy, at the cost of enlarging it, is the following (source TUB). The policy P2 above can be converted into 18 policies of the type below. This policy also has an internal identifier inside its definition ("id":1):

```
{
  // Collect location data
  "id": 1,
  "data": "http://www.specialprivacy.eu/vocabs/data#Location",
  "purpose": "http://www.specialprivacy.eu/vocabs/purposes#Telemarketing",
  "processing": "http://www.specialprivacy.eu/vocabs/processing#Collect",
  "storage": "http://www.specialprivacy.eu/vocabs/locations#OurServers",
  "recipient": "http://www.specialprivacy.eu/vocabs/recipients#Ours",
}
```

The links in the above policy elaborate on the details. For the Proximus use case, we used this practical implementation as well as the data vocabulary agreed upon in the W3C data privacy vocabulary and controls community group⁶.

An example is given here:

```
{
  "_id": "5dcc73f02f57b27eaa2b363a",
  "user": "5dcbb37136f162ceb853c580",
  "process": "5d7a1f57756fb609aa30685d",
  "data": "http://www.w3.org/ns/dpv#dpv:TVViewingBehavior",
  "purpose": "http://www.w3.org/ns/dpv#dpv:CreateEventRecommendations",
  "processing": "http://www.w3.org/ns/dpv#dpv:Collect",
  "storage": "ProcessorServers",
  "recipient": "Processor",
  "timestamp": 1572901683000,
  "instanceData": {
    "movie_title": "Failure to Launch",
    "movie_genre": "Drama|Romance",
    "watched_on": 1572901683000,
    "duration": 2390,
    "language": "English",
  }
}
```

Remaining areas for investigation are:

- Negation: The currently defined policy language does not allow negation nor exclusion. For instance, “you can track my location everywhere except on weekends at certain locations” is currently not possible.

⁶ <https://www.w3.org/community/dpvcg/>

- Linked data used in policy definition: What if the content at the link location changes? Does this require a new consent from the data subject? Perhaps policy versioning is a solution for this.

3.6.2 Sticky Policies

The sticky policies concept was in the mission statement of SPECIAL⁷ but research into it has not really materialized and there is only one proposal as to how to encode it. It is not a requirement for the Proximus use case but would nevertheless be seen within Proximus as an added value for other use cases.

One proposal for encoding sticky policies (source CeRict) in business policies and consent policies consists of attaching the sticky policy to the recipient.

For example:

Transfer Policy with Sticky Policy:

```
{
  hasData: XXX
  hasProcessing: Transfer
  hasPurpose: YYY
  hasStorage: ZZZ
  hasRecipient: {
    RecipientA
    underStickyPolicy:
      {
        hasData: XXX'
        hasProcessing: YYY'
        hasPurpose: ZZZ'
        ...
      }
  }
}
```

It does however not seem to solve the issue of keeping the policy with the data (wherever the data travels). Work is being done currently on the log vocabulary to deal with this.

3.6.3 Privacy Dashboard

The Privacy Dashboard is explained in depth in deliverable D4.5. A specific styled version was created for Proximus (Figures 3.7 and 3.8). This version will however not be maintained beyond the scope of SPECIAL, but a generic version will be available in a public repository with all the functionality. Details are to be found in deliverable D4.5.

The Privacy Dashboard developed by TUB was an important component in the second iteration of the Proximus use case and there was a high level of interaction between Proximus and TUB to establish this version.

⁷ <https://www.specialprivacy.eu/> ... harness them with *sticky* policies....

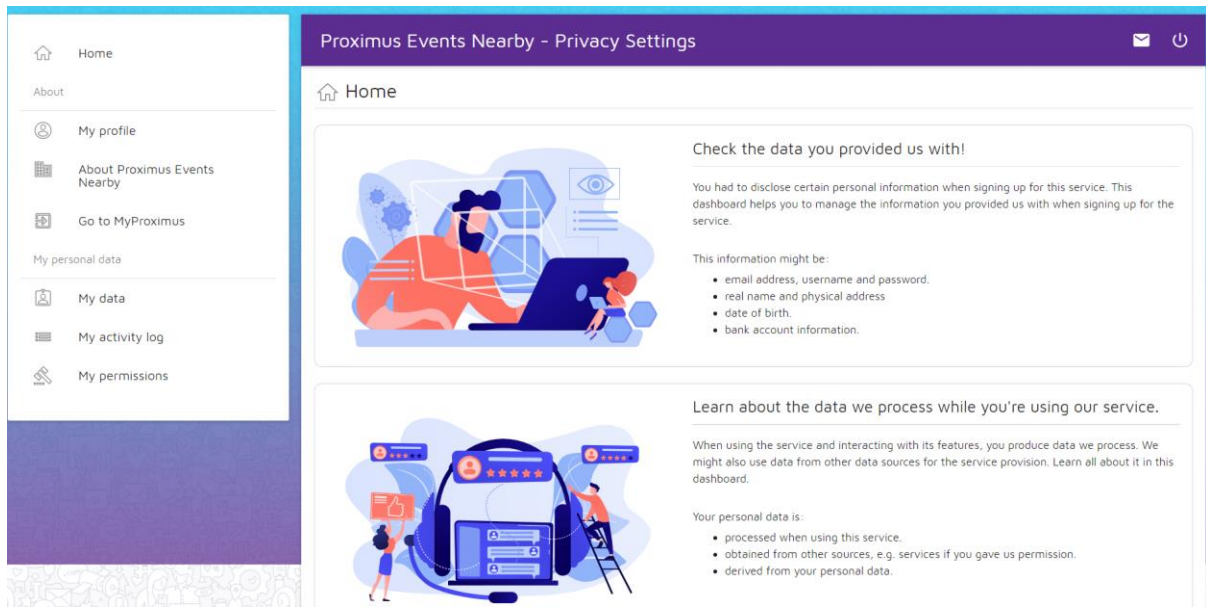


Figure 3.7 Home screen of the Privacy Dashboard

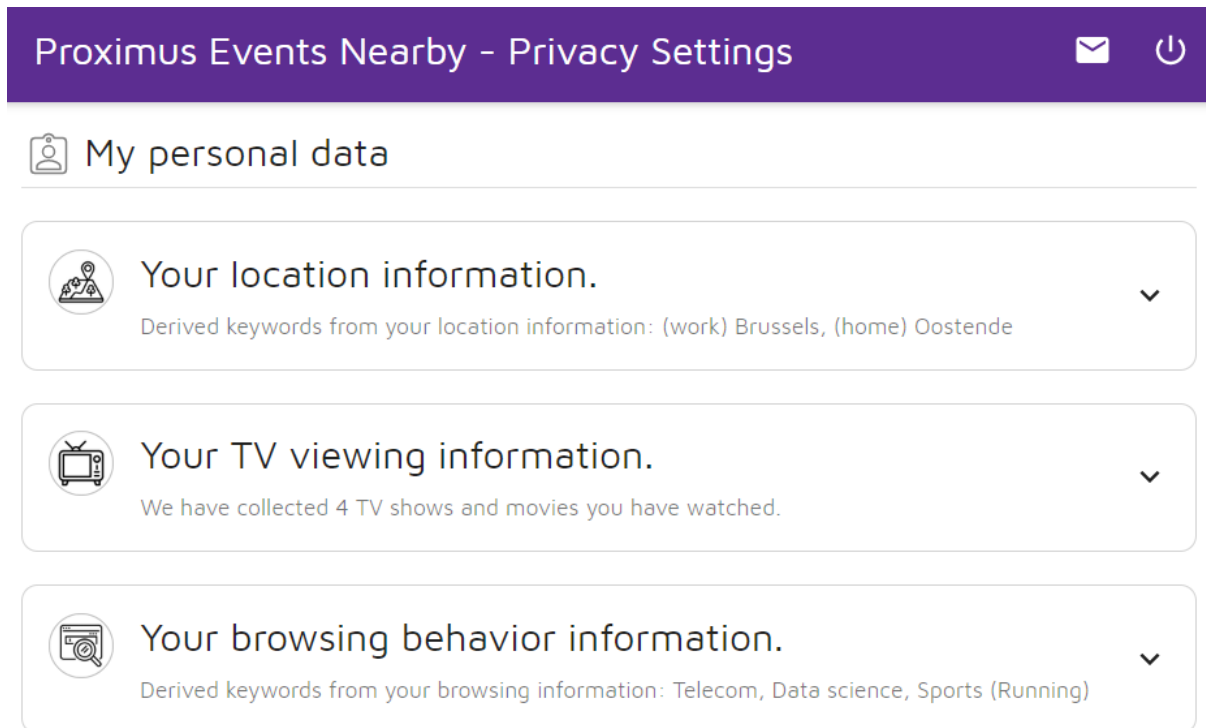


Figure 3.8 Clickable categories of the “My personal data” screen in the Privacy Dashboard.

3.6.4 Compliance checker

Proximus experimented with the ex-post compliance checker as described in D3.4 and available at [github](https://github.com/specialprivacy/compliance-checker)⁸. It demonstrates the capabilities of such a compliance checker. In the current Proximus use case, which has been built from scratch, compliance is inherent to the design and ex-ante checking is

⁸ <https://github.com/specialprivacy/compliance-checker>

done by default. Proximus did not experiment with the ex-ante compliance checker as mentioned in D3.4 as the ‘personal-data-gateway’ component was not developed at the time of doing the first experiments.

3.6.5 Event log

Proximus has not specifically deployed the SLog mechanism as described in D2.7 by WU. However, it used the content of D2.7 as inspiration to implement a more simple derived version.

For example:

```
{
  "success": true,
  "message": "Successfully loaded log entries",
  "result": [{
    "_id": "5c8913ec3c827659934b59e9",
    "user": "5c8791202be791080a464af9",
    "process": "5c6e96a849199380c2388497",
    "data": "http://www.w3.org/ns/dpv#dpv:GPSCoordinate",
    "purpose": "http://www.w3.org/ns/dpv#dpv>CreateEventRecommendations",
    "processing":
      "http://www.specialprivacy.eu/vocabs/processing#Collect",
    "storage": "http://www.specialprivacy.eu/vocabs/locations#OurServers",
    "recipient": "http://www.specialprivacy.eu/vocabs/recipients#Ours",
    "timestamp": 1552487404287,
    "instanceData": {
      "lat": 52.5129013,
      "lon": 13.3201085,
      "streetName": "TEL, Bismarckstrasse, Charlottenburg-Wilmersdorf, Berlin",
      "decision": "Agree",
      "timestamp": 1552487403050
    }
  ]
}
```

3.6.6 Dynamic Consent

Dynamic consent as proposed by ERCIM and ULD (see deliverable D4.3 for a detailed definition) is still in “research modus” and was not implemented nor experimented with in the Proximus use case. The backend architecture and frontend wireframes are still being defined.

3.6.7 Conclusion on SPECIAL components

To answer whether the effort of using SPECIAL technology pays off in the Proximus pilot and potentially in other use cases, we can separate them in three groups:

Essential: The SPECIAL components that are underlying in the Proximus use case:

- (1) The Policy language & W3C data privacy vocabulary community group
- (2) The Privacy Dashboard

In terms of the cost effectiveness for a production ready version of any application using (1) and (2), there will be higher investment for the first application, but for any following application, the existing

technology can be reused quite easily. And in fact, it should be reused to guarantee the same user experience across all applications. In this case, the Privacy Dashboard would become company centric (only one) and not application centric (one per application).

Optional: The SPECIAL components that could have made a difference if available:

- (3) Sticky Policies

To be further investigated: The SPECIAL components that require further investigation to be interesting for other enterprise wide solutions:

- (4) Compliance checker
- (5) Event log

3.7 Results of the user tests

We chose for a qualitative instead of a quantitative approach for performing the user tests. Each test was a cognitive walkthrough session with the data subject, following a predefined questionnaire and scenario. The focus of the data subjects' tests is to find feedback on the general privacy acceptance and not so much on the cosmetics of the privacy dashboard. It is therefore not a usability test, but rather an "understanding" test with the main question for the data subject being "Do you understand what happens with your data?".

3.7.1 Gender & Age

There were 12 data subjects that participated in the test. Figure 3.9 shows the gender and age distribution. No particular bias was noticed in any age, meaning that privacy is a concern for all ages.




Figure 3.9 Gender & Age distribution.

3.7.2 Expertise & daily time on the internet

We noticed that the data subjects found it a bit hard to evaluate how much time they spend on the internet daily. Many use internet for work reasons, and it might have been better to ask for "personal time" on the internet. Nobody qualified herself/himself as a beginner. (see Figure 3.10)




Figure 3.10 Expertise & daily internet time.

3.7.3 Preferred internet browsing device

It was interesting to see (Figure 3.11) that not a single data subject has a desktop as a preferred device. The trend is towards devices that can be taken along, and even more than one. This is good news, as the SPECIAL Privacy Dashboard is best viewed on a larger screen and a laptop/tablet would be more suitable than a mobile phone. For the tests, the data subjects all used their laptop.




Figure 3.11. Preferred internet browsing devices.

3.7.4 GDPR and your privacy

Almost all data subjects (10/12) try to understand the reason(s) why their personal data is needed. Only one data subject tries to read a complete Privacy Policy, while another data subject generally accepts sharing personal data without reading any further.

Similar, ten out of twelve data subjects state that “GDPR is a good thing for his/her privacy”. One data subject disagreed, while another data subject had no opinion.

3.7.5 Agreement with popular beliefs surrounding privacy




Figure 3.12. Agreement with popular beliefs surrounding privacy.

We also noticed (Figure 3.12), a strong correlation between the perception of US based companies (Google, Facebook, Apple) and that internet privacy is an illusion. This can even be extended to China. In data privacy, it confirms two distinct perceptions:

- 1) Asian-American: Data has economic value which needs to be captured and invested in, which creates new value.
- 2) European: Citizen privacy prevails over the economic value, and as a side effect also limits somewhat possible scientific research.

3.7.6 Findings after testing the application

For the test, one hour was reserved with each data subject, including the filling out of the pre- and post-questionnaire. Most data subjects (8/12) found the test to be the right amount of time. Some data subjects (3/12) thought it was a bit too long.

3.7.6.1 Events Nearby

The test application “Events Nearby” was perceived as “somewhat interesting” by five data subjects, while six data subjects found it a “fair” application, meaning that they would use it when needed. One data subject suggested that he would first look in Facebook.

The data subjects were informed upfront that the “Events Nearby” was a teaser application and that shortcuts were taken in the development and usability. Therefore, we did not focus on collecting improvement feedback particularly for “Events Nearby”.

All data subjects remembered after the test for which personal data categories they had given consent. Only one data subject did not remember any longer for which purpose consent was given.

This means that the “consent settings” within the “Events Nearby” application were clear and non-ambiguous. Seven data subjects mentioned that the consent request met their expectation for a privacy policy representation.

While one data subject preferred a traditional privacy policy, the others all preferred this consent request and they mentioned that:

“... A traditional privacy policy is always a long text that you don't read. The consent request is short and pushes you to read and to know what will happen with your data. It uses icons and simple and plain language...”

3.7.6.2 Privacy Dashboard

The data subjects were presented with a complete range of 60 adjectives to describe how they perceive the Privacy Dashboard. The top ten adjective count is given in Table 3.13:

Table 3.13. Top 10 adjective count of perception on the Privacy Dashboard.

Adjective	# times mentioned by data subjects
Easy-to-use	9
Appealing	5
Clear	5
Organized	5
Too-Technical	5
Time-consuming	4
Boring	3
Confusing	3
Friendly	3
Helpful	3

The full list of mentioned adjectives is shown in the following wordle (Fig. 3.14):



Figure 3.14 Wordle showing the adjectives used when evaluating the Privacy Dashboard.

It is impossible to show a clear percentage on such a small population in the test. It does show however that there seem to be two groups:

- A group who understands the need for the Privacy Dashboard and finds it easy-to-use.
- A group who is overwhelmed by complex and cryptical definitions and classifies the Privacy Dashboard as too technical.

During this test, a software iteration was done on both “Events Nearby” and the “Privacy Dashboard”. The results did not change significantly. Many more iterations are needed before the Privacy Dashboard will be easier to understand for a bigger group of data subjects.

A common point raised by both groups however, was the uncomfortable surprise when seeing the level of detail that was reached with the raw personal data. Specific locations, movies seen, and URL’s visited were perceived as ‘big brother’ and ‘intrusive’ to one’s personal sphere.

For some data subjects and categories, we used synthetic data. This was rather confusing for the data subject as he/she needed one more level of abstraction.

The data subjects gave valuable feedback and the following list with improvements was compiled for future iterations of the Privacy Dashboard.

Improvements:

- Difference in understanding between “Consent Settings” in “Events Nearby” (easy to understand) and “My permissions” in the “Privacy Dashboard” (difficult to understand) while both should be the same. Maybe a link to the “Events Nearby” consent settings would be better.
- A pause button for location data - not to be tracked for a while.
- Links from the heatmap to the relevant log entries.
- A message that pops up when you click erasure is confusing. Why not simply delete it immediately?
- Unchecked keywords of interest profile in “Events Nearby” are still in the Privacy Dashboard. They could be greyed out.
- “load more” on movies while there are only two.
- Controller and processor should be translated in the permission statement.
- TV films is confronting; is this detail still needed if we derived the keywords?
- Controller/processor wording is too technical.
- Why do we need the processor Microsoft?
- Map should show how the data points result in the derived keywords. It should show only the outliers as data points; the rest should be grouped as we can derive data from it.
- Map: It could be interactive: Could this be your home location? Could this be your work location?
- My permissions: “controller” and “processor” should be substituted by Proximus and Microsoft.
- Usability is very far from a real-world application reflecting legal/academic vs commercial.
- Activity log: very technical; maybe show only the last 24 hours.
- Clarify if revoking consent also means deleting all the data to which that consent applied.

3.8 Conclusion

The Proximus “Events Nearby” application was used as a teaser application to understand the privacy concerns of the data subjects. Shortcuts on functionality and usability were made to keep the focus only on the privacy aspect and less on the commercial readiness or recommender accuracy of the “Events Nearby” application.

While in the “Events Nearby” application it was found very easy to give consent for processing personal data, it proved very difficult to erase data in the “Privacy Dashboard”.

GDPR is for many internet users a new language that needs to be learned. Instead of offering a clear new innovative vocabulary, and despite the effort of the W3C community to standardize, it throws in words like “controller” and “processor” which are not easily explained.

The SPECIAL project was able to transform the legal GDPR language into a machine-readable language (the policy language) understandable by computer algorithms in the backend.

However, on the front-end, it proves difficult and it would require more development iterations before the machine-readable policies are also “human-readable”, if at all possible. There is optimism however as almost all users found the front-end more intuitive than a traditional page-long privacy policy.

Finally, we would like to mention some possibilities for further research:

- (1) A Policy Language should be able to model statements closer to a human language to capture user-defined consent like “track me everywhere except here in the evening”. This involves including negation which may affect severely the performance of any compliance checking algorithm.
- (2) Sticky policies are a unique way to pass data from one party to another respecting the consent given for the data, at least in theory. In practice, it seems that any algorithm that has access to the data after understanding the policy can simply repackage the data without the policy.
- (3) The Privacy dashboard usability is on the right track but is missing a few iterations or crowd sourced efforts to become ready for mainstream acceptance. Also, a question that was not investigated is whether the Privacy Dashboard should be application centric or controller centric, and what would be the criteria to choose one or the other.

4 Pilot 2 (Refinitiv)

4.1 Summary

The Refinitiv pilot seeks to both satisfy our Know-Your-Customer (KYC) use-case and then to generalise our learnings to support arbitrary use cases in the financial data supply chain. We've focussed our efforts in four key areas:

- Policy language expressivity and the possibility of situating the policy language within a broader business information model that governs the management and distribution of data within the financial data supply chain.
- Modelling business processes in the context of the General Data Protection Regulation (GDPR) and the user experience (UX) to support non-technical domain experts
- The efficacy of automated compliance checking of processing both against the GDPR (static checks) and against the consent statements of data subjects (dynamic checks)

We focussed our work on transfers of personal data under the GDPR. It's an issue of particular interest to us as our customers are global. But narrowing the scope also allowed us to go deeply into the intricacies of a challenging area within the regulation. Could we automate even this?

4.1 Pilot objectives

The pilot is designed to answer the following questions:

Policy Language, Business Information Models, and Standardisation

- Does the policy language and associated vocabularies, as currently specified, capture all the relevant information required to demonstrate compliance with the GDPR? If not, what use cases can we describe to add to the requirements for the next iteration of the standards incubated by the SPECIAL project.
- Can we integrate the policy language into a larger business information model that describes the rights and obligations over data in a content supply chain? (This model covers not only privacy considerations but also licensing obligations and other regulations relevant to a data supply chain).
- Can we capture the policy language within this business information model using concepts and relationships familiar to the finance industry?
- Can we import the vocabularies developed by the W3C Data Privacy Vocabularies and Controls CG (DPVCG) into the model?
- Can we then check GDPR compliance using this business information model?
- How much demand is there within the financial industry to standardise a business information model that allows automated compliance checking?

Business Process

- Can we provide a UX such that staff currently managing our Know-Your-Customer workflows can confidently specify the processes underpinning them – or will we have to rely on specialists with a technical understanding of the policy language?
- How should we surface the results of compliance testing to privacy professionals and regulators?
- How to we generalise this UX to cover a wider set of use cases?

- Do our processing policies provide a true reflection of our right to process personal information under the GDPR? If not, how should they be amended to be so?
- Are our processing policies a complete and correct reflection of the processing we actually do on personally sensitive data in our Know-Your-Customer workflows? If not, do we need to amend the policy language to provide such a complete and correct description?

Compliance Checking

- Can compliance checking be fully automated – without any recourse to manual intervention? If not, what kinds of manual intervention are needed, and when?
- Given the volume of compliance checks over consent that our live systems are likely to generate, can the SPECIAL components make decisions fast enough: we're looking for millisecond response times.
- Do smart contracts on the blockchain provide a suitable implementation environment for compliance checking?
- What is the scope of the completeness and correctness guarantees that we can offer?
- If we are to virtualise, externalise, and automate compliance decisions, how should we provide/present the resulting service?

4.2 Data protection considerations

The data used in this pilot will all be synthesised data designed to support the objectives described above. No 'genuine' personally sensitive data will be used.

4.3 Pilot components

The pilot tests the following SPECIAL components:

1. The policy language and vocabularies: are they expressive enough?
2. The compliance algorithm: does it provide the answers we expect?
3. The compliance engine provided by TenForce:
 - Does it correctly implement the compliance algorithm;
 - Are there cases that require manual intervention;
 - Can we get fast enough decisions (i.e. in milliseconds)
4. The ODRL Regulatory Compliance Profile

The pilot also develops the following components:

1. A UX proof-of-concept developed internally to explore the creation and maintenance of processing policies.

The target audience for this UX are the staff that currently design and execute our Know-Your-Customer workflows. There are many of these workflows for different clients, different countries, and different types of financial transaction.

The Proof of Concept (PoC) seeks to transform this informal knowledge into the formal processing policies specified by the Special project. If successful, we will be able to expand the scope of our coverage without recourse to technical specialists with knowledge of the underlying representation.

2. Implementations of the compliance algorithm both as smart contracts on the block chain and in serverless cloud-based environments.
3. The integration of the ODRL Regulatory Compliance Profile proposed by WU and the W3C Data Privacy Vocabularies and Controls CG (DPVCG) ontologies (both of which are still under active development) into a wider business information model that governs the distribution and use of data in the financial data supply chain.

4.4 Evaluation criteria

We are using the legal expertise both within the Special project itself and the Refinitiv Privacy Office to 'sense check' first the processing policies themselves and then the compliance decisions made by the compliance engine. Are we getting the results we expected?

The Know-Your-Customer team working with the UX-experts are evaluating the viability of using the existing Know-Your-Customer staff to create and maintain processing policies.

Software architects attached to the product are considering the response times. Given their sizing estimates, they must decide whether to pre-compute decisions or whether to run them in real-time.

We're also working with the Proposition and Product Development functions within Refinitiv to clarify the commercial potential for tools and services based on the Special research. These discussions are themselves influenced by the interest shown by our largest customers in the wider opportunity to automate compliance both to data licenses and to regulations like the GDPR.

4.5 Policy language, information models, and standardisation evaluation

We identified 19 permissions within the GDPR that allow for the transfer of personal data. Five of those permissions are based on adequacy decisions; six based on appropriate safeguards; and eight derogations.

Using the formalism provided by the policy language, we found that every permission could be expressed using the structure provided.

Then we looked at the permissions that the KYC team felt they needed to cover their workflow and processes. At the level of granularity at which they felt comfortable, there are six. Again, we could express them using the structure provided by the policy language.

The work on the policy language also specified the compliance relationship between the six processing permissions specified by our KYC team and the 19 defined by the GDPR: subsumptive containment. Using this relationship, it was easy to check which versions of the processing policies were compliant and which weren't. And for those that weren't, easy to see how to bring them into compliance.

The terms - or values - we used in the permissions fell into two buckets: those whose meaning is anchored in the GDPR and those whose meaning is domain specific. The former we could borrow

wholesale from the vocabularies developed by the W3C Data Privacy Vocabularies and Controls CG (DPVCG), e.g. the values for the legal basis on which processing is undertaken.

The values for purpose, however, seemed specific for KYC operations. This made sense, as their only role in compliance testing is to match - or be subsumed by - the purposes specified in the data subject's consent statements, which will also be domain specific - the same domain!

Some terms fell between the buckets - especially *Actions*. The KYC team has its vocabulary; the GDPR another. Further, the KYC team's terminology was often at a far more granular level than that used by the GDPR. They work with passports and lease agreements; the GDPR talks of personal data and sensitive data.

Given the policy languages guidance in using classes, however, it was easy to translate - and de-couple - the language used by the KYC team and that required for compliance testing. When they spoke of lease agreements, we could simply sub-class personal data, and the compliance tests run as normal.

The conceptual foundations of the policy language - permissions and obligations - is a natural fit with the way the financial data industry is beginning to think about compliance to licensing conditions. However, those relationships are captured at a higher level than those proposed by the policy language using a rights expression language called ODRL. Could we integrate the lower-level policy language into the more widely understood - and standardised - ODRL language?

The way forward was provided by the Open Digital Rights Language (ODRL) Regulatory Compliance Profile which described the translation we were looking for. Furthermore, its use of the "predicate constraint" inspired us to make some further simplifications of the ODRL language which we will likely recommend by included in future versions of the ODRL standard.

So now the architectures, services, and tooling under development in the financial data industry to automate rights management can also potentially be used to enforce compliance to the GDPR.

4.6 Business Process

Given that we can express our workflows as a set of permissions that we can automatically test against the GDPR, and that we can use the same architecture we use to test compliance to licenses, how can we integrate this into our quotidian business operations?

First, there is the question of vocabulary. Did our KYC team need to learn a completely new one? As discussed above - no, we can automatically translate between the terms they use and the terms of the GDPR.

How about creating the processing permissions and policies required by the KYC operations? The team can't write them by hand, but could we provide a UX that guided them through the process without presuming any knowledge of the underlying policy language?

We wrapped the processing permissions into processing steps; and processing steps into business operations. These are commonly understood concepts.

To get the permission-specific information we needed to run compliance tests, we first tried a traditional sequence of contextual controls (drop-downs, radio and edit buttons) as shown in Figure 4.1.

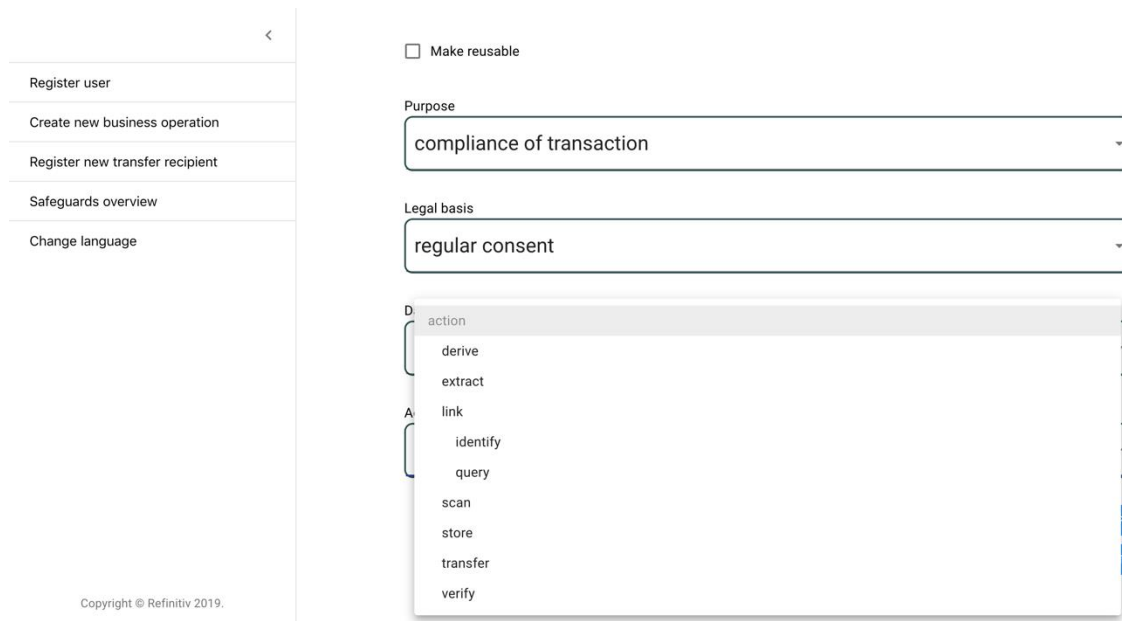


Figure 4.1 - Selecting action in traditional UX

Following instruction, it was easy for the team to generate the final permission specification, as shown in Figure 4.2.

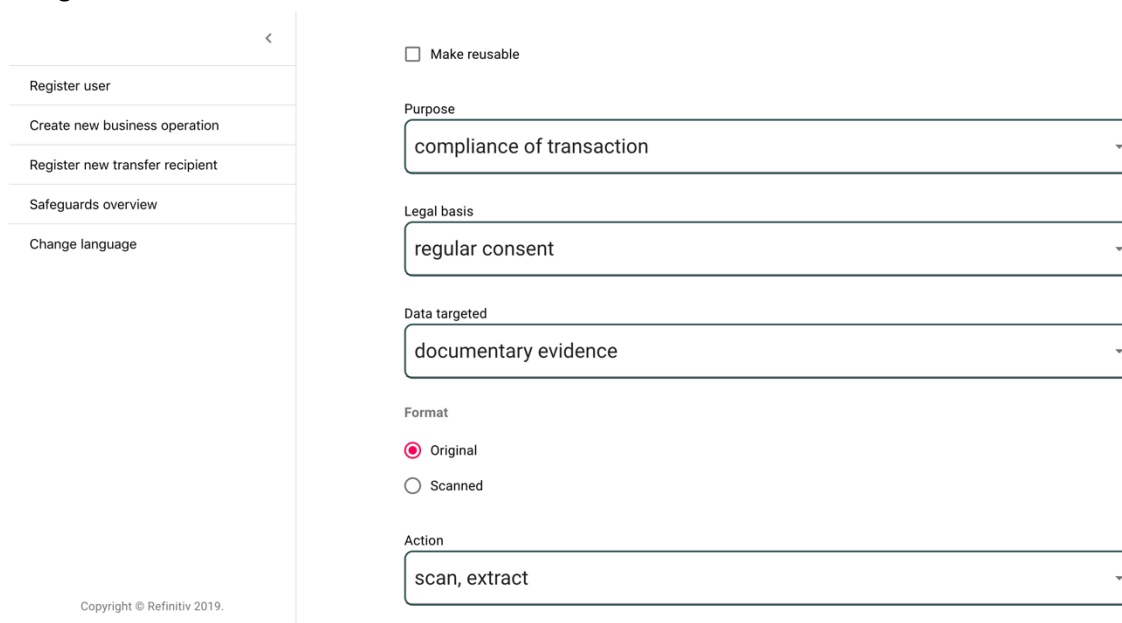


Figure 4.2 - A permission display using traditional controls

But quite what this meant, what had been achieved, was not clear. The representation is still quite abstract. A far more successful approach was to create the permission through expressing directly what you want to do - as a sentence or statement. The menus exist now within a context. Their meaning is clearer (see Figure 4.3).

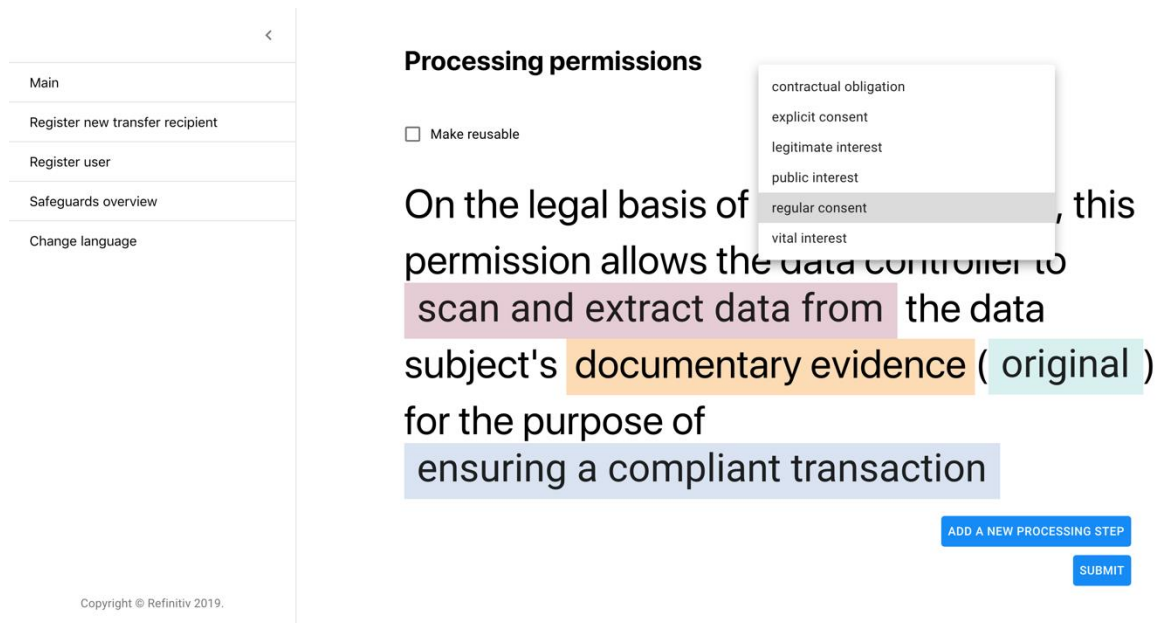


Figure 4.3 - Selecting a lawful basis with the context of a sentence

What's been achieved is also now much clearer (Figure 4.4):

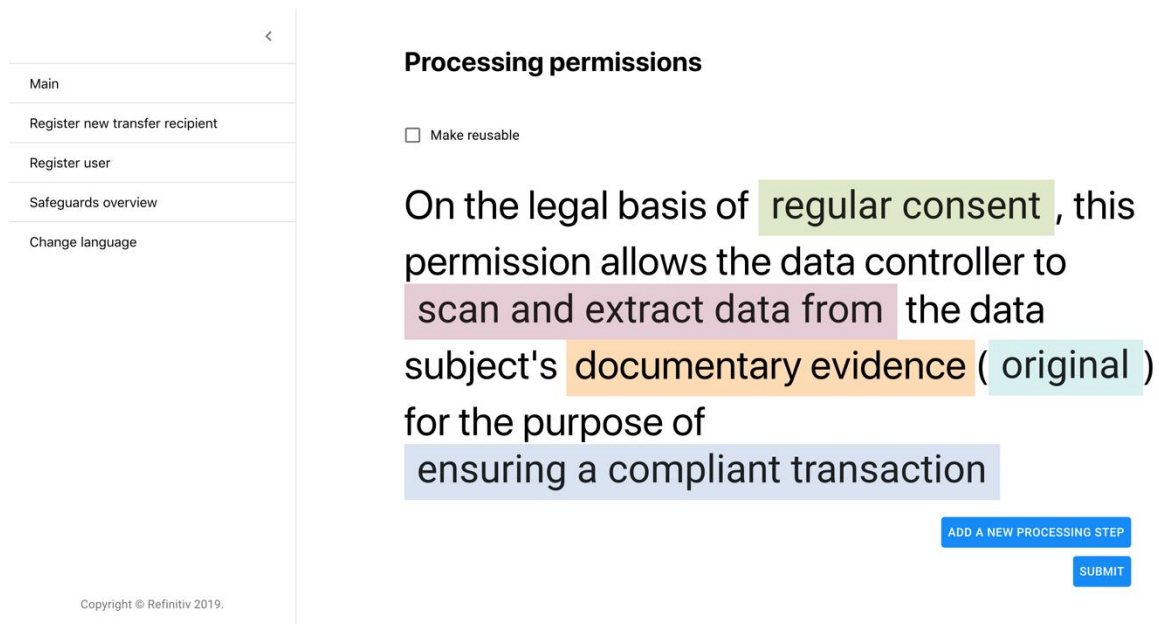


Figure 4.4 - A permission expressed as a sentence

This "sentential" approach was a breakthrough in designing an intuitive UX. People now readily understood what they were doing and why.

Of course, transfer permissions are more complex than this. The permissions controlling the transfer of personal data outside the EEA are probably the most complex permissions in the GDPR. Frequently they depend on having appropriate safeguards in place. To run compliance tests, we need some descriptive metadata about them. But that is way beyond the remit of the KYC team.

But this is within the remit of the Privacy Office. Suddenly our UX is supporting two audiences: the data operations teams that run the KYC process and privacy professionals. One valuable consequence is that the Privacy Office now has far greater access (and understanding) of the granular processes that the operations teams run.

This transparency can reach further. Exercising several of the derogation permissions that allow transfers outside the European Economic Area (EEA) require approval from a supervisory authority. Could they be a third audience, with access to real-time reporting on the processing of personal data and the permissions that control it? That is a question that requires further research.

4.7 Compliance Checking

There are two contexts in which we need to check compliance:

1. Compliance between the processing permission and the GDPR
2. In the case of consent, compliance between the data subject's consent and the processing permission

We looked at four techniques:

1. Using inferencing (along the lines suggested by the policy language specification)
2. Using closed-world rules (SHACL)
3. Using a SPARQL query
4. Using smart contracts in the context of the blockchain.

All worked. The first option, as exemplified by TenForce's compliance engine, has the performance profile that makes it the preferred option to check consent compliance in real-time.

The third option has the simplicity and flexibility to make it our preferred option to check compliance between processing permissions and the GDPR during the PoC. Whether it would retain that simplicity in production is an open question. It may be easier to maintain a discrete set of SHACL rules than a monolithic SPARQL query - or else use a generic inferencing engine.

The fourth option, smart contracts on the blockchain, worked. But designing the compliance algorithm from first principles was painstaking and the performance profile was heavily dependent on the complexity of the policies checked. Without further substantive reasons to migrate to the blockchain, this option was discarded.

Is the process fully automated? Well, yes, and no. The compliance tests pre-suppose some human judgements. These are clearly defined, are captured in the UX, and must be made by the Privacy Office. For example, if a transfer permission relies on an appropriate safeguard, is the safeguard legally binding? Is it legally enforceable? Does it provide the data subject with their rights? Does it provide them with remedies?

These are judgements we cannot automate (yet?). In our system they are captured as attestations. Without them, compliance will fail.

5 Pilot 3 (Deutsche Telekom)

5.1 Remark

This section on Deutsche Telekom's pilot is part of a "living document" in its' final (yet still vivid) version. Certain sub-chapters stayed unchanged or underwent only minor modifications and thus might be skipped by readers who are well aware of the context, plan and implementation as described in D5.1 and D5.3. However, the remarkable progress made in the pilot is covered in this chapter and deserves reader's attention. Therefore, it is recommended to focus on the following sub-chapters covering the most visible and relevant changes compared to D 5.3.

- 5.5 (updated)
- 5.6 (almost completely new)
- 5.8 (updated)
- 5.9 (new). This is the most relevant new section since it covers the evaluation of the DT pilot

The use case's and pilot's time schedule is stable since D5.3 and the present chapter 5 serves as a good overview on and summary of the last period of the project, focussing on evaluation and testing of the designed and implemented version of the pilot.

5.2 Use case description (short recap)

5.2.1 Additional use for device (client) based QoS/location data

Prior to SPECIAL (and prior to the use case pilot implementation), a DT app (Customer Network Experience/"CNE-app", by Deutsche Telekom) collected data that was (and is) used only for network quality improvement. Some components of the collected data could also be used for other purposes and by other units, e.g. by Motionlogic, a legally independent subsidiary of DT, to improve and verify algorithms in location based services. This leads to an interesting business opportunity that was impossible to implement under EU law until SPECIAL came up with an innovative solution.

5.2.2 Previous (pre-SPECIAL) situation

Users could download and install the CNE-app from the app store (Apple) or Play Store (Google). The purpose was very clear and the user interface offered the key functionality on screen. Of course, the privacy statement and terms and conditions were presented at first launch. Then, the user could initiate manual "speed tests" or change certain settings, e.g. turn on a "diagnosis mode". Figure 5.1 shows the original app.

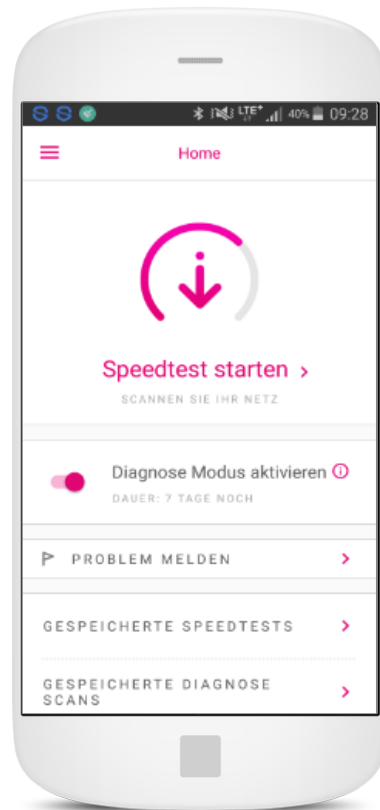


Figure 5.1 : CNE-app (customer network experience) with a slider to activate “diagnostic mode”; we intended to switch on “non-anonymous mode” similarly

In the original (pre-SEPCIAL) version of the CNE app, valuable PII (personally identifiable information, such as precise location) was collected (with user consent) and used for the original purpose (network quality measurement) by DT. Other possible uses (and users) of the data were aware of the assets but could not get access to it, e.g. Motionlogic. The following two paragraphs explain the situation and role of the two players here in more detail:

1. **Telekom Deutschland GmbH (DT)** collects Quality of Service data (QoS) for its’ mobile network service. One source of this QoS data is a smartphone app called “CNE – Customer Network Experience” (see **Error! Reference source not found.**). Each data set collected here includes (among others) geo location information measured via device GPS. Radio data (reception quality etc.) are collected, aggregated, condensed and sent to a database. Data sets are by default anonymized and/or pseudonymized to protect user’s privacy. Users give their informed consent during installation of the app. Data sets are then collected and statistically evaluated by the service quality department to improve the network quality. The department responsible for running and further developing the app and database is looking for additional use of the collected data, e.g. gain insights by applying AI (artificial intelligence) techniques to the data.
2. **Motionlogic** (<https://www.motionlogic.de/blog/de/>), an independent Spin-off company of DT, uses anonymized and time-delayed data of mobile phone usage (location data, cell-tower location) to offer Business-to-Business (B2B) location services, e.g. heat maps of population density in urban areas or traffic infrastructure. Motionlogic never exports individual data or even data sets received from DT (T-Mobile brand). Rather Motionlogic does the requested processing internally and only delivers the results (e.g. heat maps). Due to limited quality of data, the added value is limited as well. Better data, e.g. individual user tracks or even more accurate location data, would improve the results dramatically. However, due to lack of

explicit “Opt-In” by end users, Motionlogic acts way below its theoretical capabilities and capacities. Motionlogic needs more “Opt-In” users among DT customers.

5.2.3 Target scenario for pilot implementation

A new version of the CNE-app (including SPECIAL’s components and concepts) collects more frequently data, in particular more location data, and reduces (or abandons) anonymization/pseudonymization. The now much more valuable data is shared with Motionlogic (but not further) to enable better location based evaluation for business partners and monetization. To do so, much deeper consent needs to be given by users in the form of an explicit (and informed) “Opt-In”. We assumed that users choose to “Opt-In” if they keep control over their data and perceive a certain customer benefit. Policies as developed/supported by SPECIAL and the privacy dashboard are the tools to guarantee user control and transparency.

In the current implementation, the user can give consent to “additional scientific use” of his data (in German “Wissenschaftliche Zusatznutzung”) as shown in Figure 5.2.

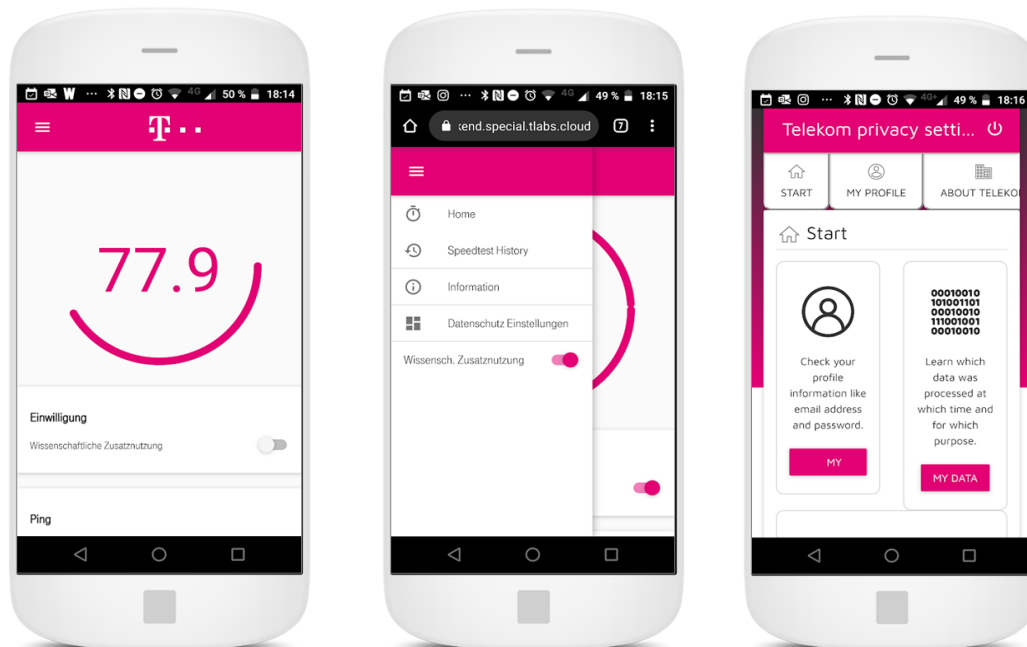


Figure 5.2 : SPECIAL version of CNE-app (customer network experience) with a sliders to activate “Scientific use” and a screenshot of the current “Telekom privacy dashboard”

Please refer to deliverable D 1.5 and D1.6 for more detailed descriptions of the use case.

As a further development of the DT use cases (especially as described D1.6), we decided to simplify the policy and minimize user’s options to edit policies in favour of general applicability in DT’s IT- and business infrastructure, and expect a larger number of users/data subjects to test scalability (in later iterations). In other words: dynamic consent is an optional part of the current initial implementation. The policy used here is more a stepwise (and explicit) approval. DT is currently assessing D1.6 and the concept of dynamic consent therein. In future versions DT might use the “dynamic consent” approach to simplify user interaction and reduce the individual consent-granting actions by users.

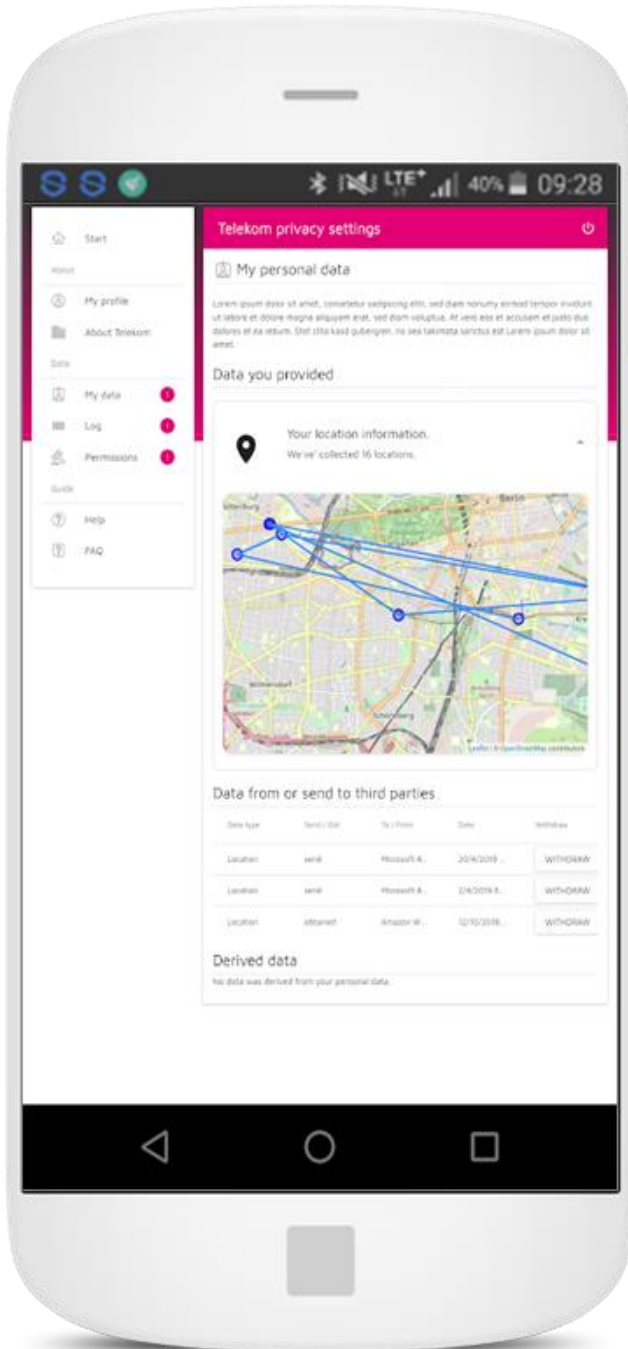


Figure 5.3 : SPECIAL privacy dashboard, launched from the CNE-app, showing the “My data” section

Since transparency is a key requirement and a core feature of SPECIAL, the use case also covers this approach by giving the end user access to his specific “privacy dashboard” directly from the app: The SPECIAL CNE app contains a link to the SPECIAL sever running the (location) data base, policy engine and privacy dashboard (transparency tool), see Figure 5.3. The user can simply follow this link on the smartphone and monitor which data has been collected and was transferred or processed to/by a third party (i.e. Motionlogic). The privacy dashboard also allows the user to remove his data from the database and/or withdraw the consent give before.

5.3 Objective of pilot implementation (in short)

5.3.1 DT's view on Big Data and AI and their relevance to SPECIAL

DT's large amounts of user data are currently only used for the purpose of (telecommunications) service delivery and optimization and are not shared with any 3rd party to avoid data breaches or other threats. DT always puts the customer and customer protection first and never risks any legal or reputation threats. Therefore, SPECIAL and the observations/experiences made in the project are of particular value for DT.

DT acknowledges the potential additional value (for the own business AND for customer) of this data and the results from (potentially AI based) data analyses.

In SPECIAL, DT wants to examine opportunities and limitations of privacy tools compliant with GDPR (and other legal regulations) in the context of a scalable, yet already "Big Data" use case: The CNE app today (prior to SPECIAL supplements) has about 20.000 users, usually about 5% of them (1000 users) are online concurrently.

Since the app was launched (in 2017) these users produced 20.000 to 70.000 datasets per day, about 1 M datasets per month. This is the base figure for the "original" CNE app. DT expected somewhat smaller numbers for non-anonymizing version of the app. However, neither the "original" nor the "SPECIAL supplemented" version of the app would be pushed in app store marketing. So user numbers are assumed to be limited on purpose. The resulting numbers (of users, datasets etc) are clearly "Big Data" but still relatively small compared to the amounts of data telecommunication operators as DT usually deal with. Therefore, a successful internal evaluation (of the "SPECIAL CNE app) is expected to open doors towards the monetization of real Big Data treasures that lie hidden in the carrier's operational data centers.

5.3.2 Objectives for the DT pilot

DT's view on the pilot implementation is that it wants to see and test a "Proof of Concept" (PoC), not a product (or even a Minimal Viable Product (MVP)). This means that DT is interested in knowledge, experience and decision support for future developments. The implementation is not intended to become part of DT's regular operational IT infrastructure.

DT's pilot is based on prior work by SPECIAL consortium members. In particular, D1.6 was used as basis for intense discussions since it contains most interesting (yet not fully assessed) concepts. D4.1 (privacy dashboard) and D4.2 (Usability testing) helped to come up with implementation decisions and formed the objectives defined for this pilot. Also WP3 (in particular D3.3, Backend scalability) helped to focus on a simple use case with relatively many users/data subjects.

DT's main objectives for the pilot implementation are:

- **Proof of feasibility** of technology developed in SPECIAL (privacy policies, policy engine, privacy dashboard/transparency tools etc.) in conjunction with DT's existing IT infrastructure and internal processes. This feasibility includes technical, legal and organizational aspects. DT as an agile and modern telecommunications service provider needs to know how and how far cutting edge data processing tools, especially privacy related ones, will affect business operations: Does DT need to change internal processes (e.g. the PSA-process, see below) to allow such technologies to be implemented?
- **Technical benefit for DT/Motionlogic (ML)** The pilot helps to optimize DT's and Motionlogic's analysis tools. The tools use anonymized user data to deduct certain insights and allow

predictions on the behavior of unknown, anonymous people based on previous experiences. The optimization and validation relies on a certain amount of confirmed (and thus non-anonymized) data sets. These data sets will be comprised of the data provided with full user consent using SPECIAL's policy mechanisms.

- **Additional user benefit:** Part of the analyses by DT will be the identification of network and coverage issues. DT currently considers to share (parts of) this knowledge of the current "health" of the network with effected customers. Thus one potential additional benefit of the pilot would be to (possibly) share relevant network quality information with users. To determine for whom the information is "relevant" (i.e. who is located close to the issue scene), only personal data can be used, in our case data, from users who gave the requested consent.
- As a **"hygiene factor"** (a feature/factor that must be present anyway, even if it does not contribute to the business value of the product) we assume (and require) that the pilot fulfils DT's corporate privacy & security standards. Especially since this is a privacy related project, DT values the protection of user data not just inside its data centers, but also protects it against external threats. A respective assessment and statement/confirmation by DT's corporate privacy officer was requested (and carried out/received in the meantime).

5.3.3 What is NOT in scope

DT does not plan to monetize user data. Neither today nor in the foreseeable future will DT sell their users' personal data. The pilot will not result in any direct marketing activities. DT does not even plan or intend to build personalized services for the end users based on the data collected in the pilot.

5.4 Pilot implementation architecture

The original CNE app and use case are described in more detail in D5.3 (chapter 5.4). The present document (D5.5) focuses on the “SPECIAL CNE app” as it was implemented and evaluated together with SPECIAL partners.

In order to improve quality of Motionlogic’s algorithms, individual data sets with concrete (individual) users are necessary. SPECIAL delivers the necessary tools and mechanisms to collect informed user consent, give the user full transparency, allow him to remove his data (and withdraw consent) if needed. Since the pilot implementation should not put any running DT process at risk, the system architecture uses a “bypass” mechanism for the privacy data (policies, log data and user interaction). This gives the user full control and does not touch mission critical processes within the company.

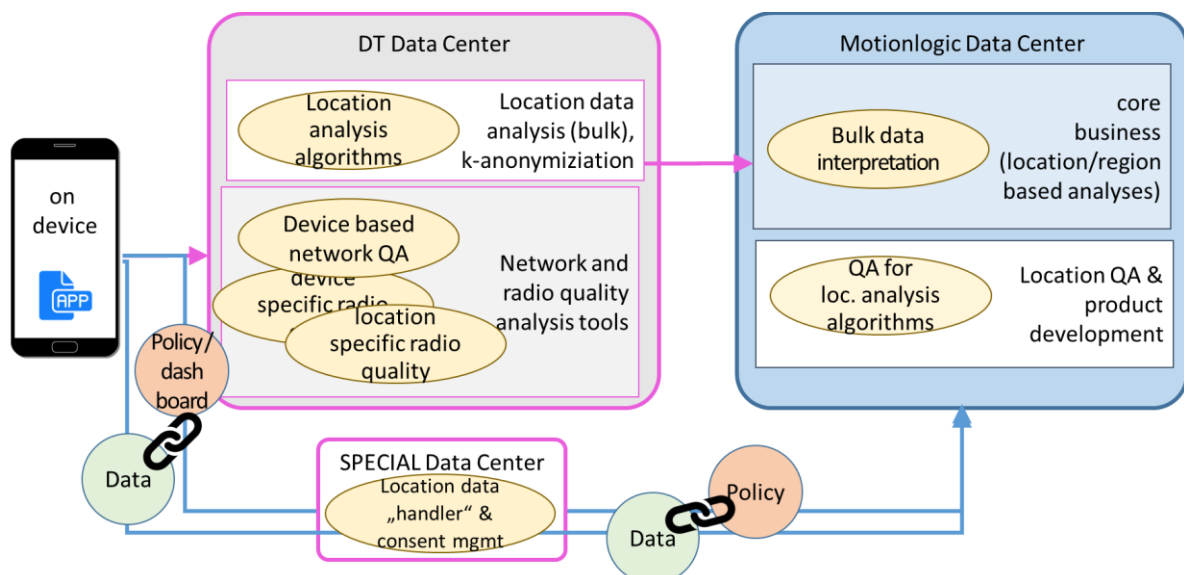


Figure 5.4 : architecture: SPECIAL tools and methods are added to the original process. Non-anonymous data will be used legally, with transparency and user control is provided

Figure 5.4 shows the pilot’s approach to demonstrate the opportunities provided by SPECIAL: A dedicated “SPECIAL Data Center” handles both, the (non-anonymous) transfer and storage of “payload”/location data, and consent/privacy/transparency management:

If the user gives his consent (in the app/on device), location data (“data/payload/content”) are transferred to the SPECIAL data center (green circle).

At the same time the consent (and the attached conditions, i.e. the “policy”) is transferred/communicated to the (same) SPECIAL data center. Thus data and policy are linked.

Motion logic can now receive the data according to the policy from the SPECIAL data center.

Transparency and control (“privacy dashboard”) for the user are conveyed via the same link between app/device and SPECIAL data center.

Even though the SPECIAL components are currently completely outside DT’s core data centers, this architecture will prove the validity of the approach and does not contradict the idea of integrating the functionality and concepts (linked data and privacy policies) directly in the core IT processes of DT.

5.5 Data flow

A dedicated analysis of the flow(s) of data was carried out to estimate the measures needed to comply to DT-internal rules, GDPR and other legal regulations. Since the original CNE app already was approved by security and DPO authorities, we could assume that this would be a stable starting point for the analysis. The pre-SPECIAL situation is described in great detail in D5.3.

Figure 5.5 shows the newly designed setting:

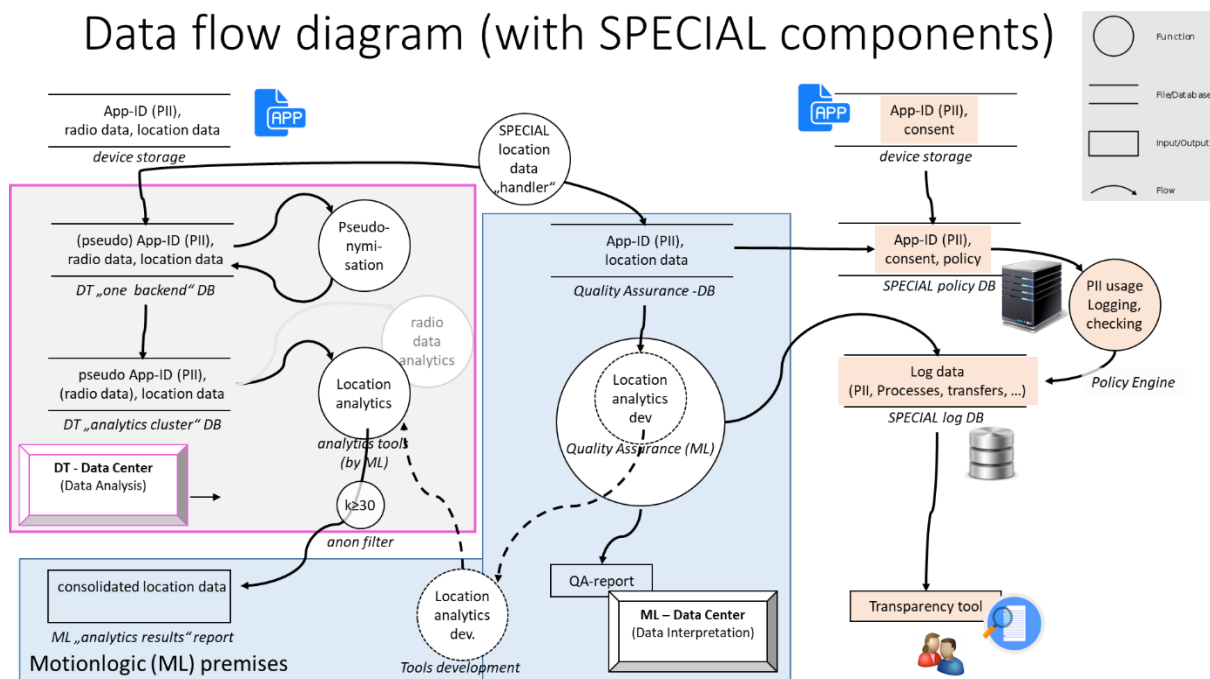

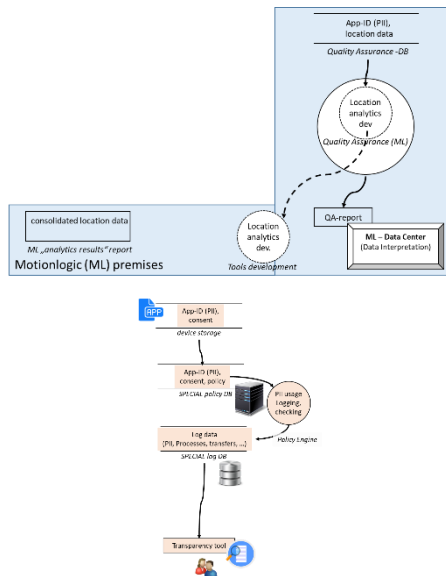


Figure 5.5 : Data flow as designed with SPECIAL components. Certain PII is transferred to Motionlogic with user consent.

For clarification, colours have been added to the drawing:



The magenta outlined grey box contains DT’s data center where analytics tools take place. These include location analytics tasks using Motionlogic’s algorithms.



The blue box depicts Motionlogic’s data center and development facilities. Utilizing certain non-anonymized location data (received from the “SPECIAL location data handler”) software development and QA can now legally improve the location analytics tools that work inside DT’s premise.

SPECIAL components (with light orange background) take care of all the privacy and GDPR constraints and also the user interaction. This is crucial since neither DT’s data center nor Motionlogic have any direct contact to end customers (the data subjects in SPECIAL terms).

The data flow designed for the new version of the overall system (CNE app, DT data center, ML data center, and privacy/transparency/control tools) ensures minimal changes to the legacy systems and provides optimal access and exploitation of the desired location data while at the same time revealing the feasibility and value of the SPECIAL approach.

The components can easily be mapped to the “SPECIAL-K reference architecture” as described in <https://www.specialprivacy.eu/platform/special-k-reference-architecture> and depicted in Figure 5.6:

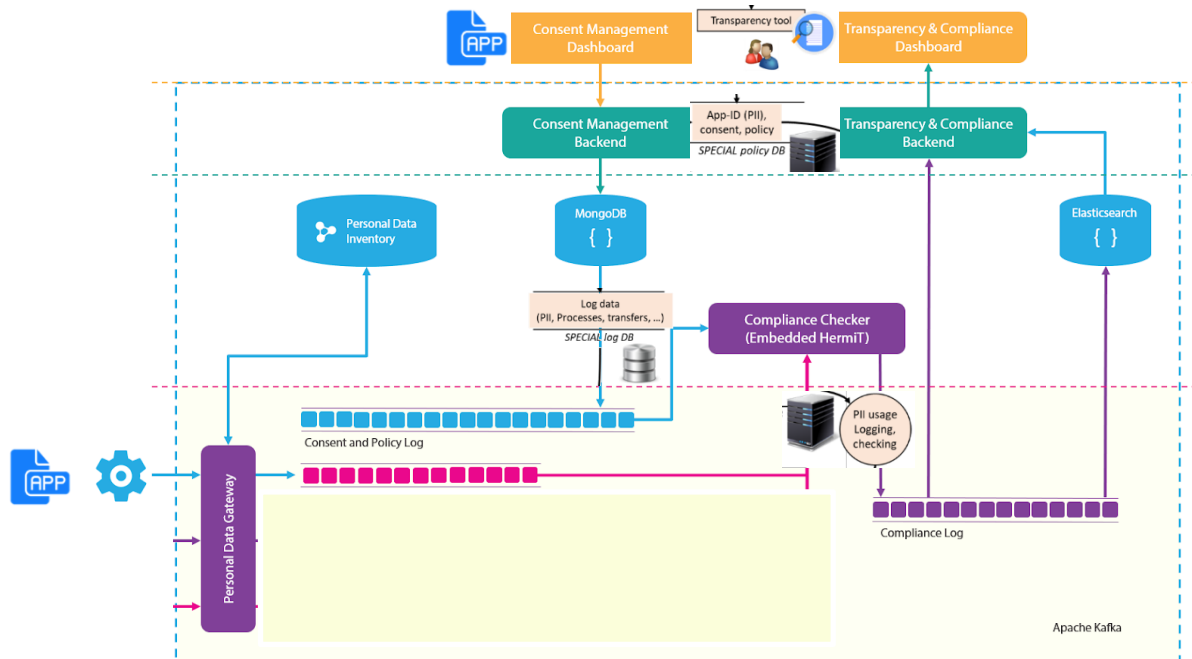


Figure 5.6 : SPECIAL-K architecture setup for ex post compliance checking (reference architecture) with overlay showing how the DT pilot mapped the reference architecture to the actual implementation

5.6 Data flow specification in GDPR terms for legal assessment

This chapter briefly explains the data flow from a more formal and legal point of view, clarifying the type of data, processes, controllers and purposes. This is done with the GDPR terminology in mind and applying it to the use case utilizing the SPECIAL policy language and vocabulary.

5.6.1 Types of user consent/policy

During Installation (and later during operation) of the CNE app, the user gives (or withdraws) certain consents for the use of their personal data. Three cases can be distinguished:

- ① **If user does not agree to any privacy terms, the app does not collect any data (and does not function at all).**

No data is collected/stored/processed/shared etc

- ② **If the user agrees to the “default” data privacy terms (necessary for the app to perform its’ core task), the following processing (etc.) takes place:**

The CNE app for measuring perceived network quality on end devices (smartphones) collects precise location data (based on GPS sensor data). The position data is transferred to DT’s network operating data center to monitor network quality together with a unique “AppID” and the timestamp and the “payload”: network data (effective bandwidth, signal strength etc).

- ③ **If the user agrees to “additional scientific use”, the following processing (etc.) takes place in addition to ②:**

No additional data is collected. However, the (precise) location data, corresponding time stamps and app-id are (conceptually) shared with a third party (“Motionlogic”) for the scientific purpose of software quality assurance. Motionlogic anyway generates assumptions about certain locations of anonymized users based on (coarse) radio-tower (cell) information; using the location data (from CNE app), Motionlogic can now verify these assumptions (or enhance the software generating these assumptions). The data is not shared any further than Motionlogic. After the QA, data is deleted, usually after 7 days (max.).

Table 5-1 Data privacy terms and instantiations as formulated above

	CNE “default” conditions ①	CNE with “scientific use” ②
Data (type of data)	AppID (1) User/device location(2), timestamp of probe (3), set of network data (from probe) (4)	AppID (1) User/device location (2), timestamp of probe (3), Declaration(Class(svd:Location)
Duration (of usage)	7 days	7 days Declaration(Class(svdu:StatedP urpose))
Locations (of processing)	EU (mostly Germany, DT data centers) (6)	EU (mostly Germany, DT/SPECIAL data centers); using ML algorithms in EU. (6)

		Declaration(Class(sv1:ControllerServers)) Declaration(Class(sv1:EU))
Processing (type of processing)	Collect, Network quality analysis (based on location info)	Software quality analysis (research) and privacy acceptance survey (7) Declaration(Class(svpr:Analyse)) Declaration(Class(svpr:Anonymize))
Purpose (of processing)	Enhance mobile network quality	Optimize analysis software that does NOT have access to personal data. (8) Contribute findings to scientific community; (8) Declaration(Class(svpu:AuxPurpose)) Declaration(Class(svpu:Feedback)) Declaration(Class(svpu:Tailoring))
Recipients (of data)	Only DT internal	DT internal & “ours”=SPECIAL DataCenter; & derived data (QA-results) will be sent to Motionlogic (9) Declaration(Class(svr:OtherRecipient)) Declaration(Class(svr:Ours))

In addition to the “timestamp” of the probe, a timestamp of each relevant processing step is kept and added to log file.

5.6.2 Background and role of Motionlogic:

Motionlogic is a spinoff of Deutsche Telekom’s R&D department. It is a 100% DT-owned company, but at the same time a 100% independent legal entity (“GmbH” / “Ltd”). In the pilot described here, Motionlogic can be regarded as a “**data controller**” since the purpose of processing (quality assurance) is defined/determined by Motionlogic. Since the algorithms applied to the (personal) data are also developed by Motionlogic and process both personal and anonymized data, Motionlogic might also be considered a “**data processor**”, however in the pilot, the “SPECIAL data center” is the most obvious actor doing “processing” tasks.

5.6.3 Detailed explanations on some of the terms/tasks mentioned above:

- (1) The CNE-App uses and “App-ID” to identify itself against the CNE-Server. This App-ID is a unique identifier, generated during App-Installation. The App-ID identifies a specific instance of the App, i.e. it is unique and connects each data-set with the device (smartphone) that generated the data-set. Since the Smartphone can be assigned to a specific person (at least when using additional information), the App-ID is considered “PII” (personally identifiable information).
- (2) User/device location: location is the most exact position available on the phone. Usually it is the (advanced) GPS position. Accuracy of position may as good as +/- 10m. It the primary purpose of the CNE app to measure network quality on the end device (smartphone) together with the exact position to white (or “grey”) spots of mobile phone connectivity. Therefore, knowing and sharing the position is necessary for the CNE-App. At the same time, this exact position is needed for the “additional use” by/for Motionlogic.
- (3) The CNE-app takes probes based on three triggers/events:
 - a. Every hour a scan of network features is carried out
 - b. Whenever the device’s (and thus user’s) mobile network-context switches,
 - c. When the user triggers the network-scan manuallythe resulting probe-data is transferred (together with the location and the current time) to DT’s backend data center. If the user consented to the “additional use”, the AppID, position and timestamp data are also transferred to the “SPECIAL” server (ignoring the network probe).
- (4) The data collected on the device mainly is network-quality data (cell-id, ping-time, frequency-band, wifi-availability, upload-time/bandwidth, download-time/bandwidth, context-info (device type, SIM-type etc) and identification-info (MSISDN = User-ID, IMSI = device ID) as well as the already mentioned App-ID (unique per device and thus user) and position. Of this data set, only position, timestamp and AppID are of interest for the “additional use”. Thus only these three fields of the data set are transferred to the SPECIAL server (for analysis by Motionlogic).
- (5) A data point that is transferred for “additional use” to the SPECIAL data center has the following fields:
 - a. AppID (integer) e.g. 4294967295
 - b. Longitude (WGS84 coordinate) e.g. 33UUU E: 91776.15689 (= 13.404954 E)
 - c. Latitude (WGS84 coordinate) e.g. 33UUU N: 20073.007 (= 52.520007 N)
 - d. Timestamp (dd/mm/yyyy hh24:mm:ss) e.g. 25/05/2018 01:23:45
- (6) All processing happens in EU. In the “standard” CNE-case, all processing happens in Deutsche Telekom data centers. No third parties involved here. The additional processing for “software quality analysis” (see (8)) is carried out on SPECIAL servers. These are servers run and maintained by DT in Berlin, Germany, but not inside DT’s internal corporate network. Rather, these servers run in a separate, protected network with strict access control.
- (7) Processing:
 - a. measuring and mapping network quality based on location and device features (core CNE-functionality) not covered here.
 - b. Software quality analysis: Motionlogic usually works with very coarse data (location accuracy > 500m, no individual positions (only position of groups of 30 or more) and no “tracks”, i.e. sequences of positions or “paths” a user followed). Nevertheless, Motionlogic is able to produce very valuable location based statements/predictions (e.g.: How many people leave the subway station X at 8 a.m. and proceed to the

shopping center Y). It is hard to prove these statements and even harder to enhance prediction quality. The exact position data together with AppID and timestamps allow a verification and thus optimization/enhancement of Motionlogic's analysis tools (that do NOT use individual position data). This software quality analysis requires user consent or "OptIn" to be accomplished.

- c. The second research aspect is SPECIALs approach itself: How do users react to the consent-request and the transparency offer? To answer these questions, the app needs to be designed and the usage needs to be observed and evaluated scientifically. No explicit large scale survey is intended. Rather a "brd-eye view" on selected pilot users. Thus the data processing is not part of the research but rather the subject of observation. The result could be called a "privacy acceptance survey".

(8) The purpose of processing is mentioned above:

- a. Motionlogic's standard software, described in (7)b., (that works without PII) shall be improved by comparing its' analysis results/predictions with "real world data", i.e. data that originated from identifiable users/devices. This is a part of "software quality analysis" in Software Engineering.
- b. The other purpose ("privacy acceptance survey") is fulfilled by the work of the SPECIAL project, in particular the dissemination work.

(9) Recipients: Personal data is collected on the device and transferred to DT datacenters (for standard CNE purposes; not be discussed here). With user consent ("OptIn") personal data is also transferred to our (SPECIAL) pilot server in the DT domain. Our "SPECIAL" servers host the personal data, the policy-engine and the transparency tool to monitor the use of the data. For the processing described in (7), motion logic transfers its analysis tools (software/algorithms) to the SPECIAL data center. Here, inside SPECIAL datacenters, the quality analysis takes place: Motionlogics software compares its' standard results with the data from real world users. The result of this comparison is then transferred back to Motionlogic as "derived data". It is not intended or possible to re-construct personal data from such "derived data". Results of the comparison would rather be of the kind: "prediction and real world match" or "prediction and real world differ by 25m east" or "prediction and real world diverge completely".

5.7 Use of SPECIAL components

5.7.1 Policy Language

Since the use case of DT's pilot is rather "straight forward" with a limited number (and sequence) of possible consents, the SPECIAL policy language is not fully exploited but used in a pragmatic way. We experimented with formulating the current consent statements in the form of "policies" and succeeded easily.

In addition, DT is considering to use/define more complex policies for other use cases. To be able to judge the applicability of SPECIAL's policy language, complex and partially "synthetic" requirements have been contributed to WP2. As a result, the policy language was found to be a full match of even the most advanced future requirements.

In the current pilot, DT only scratches the surface of the policy language's potential. Even with the (present) limited use, DT is able to cover a very large portion of "real-world" requirements. Plans for follow-up activities, using the policy language as well as the vocabulary and (certain) tools are maturing

and are intended to be implemented after the END of project SPECIAL building on top of SPECIAL's results.

5.7.2 Sticky Policies

“Sticky policies” are a charming feature and an attractive technical concept. On this (conceptual) level, the pilot uses sticky policies in the SPECIAL server. However, there is no end-to-end implementation of the concept available at this time. DT's use of the concept is meant to be a step towards the subsequent implementation in the productization phases.

5.7.3 Privacy Dashboard

Since transparency and control for the end user are of utmost importance for DT in this pilot (and beyond), special attention was given to the implementation of the privacy dashboard. This tool grants not only the legally required data transparency, but also is a convincing proof of DT's reputation of being not only responsible and protective with user's data but also innovation friendly and “leading edge” in terms of technology. DT puts the user/customer in the front and takes care of the users security and wellbeing while using DT's products.

The privacy dashboard is used to underline this reputation and as an argument to gain user consent. In fact, the dashboard was the main argument to get DT involved in the project. This statement also holds for follow-up activities (as mentioned above).

5.7.4 Compliance Checker

DT evaluated the various versions of the ex-post compliance checker (see D3.6 for details on this tool). Since the used policies in the current DT pilot are very simple and the type of processing/transmitting data is at the same time simple and conceptually happening outside the SPECIAL server scope, no explicit ex-ante compliance check was implemented. However, the method of collecting and storing relevant information (in the event log) is a perfect precondition for formal compliance checks.

5.7.5 Event log

The current version of the pilot uses SPECIAL event logs to conceptually check compliance (ex-post) and to feed the privacy dashboard. The event log plays a central role for all three main purposes of the privacy (sub)system of the CNE-app and potential implementations in future products:

- The event log provides the data used in the transparency tool (dashboard)
- The compliance checker uses the event log for checking process steps (events) versus policies
- Audits (by corporate or external DPOs) use the event log as a central source of information

In the current form, the event log fulfils its purpose (or purposes, see above). However, DT internal discussions and evaluations suggested that a reworked version/implementation using either a simple data base or block chain technology might be more appropriate. For the current pilot the used implementation was the optimal choice.

5.7.6 Dynamic consent

Dynamic consent (<https://epub.wu.ac.at/6494/>) as a concept would fit nicely to the use case, in particular if you consider ongoing and extended use of PII by DT or at least under DT's management. DT considers dynamic consent an extremely valuable concept and a key reason for DT to continue evaluation and work on SPECIAL results even after the official termination of the project (in 2020 and later).

More concrete: The concept of “Dynamic Consent” will be studied in a new context of a follow-up-activity planned around a European “data market place”. Here, dynamic consent might be the tool of choice to keep user experience good and at the same time grant full user control over personal data used in different contexts, by different controllers and for different purposes.

5.8 Testing

In addition to the usual software tests (for meeting software quality requirements) and security and privacy tests, T-Labs ran qualitative tests based on user questionnaires and expert interviews:

All those tests were executed by experienced test experts, collaborating with experts from all relevant DT-business units (CNE-app, Motionlogic, T-Labs) and SPECIAL experts.

The tests were carried out during a dedicated test period. The new (SPECIAL supplemented) version of the CNE app was made available for download and installation by designated test users.

During the first use of the new CNE app version, the users were asked to renew the acknowledgement of the (updated) privacy statement (by changing a switch in the interface).

In addition, the users were asked if they were willing to participate in the non-anonymized procedure for further quality enhancement. The users then gave consent via checkbox/slider and were provided a link to the privacy dashboard (transparency tool). The whole setting was clearly marked as a test procedure.

The test run was planned for one month with an option to extend the period. Of course, the test could be aborted anytime by stopping the data collection on the device or on the server. Users were informed immediately about the stopping of data collection via the transparency tool.

During the test period (October and November 2019) the mentioned expert interviews and survey took place. Results are described in the present deliverable (D5.5) below.

5.9 Results and Evaluation

5.9.1 Results and Observations

According to the plan described in D5.3, DT experts and pilot users have been interviewed informally. The interviews relied on four predefined questions. A representative summary of typical answers can be found below. The results can be summarized as answers to the following questions:

1. **Does it work?** (Does SPECIAL enable DT and business partners collecting, storing and processing personal data?)

Yes.

All three aspects of the pilot implementation fulfil their purpose:

- a) The CNE-App still delivers the originally intended network-quality data, and users are still using the app in about the same intensity as before (i.e. No harm to existing business).
- b) The collected location data enriches the data pool of Motionlogic. Being able to compare “synthetic” assumptions with real world observations helped to optimize the algorithms that generate the “synthetic” results
- c) The privacy related components (dashboard, policy language and compliance checker) worked as expected

2. **Does it pay off?** (Is the effort in a reasonable ratio to the revenue?)

Yes (mostly). Since the desired result was very much limited, the purpose (software quality assurance) would be achievable and has perceivable results. In the current context of a research prototype (internal pilot) the effort was relatively high, in particular on the project side. If the business partner would need to pay for (at least) the operational work, the effort would probably considered “not be worth” the results. Developing the concepts and components of SPECIAL for just one use case (i.e. CNE-App data for Motionlogic) would just be too expensive. However, the tests made “appetite for more”: if working with “real (non-anonymized) data” is possible, than this should be the “default” processing option, and the effort would for sure pay off.

3. **Are real world users satisfied?** (user satisfaction regarding usability and functionality of SPECIAL tools)

Yes, but users are not completely happy with the UI. They still prefer “simpler” or even “invisible” user interfaces for a smooth user experience. This may be a hint towards the next steps: hide the privacy settings even more and optimize ease of use for transparency. Possibly an “implicit dynamic consent” would be a user oriented solution. Now the question with respect to legality remains. Transparency tool was perceived as a great idea and opened users’ eyes in more than the current context – users assumed that other apps/services have at least the same information “somewhere on the servers” but do not allow users to check or delete this data.

4. **Is it legal?** (This is not the least question, but rather the most critical one since it is directly relate to DT’s reputation. We assume that the tools are “legally acceptable”, but we expected more: the tools were meant to be accepted by end-users as legal and justifiable, not only (but also) by lawyers.

Yes, absolutely. It is not completely clear if the users expected a DT tool to be legal and just trusted DT due to its reputation and own expectation, or the users judged according to their own (layman)

knowledge of legal context. However, the pilot was perceived “even more legal than most current apps and websites using my data” because of the transparency given.

5.9.2 Evaluation

Evaluation along the criteria listed in D5.3 resulted in the following findings:

5.9.2.1 *Hygiene factors (criteria that need to be fulfilled in all products/prototypes anyway)*

Security:

Since the app was built on top of an existing one and did not impose many critical attack vectors, security was ensured at all times. In addition the “closed world” nature of the pilot helped to prevent any security breaches.

Using advanced authentication methods between the app, the SPECIAL data center and other potential communication or business partners made sure that no PII leaked.

Data Privacy:

In our context, this is more than a hygiene factor. The concept, architecture, data flow and all processing steps have been discussed and agreed with several experienced experts and lawyers inside the project and from Deutsche Telekom. Data Privacy was guaranteed at all times according to all applicable laws and regulations as well as according to company rules.

5.9.2.2 *Value of Insights gained*

Usefulness of gained information:

Payload data was useful. However the key outcome for Motionlogic was: “If we can get SOME non-anonymized data from DT data centers, why not ALL the data (user consent assumed)”.

The more valuable insights regarding consent, in particular the user interaction to get and control consent: Users don’t want to be bothered. User interface must be simple, choices must be minimal, “policies” need to be explained in simple terms and transparency needs to be given in short and easy-to-grasp format.

Plausibility of gained information:

Plausibility is given for both “payload” data/results and privacy-metadata. No real surprises observed.

Corporate fit:

A key finding of the project was that SPECIAL technology, in particular consent management, is a recurring task in many current and upcoming products, services and offerings by Deutsche Telekom and its business partners (customers). A solution for this challenge is needed, and SPECIAL gave at least valuable hints on the way to structure this task.

Another finding was that consent management is not a business by itself (alone) but rather a component of other products. Consent management needs to be “built in” and pretty much invisible for the end user. Thus UI issues can and should be handled by actual product managers, not by privacy officers. One way to gain this would be to re-shape the components as an “SDK” that other products include and use. This plan is currently under discussion in DT.

Basic technical criteria:

Computational load is moderate. The limited computational resource that we assigned to the test runs never got even close to their limits.

Network load depends on actual implementation and “real-time” requirements. In the case of this pilot, network load could be neglected since the core purpose of the app was network measurement. A few more bytes to be transferred did not affect the load.

Storage is an issue: Not only can the sheer number of transactions fill up a database quickly but also administrative issues arose: when and how will the transaction data be deleted? What happens with old policies? There are solutions for these (and similar) issues available, but they all consume storage capacity and admin effort.

Scalability for Big Data applications: The relatively simple use case, applied to a relatively large number of users and datasets is expected to bring up findings regarding scalability. DT's mid term business interests lay in legally monetizing its Big Data assets.

All our test indicated that scalability may be issue. The approach scales linearly, but the technology requires a good stream-processing-performance and (potentially) huge storage. Both aspects can be covered but need attention and might require additional (financial) resources in an industrial context.

5.9.2.3 *User acceptance criteria*

(Relative) **number of "Opt-Ins"** made with the SPECIAL / pilot.

Due to the "DT internal" nature of the implementation, test and evaluation, the overall number of participants was small. Thus it was decided to forgo statistical evaluation. Finally, all participants used the "opt-in" in order to try out the new features and to participate in the scientific re-use of their location data. This could be seen as a "100% Opt-In rate" which is of course not representative for the average user/customer group. Nevertheless, the feedback given by those internal users was very encouraging and gave reason to believe that the opt-in rate would indeed be much higher than without SPECIAL's user-centric concepts and tools. See next paragraph for details.

User response/acceptance. In addition to the (relative) number of users who opted in, we want to collect reasons for opting in (giving consent), or rejecting the requested consent. These statements will be a valuable evaluation result that will be considered in future products (even outside the domain of the pilot). Thus, these user opinions will (probably) form some of the most valuable qualitative insight for future privacy related products/components in DT's user facing service offering

Even though all test users (most of them experienced and somewhat affine to technology) opted in (gave full consent) to the scientific additional use, several of them also shared critical opinions about the opt-in in general and about the overall pilot:

- Some users asked about more details on the "scientific use". They then said the reason behind it is ok, but they would like a more open and clear explanation in the first place.
- All users appreciated that they were not bothered by any additional interaction/attention requested by the app: No "location based advertisement" or any other form of intrusion. The CNE app remained what it was before (a sort of "speed test") for the user.
- The app had three places to give/withdraw consent: during installation, before speed test execution and in the settings menu. These carefully chosen, sparse places that allowed (but not requested) user interaction (turn "scientific use" on or off as a default and on a case-by-case-basis) was generally acknowledged but lead to confusion for some users. The case-by-case setting was rarely used.

All users liked the transparency tool (dashboard) very much: Being able to review the data collection and even trace the share and use of one's data was most interesting for our users (keep in mind that those users were somewhat biased "technology affine" corporate internal users). By "liking the tool" we mean that they liked the concept and the options the tool provides them. The current implementation was for some users still confusing and "too much" information.

One particular aspect of interest was the opportunity to delete data points from the location data base: This feature was very much appreciated but sometimes not fully understood. However, the simple "on-click experience" caused users to ask why such a feature was not present in other tools. At the same time it was not clear when the actual deletion would happen and what happened to the derived information based on the old dataset (with the now deleted data point still included).

Interestingly, some of the (expert) users asked for some sort of “intelligent auto-fill” for the opt-in: Basically they wanted the “intelligent dynamic consent” that learns from their behaviour what decisions the user would make (mainly in the case-by-case opt-in) and then at least propose the assumed user decision or even make it for the user. Of course, users want the ability to over-rule any decision made for them.

Finally, users raised the question how the SPECIAL concept would gain a wide scale adoption around the online industry. Using the SPECIAL concepts (but not all its technical components) could then result in a huge number of user interfaces and a rise of the number of decisions to be made by users and monitoring needs. This would be a challenge and would rise the need for some sort of standardization. Transparency and control are very much appreciated, but many users were afraid of the additional workload they would face if the concept would be adapted by a large number of apps. Some assumed that a similar effect to the “accept cookies request” would be possible: Users just opt-in (or opt-out) immediately without thinking about the consequences.

6 Conclusion

This is the last version of SPECIAL's deliverable on pilot implementation and testing, D5.5. While this version ("V3") now has all the final observations and conclusions gained during the project period, we would like to point out one finding in particular: All pilot partners worked on pilot implementations, not market level products. The different nature of the pilots contributed to the diversity of experiences made and thus indicates the diversity of possible applications and implementations in future products. At this point of time, at the end of the R&D project, all results and observations, results and conclusions can only be based on pilot implementations. Hence, further work is necessary to make the SPECIAL approach a true success not just on an R&D level but also in the industrial real world of data driven business.

Some conclusions can be met as "overall conclusions" by all three pilot partners which are drawn at the end of this chapter. However, it seems to make sense and provide more specific insights to also give partner-specific conclusions:

The Proximus "Events Nearby" application was used as a teaser application to understand the privacy concerns of the data subjects. While in the "Events Nearby" application it was found very easy to give consent for processing personal data, it proved very difficult to erase data in the "Privacy Dashboard". The SPECIAL project was able to transform the legal GDPR language into a machine-readable language (the policy language) understandable by computer algorithms in the backend. However, on the front-end, it proves difficult and it would require more development iterations before the machine-readable policies are also "human-readable", if at all possible.

Refinitiv focussed on the policy language and its' relation to legal texts, namely the GDPR and on the business aspects of this formal description. Business value would mainly result from the automatic compliance checking. This approach was a complete success. In addition our participation in SPECIAL enabled several teams in the company (KYC, compliance and others) to "speak" the same language, i.e. SPECIAL's policy language and thus collaborate more effectively.

The DT "CNE-App" showed that adding a GDPR-compliant workflow utilizing personal data (location) next to an already existing processing without any GDPR relevance (no PII affected) is possible. At the same time the value of an overarching "privacy and transparency" approach is now very clear. The pilot unveiled opportunities due to a GDPR compliant consent management. As with the other pilots, the UI/UX (or rather the user capabilities) form a major obstacle. DT therefore considers follow-up activities on the lower levels of applications/services, i.e. "Software Development Kits" with GDPR compliant consent management built in.