SPECIAL

Scalable Policy-awarE Linked Data arChitecture for prIvacy, trAnsparency and compLiance

Deliverable 5.1

Pilot implementations and testing plans V1

Document version: 1.0

SPECIAL DELIVERABLE

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

Mr Philippe Rohou t: +33 4 97 15 53 06 f: +33 4 92 38 78 22 e: philippe.rohou@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-awarE 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 – M19
Deliverable number:	D5.1
Deliverable title	Pilot implementations and testing plans V1
Contractual Date of Delivery:	31.07.2018
Actual Date of Delivery:	31.07.2018
Editor (s):	Martin Kurze (DTAG)
Author (s):	Benedict Whittamsmith, Rudy Jacob, Martin Kurze
Reviewer (s):	Sabrina Kirrane, Eva Schlehan, Piero Bonatti, Olha Drozd
Participant(s):	TR , Prox, DTAG/TLABS
Work package no.:	WP5
Work package title:	Use Case Implementation & Evaluation
Work package leader:	TR
Distribution:	PU
Version/Revision:	1.0
Draft/Final:	Final
Total number of pages (including cover):	26

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 5

2 Introduction to pilots and testing plans in SPECIAL 6

2.1 Types and purposes of pilots, business/industry context 6

2.2 Purpose and requirements for tests/testing plans 6

3 Pilot 1 (Proximus) 8

3.1 Summary 8

3.2 Objective and evaluation for Proximus 8

3.3 Data protection considerations 9

3.4 Dissemination within Proximus 9

3.5 From the first proof of concept iteration towards integration within Proximus 10

3.6 User interface 11

3.7 Recommender engine 13

4 Pilot 2 (Thomson Reuters) 15

4.1 Summary 15

4.2 Pilot objectives 15

4.3 Data protection considerations 15

4.4 Pilot components 15

4.5 Evaluation criteria 16

5 Pilot 3 (Deutsche Telekom) 17

5.1 Use case description (short recap) 17

5.1.1 Additional use for device (client) based QoS/location data 17

5.1.2 Current (pre-SPECIAL) situation 17

5.1.3 Target scenario for pilot implementation 18

5.2 Objective of pilot implementation 18

5.2.1 DT's view on Big Data and AI and their relevance to SPECIAL 18

5.2.2 Objectives for the first iteration of DT pilot 19

5.2.3 What is NOT in scope 20

5.2.4 Public challenge 20

5.3 Pilot implementation architecture 21

5.4 Testing plans 23

5.5 Expected results & evaluation criteria 23

5.5.1 Expected Results 23

5.5.2 Evaluation Criteria 24

6 Conclusion 26

1 Summary

Three industry partners (Proximus, Thomson Reuters 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 D1.1, D1.2, D1.3 as well as D1.4 and subsequently D2.4 and D3.2 , D3.3)

This document was designed for "staging", i.e. it will be reworked on a regular basis. The present version "Deliverable 5.1 V1" is the first iteration and includes implementation plans, first results of the (already implemented) pilots and the test plans. Later iterations of the document will be found in D5.2 and D5.3. 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 is "V1", neither a profound summary nor ultimate conclusions can be expected. The reader is kindly asked to take advantage of the current (preliminary) state of findings and wait for the subsequent versions of the deliverable (D5.2, D5.3) for deeper insights.

The present document aims at giving the reader a comprehensive view into the current state of implementation of the pilots, testing plans and expected 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 first phase of the product development, the prototype aims at implementing and testing a relatively complex use case with relatively small amount of data subjects.
- **Thomson Reuters** 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** (and its' R&D unit T-Labs) is aiming more at applications that use the huge amount of (personal) data that exists in DT and cannot currently 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 straight-forward (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 expected users is relatively high which gives 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 style 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 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 (Security, Privacy and Usability; see below).

Iteration 1 (this document, deliverable D5.1)
Proximus evaluates mainly the overall architecture and the composition of building blocks, general feasibility needs to be checked, and first experiences with corporate internal processes need to be collected. First tests in conjunction with Proximus' recommender engine will be carried out.
Thomson Reuters has a strong focus on the policy language and checks its' expressivity and compliance in detail.

Deutsche Telekom uses this phase for a complementary approach: internal corporate (IT) processes are checked or revised to implement core aspects of SPECIAL. This leads to an in-depth evaluation of SPECIAL's modularity. Also scalability will be tested relatively early (many users) with a "minimal viable prototype" (put in quotation marks because we do 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 (deliverable 5.2)
Proximus will extend the user group and thus plans to focus on scalability of the approach. First business expectations will be verified and the business exploitation plan refined.
Thomson Reuters in this phase will test the implementation details and challenges resulting of applying SPECIAL technology inside TR's IT systems.
Deutsche Telekom will focus on technical ease implementation and applicability of DT's corporate design and corporate identity rules on the tools. In addition DT will also do a first evaluation of business relevance.

Iteration 3 (deliverable 5.3)
All industry partners will evaluate the applicability of their respective pilots in their organizations. Also the overall concepts of SPECIAL (linked data, privacy 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 "sanitary features" of usual products (security and privacy just 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 apply 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 1.5 (2nd version of the requirements definition). In a nutshell, the use case concentrates on the following concepts:

- Recommender Engine for events at the Belgian coast
- Mobile user interface
- Data subjects of Proximus
- Data requested from data subjects:
  - Location
  - Television viewing
  - Browsing history

The first iteration, an initial end-to-end version, albeit for only 5 data subjects, is ready for beta testing. In 2019, a second and third iteration is planned.

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 at this stage, 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 data subject (This is also being investigated in D4.2)
- Privacy awareness of the data subject (This is also being investigated in D4.2)

Specifically, in iteration 1 (with 5 data subjects), we are looking to technically test whether the solution works end-to-end, and focus highly on the intuitive interface, in order to “attract” enough data subjects for the second iteration (with 50 data subjects).

For the SPECIAL project itself, Proximus is asked to evaluate the usefulness of the SPECIAL components (called building blocks) for its use case.
We will answer the following questions:

- Do the SPECIAL components and policy languages help or create more overhead in the use case?
- Are the SPECIAL components a "conditio sine qua non" or are there easy alternatives?
- Are the SPECIAL components putting a measurable extra load on the system (CPU/diskspace/speed of execution)?

3.3 Data protection considerations

On July 19th, 2018, the current state 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 is recommended to avoid ending up in the more complicated governance of the Proximus project methodology with a release calendar beyond our control.

The project team will have to go back to this Privacy Council to ask for more approvals when going to a higher number of data subjects (50 or 500) or when adding any of these two other data sources.

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.
  - Google Firebase (art 4.1 Application of EU Legislation) as the Data Processor for the Consent & Identity/Access Management.

  (Please note that this will be different when running in production)

3.4 Dissemination within Proximus

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

- Alignment with the Proximus project team working on GDPR compliance for all IT/CRM systems
- Alignment with the Enterprise Architecture team: One SPECIAL project team member is part of the Enterprise Architecture team
- Alignment with the Proximus Privacy Council
- Future marketing on the internal intranet when going to the second (50 data subjects) and third iteration (500 data subjects). A SPECIAL project video would facilitate this task.
- Alignment with the Consumer Business Unit with its innovation initiatives.

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

3.5 From the first proof of concept iteration towards integration within Proximus

An estimated path to integrate the SPECIAL use case into the existing Proximus Ecosystem of applications is shown in Table 1 below.

Table 1 Iterations of the SPECIAL use case at Proximus

Building block	**SPECIAL M20: Iteration 1 (End-to-End Proof of Concept with 5 data subjects)**	**SPECIAL M25: Iteration 2 (limited pilot with 50 data subjects)**	**SPECIAL M31: Iteration 3 (companywide pilot with 500 data subjects)**	**POST SPECIAL: Integration within Proximus**
Identity Access Management	Firebase (manual)	Firebase (automated sign up)	Firebase (automated sign up)	Proximus Enco IAM or MyProximus
Frontend UI	Version 1 - one basic vertical page (React.js) developed by TUB	Version 2 - fully GDPR compliant interface (incl T&Cs) (React.js)	Version 3, including GPS location tracking, user feedback (React.js)	MyProximus libraries with input from Version 3
Privacy Dashboard (TUB)	High interest, but (1) too little PXS capacity to implement (2) It is not yet built for a mobile screen	Planned for implementation	Implemented	Evaluate if useful for MyProximus
Layered Privacy Statement/Notice	Suggested by SPECIAL team but still in ideation phase	Depends on whether "pilot ready"	A/B testing (linear vs dynamic) could be a possibility	Evaluate
Consent Datastore	Firebase Database	Firebase Database	Firebase Database	Proximus IAMS
Event Datastore	MS Azure	MS Azure	MS Azure	MS Azure
Profile Datastore	MS Azure (MySQL / MongoDB)	MS Azure (MySQL / MongoDB)	MS Azure (MySQL / MongoDB)	Proximus 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	Kafka + datastore (MySQL / ElasticSearch)	Kafka + datastore (MySQL / ElasticSearch)	Kafka + datastore (MySQL / ElasticSearch)
Machine learning	Python	Python	Python	Python

Building block	SPECIAL M20: Iteration 1 (End-to-End Proof of Concept with 5 data subjects)	SPECIAL M25: Iteration 2 (limited pilot with 50 data subjects)	SPECIAL M31: Iteration 3 (companywide pilot with 500 data subjects)	POST SPECIAL: Integration within Proximus
SMS / email sending	Proximus enco	Proximus enco	Proximus enco	Proximus enco

3.6 User interface

The user interface was designed and developed by TUB. It is a mobile responsive HTML page, developed in react.js following the Proximus branding style. It is currently hosted on MS Azure and the consent settings of the data subjects will be stored in Firebase database. Fig 1 to 3 show the UI.

Authentication is done via an email/password method stored in Firebase. For the current version, the 5 data subjects (email and initial password) have been added manually via the Firebase admin console into the Firebase database. In the next version, signup will be possible via the user interface.

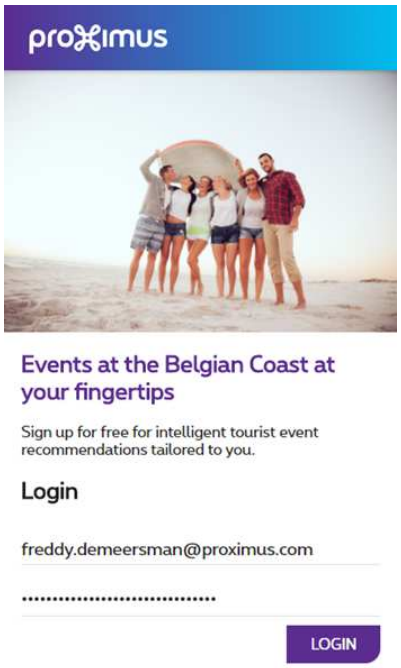

Figure 1 Login Screen

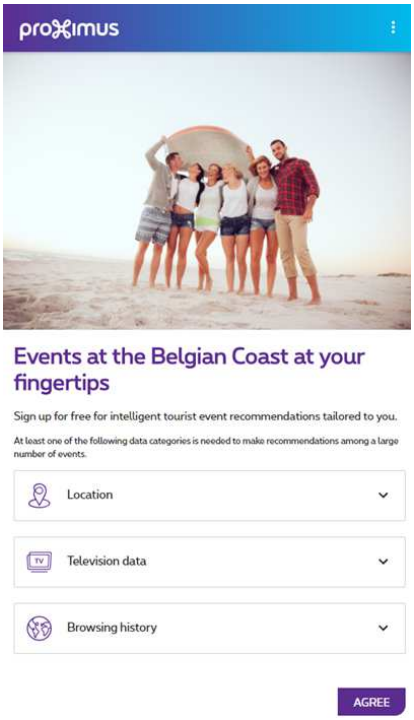

Figure 2 Consent overview

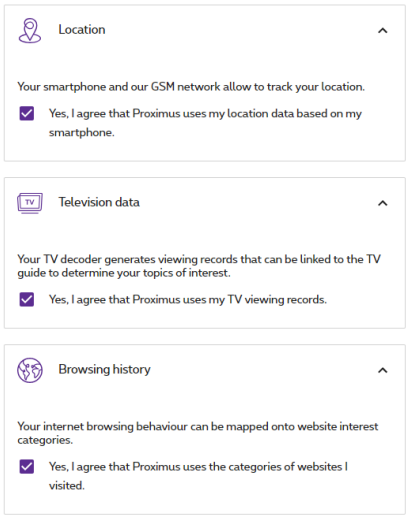

Figure 3 Consent Detail

3.7 Recommender engine

The Proximus Recommender engine is shown in Figure 4 Proximus Recommender engine.

As a side (internal) deliverable, Proximus will develop a NLP (Natural Language Processing) algorithm for creating customer profiles based on keywords.

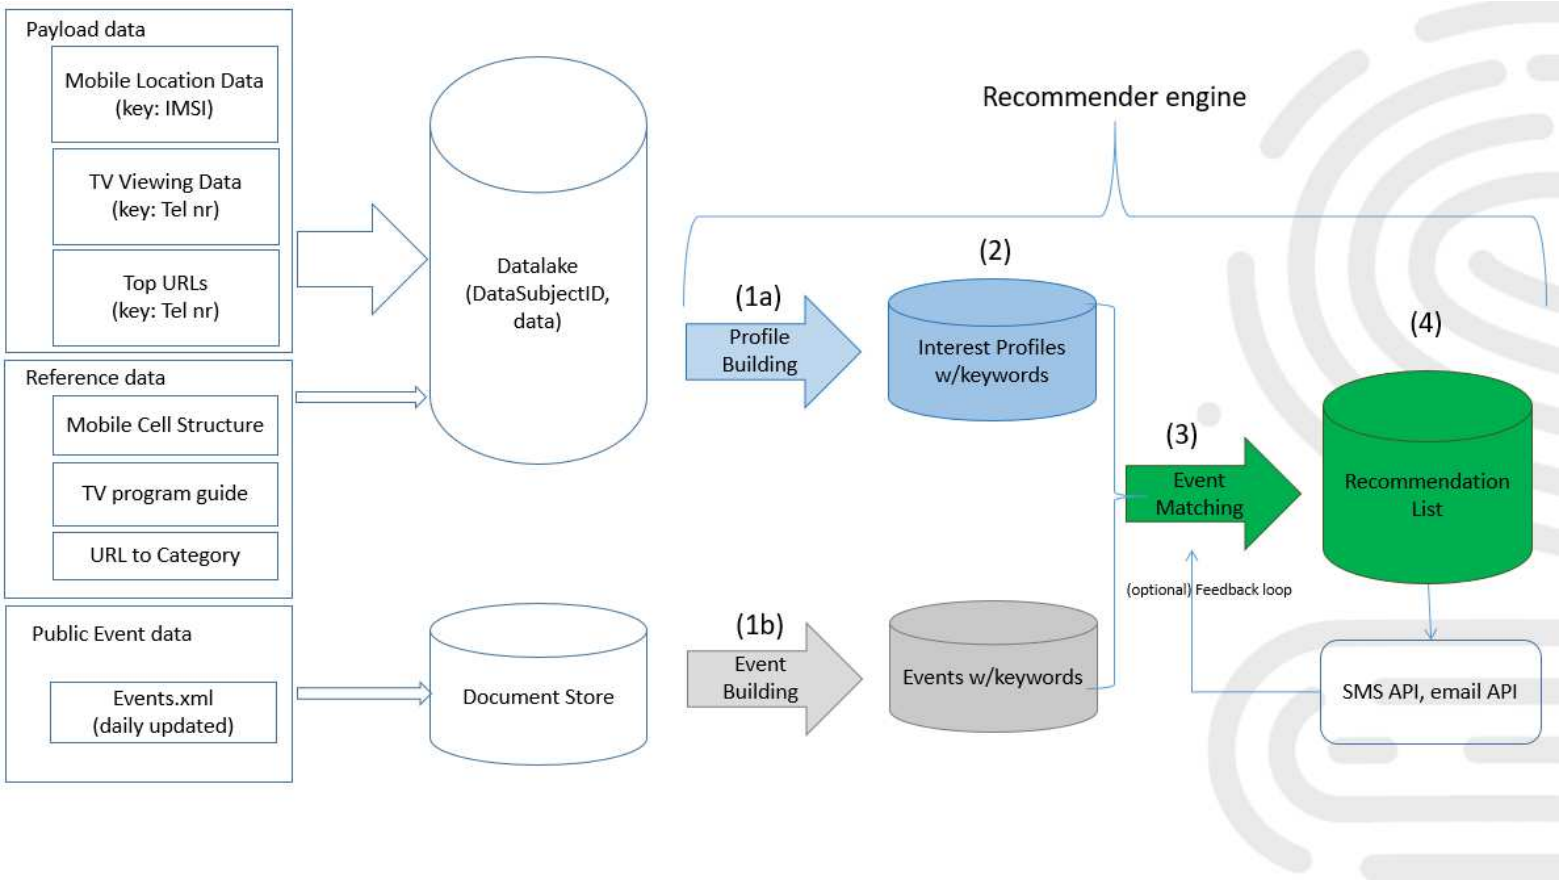

Figure 4 Proximus Recommender engine

Keyword and category matching was done to obtain a first level of recommendations. Figure 5 Keyword and category matching shows examples of the matching that occurs between TV programs (left side, blue) and Tourist Events (right side, gray).

Two methods are used: (1) Category and (2) Keywords, which are combined and calculated into a match strength score.

Channel Name	Program Name	Program descriptionWords	Program Category	Program Keywords	Strength	Event Keywords	Event Category	Event DescriptionWords	Event Title
Food Network	Diners, Drive-Ins	['road', 'guy', 'Search', 'tough', 'b	C.Documentary: C.Gastronor	{'Education', 'Cooking', 'Ga	0,855	{'Education'}	Begeleide rondleiding	['late', 'seduce', 'delicious', 'cc	Bezoek koffiebranderij
AMC	The Brothers Wa	['intimate', 'portrait', 'sage', 'film	C.Documentary: C.Biography	{'History', 'Education'}	0,879	{'Education'}	Begeleide uitstap of rond	['arnout', 'co-author', 'native'	fietstocht: merkwaardig
Sporza	Wonderland	['music', 'wonder']	C.Music: C.General	{'Music'}	0,945	{'Music', 'Lifestyle'}	Concert	['performance', 'quirky', 'brid	Zomerbar: Ambrazar
AMC	Hollywood Singin	['delivery', 'years', 'bid']	C.Music: C.Dance	{'Music'}	0,942	{'Music', 'Lifestyle'}	Concert	['whitesberge', 'host', 'biggest	Radio 2 Zomerhit 2018
VH1 Classic	The Rock Show	['know', 'love', 'hurtle', 'around',	C.Music: C.General	{'Music'}	0,944	{'Music', 'Lifestyle'}	Concert	['whitesberge', 'host', 'biggest	Radio 2 Zomerhit 2018
MNMH	Muzieklus	['radio', 'broadcast', 'chock-full',	C.Music: C.General	{'Music'}	0,946	{'Music', 'Lifestyle'}	Concert	['whitesberge', 'host', 'biggest	Radio 2 Zomerhit 2018
één HD	Clips	['clips']	C.Music: C.General	{'Music'}	0,946	{'Music', 'Lifestyle'}	Concert	['whitesberge', 'host', 'biggest	Radio 2 Zomerhit 2018
Radio1	De nacht van Rad	['news', 'least']	C.Music: C.General	{'Music'}	0,862	{'Music', 'Lifestyle'}	Concert	['dirk', 'cassiers', 'erik', 'gooss	Wednesday Night Jams
TLC	Say Yes to the Dr	['niki', 'budget', 'dollar', 'phrases'	C.Documentary: C.Fasion	{'Education', 'Fashion'}	0,923	{'Education'}	Cursus of workshop	['graffiti', 'busy', 'opinion', 'm	SWAP - Graffiti art
Disc World	Extreme Enginee	['azerbaijan', 'invented', 'decade	C.Documentary: C.AppliedSci	{'Education', 'Technology',	0,845	{'Education'}	Cursus of workshop	['fifth', 'sixth', 'love', 'science',	Wetenschapsacademie
Disc Science	Nextworld	['nextworld', 'future', 'statue', 'te	C.Documentary: C.Science	{'Education', 'Technology',	0,737	{'Education'}	Cursus of workshop	['fifth', 'sixth', 'love', 'science',	Wetenschapsacademie
Disc Science	How It's Made	['world', 'life', 'state', 'objects', 'p	C.Documentary: C.Science	{'Education', 'Technology',	0,798	{'Education'}	Cursus of workshop	['fifth', 'sixth', 'love', 'science',	Wetenschapsacademie
BRUZZ	BRUZZ 24	['bruzz', 'daily', 'appointment', 'br	C.Magazine: C.Documentary	{'Education'}	0,95	{'Education'}	Cursus of workshop	['local', 'service', 'butterfly', 'p	Aanbod dienstencentru
NPO 2 HD	De rijdende recht	['jetske', 'Pine', 'elsen', 'impact', '	C.Magazine: C.Documentary	{'Education'}	0,992	{'Education'}	Cursus of workshop	['Ostend', 'association', 'kolor	Arabische taallessen vo
VIER	BBQ Vanavond?!	['naanbread', 'yogurt', 'cucumber	C.Magazine: C.Gastronomy	{'Cooking', 'Gastronomy'}	0,738	{'Gastronomy'}	Eten en drinken	['women', 'kvlv', 'runway', "fa	Ontbijtpakketten vaderd
VTM HD	Sofie in de Keuke	['week', 'sofie', 'kitchen', 'dina', '	C.Magazine: C.Gastronomy	{'Cooking', 'Gastronomy'}	0,73	{'Gastronomy'}	Eten en drinken	['women', 'kvlv', 'runway', "fa	Ontbijtpakketten vaderd
Kadet	Star Wars Rebels	['ezra', 'zeb', 'want', 'relationship	C.Kids: C.Cartoon	{'Family'}	0,845	{'Gastronomy'}	Eten en drinken	['yummy', 'healthy', 'breakfas	Gezond en fit ontbijt vo
ms N	The Rewrite	['career', 'british', 'scriptwriter', '	C.Movies: C.Romance	{'Movies'}	0,956	{'Movies'}	Film	['time', 'gold', 'bazin', 'Dick', 'a	Opera Live 2018: La Fa
VTM HD	Sahara	['explorer', 'dirk', 'pitt', 'experien	C.Movies: C.Adventure	{'Movies'}	0,983	{'Movies'}	Film	['twin', 'siegmund', 'sieglinde',	Opera Live 2019: Die W
nick jr N	Dora	['Dora', 'surprise', 'new', 'bicycle'	C.Kids: C.Cartoon	{'Family'}	0,843	{'Tourism', 'Family'}	Kamp of vakantie	['piet', 'pirate', 'studio', 'figure	Kidskriebels – Studio 10
DisnJr N	Puppy Vriendjes	['puppy', 'friends', 'bingo', 'rolly',	C.Kids: C.Cartoon	{'Family'}	0,85	{'Tourism', 'Family'}	Kamp of vakantie	['trial', 'salt', 'sand', 'bustling',	Expeditie Noordzee
DisnJr N	Puppy Vriendjes	['puppy', 'friends', 'bingo', 'rolly',	C.Kids: C.Cartoon	{'Family'}	0,85	{'Tourism', 'Family'}	Kamp of vakantie	['trial', 'salt', 'sand', 'bustling',	Expeditie Noordzee
Kadet	Dinotrux	['scraper', 'steal', 'garby', 'compo	C.Kids: C.Serie	{'Family'}	0,884	{'Tourism', 'Family'}	Kamp of vakantie	['late', 'adrenaline', 'body', 'vu	Waterpretparkweek
VTM	Barbie Dreamtop	['dreamtopia', 'swimming', 'Barb	C.Kids: C.Serie	{'Family'}	0,874	{'Tourism', 'Family'}	Kamp of vakantie	['late', 'adrenaline', 'body', 'vu	Waterpretparkweek

Figure 5 Keyword and category matching

4 Pilot 2 (Thomson Reuters)

4.1 Summary

In this first iteration of the pilot Thomson Reuters is focussed on testing expressivity (of the policy language), compliance, and the user experience of capturing our processes as policies.

Effectively it is an internally focussed pilot to test the complexity (or not) of integrating Special into our existing Know-Your-Customer workflows and the confidence we can have in its results.

4.2 Pilot objectives

The first-phase pilot is designed to answer the following six questions:

- Does the policy language and associated vocabularies, as currently specified, capture all the relevant information required to decide compliance to the GDPR? If not, what use cases can we describe to add to the requirements for the next iteration of that language/vocabularies.
- 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?
- Can compliance checking be fully automated – without any recourse to manual intervention? If not, what kinds of manual intervention are needed, and when?
- 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?
- Given the volume of compliance checks our live systems are likely to generate, can the Special components make decisions fast enough: we're looking for millisecond response times.

4.3 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.

4.4 Pilot components

The first-phase pilot will test 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)

The pilot will also evaluate 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.

4.5 Evaluation criteria

We are using the legal expertise both within the Special project itself and the Thomson Reuters 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 within the Thomson Reuters Innovations Labs are evaluating the viability of using the existing Know-Your-Customer staff to create and maintain processing policies.

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

5 Pilot 3 (Deutsche Telekom)

5.1 Use case description (short recap)

5.1.1 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 a potential solution.

Since we use an existing app and thus have an existing user base, the new version of the app, and thus SPECIAL, will use real personal data in this pilot.

5.1.2 Current (pre-SPECIAL) situation

Users can download and install the CNE-app from the app store (Apple) or Play Store (Google). The purpose is very clear and the user interface offers the key functionality on screen. Of course the privacy statement and terms and conditions are presented at first launch. Then, the user can initiate manual "speed tests" or change certain settings, e.g. turn on a "diagnosis mode". Figure 6 shows the current app.

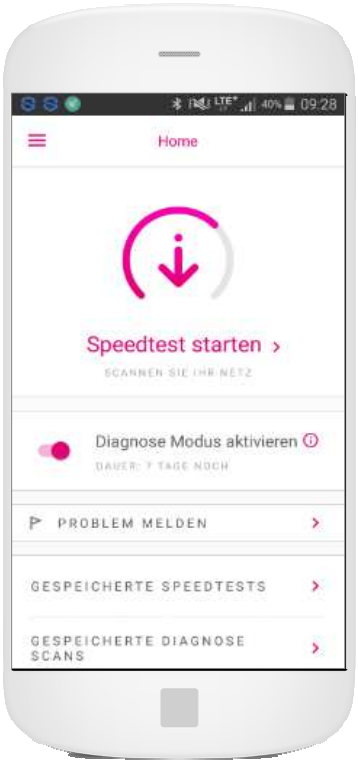

Figure 6 : Current CNE-app (customer network experience) with a slider to activate "diagnostic mode"; we intend to switch on "non-anonymous mode" similarly

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). 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 their theoretical capabilities and capacities. Motionlogic needs more "Opt-In" users among DT customers.

5.1.3 Target scenario for pilot implementation

CNE-app (or new versions of other DT-apps incorporating the CNE functionality) 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 assume 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 would be the tool to guarantee user control and transparency.

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 or users/data subjects to test scalability (in later iterations). In other words: dynamic consent is not part of the current initial implementation. The policy used here is more a one-step approval. DT is currently assessing D1.6 and the concept of dynamic consent therein. In future versions DT might use this to stepwise increase user's consent-level.

5.2 Objective of pilot implementation

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

DT owns a large amount of valuable user data, which it needs, to deliver the high quality services the users expect. Since DT has a first class reputation in terms of security and privacy protection, DT is

more than reluctant to experiment with user data, potentially breaking the valuable trust that was built up over time.

Nevertheless, DT acknowledges the potential additional value of this data and the results from the data analyses. In particular, when comparing with OTT ("over the top") players such as Amazon, Facebook or Google, DT sees huge opportunities in adding more and better services to customers based on "Big Data Analysis" or "AI".

DT always puts the customer and customer protection first and never risks legal or reputation threats. Therefore, SPECIAL and the observations/experiences made in the project are of perticular value for DT.

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.

These users currently 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 expects somewhat smaller numbers for non-anonymizing version of the app. However, neither the "original" nor the "SPECIAL supplemented" version of the app will be pushed in app store marketing. So user numbers will 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 pilot and internal evaluation is expected to open doors towards the monetization of real Big Data treasures that lie hidden in the carrier's operational data centers.

5.2.2 Objectives for the first iteration of 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 will most likely not 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 (Transparency 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.

T-Labs (and thus DT's) main objectives for the pilot implementation are:

- **Proof of feasibility** of technology developed in SPECIAL (privacy policies, policy engine, 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 will help 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 additional objective 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 **"sanitary factor"** 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.

5.2.3 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.

5.2.4 Public challenge

The main challenge for the pilot will be to convince as many users as possible to "Opt-In" to the use of their personal data (i.e. location data). This is not so much a "hacking challenge" but rather a proposed challenge to formulate the privacy information accordingly and let the user make a truly informed decision about the use of their data.

A resulting challenge (a core feature of DT's pilot) is the use and acceptance of the "transparency tool". We assume that the more users check and control their settings using this tool, the better is the overall acceptance. Of course, we assume that people will not only use the tool to "opt-out". So the intensity and type of use of the transparency tool will be a key aspect of this pilot. This also applies to the next iterations of the pilot/deliverable.

5.3 Pilot implementation architecture

The original CNE app and Use case (prior to the advanced version to be piloted in SPECIAL) is based on a straight forward data forwarding schema: Network Quality Data gained on user's/customer's device are uploaded to buffer server and then forwarded to the DT data center where they are evaluated in an anonymized form.

The analysis results are then transferred to Motionlogic for further processing and presentation. See for more details. The analysis results are completely anonymized and individual data sets cannot be reconstructed by Motionlogic. Obviously, Motionlogic also cannot assign any data set to a particular user/customer.

Motionlogic's operation and business model does not require any individual data set or any link to individual users. Therefore, the current implementations and data quality seem to be "good enough" for Motionlogic.

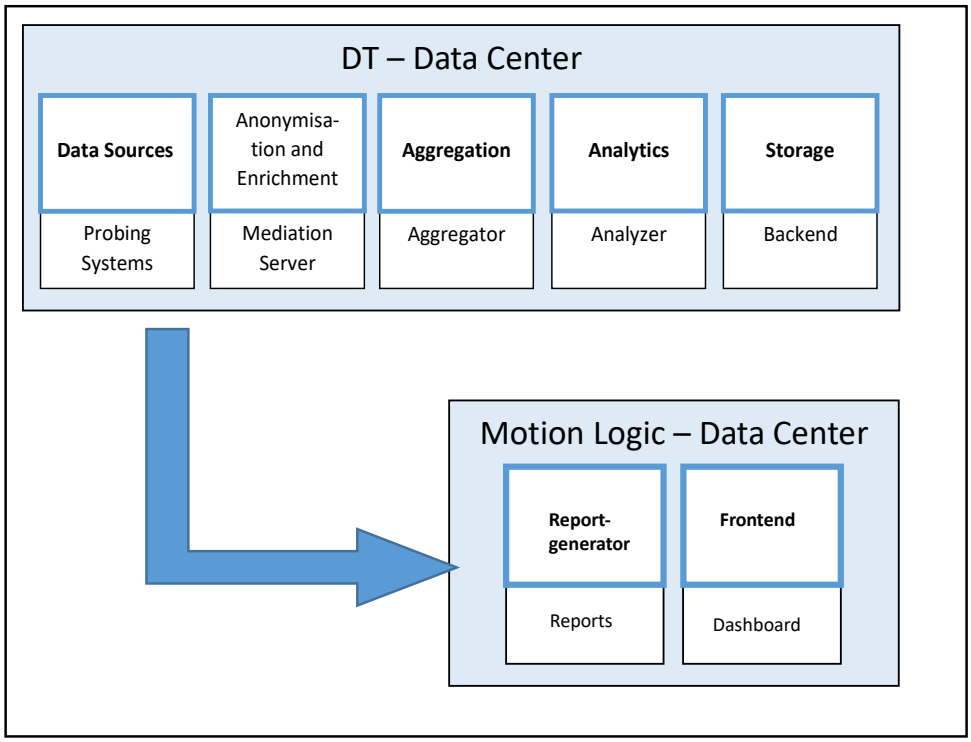

Figure 7: Current situation: data is collected and analysed in DT data center. Anonymized results are exported to Motionlogic for further processing/presentation

However, to improve quality of Motionlogic's algorithms, individual data sets with concrete (individual) users are necessary. SPECIAL delivers the necessary tools an 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. depicts the architecture with the "bypass" around the DT data center.

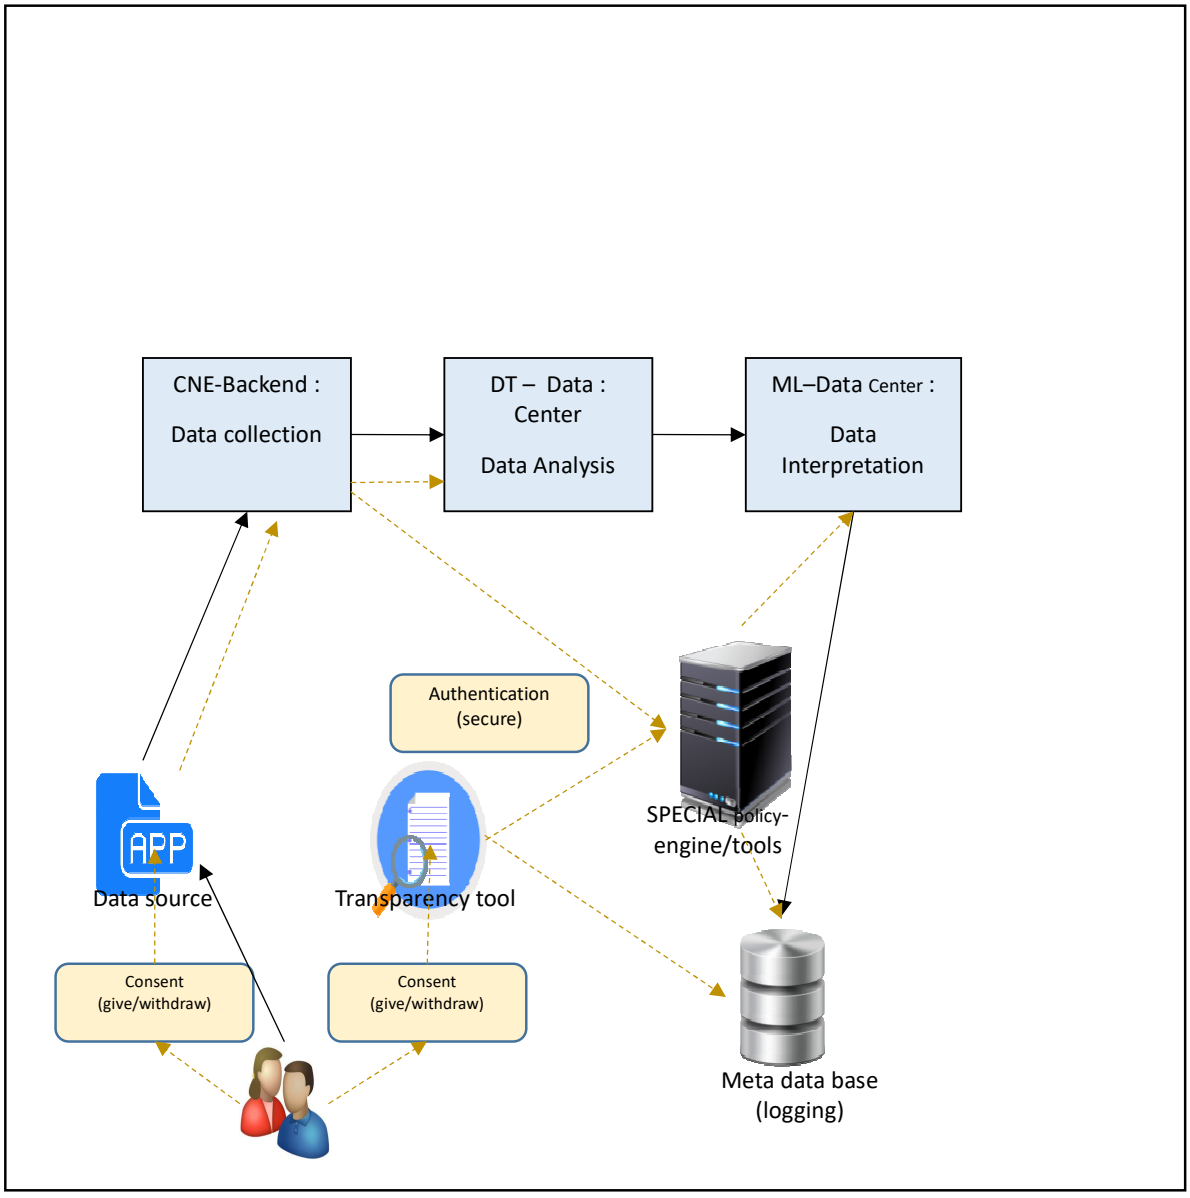

Figure 8 : Planned architecture: SPECIAL tools and methods are added to the original process. N on-anonymous data will be used legally, with transparency and user control is provided.

The following chapters will detail the implementation and explain certain design decisions as well as the planned test criteria and procedures.

5.4 Testing plans

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

All those tests will be executed by experienced test experts, collaborating with experts for both DT-business units (CNE-app, Motionlogic) and SPECIAL experts.

The tests will be carried out during a dedicated test period. The new (SPECIAL supplemented) version of the CNE app will be made available for download and install in the app store / play store. User devices will update automatically (as usual) and new instalments will use the new version.

During the first use of the new CNE app version, the user will be asked to renew the acknowledgement of the (updated) privacy statement (by changing a switch in the interface). This can be considered a "minimalistic version" of the dynamic consent mechanism described in D1.6. Due to the earlier decision to start with a minimal feature set but a relatively larger number of users, it is assumed to be fine starting with the smallest non-zero dynamics at this stage (i.e. no consent or minimal consent).

In addition, the user will be asked if he is willing to participate in the non-anonymized procedure for further quality enhancement. The user then gives his informed consent via checkbox and is provided a link to the transparency tool. The whole setting will clearly be marked as test procedure.

The test run is planned for at least one month with an option to extend the period. Of course, the test can be aborted anytime by replacing the app in the playstore and/or by stopping the data collection on the server side. Users will be informed immediately via the transparency tool.

During the test period and after the first month, the mentioned expert interviews and survey will take place.

5.5 Expected results & evaluation criteria

5.5.1 Expected Results

Given the ***qualitative objectives*** of the pilot and the tests, DT expects pretty simple, almost binary (Yes/No), answers to a few strategic questions. All of these are based on the situation of a large European telecommunications provider with an excellent reputation in many aspects including privacy trustworthiness.

Thus the expected results can be summarized as answers to the following list of questions:

1. Does it work? (Does SPECIAL enable DT and business partners collecting, storing and processing personal data?)
2. Does it pay off? (Is the effort in a reasonable ratio to the revenue?)
3. Are real world users satisfied? (user satisfaction regarding usability and functionality of SPECIAL tools)
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 expect more: the tools are meant to be accepted by end-users as legal and justifiable, not only (but also) by lawyers.

These four questions need to be asked and evaluated using the criteria described in the next chapter where we explain the evaluation criteria that ultimately lead to the four answers to the four questions.

While these questions might be answered on the basis of the pilot implementation with a limited (yet "big") number of users, the real value of the approach will be extrapolated to the assumed application to DT's "real" Big Data assets.

5.5.2 Evaluation Criteria

Evaluation of the pilot implementation will be done on a qualitative level (no exact quantitative measurements planned). Since the overall goals (objectives, see chapter 5.2) are feasibility and effects on existing IT infrastructure and user acceptance, it is considered to be of secondary relevance to collect any numerical data.

Our evaluation criteria can be divided according to objectives in four groups:

5.5.2.1 Sanitary criteria (criteria that need to be fulfilled in all products/prototypes anyway)

Security: Does the pilot implementation fulfil basic security requirements, e.g. is the app secured appropriately and protected against external threats. Are all databases protected against a moderately serious attacker? Is network communication encrypted and is the encryption sufficient?

Data Privacy: Does the pilot and all its' components comply with all applicable privacy [1]laws, in particular with the GDPR?

5.5.2.2 Value of Insights gained

Usefulness of gained information: Two types of information will be gained with the pilot: privacy/consent related data, and "payload" – data that is useful for Motionlogic's quality assurance. Both categories of information will be evaluated separately (simple yes/no decision, met by a subject matter expert)

Plausibility of gained information: Since the core findings (from payload data) will be based on machine learning style algorithms, they might contradict proven knowledge in this field. Even if we hope for "surprising" results with big innovation value, we will double-check the most suspicious results and thus rate the overall plausibility of the information.

Corporate fit: Organizational effort for implementation (minus re-usability effect for production version) is expected to be significant (as is usual in large organisations). Barriers for a pilot may be lower than for a fully-fledged operative system, however if the effort to get the pilot confirmed is protectively high, critical assumptions can be made regarding organizational acceptance of the final product. On the other hand, we assume that the pilot (if successful) will open doors for similar systems products in the times to come.

Basic technical criteria: Computational, storage and network load, scalability. Based on experiences with the pilot, DT needs to rate the technology (not the pilot) regarding potential impact on internal IT infrastructure. This can be done with standard IT performance measurement tools but will also be based on expert opinion. The rating will be qualitative only.

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

5.5.2.3 User acceptance criteria

(Relative) **number of "Opt-Ins"** made with the SPECIAL / pilot. Giving consent to a collect and process personal data such as location is a specific act that a user needs to perform to allow DT/Motionlogic to process his data. We assume that not all users will give this (binary) consent. This is the minimal version of the "progressive" consent granting described in D1.6. More stages/degrees of consent (using more personal data and/or more purposes) are foreseen for later stages (not within SPECIALs current time scope). Thus, the portion (relative share) of users that give consent in a very limited

[1] The goal of SPECIAL is not to demonstrate compliance to *all* privacy laws, but an assessment is necessary nevertheless.

experimental environment will be a most valuable insight for other situations where user consent might be asked.

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 valuable evaluation results and 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.

6 Conclusion

This deliverable is intentionally labelled “V1”, therefore final conclusions cannot be expected here. As mentioned above, D5.2 and D5.3 will reflect on later stages of the implementation and testing.

However, based on the experience gained prior to and while writing this document, several observations have been made that could be seen as “preliminary conclusions”:

1. Despite the different use cases and different stages of implementation, all three pilots came up with comparable implementation objectives and testing plans:
   a. Check feasibility/effort,
   b. Check functionality,
   c. Check security and privacy.
2. The general concept of SPECIAL (using linked data, privacy policies and the established prior work of previous EU projects) succeeded in convincing the industry partners.
3. All industry partners use the implemented use cases/pilots to test “corporate compatibility” with internal processes and IT infrastructure.
4. Usability is crucial, mainly for end customers but also for admins and backend operators.
5. Big Data and AI applications require techniques beyond SPECIAL’s scope. The pilots and testing plans are designed to deliver proven experience on feasibility of respective (Big Data/AI) analyses under the governance of modern privacy regulations.

First experiences (before the actual tests run) indicate that SPECIAL’s concept and tools are very well applicable, especially since they are designed in a way to be implemented in various ways including different corporate IT infrastructure environments.