SEVENTH FRAMEWORK PROGRAMME
Challenge 1
Information and Communication Technologies

Trusted Architecture for Securely Shared Services

Document Type:	Deliverable
Title:	**Pilot specification and use case scenarios**
Work Package:	WP 9
Deliverable Nr:	D9.1
Dissemination:	Public
Preparation Date:	31 December 2011
Version:	3.0 Resubmitted

The TAS³ Consortium

Benef. Nr	Beneficiary name	Benef. short name	Country	Date enter project	Date exit project
1	Katholieke Universiteit Leuven (Proj. Coord)	KUL	BE	M1	M48
2	Synergetics nv/sa	SYN	BE	M1	M48
3	University of Kent	KENT	UK	M1	M48
4	University of Karlsruhe	KARL	DE	M1	M48
5	Technical University Eindhoven	TU/E	NL	M1	M48
6	Consiglio Nazionale delle Ricerche	CNR	IT	M1	M48
7	University of Koblenz-Landau	UNIKOLD	DE	M1	M48
8	Vrije Universiteit Brussel	VUB	BE	M1	M48
9	University of Zaragoza	UNIZAR	ES	M1	M48
10	University of Nottingham	NOT	UK	M1	M48
11	SAP research (S&T Coord.)	SAP	DE	M1	M48
12	Eifel asbl	EIF	FR	M1	M30
13	Intalio Ltd	INT	UK	M1	M18
14	Risaris Ltd	RIS	IR	M1	M48
15	Kenteq	KETQ	NL	M1	M48
16	Oracle	ORACLE	NL	M1	M48
17	Custodix nv/sa	CUS	BE	M1	M48
18	Medisoft bv	MEDI	NL	M1	M12
19	Karlsruhe Institute of Technology (KIT)	KARL	DE	M1	M48
20	Symlabs SA	SYM	PT	M18	M30

Contributors

	Name	Organisation
1	Sandra Winfield	Nottingham
2	Thomas Kirkham	Nottingham
3	Dries Pruis	Kenteq
4	Lex Polman	Kenteq
5	Louis Schilders	Custodix
6	Brecht Claerhout	Custodix

List of Figures

1 Executive Summary

The objective of WP9 'Employability and Healthcare Demonstrators' is to prove the generic applicability of the TAS³ trust infrastructure for exchanging and managing personal information in different domains, in particular in the areas of employability and healthcare.

Building on the initial UK employability integration trial from Year 2, described in the previous iteration of this document, the updated demonstrator looks further at ways in which the process can be both captured and automated as much as possible through the use of TAS³ services. Additional TAS³ components have been integrated into the upgraded version of this integration trial, i.e. delegation of authority, Online Compliance Testing and the Ontology Service. Other TAS³ components, used in the previous version, have been updated mainly with a view to improving user centricity and usability.

The new demonstrator in the Netherlands explores the use of TAS³ in a legacy application setting. The TRIPOD consortium consisting of existing actors in the employability field has been established and full collaboration of the different service providers has been obtained. The scenario chosen is of high likelihood in the current difficult economic situation. The mass layoff scenario represents a typical use case when a factory closes down or a significant downsizing of activities takes place: in such situations a large number of employees are affected. Many European countries have already provided a legal framework for these unfortunate situations: see, for example, the recent closure of the General Motors plant in Antwerp affecting 2600 workers.

Finally, the Healthcare scenario has been updated and partially modified as a result of the withdrawal of the original Dutch Healthcare Partners. However the new scenario builds further on the Healthcare integration trial from Year 2, described in the previous version of this document, and focuses on enabling user-centricity in the Healthcare environment using the TAS³ architecture. It also supports the type of domain-specific implementation of TAS³ in the exploitation plan described in Deliverable 11.6.

Simplifying the connectivity to data security, privacy protection and trust ranking components within the same architecture is becoming reality. The main added value of TAS³ consists of integrating existing and new components in one architecture, yet still to be translated in one infrastructure.

Until today enhancing existing applications with privacy and security components has been a tedious and almost impossible task. Integrating components within existing (legacy) applications often results in significant development time and cost and without any guarantee that the different components are correctly integrated within the network infrastructure. In many cases, workarounds and interfaces between components are developed at the Service Provider or User application end. In addition, the TAS³ architecture benefits from a comprehensive legal and contractual framework which evolves into a complete ecosystem, providing accountability, the capacity to demonstrate compliance via two different intake process specifications (end user and Service Provider) and the provision of a high level of auditability and visibility to access and use information at both system and individual levels. More detailed information can be found in deliverables D6.1 and D6.2.

Moreover the User Centricity by design which is the cornerstone of TAS³ provides the individual user with the capacity to manage his/her personal information sharing.

2 UK Employability Integration Trial

2.1 Non-legacy application, 'bottom-up' development

2.1.1 Introduction

Graduate employability, as discussed in deliverables D1.1 and D1.4 as well as in the two previous iterations of this document is a topic that has come under increased focus in the recent economic climate, and especially in the UK. Recent national spending reviews and cuts have brought into focus the issue that graduates want to see a return on their educational (and financial) investment and are seeking ways of improving their position in the competitive market for jobs. The recent Browne report[1] in the UK, while proposing changes in funding for UK Higher Education, has increased the emphasis upon potential students, as paying customers, considering how their employability will be improved when choosing a university course.

In the UK an increasing number of institutions are addressing this by offering students practical workplace experience through work placements. There is a myriad of such schemes, varying in duration from 6 weeks to one year, some associated directly with courses, others allowing students 'time out' from their academic trajectory. Students are placed in UK companies, SMEs and, increasingly, overseas. There is also a growing number of agencies and cross-institutional schemes operating in this field, offering services of varying quality. Many UK Higher Education Institutions (HEIs) are contracting specific agencies as approved suppliers in order to attempt to guarantee an appropriate level of service for their students.

Building on the initial UK employability integration trial from Year 2 (described in the previous iteration of this document), the updated demonstrator looks further at ways in which the process can be both captured and automated as much as possible through use of TAS[3]. The exchange of sensitive and private data between parties (university, student, placement provider/agency and organisations offering employment) remains key to being able to offer the best matched solution. The previous integration trial is extended to show pathways for more than one student user, and incorporates the integration of TAS[3] technology into a Personal Data Store (PDS), owned and managed by the student and able to be used to provide data to enable more personalised services.

Efficient flow and exchange of information about the students themselves, the programmes they are eligible for, and the vacancies available, is needed to support the matching of learners to appropriate placements. There are elements of choice at both ends of the process: learners want to be able to choose from a selection of suitable placements, while employers wish to choose from a selection of suitable students. Preservation of anonymity and gradual ('staged') release of data, both from the students about themselves and the placement provider (employer) about the placement, help to ensure fairness and impartiality throughout the process.

2.1.2 Summary evaluation of previous integration trial

The main aim of the Year 2 integration trial was to show integration of a collection of TAS[3] components in a 'green field' situation where no previous system existed.

---

[1] http://www.bis.gov.uk/assets/biscore/corporate/docs/s/10-1208-securing-sustainable-higher-education-browne-report.pdf

We outlined a storyboard based on a realistic scenario: this involves two students seeking work placements who are interacting with TAS³, one coming from within the institution and with previous experience of TAS³ for other processes, the other as an external student who is new to TAS³. From this we developed an end-to-end business process, where the user was able to select a matching service on the basis of chosen levels of trust and security. Development effort concentrated on reliable integration of a number of TAS³ components, taking initial prototypes developed independently and incorporating them into an overall process so that meaningful results were generated and processed by each. This choreography integrated work by six different WPs (Single sign-on, use of security policies, service discovery, service selection, use of Trust settings, Workflow management and use of the Audit Bus and Dashboard, as well as interaction with the Legal workpackage) while paying consideration to user centricity.

We broke the process down into five phases: joining, registration, service discovery and selection, service execution and an end phase. The trust database was populated with default values, the student user authenticated to the system using ZXID[2] and logged in using her University credentials. Placeholder terms and conditions were used to show introduction of legal content. A SAML[3] token returned basic user data from the IDP, which was viewed and accepted for the process before being used for programme matching and selection. Further registration data was collected, the user set her policy levels for trust and security, and the Discovery service, invoked by the workflow, located Service Providers who could offer the correct programme. The Discovery Service called the Trust PDP to ensure that the list of services returned to the user matched the user's settings, and the list of services and trust ranking were returned to the workflow PEP, which delivered it to the workflow.

We were able to show how an empty set of results could be returned if the user preferences could not be met by any of the available Service Providers; in this instance the user was given the opportunity to refine her policy settings and return to the service discovery phase.

We were then able to execute an exception where, by manually manipulating the trust ranking of a service in the trust database, the user was warned that her choice no longer matched her set parameters and the process looped back to repeat the Discovery Service.

On successful repeat of the Discovery service the chosen service was invoked and the invocation call, including personal data and policy information, was passed through the PEP, checked with the PDP and the service executed. We were able to show all transactions in the Dashboard, and finally the user was asked for feedback and issued with a receipt consisting of a user-friendly handle on the log information collected during the transaction.

Following a significant amount of development and integration effort, we were able to show successfully that it was possible to integrate remote components based on different machines to perform this end-to-end demonstrator, which was shown at the project review in March 2010. However we were aware that at this stage in the integration process, we were using placeholder, very coarse-grained, policies and minimal personal data gathered dynamically from an application form. To build on this, we saw the need

---

[2] http://www.zxid.org/

[3] http://www.oasis-open.org/specs/#samlv2.0

to progress to finer-grained policies and data, and to incorporate use of an existing user data store in order to move closer to a full pilot situation.

2.1.3 Outline of new scenario with new features highlighted

This Year 3 demonstration develops further the scenario described in the previous release of this deliverable. As before, it aims to be flexible and focuses on showcasing the integration of TAS³ technical components within the context of a real-life situation.

We aimed to continue the integration theme and to encompass the full range of TAS³ components (in particular use of the Ontology Service via the Credential Validation Service and use of OCT to assess a service provider's reliability and therefore contributing via the KPI to its trust ranking) but also to include multiple user workflows to show an episode of interaction with TAS³ from multiple perspectives, and to show use of personal data (represented by a placeholder in the previous demonstration) from a TAS³-enabled version of an ePortfolio of the type that is currently used by students at the University. The TAS³-enablement of the ePortfolio also has led the project into the direction of presenting a Personal Data Store (PDS) as an unexpected output.

Technically the demonstrator involves the integration of services with existing web-based sources of data, specifically with a TAS³-enabled PDS. As before, the central processes involved in the real-life scenario are held largely within office procedures and are not automated, with significant use of integrated computing applications. The lack of systems supporting existing processes in the scenario has enabled us to focus on the user's needs in terms of planning user-centric service development and integration of the core TAS³ components.

We have also broadened the scenario so that instead of just one, it includes the perspective of two student users: Learner 1 is new to TAS³ and has no previous experience with using a PDS; Learner 2, on the other hand, already has a TAS³-enabled PDS, has used TAS³ before and is familiar with its functionality, but has never used it before in the context of seeking work placement. We are therefore able to show how both new and existing TAS³ users are able to interact with the system.

The Trust service (as described in deliverable D5.3) involves a combination of three engines: reputation, KPI and credential based. Reputation is based on feedback from the business process, ranked in the Trust database. Trust criteria can therefore consist of asking for services with different types of reputation scores, which meet boundaries for testing scores (e.g. no failures in the previous month) or for specific credentials (only those services certified by a certain party). Users can generate trust policies according to these criteria or combinations of these criteria. For example:

1) Only accept services with a reliability score of:

- >90%
- >75%
- >50%
- accept any, this is not important to me

2) Only accept services with an average feedback score of:

- >9/10
- >7/10
- >5/10

- I don't mind what the average feedback rating is

3) Only accept services which have been certified by the University of Nottingham):

- require University of Nottingham certification
- accept any, this is not important to me

We anticipate that later versions of the design might also allow the following:

4) Only accept services which have a performance score that is

- very high
- high
- medium
- I don't mind

(Performance score is computed from stock price, frequency of use, etc.)

The OCT component of TAS³ (see deliverable D10.3) publishes messages about a service's performance quality to the audit bus; the Trust PDP uses a listener client to monitor this. Ultimately, however, the result returned by the Trust PDP will be Accept or Deny for a service. The use of the OCT service to automatically test services used in the TAS³ set of applications allows the demonstrator to show automatic reactions of the system to changes in trust ranking of specific services.

The Ontology Service communicates with the Credential Validation Service (CVS), which is part of the SP authentication process in the Matching Service (see the process illustrated in the sequence diagram in Figure 5 below). For example, if the language used by a Service Provider and the learner's policy are different, the Ontology service is called to see if there is a valid match using the structure of subject/attribute/role. The Ontology service communicates with the CVS, which in turn talks to the PDP, then the PEP, and finally to the Matching Service. If the Ontology Service fails, the Matching Service fails in the same way as policy denies access but with a specific error linked to the failure of the system to link to the data specific policy to service provider rules in the PDP.

2.2 Integration trial demonstration environment and scenario

2.2.1 Basic Storyboard

Actor	Person	Description
New TAS³ user (external)	Learner 1	Accesses TAS³ via the application which supports it when looking for a job (Placement Co-ordinator website). She creates an account with the SP that is TAS³ enabled.
Established TAS³ user (internal)	Learner 2	A more long-term user whose institution has already provided him with a TAS³-enabled personal data store in the form of an ePortfolio
Placement Co-ordinator	World of Skills	Has a contract with an institution to manage student placements using TAS³, but also has an external

Actor	Person	Description
		presence and will offer similar services to external candidates from other institutions
Service Provider	Matching Service	Provides a TAS³-enabled matching service
TS3-enabled PDS	Mahara[4] PDS	ePortfolio service with additional TAS³ functionality

2.2.2 Use of the architecture

This scenario uses the following components from the TAS³ architecture:

- Web browser
- Front end Web GUI
- Business Process Engine
- Web services
- User Audit/Dashboard
- Policy Editor and Consent Management
- Delegation settings
- Identity Provider
- Trust and Reputation
- Authorisation
- Delegation Service
- Ontology Handler
- Discovery Registry
- Trust Network Management Processes
- Business Process Models
- Policies
- Modelling Tools
- Organisation Level Ontology
- Audit Events
- Online Compliance Testing

The components used are highlighted on the latest version of the project architecture diagram in Figure 1.
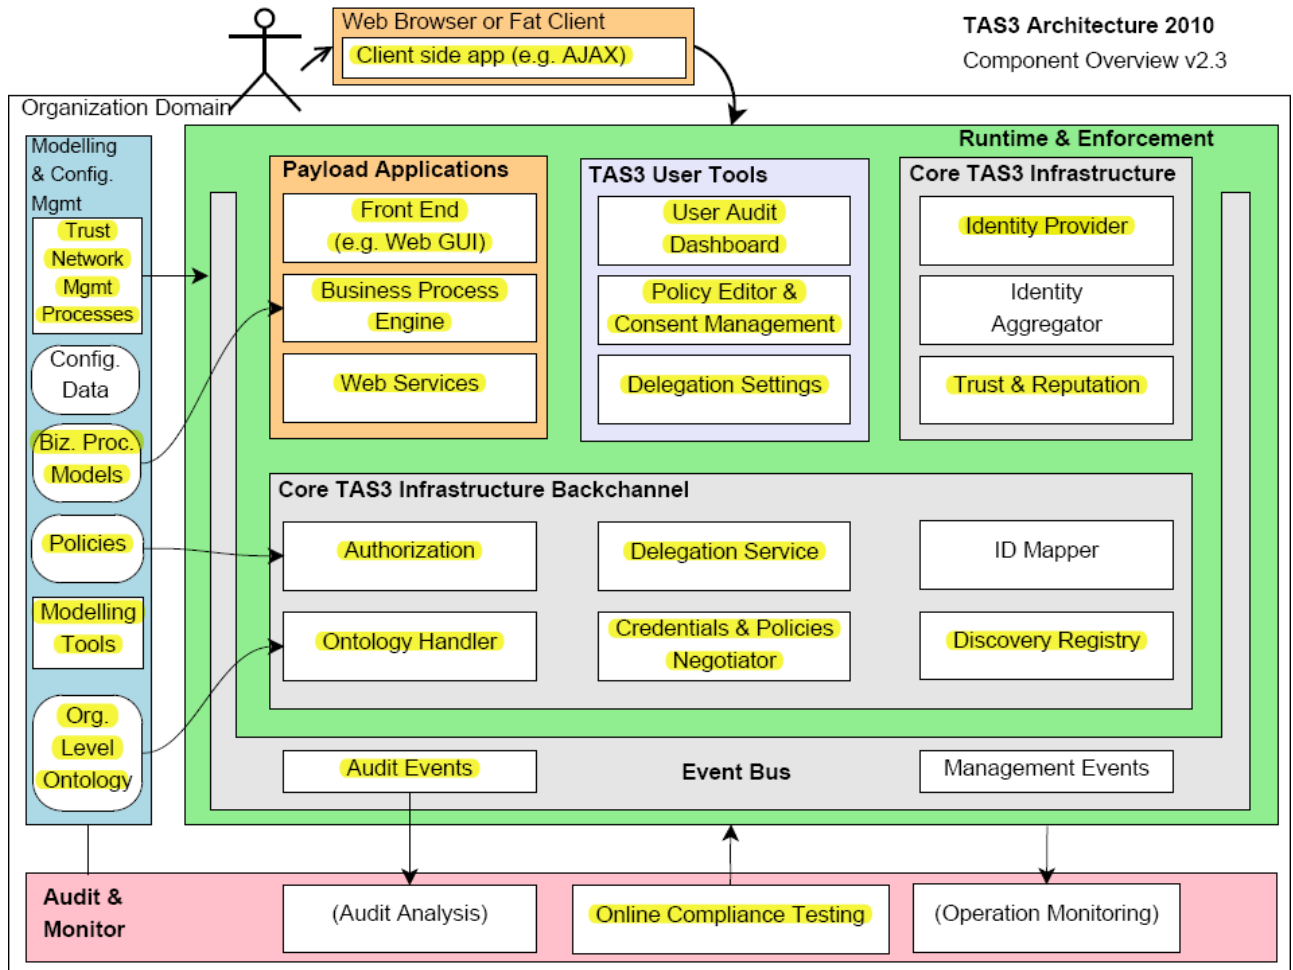
---

[4] http://mahara.org/

Figure 1 TAS³ architecture diagram showing components used

2.2.3 Prerequisites

All actors (placement administrators, placement providers and any application-specific service providers) are already registered with a TAS³ network and have agreed to any contractual obligations; vacancy profiles are available to the system, and have previously passed through a registration process similar to that for service providers.

- University and Placement Co-ordinator have an agreement and are registered with TAS³ IdP

- There is a TAS³-enabled PDS available: the Placement Co-ordinator is able to trigger automatic generation of new accounts and Learner 2 already has an account

- Matching Service Providers have signed up to the TAS³ network and conform with any overarching contractual obligations

- Placement Providers have registered placements with Matching Service Providers and these are categorised accurately using a profile that can be matched to learners

- Trust network is in place

2.2.4 Basic storyboard

1. Learner 1 (an external user) is seeking a placement and decides to register with a placement administrator who is able to offer places on a suitable programme. The learner enters the placement administrator system via the website and is asked for authentication. She is able to do this by selecting her preferred IDP and using SSO: if she is successfully authenticated she is able to log in.

2. If login is successful, Learner 1 is asked to accept both general TAS[3] terms and conditions as well as the specific terms and conditions for the service, and is asked to provide registration information. This is the first point at which this user actively submits her data to the TAS[3] infrastructure. This data is collected using a registration form and is used by the Placement Co-ordinator to check eligibility (most such programmes are only open to those who are current registered students with a UK Higher Education Institution). Previously we used application form data to represent this information. By collecting it in a form marked up using the UK Leap2A[5] ePortfolio interoperability standard, it becomes possible to import this data into a Mahara ePortfolio/PDS.

3. The Placement Co-ordinator triggers the process to set Learner 1 up automatically with a Mahara PDS and the application form data is transferred to it with a base default policy attached to it. Learner 1 can edit this policy within Mahara; she can also now use this data for other purposes, should she so wish. She can also use Mahara to view her own personal Dashboard to track use of her data within TAS[3]. As in the previous demonstrator, PERMIS policies will be used

4. Meanwhile Learner 2 (an internal user) has reached a stage in his course where his institution says he must carry out a work placement. His institution already has a contract with the Placement Co-ordinator to manage learner work placements for this programme, and has already provided all students on this course with TAS[3]-enabled Mahara PDS accounts, which learners can log into using SSO and have been using for other TAS[3]-enabled transactions. Learner 2 therefore already has the necessary placement matching data within his PDS, and it already has policies attached to it. He also has a personal TAS[3] Dashboard with records of his other transactions (business processes and sets of audit data). Learner 2's institution has enabled the functionality within the system that allows him to click on a button that activates the placement process; this functionality is now pushed to the Learner.

5. Both learners now specify trust policies for the placement service and create a profile using a view within Mahara which they share with the Placement Co-ordinator.

6. The Placement Co-ordinator uses each learner's profile to match him or her to a programme. Specific programmes may require additional data, or specific policies to be attached to data. Both learners can add data to their view, fine tune their policies and agree to any additional terms and conditions that may be attached to their delegated programme.

7. Learners then specify their delegation criteria: they can choose to delegate responsibility for finding matches to the Placement Co-ordinator.

8. The Placement Co-ordinator also has a set of policies for dealing with Matching Service providers. These are added to the user data and the Service Provider must comply with both.

---

[5] http://wiki.leapspecs.org/2A/specification

9. The Placement Co-ordinator triggers the Discovery process, which then selects or rejects appropriate candidate services in the network according to their compliance with trust and security policy levels. These are partly determined by the results of routine testing using OCT and the current status of the SP's KPI. The ontology service is used to resolve any mismatch in policy terminology (see process in Figure 2).
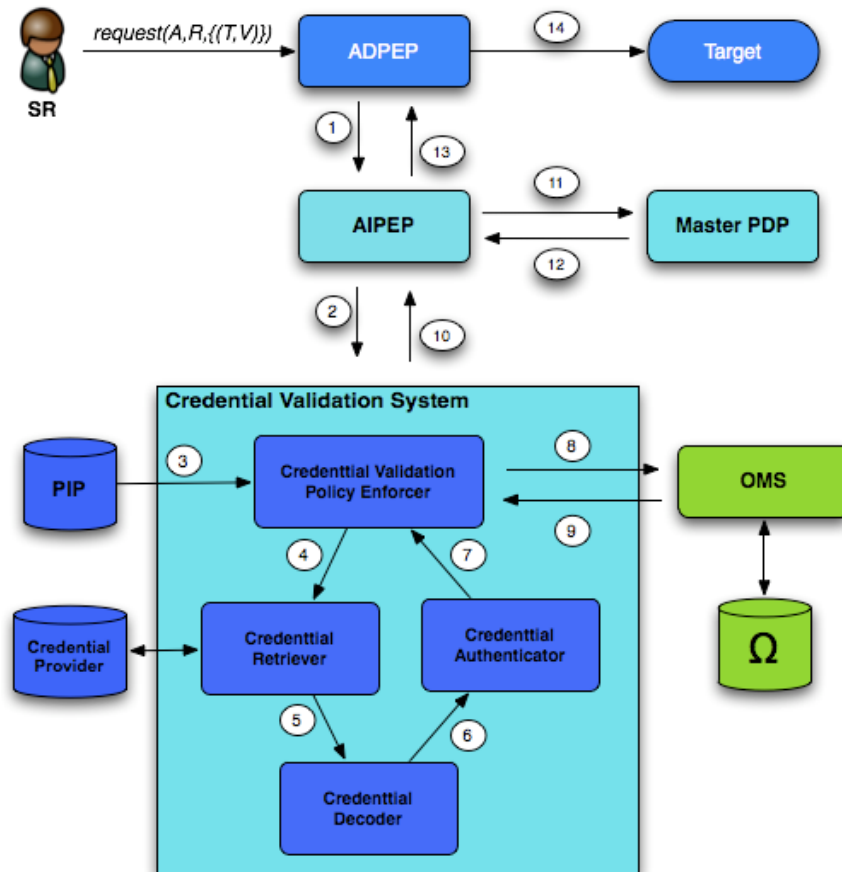

Figure 2 Position of Ontology Management Services (OMS)

10. Following the negotiation phase, the learner or Placement Co-ordinator (depending on delegation criteria) is presented with a list of suitable Service Providers: this can be none, one or many. (For current demonstration purposes we do not need to limit this as we will be using dummy data; at a later stage we will consider an option to limit results if a very large number of matches is returned.)

11. If a trust ranking changes during service discovery, the discovery service loops back.

12. The chosen matching service is executed: learner data is released to it and the execution process involves interaction of the main policy decision and enforcement points to ensure compliance with the learner's policies. The request to invoke the matching service is authorised, and this invocation is achieved via the policy framework. The results are also passed back through the policy management services to ensure security of the data. This phase of the demonstrator is a thin slice at this stage but is nevertheless able to illustrate how the policies are used within the system, how data is protected by the binding terms of the TAS[3] infrastructure and how these can be expanded further to include service providers.

13. The resulting list of matching vacancies is written to the learner's PDS with an associated policy that includes rights for the learner only. The learner may choose to grant additional access rights. The Placement Co-ordinator receives notification that results are available, and may choose to email learners to tell them they can be viewed.

14. Once the match is complete the learner makes specific applications for the placements he or she is interested in. If the learner does not receive any matches or rejects the choices offered the whole matching process can be repeated with the learner changing his or her trust requirements.

15. The TAS³ monitoring and auditing services are operational throughout, so at any stage each learner is able to interrogate the Dashboard to see what has happened to his or her data. At the end of the process the learner, Placement Co-ordinator and Service Provider are each issued with a receipt that gives information about data has been used and the actions performed.

16. All data (apart from the audit data) associated with the transaction is destroyed: any new interaction with the system is via a separate workflow. The log events are never deleted.

2.3 Demonstrator status and objectives

This process has at its heart the user selecting and securing personal data for use in a specific application provided by a TAS³-compliant framework. The interaction in this framework shows how policies are used in the system in order to protect the user, the service provider and the wider integrity of the application framework.

In practical terms, the integration of TAS³ systems and data remains a key development challenge presented by the scenario. In terms of the wider scenario and real-life application, the demonstrator puts the user at the centre of the process and in control of who provides job matching, what personal data they can use, and how long they can use it for.

This control by users over their personal data highlights practical steps to show how user-centric privacy can exist in distributed computing applications. This privacy extends to allow both employers and service providers to secure any information that they present to the system.

Overall, the following issues are addressed:

- The need to integrate systems and transfer data

- Provision of a better choice of matching facilities for learners

- Giving learners control over access to their data and the ability to check who has accessed it and in what context

- Preservation of privacy and anonymity within the process.

The management of policies is at the centre of the work. Further aims are to demonstrate:

- authentication/authorisation of access to sensitive data secured by both users and service providers

- policy setting/tuning managed in a user-friendly way but also secure enough for use in a complex and distributed system

- trust negotiation using metrics that the user can understand and relate to (negotiation must present meaningful results and ensure that the trust levels of service providers can be readjusted in real time, and in extreme cases re-negotiations can take place mid-process)

- some key service integration between application level services and the TAS³ trust infrastructure to demonstrate external service provider interaction with the trust framework

- that the user is at the centre and in control of use of personal data (achieved by feeding significant events in the process back to the user for approval; clear and understandable interfaces need to be presented to the user, particularly for the policy-setting process and key decision points in the workflow)

- a sizeable part of the complete integrated TAS³ trust infrastructure in action (the flow will illustrate how a user can begin interaction with TAS³ and continue through to the execution of an application in the framework and the retrieval of the result; this presents interactions with the key security and user-specific components in TAS³ and presents a strong basis for future developments).

The main achievements in this demonstrator are the integration of a variety of components generated by a number of partners. Technically the cross-domain management of user-secured data in the TAS³ policy framework is an innovation that the project will build upon in further development phases. This is also reflected in the logging framework, with Dashboard interaction and the receipt as an additional log-based innovation, giving users the wherewithal to investigate application execution once it has completed, should they wish to do so.

The sequence diagrams in Figures 3-6 below show how the components built by project partners interact in this demonstrator. This extends and expands the model used in the previous demonstrator: new services not included then are shaded in blue.
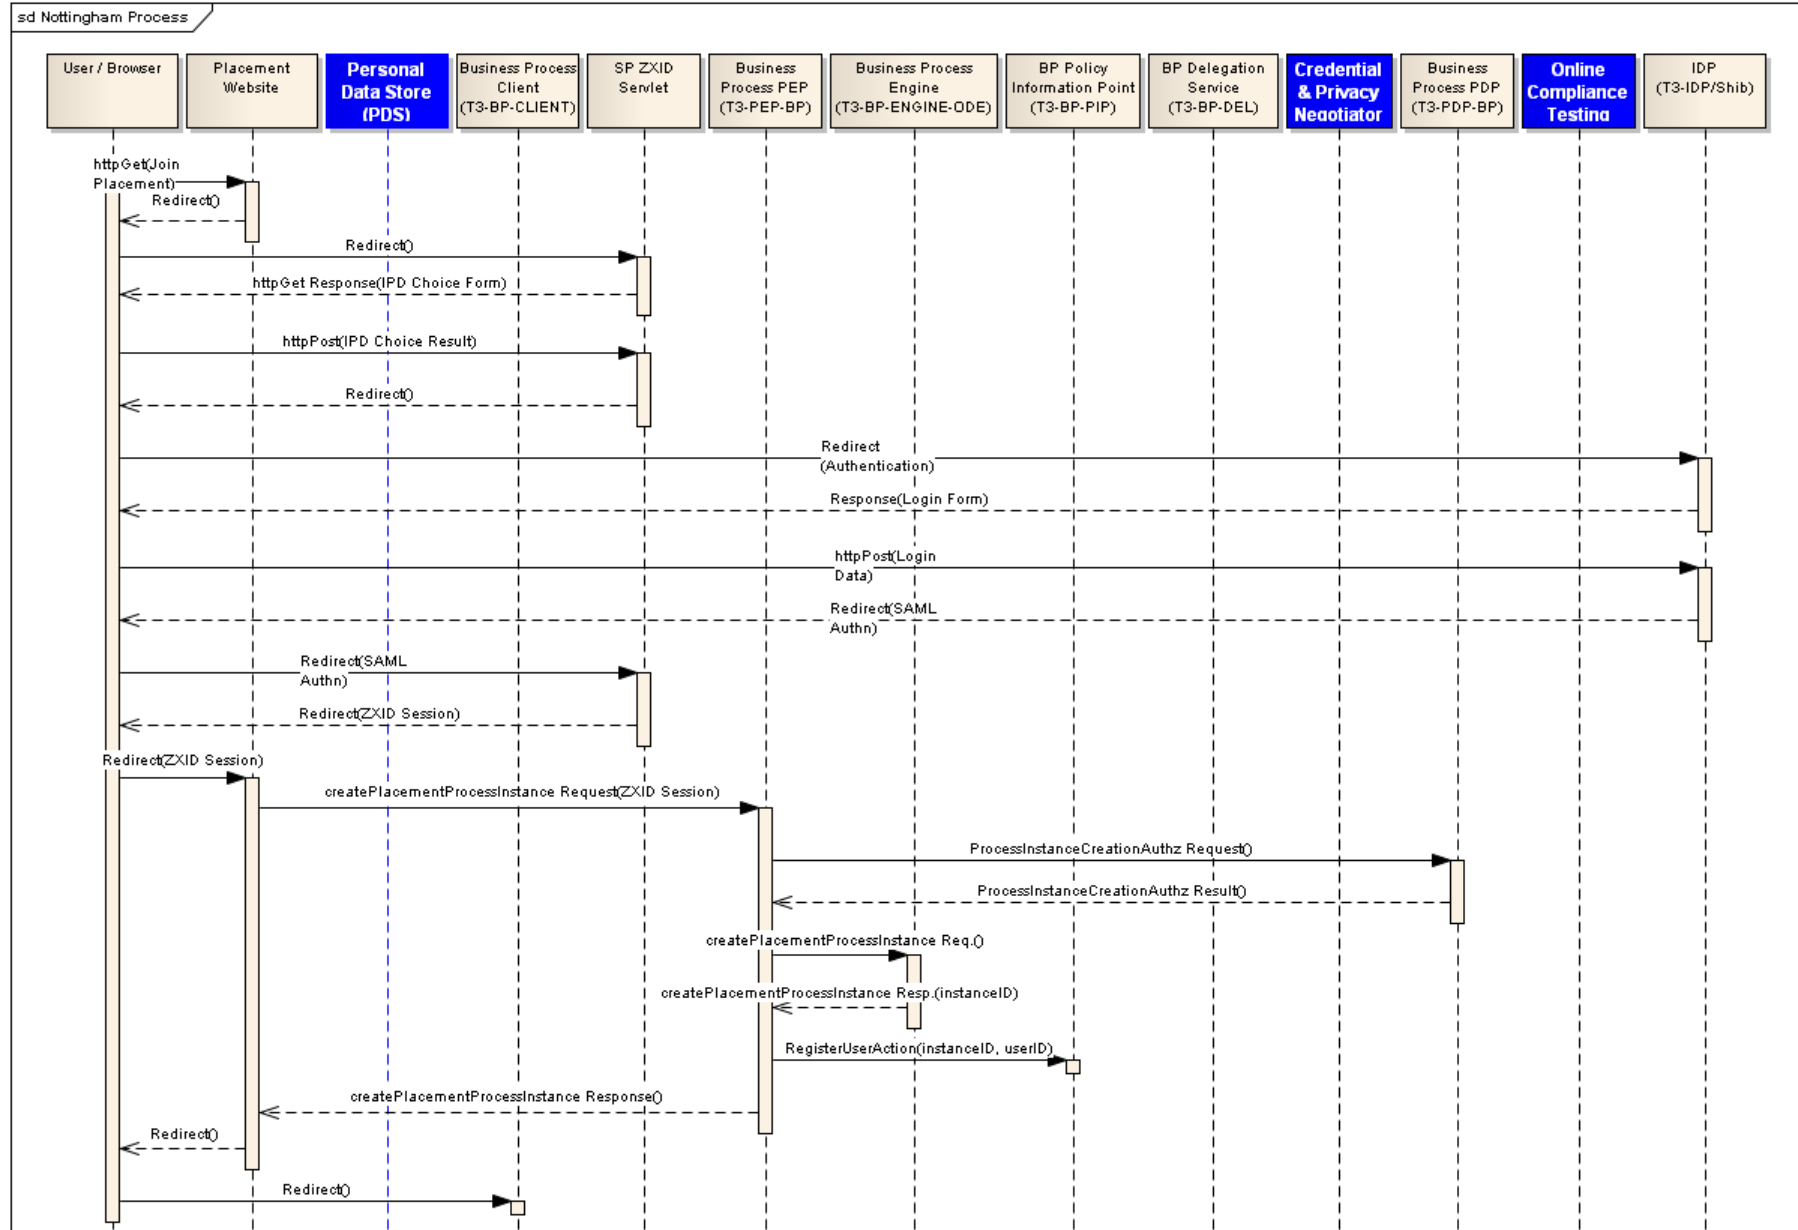
sd Nottingham Process

| User / Browser | Placement Website | Personal Data Store (PDS) | Business Process Client (T3-BP-CLIENT) | SP ZXID Servlet | Business Process PEP (T3-PEP-BP) | Business Process Engine (T3-BP-ENGINE-ODE) | BP Policy Information Point (T3-BP-PIP) | BP Delegation Service (T3-BP-DEL) | Credential & Privacy Negotiator | Business Process PDP (T3-PDP-BP) | Online Compliance Testing | IDP (T3-IDP/Shib) |

httpGet(Join Placement)

Redirect()

Redirect()

httpGet Response(IPD Choice Form)

httpPost(IPD Choice Result)

Redirect()

Redirect (Authentication)

Response(Login Form)

httpPost(Login Data)

Redirect(SAML Authn)

Redirect(SAML Authn)

Redirect(ZXID Session)

Redirect(ZXID Session)

createPlacementProcessInstance Request(ZXID Session)

ProcessInstanceCreationAuthz Request()

ProcessInstanceCreationAuthz Result()

createPlacementProcessInstance Req.()

createPlacementProcessInstance Resp.(instanceID)

RegisterUserAction(instanceID, userID)

createPlacementProcessInstance Response()
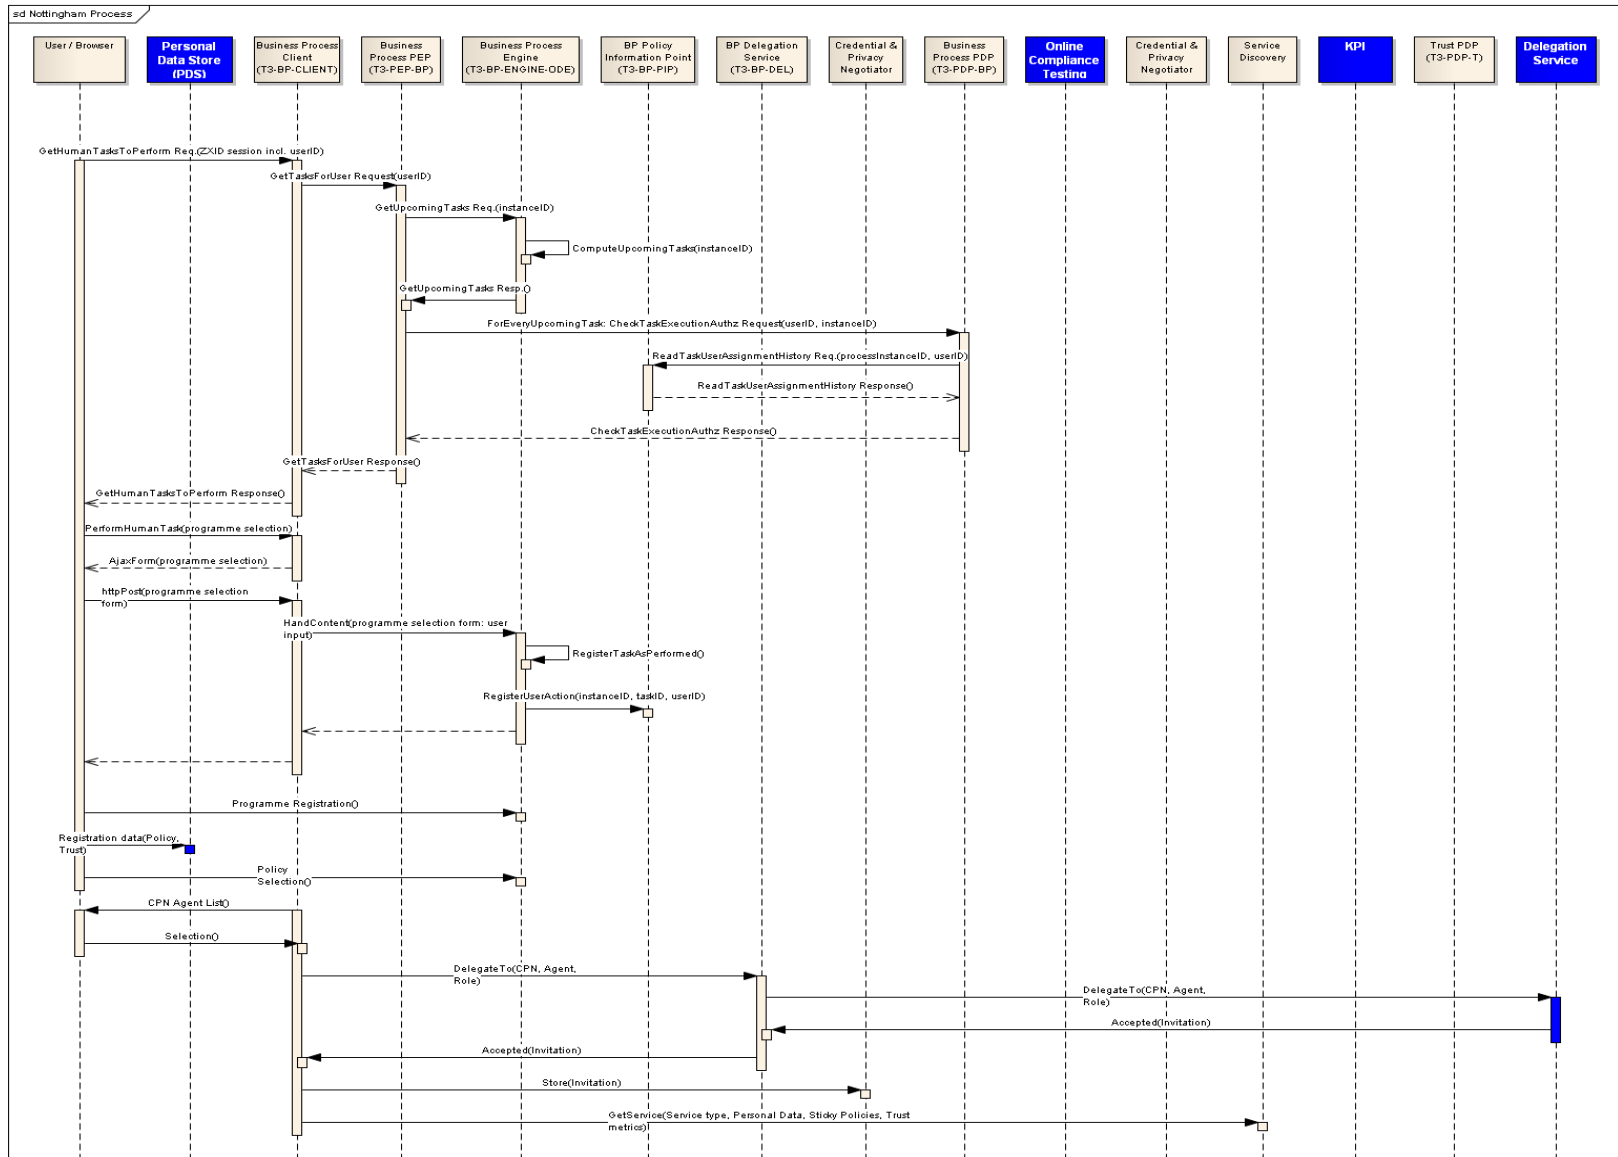
Redirect()

Redirect()

Figure 3 SSO phase

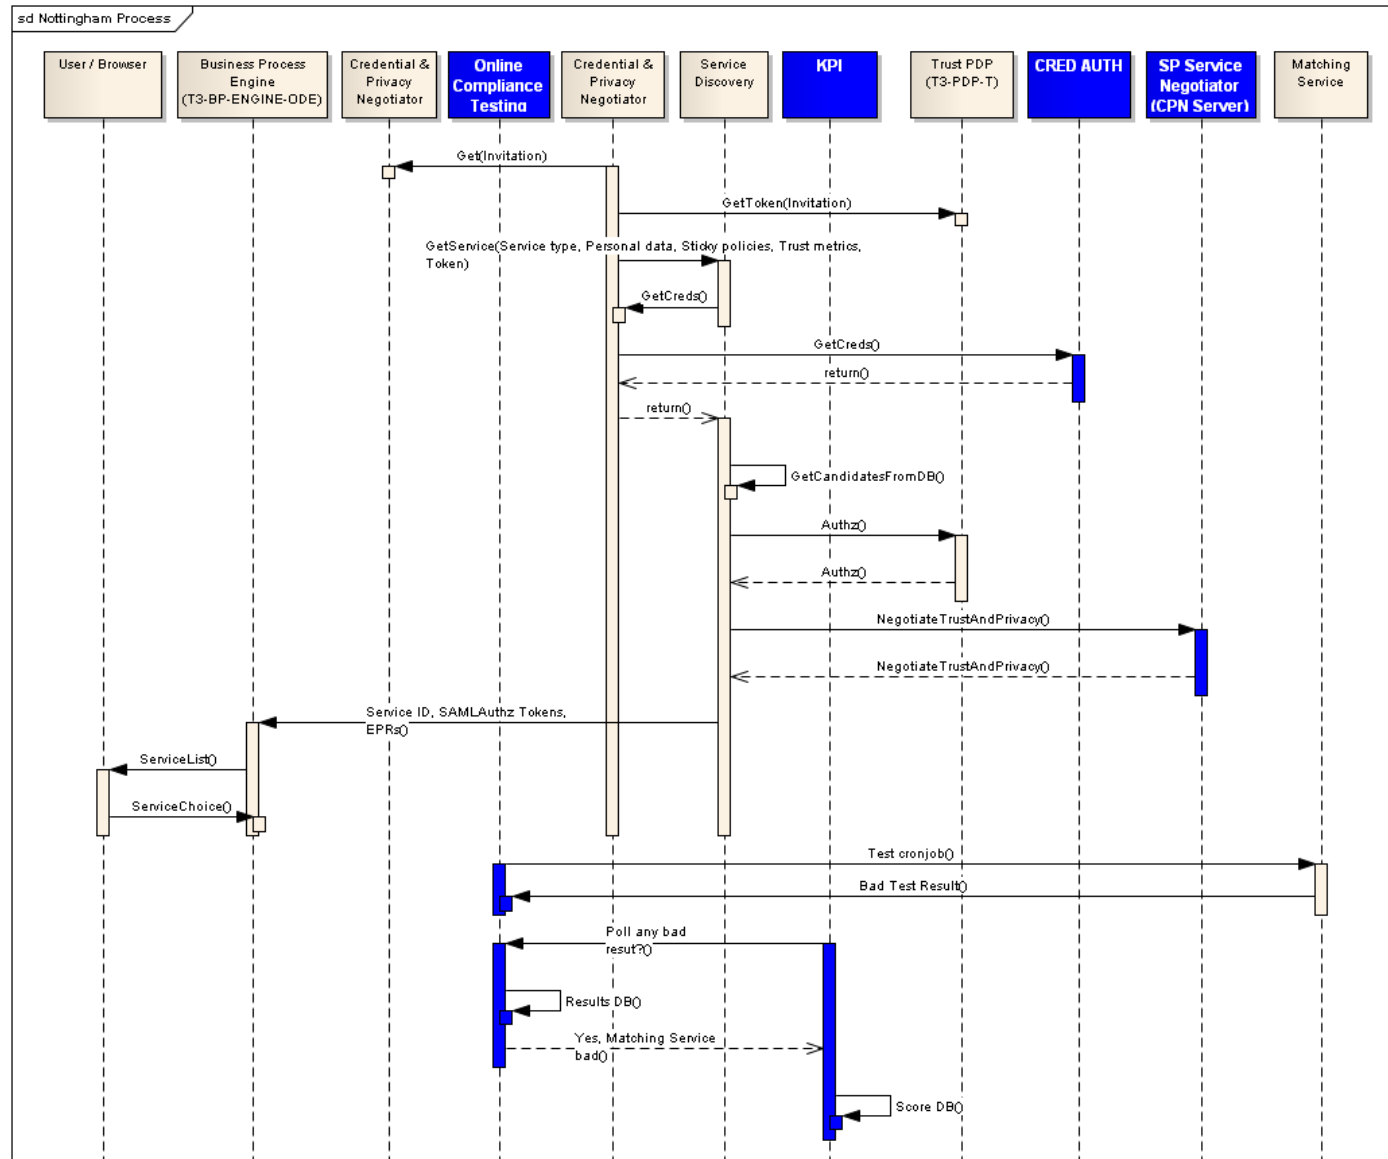**Figure 4 Service selection and delegation**

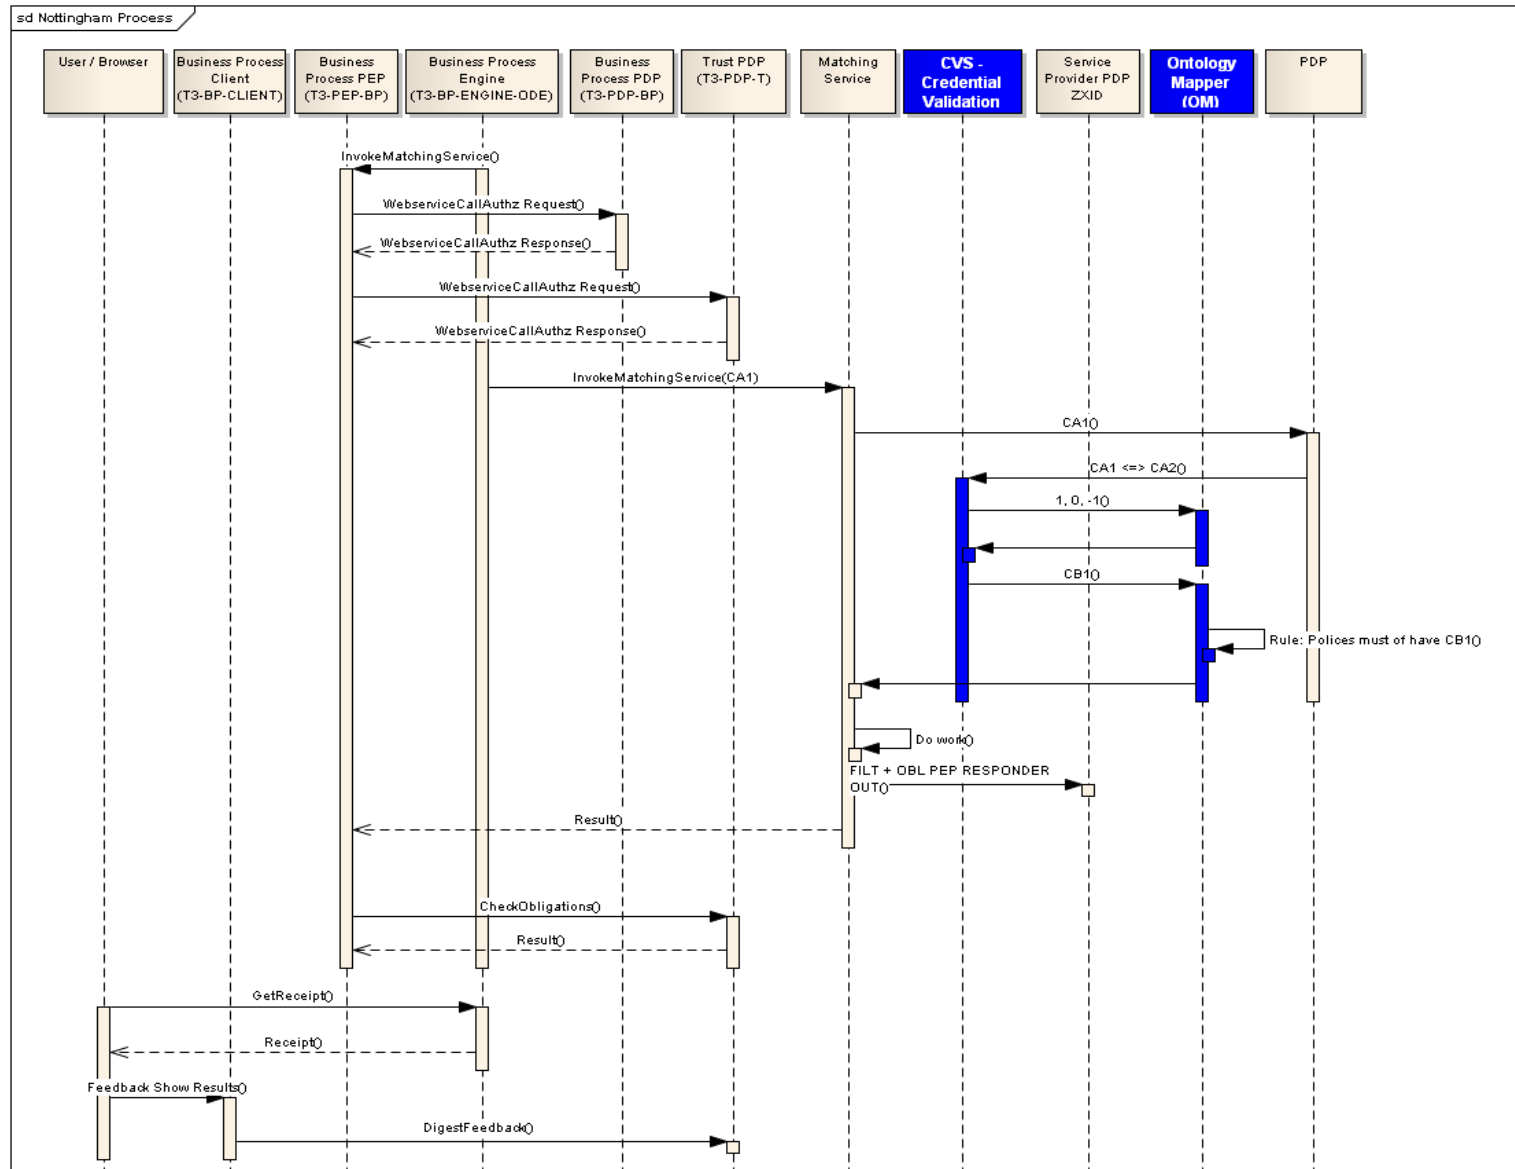**Figure 5 Negotiation and OCT**

Figure 6 Web service prime, call and end phase (including ontology use)

Overall the sequence is divided into four main phases, shown on the three diagrams above. The first is the SSO phase, where the user requests, and is either granted or denied access to, the application that is running on the TAS³ framework. The user at this phase is represented by information held by their home organisation in SAML format. This basic limited data will be used to identify the role that the student has at that organisation (e.g. undergraduate student, postgraduate student, etc). Once granted access, the student will present personal data and select his or her data security policies in the service selection and delegation (SSD) phases.

The SSD phase will involve the user being presented with a registration form which acts as an application form. As with a traditional application form, this will request personal data. This data can be input directly via a web form, or alternatively the user could give pointers to data held remotely (e.g. in an ePortfolio or PDS). As users are presenting personal data at this phase, this is the point where they select policies to protect it. This offers an improvement on traditional application methods as it offers finer granularity: for example different sections of a CV or ePortfolio view can be given different degrees of security.

If delegation has taken place, the CVS checks that the presenting user actually has authority to handle data in this context. The Negotiation phase takes the user's settings and checks the available service providers to see who can fulfil the requirements according to the user's preferences. It presents the user with a list of potential services to provide a job match based on the preferences and security/trust settings declared in their registration/application form. If no services can be presented at this point, the process can loop to allow the user to modify or add new security. Note that this process involves use of OCT to check the reliability of providers' services. This check continues after service selection to ensure the system is providing accurate services based on the user's preference. It presents the user with a list of potential services to provide a job match based on the preferences and security/trust settings declared in their registration/application form. If no services can be presented at this point, the process can loop to allow the user to modify or add new security settings and try again.

The final phase is the actual call to get the placement match. This consists of the transfer of personal data across domains and involves obligations and policy checks. As in other phases, calls to policy enforcement and decision points and application-specific calls will all be logged and directed to the user Dashboard. The results of the match will be returned to the user's PDS. Finally a receipt will be passed to the user giving a summary of the Dashboard logging data with links to log sources for more detail.

2.4 Future Work

While the eHealth demonstrator is trialling integration of TAS³ components into an existing system, the UK employability demonstrator continues to demonstrate how a new system and new approach to existing ways of online data sharing can be developed and demonstrated using TAS³ as a catalyst.

Future work in the final year of the project will extend this work into a pilot and focus on expanding the range of functionality in all areas.

This phase of integration effort builds on that from phase 1, which concentrated on standards compliance and testing of basic TAS³ components: the emphasis now is on demonstrating more advanced functionality as components develop and mature, and integration with a user-owned PDS.

The final phase of development will extend this further to incorporate a wider range of service providers, a larger number of end users and, if feasible, investigate application of TAS³ to international data exchange involving the NL employability demonstrator. In terms of storyboard, for this final phase the process will be demonstrated for a vacancy provider joining the system and securing vacancy information prior to it being matched. This will illustrate how other users can take part in the application, thereby bringing in user data from various sources and different groups.

As with the healthcare demonstration, the scope of the final integration phase and its technical implementation is heavily dependent on the TAS³ components which will be available at the time of planning. In terms of user centricity, in phases 3 testing and further development on the interfaces will take place following rigorous engagement with users. This involvement will often consist of monitoring, but in cases such as trust negotiation we expect to explore how interested users can influence the process in real time.

To support this increased user centricity the flexibility of both the Dashboard and other audit/event notification mechanisms is being improved upon, alongside further improvements in workflow adaptability, will all be provided in consultation the project's new usability expert. As in the case of negotiation and OCT testing we will aim to make the user aware in real time of what is happening with the workflow, services selected and user data in the system.

In terms of technical development around policy and data security efforts to increase the level of technical integration between standards and software will continue. The work using ontologies is critical to this development as is the implementation work using ZXID and PERMIS.

3 Employability NL Mass Layoff

3.1 Legacy Application integration

3.1.1 Introduction

Mass layoff is the term used to describe the process of ending employment contracts of more than 50 people within one organisation.

The process starts with an announcement by the employer. The information has to be sent to public authorities and representatives of the workforce. Procedures are regulated in a European Directive and in national law, jurisdiction and collective labour agreements. Typical reasons for terminating a large number of (long-term) contracts at once are changes in strategy and markets followed by a decision to reorganise, economical decrease, or shut down.

A mass layoff is not just an economic action. It also breaks a psychological contract between employer and employee. Labour relations can be described as partial gift exchange. Not only do employers invest in employees and pay salary, contributions for social security and training: employees also invest a lot in a job to build up their careers. Career perspectives within the company are drivers of motivation and performance. A typical outcome of a mass layoff is abrupt ending of career perspectives and income security. Damage control can be a lever for negotiating with public and private stakeholders about mobility and outplacement services.

Mass layoff seen as forced mobility from job-to-job or a special type of outplacement can be more or less effective depending on the pace and quality of information exchange including the assurance of legal aspects and securing privacy and trust.

Most important is to provide valid and reliable career information and job descriptions on which a personal competency profile and mobility plan can be based for all employees involved. The translation of these job descriptions and core tasks is the key in successful searching for new jobs and matching on competencies. In the Netherlands these Employability Services are provided by the specialists of Centres of Expertise like Kenteq.

In the Netherlands, both private and public employability providers work together to move redundant employees from job to job without having to rely on the social security funds. In the use case Job seeker we follow an employee who must find a another job in six months. In the search for another job, the job seeker meets different organizations; each organization needs his personal information to conduct their services. In this chain of services personal employability information is supplemented each time the information is exchanged between the service providers.

The TAS[3] trusted infrastructure consists of integrating the individual components, to ensure that the information is transferred securely and exchanged reliably.

3.1.2 Objectives of the use case

The objective of the use case is to prove that the TAS[3] trusted infrastructure can perform in a realistic scenario for which no current system exists. We aim to demonstrate that TAS[3] can work with a variety of existing national legacy systems to

address a real problem. Reliable access to transferred of legacy data is therefore a major driver for this scenario.

The scenario demonstrates the following areas:

- Usability of the graphical user interface
- Single sign-on across multiple real-life legacy systems
- Policy management
- Data discovery
- Involvement and audit of multiple Service Providers
- Legacy data integration
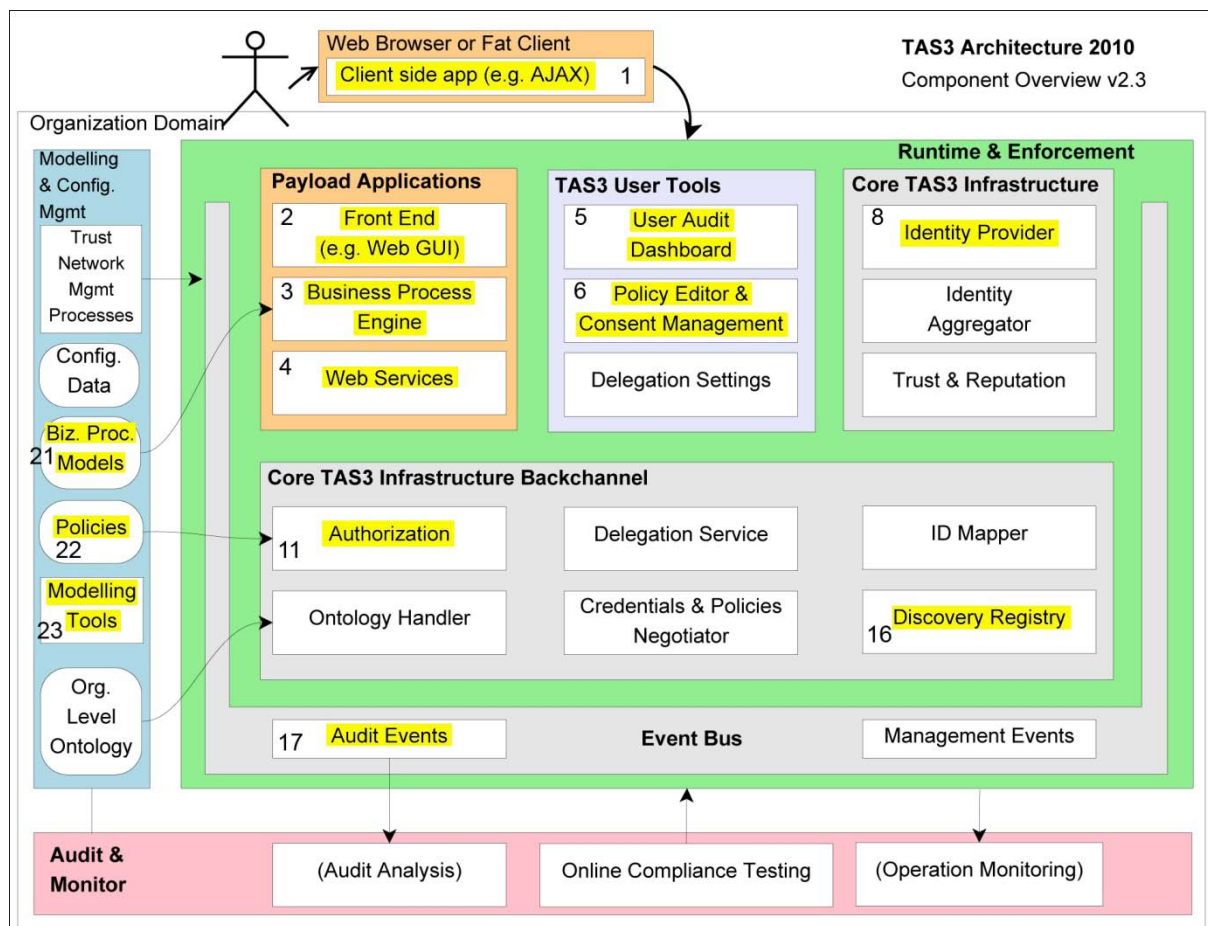- Improved Dashboard functionality.

3.1.3 TAS³ Components

Figure 7 TAS3 architecture diagram showing components used

The following components are used in this use case (numbers correspond to highlighted areas on the diagram in **Error! Reference source not found.**):

1. Web browser

2. Web GUI (Competent as skill and attitude measure instrument, the ePortfolio and the vacancy service)

3. Business Process Engine (To orchestrate the job seeker process)

4. Web services (SOA Gateway used to connect legacy applications)

5. TAS³ Dashboard (Job seeker can see audit information)

6. Policy viewer (Job seeker can view default policies)

8. Identity Provider (We use TAS³ IDP)

11. Authorization

16. Discovery service (To discover the Service Providers that hold relevant personal data)

17 Audit Events (To audit actions)

23. Modelling tool

3.1.4 Use case components

4. SOA Gateway

The SOA Gateway is the connector to create the web services between the legacy applications and the TAS³ infrastructure. This component is outlined in section 3.2.3.

8. Identity Provider (IdP)

The IdP provides Single Sign On capability (SSO). SSO provides access control to multiple software systems who are all part of an agreed security federation. The TAS³ IdP is provided by ZXID. Components who wish to implement SSO must redirect to the IdP website, where users will be requested to provide sign-on credentials. Once signed in successfully, they are redirected to the calling components where the application specific login procedures can progress. Further details on SSO and the ZXID IdP can be found in deliverable D2.1.

16. Discovery Services

The Discovery Service allows clients to discover which Service Providers hold the job seeker's personal data. Each SP must register its service with the IdP. Once the client has performed SSO, there is the ability to list all the services available at this IdP. Once the client has found the service with the appropriate data, they can use ZXID interfaces to retrieve this.

17. Audit Events

Audit events are handled by the Audit Service. This allows clients to relay messages regarding the trust fabric to a central place within the TAS³ network. More details on the Audit service can be found in deliverable D8.2.

3.2 Description of the Use Case

3.2.1 Basic storyboard

The employability NL pilots are defined as employability solutions for the individual career of an employee or f group of employees in a threatened situation. In all pilots we use employability services from Kenteq, the NL Centre of Expertise in the technical sector and test trust and security of shared services based on the use cases. In the pilots the employee is central.

TRIPOD is a Dutch cooperation of employability partners, with the goal of improving service and user centricity. TRIPOD has agreed on using the NTA 2035 standard for data exchange from ePortfolios (http://www.nen.nl/web/Werken/NTA-2035-Eportfolio-NL.htm) and the use of the same tables and vocabularies by all partners. TRIPOD includes the following service providers:

- NewCarFactory   : HR information
- Kenteq   : PCP/APL provider
- Paragin   : ePortfolio provider
- UWV-werk.nl   : Vacancy provider

The following actors are involved:

- Job seeker: redundant employee
- TRIPOD: consortium organizing employability services
- Kenteq: partner in TRIPOD which provides assessment
- Paragin: partner in TRIPOD which provides the ePortfolio
- Werk.nl: partner in TRIPOD which provides job vacancies.

A precondition in the use case is that the TRIPOD organizer has already registered all actors at ZXidp for Single Sign On services.

A brief overview of the process in nine steps (NB these are linear, each step needs to be completed before the next one can begin):

1. The company ceases production and the employees are faced with redundancy in 6 months
2. The job seeker does an intake at TRIPOD (consortium of employability providers)
3. His HR data from the company is exchanged with the employability provider (Kenteq)
4. He does an assessment which results in his Personal Competency Profile (PCP)
5. His PCP from the employability provider is exchanged with his ePortfolio
6. He gets access to a Vacancy data provider (database with Vacancy Competency Profiles -VCP)
7. He searches for a suitable vacancy (Match between PCP and VCP)
8. He applies for a job
9. He checks whether his data policies have been complied with

3.2.2 Use case actors

Table 1 describes all actors involved in the Job seeker use case.

Scenario Actor	Person	Task
NewCarFactory	HR Manager A	Provides HR data
Job Seeker (End user)	Dirk Brown	Find a new job
ID provider (ZXidp)	ID coordinator	Provide an ID for SSO to

Scenario Actor	Person	Task
		Dirk Brown
TRIPOD	Organizer	Organize TRIPOD services
PCP/APL provider (Kenteq)	Administrator (Kenteq) Assessor (Kenteq)	Provide PCP services Execute the PCP assessment Provide PCP data
ePortfolio provider (Paragin)	Administrator (Paragin) Job Coach (Paragin)	Provide ePortfolio services Coach job seeker
Vacancy provider (Werk.nl)	Administrator (werk.nl)	Maintain and provide Vacancy database
New employer	HR Manager B	Search for new staff

Table 1: Use case actors

3.2.3 Legacy Architecture

Figure 8 shows the legacy architecture of the Job seeker use case. The end user has access to the TRIPOD portal, and can access TAS³ facilities via the web front end. He can set his policies, view information on the Dashboard and access TAS³ certified services. Authentication is done by the Identity Provider t ZXidp. Through a single sign on, he has direct access to services of TAS³ external applications such as the PCP assessment system, the ePortfolio system and the Vacancy Management system. All the employability service providers exchange information using web services.
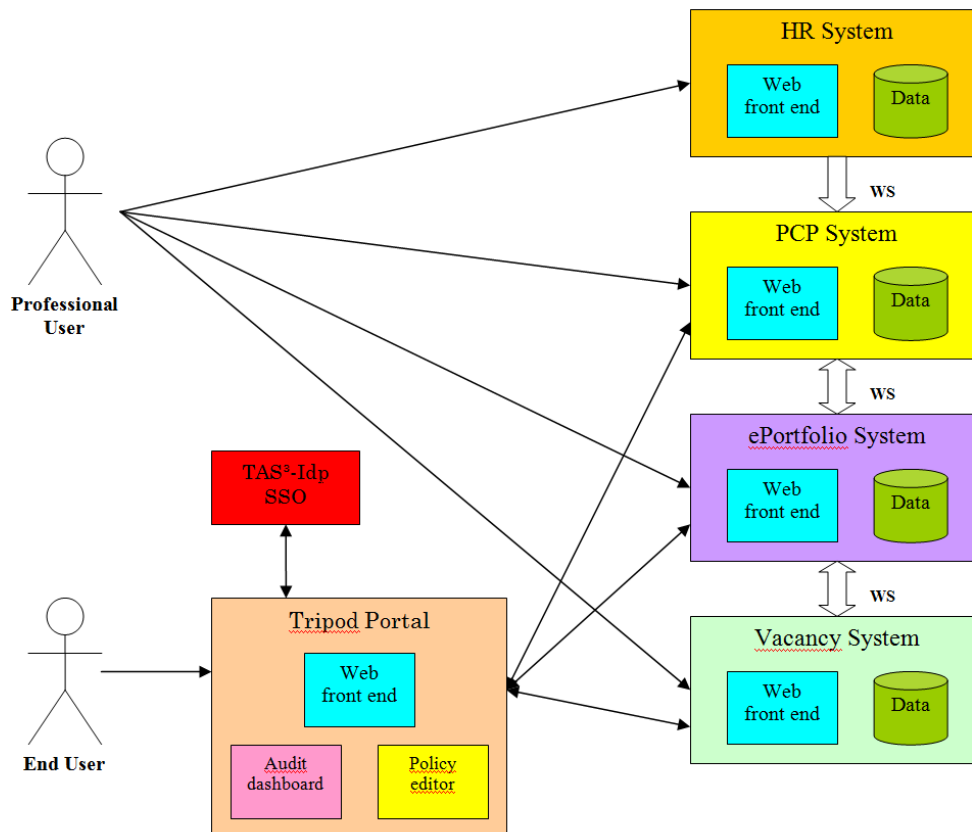
Figure 8: Legacy architecture

3.2.3.1 SOA Gateway

The SOA Gateway is a software tool that makes it possible to expose data to new or existing applications. It enables access to data from a wide range of database languages without the need for server-side code or middleware. The SOA Gateway uses multiple different drivers to create web services: depending on the amount of metadata available for a given database or programming language, the SOA Gateway Resource Discovery wizard can automatically define the components required to wrap a database or application program as a web service.

The SOA Gateway works by inspecting the backend catalogue or language and creating a Resource Definition file. This file is used to map back end columns/parameters internally with upfront web service definition parameters. The resulting file is highly flexible and can be edited by administrators to suit their particular needs and requirements. No changes are required to existing database schemas or language structures, thereby making integration with existing legacy applications a completely transparent process.

The SOA Gateway integration as a Service Provider within the TAS³ architecture is at the web service level. An additional T3-SG-WSP component which uses ZXID has been developed by the TAS³ project. T3-SG-WSP allows clients to retrieve the metadata for any web service by providing the 'o=B' parameter as an argument on the web service URL. T3-SG-WSP also ensures that Web Service Client requests are validated to ensure 100% compliance with TAS³ specifications and security credentials. On validation, the raw request is provided to the SOA Gateway for processing. T3-SG-WSP can accept a security policy definition and ensure that this policy is applied to the data by making a request to an internal or external PDP. Finally, T3-SG-WSP adapts the web service

response to include the necessary TAS³ security credentials, thereby allowing the Web Service Client to validate the response.

The full process is as follows:

1. Use the SOA Gateway Discovery Wizard to create web services from existing data

2. Retrieve the metadata of these services, e.g. http://host/myService?o=B

3. Register the metadata in the Circle of Trust area via one or more Identity Providers

4. Provide the Web Service Client with the Entity ID of the registered service. This entity ID can be used directly or in conjunction with the IdP for ID-WSF 2.0 service discovery

5. Web Service Clients can use this IdP to perform SSO

6. When the T3-SG-WSP receives a web service request, it validates this using the TAS³ security protocols

7. If the request is valid and policy checks are required, request the PDP to validate the request against the policy for this service

8. If the PDP confirms the request is valid, request the SOA Gateway to handle it

9. If the request has been handled successfully and policy checks are required, request the PDP to validate the response against the policy for this service

10. If the PDP confirms the response as valid, wrap the response with TAS³ security protocols

11. Send the response to the client.

Further technical details on the SOA Gateway can be found in deliverable D8.1.

3.2.4 Use case sequence diagram

The following diagram shows the collaboration between the systems over time. This co-operation is represented by sending messages. The nine major steps described in 3.2.1 are represented in the major sequences in figure 9.
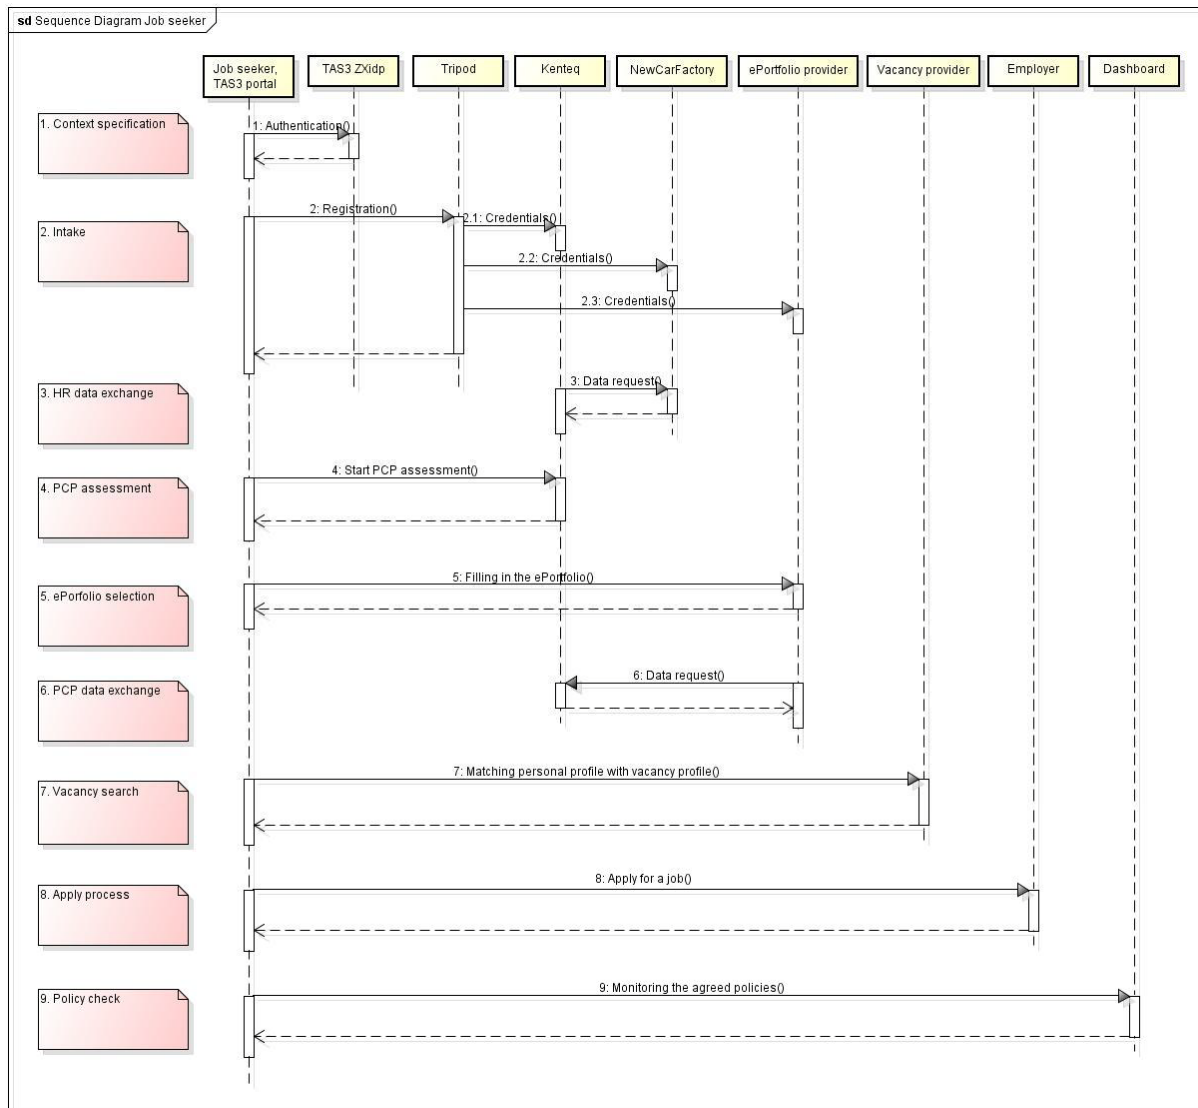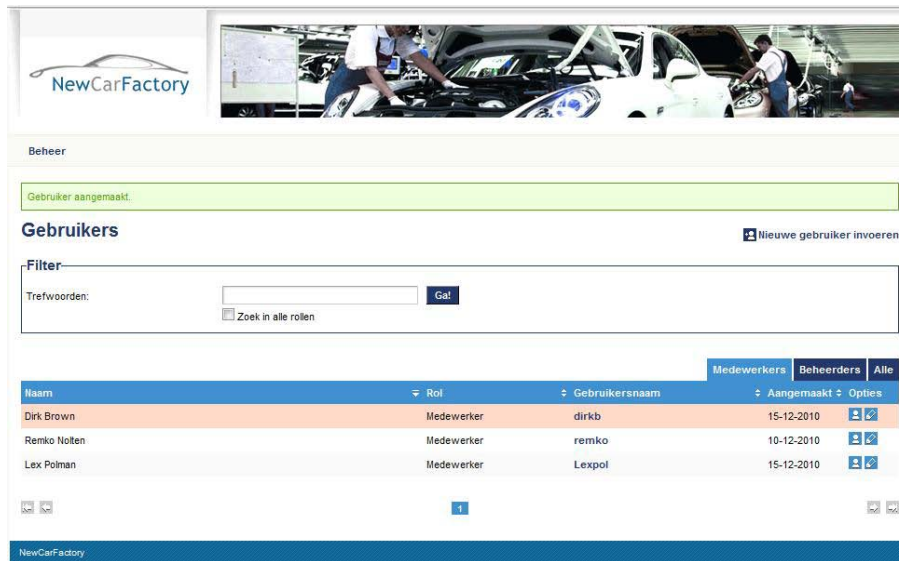
Figure 9: Overall sequence diagram for Job seeker process

3.3 Use case scenarios

From the storyboard we describe scenarios which the TAS³ infrastructure should be able to support in order to have a trustful data exchange of personal employability information. Added value from TAS³ in this use case is derived from trustful exchange of personal data, SSO for all services and seamless steps between services where personal data can be reused.

Dirk is 41 years old and has a basic education.

- He has 15 years' experience at NewCarFactory in various positions

- NewCarFactory stops production of a certain model of car and Dirk is faced with redundancy in six months.

- The HR manager of NewCarFactory sends basic information (names and email addresses) about redundant staff to the TRIPOD organizer.

- The TRIPOD organizer completes the contract data and basic information about Dirk on the Competent system.

- IdP accounts are created for all redundant staff.

Figure 10: NewCarFactory - Dirk's current employer

3.3.1 Context

Figure 11: TRIPOD portal

- Dirk is made redundant by NewCarFactory
- He goes to the TRIPOD website
- He authenticates via TAS³ Idp SSO

3.3.2 Intake

Figure 12: Intake process for new users

- Dirk starts on the TRIPOD portal

- Dirk accepts TAS³ terms and conditions

- He selects a PCP Provider (Kenteq)

- Dirk accepts terms and conditions from Kenteq

- Dirk also accepts the default policies for the use of his personal data by Kenteq

- Dirk is asked if he wants to search for his existing personal data in the TAS³ system

- He gives consent to transfer his HR data from NewCarFactory to Kenteq

3.3.3 PCP Assessment

Figure 13: Competent application for PCP assessments

- TRIPOD forwards Dirk to the Kenteq PCP system (SSO action) to carry out his assessment

- Dirk corrects or completes his data in the PCP system

- Dirk completes the PCP assessment

- The result of the assessment is a Personal Competency Profile (PCP)

- The assessor provides the PCP results to Dirk through the Competent interface

3.3.4 ePortfolio

Figure 14: Paragin ePortfolio

- Dirk receives an email from Competent informing him that his PCP assessment is complete

- He selects an ePortfolio Provider (Paragin)

- Dirk accepts the Paragin (Toegang tot Techniek) terms and conditions

- Dirk also accepts policies from Paragin

- Dirk is asked if he wants to search for personal data in the system again

- He gives consent to transfer his PCP data to his ePortfolio

- Dirk is forwarded to Paragin (SSO action) to check his ePortfolio data

3.3.5 Vacancy search
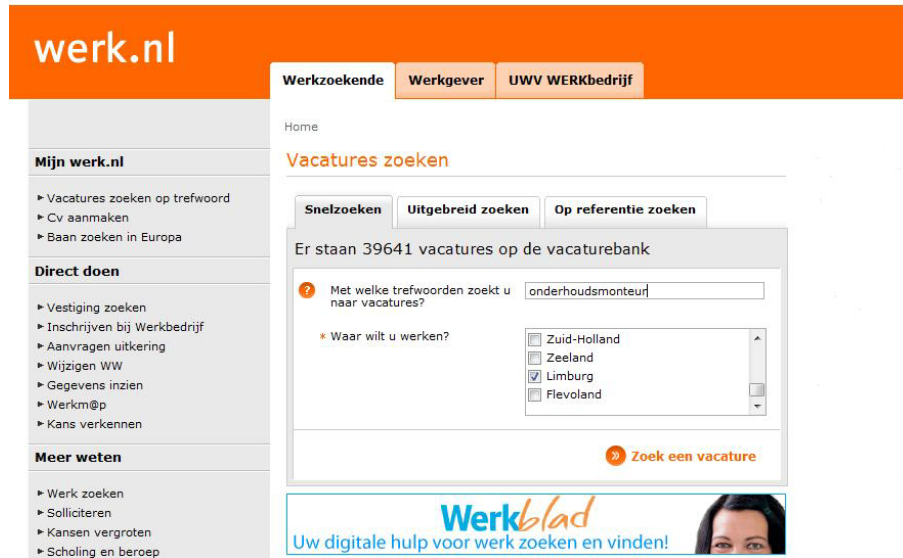
Figure 15: werk.nl website

- Dirk selects a Vacancy Provider (UWV)

- Dirk accepts UWV-werk.nl terms and conditions

- To proceed with the TRIPOD business process Dirk has to accept specific policies on his data required by the vacancy service

- He completes his search data

- Werk.nl matches his PCP with the Vacancy Competency Profile (VCP)

- Dirk gets the results of the vacancy search

- The results consist of well-matched vacancies

3.3.6 Applying
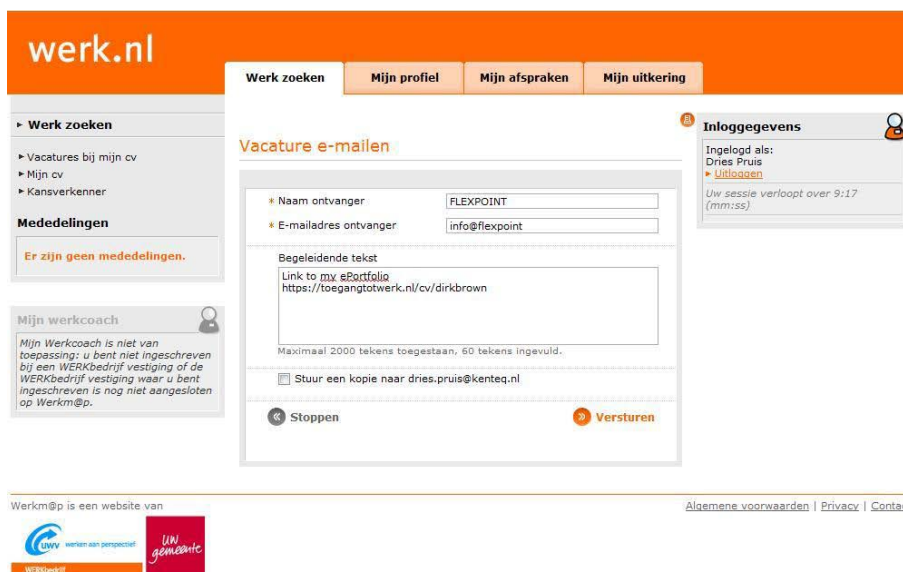
Figure 16: Job application

- Dirk selects a vacancy and applies for that job

- He sends a PDF of the showcase from his ePortfolio to his new employer
- He accepts a suitable job at a mechanical company and signs a contract.

3.3.7 Policy check
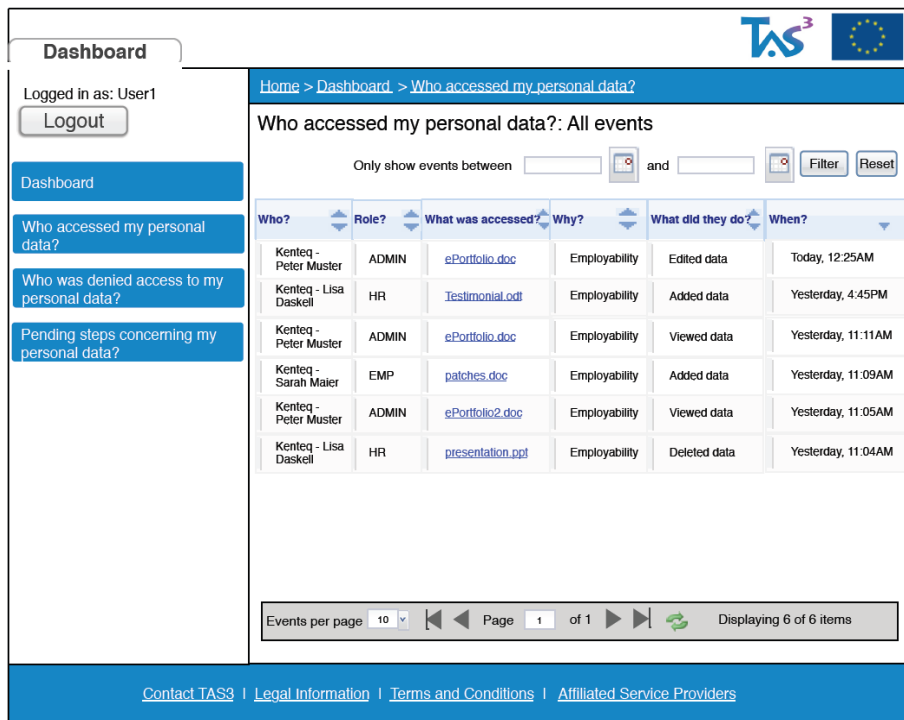
Figure 17: TAS³ Dashboard

- Dirk checks whether the service provider has complied with his policy setting at the TAS³ dashboard

3.4 Future work

In the final year of the project we will expand the use case to encompass multiple simultaneous end users and focus on expanding the range of functionality in all areas of employability. We will integrate further components developed in the project and we test these in a variety of use cases. We will demonstrate more advanced functionality of the TAS³ architecture while continuing to focus on the needs and wishes of the end user.

TAS³ features we expect to demonstrate include

- Compliance testing for service providers.
- Trust negotiation
- Discovery services
- Policy management
- Dashboard functionalities

We also plan to create a practical demonstration of a use case focusing on Accreditation of Prior Learning (APL) as described in the previous version of this deliverable.

We expect to expand our activities to include a wider range of service providers and a larger number of end users, with the potential for exploring international exchange of

personal information relating to the UK employability demonstrator and combining the current UK work placement pilot with the method used in the Netherlands.

Furthermore, we will improve the usability of interfaces in consultation with a diverse group of end users. The involvement of the end user is particularly important for the technical design of the Dashboard functionalities.

4 Healthcare Integration Trial

4.1 Introduction

4.1.1 Phase 1: Improved Information Exchange

The drivers behind the current eHealth evolution have been sketched in detail in D1.1 and D9.1 v1. In particular, the vision of providing 'Continuity of Care' as a cornerstone for improving quality of care has been explained. Continuity of Care requires Health Care Professionals (HCPs) to have at any point in time access to the relevant medical history of a patient. This requires a comprehensive follow-up of patients from cradle to grave.

In year two of the project, the healthcare integration trial has shown TAS[3] enabling improved information exchange between HCPs in the eHealth domain (cf. D9.1 v2). The trial scenario in deliverable D9.1 v2 was staged in the Belgian healthcare environment (which is comparable to that of many other EU countries). In Belgium, a number of hospitals provide access to their Hospital Information System (HIS) or a results server to professionals (essentially GPs). Patient access is currently still very rare (except possibly for downloading medical images). Exchange of medical data is mostly based on 'documents', which is reflected in many electronic data formats and in the use of result servers which contain episode reports (such as a lab outcome, a radiology report, a medical consultation report, etc). In the trial, data was exchanged using the KMEHR (Kind Messages for Electronic Healthcare Record) XML standard[6].
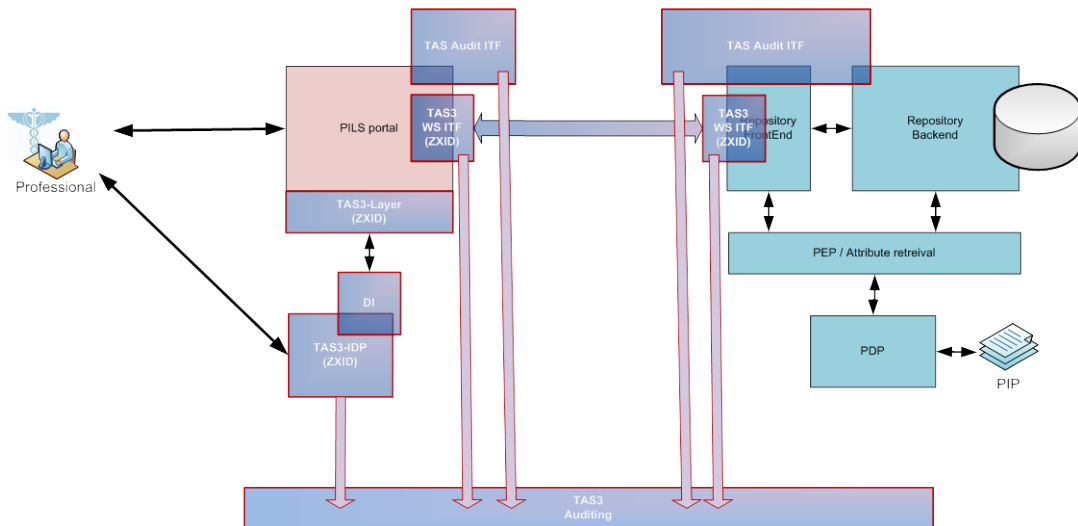
Figure 18: Phase 1 integration trial services (D9.1 v2)

The legacy PILS/Repository application used for the phase 1 integration trial (cf. D9.1 v2) demonstrated at the year 2 Review meeting consisted of four independent services deployed in a test setup at the Custodix Data Centre:

- PILS portal

___
[6] See http://www.chu-charleroi.be/kmehr/htm/kmehr.htm for specification. This is a Belgian national standard, it is not used elsewhere.

- Identity Provider Service

- Two repository Services, filled with dummy data

In phase 1, the aim was to replace basic existent security functionality with TAS³ compliant implementations (cf. Figure 18). The focus was on reaching compliance with the TAS³ protocol stack, rather than on integrating the extended TAS³ security functionality (for which implementations are not yet available) into the application workflow.

4.1.2 Phase 2: Patient Empowerment

The vision of TAS³ is to deliver an environment where end users have full control over their personal data. Whereas in the first demonstration phase TAS³ has been mainly demonstrated as an enabler for building secure architectures, phase 2 focuses on the introduction of user-centricity in the management of personal information. As a consequence, the centre of gravity for this deliverable shifts towards the authorisation part of the architecture.

The roadmap for introducing user centricity in the health environment and demonstrating full TAS³ capabilities was presented in the second project review in March 2010 using Figure 19.
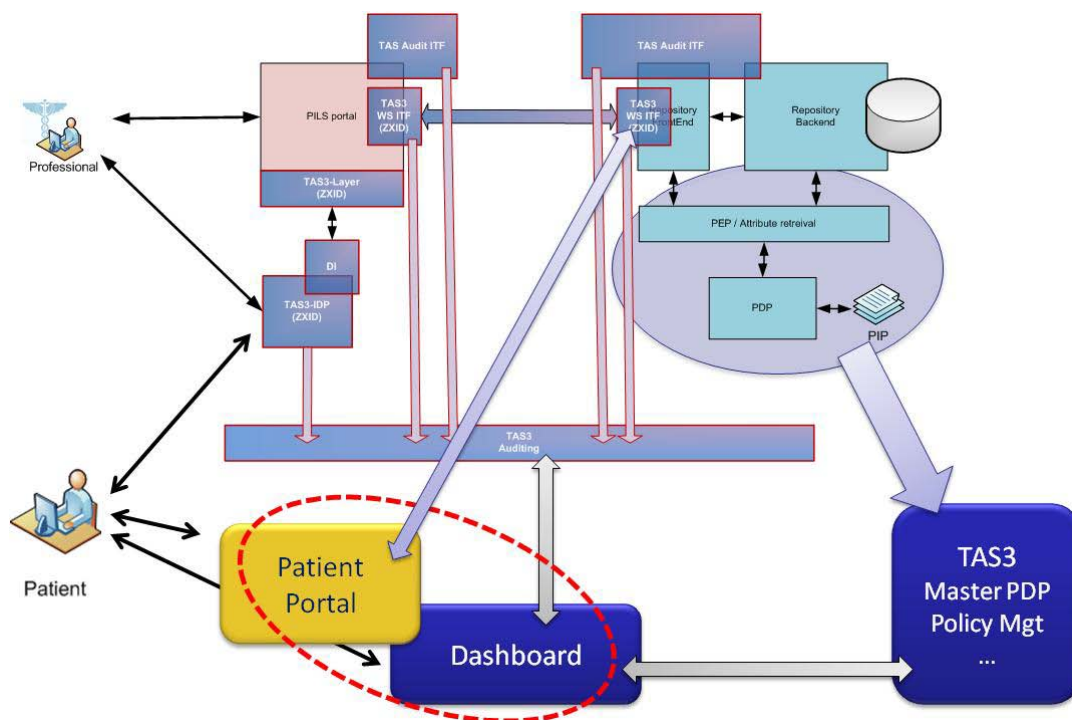
Figure 19: TAS3 eHealth ecosystem roadmap (phase 2)

4.1.2.1 User-centric Data Management

The first question that needs to be answered is: what exactly does TAS³-enabled user-centric personal data management mean within the eHealth environment?

The main reason behind an interest in user-centric data management is the fact that it is an important aspect of patient empowerment (cf. D1.1, D9.1 v1). Obviously, data protection in eHealth is not a green-field area. Common data flows that are required for

good functioning of healthcare are well defined and regulated by sets of laws and regulations. These laws and regulations (are supposed to) represent a least common denominator of what society expects of healthcare organisations and practitioners when they are treating their patients' sensitive data.

With the introduction of user-centric personal data management, patients should be able to express their personal preferences (differentiating from the default policies) as long as they are aware of the impact of their own choices. This end-user control is, however, not (and should not be) absolute: the societal and ethical framework puts some limits on the personal decision power concerning data exchange within healthcare (as in all domains). For example, one cannot expect a physician to be responsible for a treatment if a patient refuses to disclose relevant medical information. Or, as another example, in many cases community benefit is deemed higher than the personal benefit and disclosure of data for research is mandatory (e.g. country disease registers).

Finally, it should be noted that empowering people (with respect to managing their own data privacy) certainly does not mean that they will restrict access to their information as much as possible, and in such a way hamper established business processes. On the contrary, if properly implemented, a system for personal privacy policy management could facilitate requesting consent (by making it more convenient, by establishing more trust) of patients for more elaborate data sharing for specific purposes (e.g. research projects).

In summary, user-centric personal data management in the eHealth ecosystem requires:

1. A convenient way for patients to adjust the default eHealth domain policies that are determined by legislation and ethical guidelines, so that their own data is governed according to their personal preferences on data protection.

2. A system capable of dealing with specific (i.e. a limited domain) requests for data processing (consent). This should bring benefit to both consent requestors and patients.

Key to the success of user-centric personal data management is the usability of the system. Experience has shown that even people that take genuine interest in privacy and care about controlling their own data, easily reach the point of configuration-fatigue (a point where one gives up on configuration because the number of options offered is simply too high and expressing personal preferences becomes too cumbersome and complex). This is especially a problem with security and privacy settings, because for one these settings are not immediately related to functionality and thus not of primary interest. Secondly, there has not been sufficient research on facilitating authoring and enforcing of fine-grained access policies. In this integration trial, the risk for configuration fatigue is reduced by the fact that policies are specified through modifications of a default, and that options for modification are limited.

Further, usability is augmented in the trial by using the TAS³ environment to construct a central point for managing personal preferences for data in the health domain. On this central Dashboard an end user can both manage the 'static' personal preferences and specific consent requests and at the same time get an overview of who has accessed and processed his data.

It should be noted that although management is centrally orchestrated from the end user's perspective, it does not need to be implementation aware. In the same sense, the Dashboard does not need to be a unique instance. First of all, the Dashboard can be a portal uniting all functionality concerning personal data policy management

implemented by different service providers. Secondly, there could be several Dashboard providers (possibly with their own look and feel) from which end users could choose.

4.2 Integration Trial Environment Scenarios

4.2.1 Setting

The environment for the integration trial is illustrated in Figure 2020. The setting extends phase 1 of the integration trial that was presented in D9.1 v2 (and during the 2010 project review).

In the healthcare integration trial environment, there is a service Provider hosting a 'PILS Portal' which can be used by professionals to look for patient information on repositories which are registered in a Medical Circle of Trust (i.e. a CoT formed by medical service providers who want to give uniform access to HCP). In this demonstration two repositories have been connected:

- One hospital repository (containing mainly discharge information)

- One Summary Repository (containing summary records originating from primary care)

There can be multiple Identity Providers. In this trial, they are all authoritative with respect to unique user identities and unique healthcare professional identifiers. This simplified approach to identity managed is largely modelled according to the Belgian situation (cf. D9.1 v2 for details).
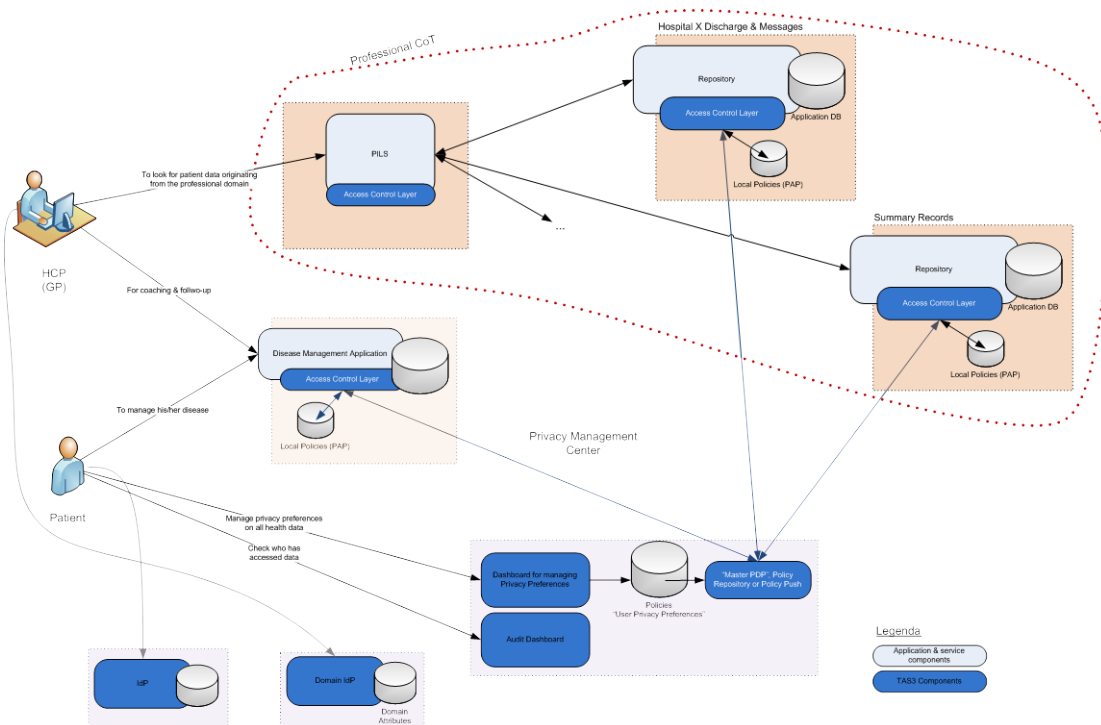

Figure 20: 2011 integration trial setting

This trial scenario aims to demonstrate TAS³ as user-centricity enabler, hence there is a central point ('privacy management centre') where the end-user (i.e. the patient) can control his personal preferences regarding data protection and check audit logs of whom accessed his data.

A central point for managing policies in a distributed environment can be implemented in many different ways. Policies (expressing personal preferences) can be generated in the central dashboard and made available for download by the different distributed service providers or the policies could be pushed towards them. The privacy centre could implement a PDP (Policy Decision Point), meaning that service providers would forward access requests to the central PDP. Equally, the dashboard could be a portal towards local privacy preference 'configurators' at the local service providers. For sake of simplicity and easy of demonstration, the dashboard infrastructure acts a PDP in this integration trial (see further).

The presence of a third party application that is not part of the 'professional healthcare environment' is new in the phase 2 scenario. This application represents one of the many independent service providers which offer health-related services to patients. In this particular trial, the service provider offers a disease management application for diabetes type 1 patients (explained below). In part, as it also provides guidance and coaching, the application serves a partial, disease-specific, Personal Healthcare Record.

4.2.2 Scenarios

The overall demonstration scenario is an extension of the 'information retrieval' scenario explained in the earlier D9.1 deliverables. In the phase 2 trial, technical aspects of the TAS³ components and the enhanced user-centricity provided by the TAS³ combined architecture is demonstrated.

The phase 2 scenario consists of three use cases.

4.2.2.1 Use Case: Phase 1 Revisited – User-Centricity

Storyline: *'Homer visits his general practitioner (GP) after an episode at the hospital where he was admitted because of breathing problems (emergency). His GP wants to find more information about Homer's medical condition, the diagnosis and treatment received at the hospital, medication prescribed and further follow-up needed.'*

Homer's GP accesses Homer's discharge information through the PILS. Governing policies are now a composition of the third party policies (law, local hospital regulations) and Homer's preferences.

- This demonstrates basic policy interaction and dashboard functionality for configuring user preferences.

4.2.2.2 Use Case: Central Control – User-Centricity

'Homer is a diabetic patient. Recently he discovered a disease management application on the internet which allows him to keep a diary about the disease. Homer uses this site as a diabetes diary as it provides benefits over his paper diary (cannot get lost, provides more detailed logging, provides relevant personalised information, etc.). As with the paper diary, Homer wants to share the recorded information with his diabetologist during consultations.'

The diabetes management site integrates with the eHealth policy platform (as it is TAS³ compliant and speaking the correct 'eHealth domain vocabulary' for access requests). Therefore access to the platform by healthcare professionals is automatically governed by the personal preferences of patients, without them needing to configure sharing rights specifically for this one site. Moreover, healthcare professionals do not even need to be registered to the platform.

- This demonstrates the benefits delivered by the TAS³ architecture for making individual control over personal data work in a large ecosystem with minimal effort.

4.2.2.3 Use Case: Consent Request

'A new clinical trial on diabetes is started. The clinical researchers have difficulties finding patients which fit the inclusion criteria. They believe that looking for patients through the different diabetes disease management service providers could offer a solution. One of the biggest concerns in executing this plan is ensuring that their actions are in line with the data protection legislation.'

The TAS³ framework offers a solution to automatically handling the complete process of obtaining consent for scanning the data and for accessing it for research, within the boundaries of legislation and governing ethical frameworks.

- Illustrates the possibilities of centralised privacy preference management for dynamic consent questions (i.e. the basis of an e-consent system).

4.2.3 Demonstration Ecosystem

For the integration trial a dummy health ecosystem has been created. It is in principle similar to the existing health environment in Belgium (and many other countries). A number of assumptions have been made which make the practical implementation easier. The implementation is for demonstration purposes and does not represent the best practical solution in terms of architecture.

4.2.3.1 User Privacy Preferences

In the ecosystem, a central PDP produces access control decisions based on default governing policies and patient privacy preferences. The central PDP deals with granularity at document level. These 'documents' are defined according to the default data exchange that takes place in the domain. In the trial the set of documents used is very limited, but in reality this could be very broad[7]. Every service provider willing to be part of the eHealth information exchange ecosystem needs to follow a number of rules:

- Service providers need to be aware of the scope of the central PDP access decisions (defined data sets on which the PDP needs to be contacted). For every access request that falls within this scope, they need to forward the request to the central PDP.

- Service providers are responsible for translating their local context into a request that uses the common policy vocabulary[8] of the eHealth domain (which the central PDP uses)[9].

---

[7] Cf. several medical terminologies.

[8] 'vocabulary' refers to how the request metadata is communicated. E.g. in XACML terms: what the subject, resource, action, etc. attributes used are called and what their content is.

- When the central PDP does not provide a decision, local policies govern.

In order to fully specify the ecosystem in accordance with governing legislation and best practices, it would be necessary to determine the detailed rules on when (and with what conditions) service providers can override central PDP decisions. However this is beyond of the scope of this demonstration.

Access requests formulated by a service provider need to contain the following information in order to be accepted by the central PDP:

Information concerning[10]	Type of Information	Value set	URN reference[11]
Subject	Identity of the subject requesting access	Social Security ID (cf. D9.1 v2)	:person-id:insz-niss
Subject	Healthcare identification number of the subject requesting access	Healthcare Professional ID (cf. D9.1 v2)	:hcp-id:riziv-inami
Subject	Type of HCP	HCP type KMEHR code	:hcp-type
Action	CRUD operation	create read update delete	:action
Action	Purpose for which the action is executed.	e.g.[12] TREATMENT, PAYMENT, OPERATIONS, EMERGENCY, SYSADMIN, RESEARCH, MARKETING, REQUEST, PUBLICHEALTH	:purpose
Resource	Identity of the data subject	Social Security ID (cf. D9.1 v2)	:patient-id
Resource	In principle all HCPs directly involved in	Healthcare Professional ID (cf. D9.1 v2)	:treating-hcp-id:riziv-inami

---

[9] This translation could be generically covered by the TAS3 ontology services (which are not yet available).

[10] The XACML vocabulary is used. An access requests means that a 'subject' requests permission to perform an 'action' on a 'resource'.

[11] URN reference in the domain wide vocabulary used in the trial. Prefix is urn:custodix:tas3:ehealth.

[12] From 'PurposeCode Value Set Definition' (OASIS XSPA Profile of SAML)

	the treatment this resource is about. However, it is only relevant for the policy if the request 'subject' is involved or not.[13]		
Resource	The resource type (broad classes as specified by the default policies)	e.g. discharge, summary, medical diary	:resource-type
Resource	Note that the resource contains information on specific sensitive subjects	ETH HIV PSY SDV	:sensitivity
Other	Specific situations	emergency out-of-office	:extraordinary-circumstance
Other	Date and time of the request		

Table 1: Content of an access request to the central PDP

The central PDP itself has access to several attribute information points that provide authoritative information for the complete domain. Two of these are:

- The GMD holder[14] database which lists for each patient who maintains their combined primary care record or GMD (if they have one).
  In the central PDP vocabulary this is represented by an attribute attributed to a physician 'urn:custodix:tas3:ehealth:hcp-relation:gmd-holder-of' containing a patient ID.

- The therapeutic relationship database[15], which records which physicians have a known therapeutic relationship with a patient.   In the central PDP vocabulary

---

[13] This attribute is only used to see if there is a direct relation with the requestor. It is undoubtedly better to define an attribute in that respect. However, the current implementation of the demonstration repositories makes this more difficult.

[14] The concept of the combined Belgian primary care record or 'Globaal Medisch Dossier' was explained in D9.1 v2.

[15] How such a database is maintained is outside the scope of this project, but relevant topic of discussion in real world setups.

this is represented by an attribute attributed to a physician 'urn:custodix:tas3:ehealth:hcp-relation:therapeutic-relation-with' which lists a patient ID.

The central PDP will base its access control decisions on a mixture of the default policies that are supposed to govern the eHealth domain and personal preferences. Policies can be defined based on the request content of Table 1 and the attributes that can be obtained from authoritative sources. As mentioned earlier, end-users cannot be given full liberty in defining their personal policies (see also further).

4.2.4 Consent Directive Requests

The above describes how the TAS³ infrastructure assists in enforcing user-defined privacy preferences over the complete health domain. However, these settings are static and define privacy policies for common data flows. There are many occasions on which personal data is needed outside these. One example is given in the 'clinical trial use case' (cf. section 4.2.2.3).

The TAS³ infrastructure allows dealing with dynamic requests to end users for specific consent directives (i.e. privacy settings). The mechanism that is used in the healthcare integration trial is illustrated in Figure 21.

When a candidate data user (or data provider) wishes to obtain specific consent directives, he will need to send a request to the central privacy management service. After validation, this request is forwarded to the end users concerned through their dashboard. The dashboard offers the end users the possibility of answering the consent request by stating their preference (possibly giving consent under specific conditions). The central management component can then transform the consent directive into a policy that is stored in the PAP of the central PDP. One of the difficulties is the fact that the 'vocabulary' used for writing the policies of the central PDP needs to be the same as (or be translatable in) the one used by the PEPs (or their context handler or local PDPs[16]) that forward the access request to the central PDP.

At present, TAS³ has not yet specified the mechanism of obtaining dynamic consent. The method by which a request towards the dashboard should be formulated and how this can be translated into a policy is not yet formalised. The following points of note are relevant:

- The consent request needs to contain the scope of users affected (so that the dashboard can determine who should be presented with it).

- Requests for consent need to be sufficiently specific (minimum requirements) and must contain a definition of the purpose of use of data, validity period, etc.

- The policies resulting from the consent directives can only be used by all data providers in the domain when they are described using a domain-wide vocabulary (i.e. so that local access requests can be translated by the providers into requests understandable by the central PDP). Nonetheless, in many cases consent directives are only relevant to a limited set of parties active in the domain; hence they could use their own vocabulary for specifying more complex rules than possible by default.

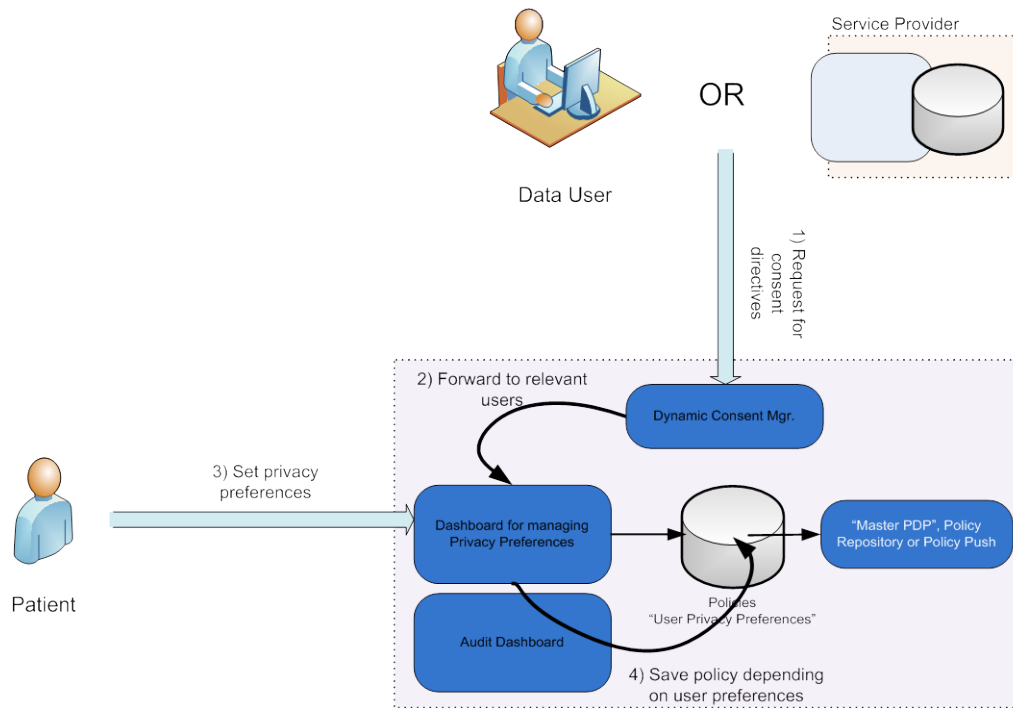[16] Depending on the implementation of the request forwarding mechanism.

Figure 21: requests for specific consent directives

In the integration trial, a simplified approach is taken. The complexity of the consent directive request is avoided by including the policy, that reflects the consent request, in the request itself in a form so that it is directly interpretable by the central PDP. The end user's decision on the consent directive request then simply determines the policy outcome (details, such as validity period, could be configured through the dashboard).

Although this approach seems naïve, it is possibly a good starting point for a generic solution.

4.3 Integration Trial Environment

4.3.1 Legacy Applications

4.3.1.1 PILS Portal and Repositories

The background and architecture of the PILS (Patient Information Location Service) portal and associated repositories have been described to great extend in D9.1 v2. Phase 2 of the trial uses the same legacy implementations of portal and repositories.

In summary the PILS provides federated access to patient information stored in different document repositories. This information locator is not an index server, but rather a search engine supporting distributed searching. The PILS application is oriented towards professional use, in typical information storage and retrieval scenarios. It has been demonstrated in the professional context in phase 1 of the integration trial.

With respect to the information repositories, two typical applications are use as a hospital results server and use as a summary record repository. A hospital results server would contain medical reports on episodes of hospital care. A summary server would contain 'summary records' created by primary care physicians. A summary record is, roughly speaking, a set of data that a physician needs in order to understand the medical status of the patient in just a few minutes, and to ensure continuity of care. The

main difference between the two from a security perspective is that a results document does not usually need fine-grained access control, whereas a summary record does.

The two demo repositories in the phase 2 trial are in fact a health record summary and hospital result server. In the phase 1 integration trial, access control in the repositories was handled locally, through an internal PDP. In phase 2, these repositories have been modified so that they interoperate with the central PDP.

4.3.1.2 Disease Management Application

4.3.1.2.1 Introduction

Empowering people in healthcare means putting them in charge of their own health[17] and implies that a large part of the responsibility regarding health management shifts towards the patient.

The need in healthcare to reduce cost on the long run (in view of the ageing population) and the desire to further increase quality of life has been explained in D1.1. Making a patient more responsible for his own health is one of the approaches which aim to save cost and at the same time provide better quality of life.

An obvious starting point in the empowerment process is to make patients important actors with respect to preventive medicine[18]. Involving people in the preventive care process by informing them better and giving them more decision-making power makes them feel responsible and more inclined to participate actively in managing their health status (e.g. adhere to a healthy lifestyle). It needs little explanation that the cost of coaching a person in such a preventive care programme is small compared to that incurred when this person needs to rely on the healthcare system for surgery or medication, or eventually becomes dependent on the social security because of inability to work.

Note that preventive care is certainly not restricted to primary prevention (avoiding the development of disease) but also, for example, deals with diagnosing and treating diseases in an early stage to reduce the negative impact of the disease or disease-related complications (disease management). The patient's responsibility in the preventive care process mainly relates to lifestyle management (typically food and exercise), therapy adherence and monitoring (either through devices or manually). Diabetes is a typical chronic condition in which proper disease management can seriously reduce the impact of comorbidities.

A common task in disease management processes for patients is monitoring and recording of relevant medical information (vital parameters, medication taken, etc.) for evaluation by physicians. When patients start to store this health status information electronically in Personal Health Records (PHR), this becomes a particularly interesting use case with respect to personal data exchange.

4.3.1.2.2 Trial Legacy PHR

The legacy application used in the TAS³ demonstration trial is a health status recording site for diabetes Type I patients.

---

[17] It does not mean assigning patients as quality controllers of the work of professionals. It is often misunderstood that way, both by professionals fearing review and by patients wanting to break free from a feeling of powerlessness.

Diabetes is a group of metabolic diseases in which the body does not produce enough insulin[19], or cells do not respond sufficiently to the insulin that is produced. This leads to high blood sugar levels, which on its turn leads to (when not adequately treated) damage of blood vessels, heart, kidneys, eyes, nerves, and other tissues or organs. In general, diabetes is a chronic condition that cannot be cured.

Type 1 diabetes is the form of the disease where the body does not produce enough insulin. Type 1 patients are thus all insulin dependent (unlike type 2 patients). Patients need to maintain their blood sugar at levels as close as possible to normal. They do this by injecting insulin to compensate the shortage of natural body insulin. The dosage is adjusted based on results of blood sugar tests (using small blood glucose monitor devices). Obviously, food heavily influences the blood sugar level: diabetic patients therefore need to take into account what they eat.

Diabetic patients are under lifelong treatment by a whole team of experts (diabetologist, podologist, dietician, etc). The diabetologist is the main actor in keeping the diabetes itself under control (not the comorbidities), i.e. helps with the disease management. In order to assess the disease evolution correctly, they require daily recordings of glucose level and insulin taken. This is typically recorded by the patient in a so called 'diabetes diary'.

The diabetes diary is often a paper diary, but off course there also exist a large number of electronic recording means. The legacy application chosen for the TAS[3] integration trial is a website implementation of such a diary.

---

[19] Insulin is a hormone that helps in the conversion of food into energy for the body. Without insulin, glucose (sugar) from food cannot enter cells, so it builds up in the blood while the body tissue becomes starved for energy.
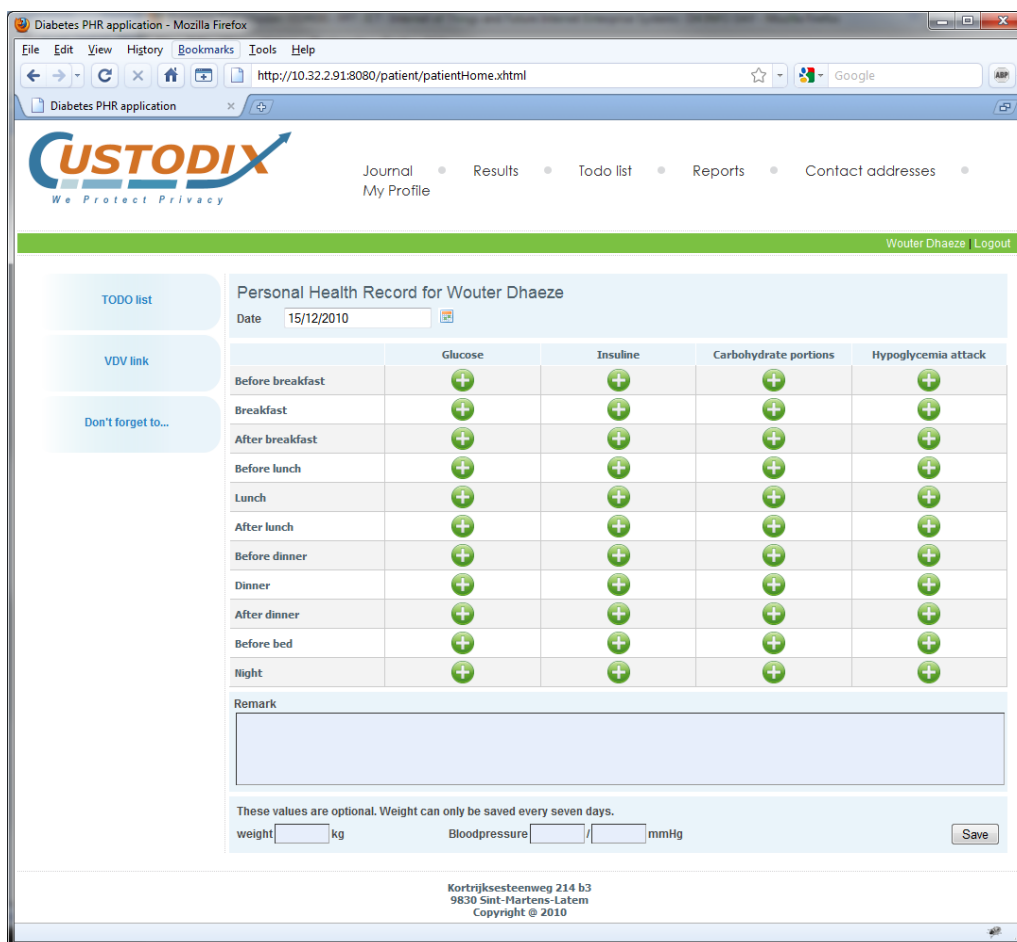
Figure 22: Diabetes diary application screenshot

The diabetes diary application used in the integration trial is limited in scope and focuses on the essential functionality of recording important parameters. The main component of the application is the diary itself (see screenshot on Figure 22) where patients can log information about their measured glucose level (blood sugar), the insulin that they have been taken and the carbohydrate content of the food they have eaten. Finally they can register all hypoglycaemia attacks (acute glucose shortage) that they have experienced. The application allows patients to share the information in this diary with their physician.

4.3.2 Dashboard

Figure 233 shows an example of how the dashboard could look for the integration trial. As explained, personal privacy policies are defined as modifications of the default settings that are accepted as commonly acceptable in the health domain. These personal settings refer to read operations only, as setting access rights for the other CRUD-operations make no sense in this context.

As one can see in Figure 23, an end user can specify user preferences for different types of information exchange in specific cases. In the example environment these (realistic) cases are:

- Data accessible by the GMD holder

- o The 'greyed-out' selection boxes are default policy settings that cannot be modified. One cannot restrict access of the GMD holder to most official communication as he is supposed to be the central person managing it.
- Data accessible by a physician with whom the patient already has a therapeutic relation. This is only relevant for data outside of that therapeutic relationship.
- Data accessible by random physicians on a first encounter or sporadic encounter i.e. when a real therapeutic relationship has (not yet) been established and registered.
- Data access in emergency and out-of-office situations.

Next to defining privacy settings on document types, a user can also specify his preferences depending on particular sensitive information that is recorded in these documents. The sensitivity configuration takes priority over the other settings: hence with respect to Figure 23 this patient's GMD holder has access to all data except for psychiatry related information.

It is important to note that the background mechanics are generic. Thus although the dashboard user interface limits the possibility of detailed access configuration, there is no technical limitation in the policy management system itself. One could for example add an 'advanced' view on the dashboard for people with an interest in more control where they could specify their preferences in more detail (e.g. enhancing the access settings for specific individual HCPs).
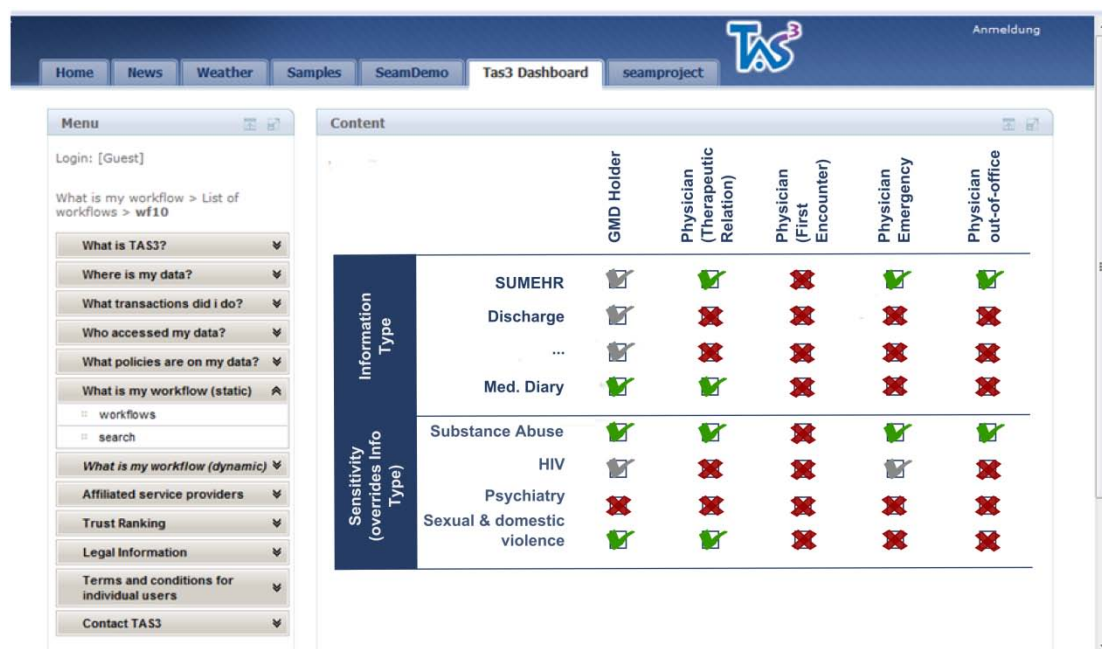

Figure 23: Example dashboard for the healthcare demonstrator

4.4 Integration Trial Status

4.4.1 Summary Evaluation of Phase 1

The phase 1 integration trial had as major objective to achieve TAS[3] compliance in a realistic setting. This meant full integration of the TAS[3] communication protocol stack, thus demonstrate TAS[3] compliant Single Sign On and Service Provider communication. Demonstration of the TAS[3] logging mechanism was a secondary objective.

These objectives have been reached as demonstrated on the year 2 review. The following has been demonstrated:

- Login to PILS portal using TAS3-ZXID-IDP
    - SSO capability (Single Sign On)
    - SLO capability (Single Logout)
    - A functioning ID-FF Discovery Service
- PILS – Repository communication using the TAS3 communication stack.

Integration with the TAS³ audit bus has also been demonstrated, although the main components of the TAS³ audit infrastructure (management and user interface) were not completed.

4.4.2 Status and Objectives of Phase 2

The objective of the phase 2 trial is to show the possibilities of TAS³ for building environments for information and service sharing where the end user remains in control of his data.

This goes beyond demonstrating technical compliance (phase1) and focuses more on the vision of what TAS³ can offer. Phase 2 can be considered successful if the three use cases of section 4.2.2 can be successfully demonstrated. A number of technical shortcuts and simplifications will be taken in the healthcare integration trail due to the current implementation state of the different TAS³ components (relevant components used in the trial are shown on Figure 24).

However, the phase2 integration trial should succeed in showing the possible advantage that TAS³ can offer for building complex environments which allow for user-centric personal data management. The concrete trial scenario will be available at the time of demonstration (review).
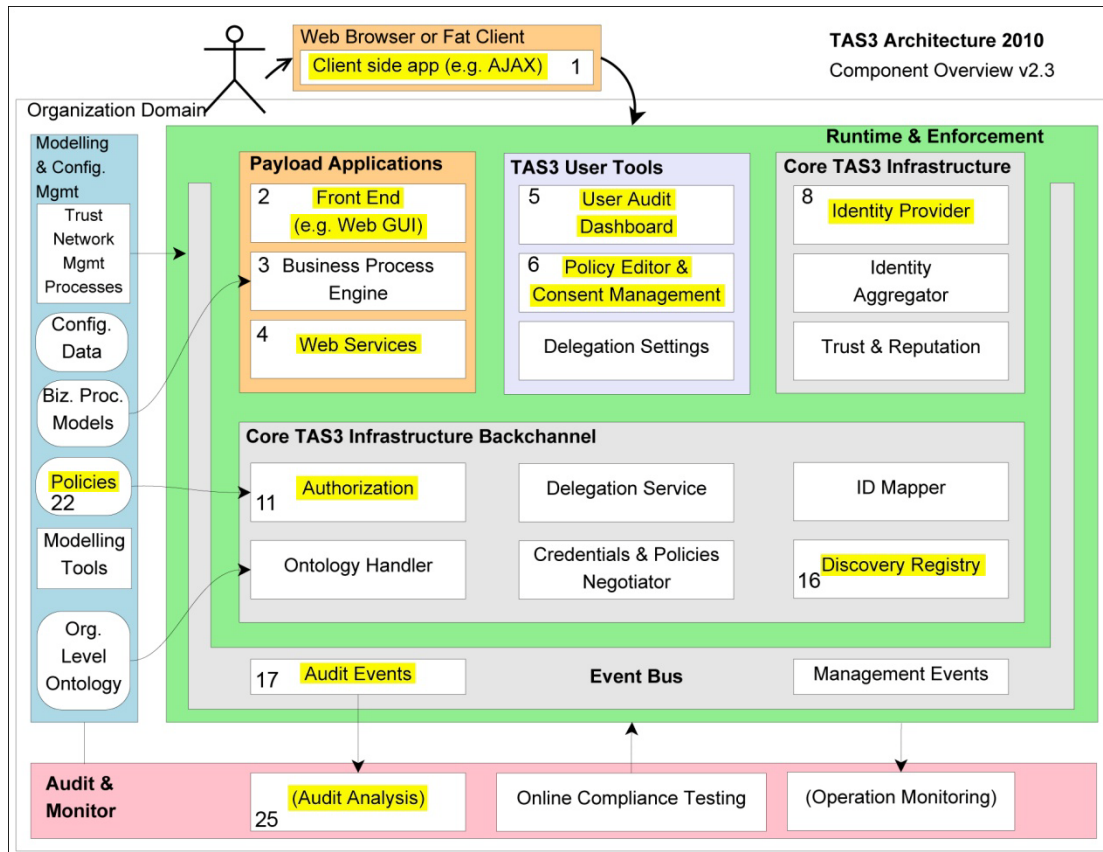
Figure 24: Healthcare Integration trial, demonstrated TAS3 components (highlighted parts integrated in the trial)

A final remark that should be made is that although the integration trial is staged in a realistic setting, there are a number of domain-specific issues that have not been worked out in the demonstrator because they are not within the scope of this work.

One issue is that some documents (e.g. a SUMEHR) would always include sensitive information (if that is registered for a patient). Hence the combination of policies on the document level with respect to 'information type' settings and 'sensitivity settings' can only work if the access control is finer grained than at the document level. If not, the unwillingness to share some kind of 'sensitivity settings' would (for some patients) always block such documents, which is not the intended behaviour. The mechanisms required to make such cases work have not been technically elaborated, however the following is a possible solution within the presented framework (requiring more complexity at the SPs): as a rule, the central PDP should comment on the reason why an access request was denied. If the SP believes that it is possible and relevant to 'remove the barriers' that caused the access denial (e.g. filter substance abuse information from a SUMEHR), then it should do so and resubmit an adjusted request to the central PDP. Depending on the central PDP re-evaluation of the request the information should or should not be delivered.

5 Table of Acronyms

CVS	Credential Validation Service
DAO	Data Access Object
GMD	Globaal Medisch Dossier (Belgian primary care health record)
GP	General Practitioner
HCP	HealthCare Professional
HEI	Higher Education Institution
HIS	Hospital Information System
IdP	Identity Provider
IDP	Identity Provider
KMEHR	Kind Messages for Electronic Healthcare Record
KPI	Key Performance Indicator
OCT	Online Compliance Testing
PDP	Policy Decision Point
PDS	Personal Data Store
PEP	Policy Enforcement Point
PERMIS	PrivilEge and Role Management Infrastructure Standards
PHR	Personal Health Record
PII	Personally Identifiable Information
SAML	Security Assertion Markup Language
SME	Small or Medium Enterprise
SP	Service Provider
SSO	Single Sign On
XACML	eXtensible Access Control Markup Language

Amendment History

Ver	Date	Author	Description/Comments
V1.0	20/12/2010	Nottingham Kenteq Custodix	Version ready to be reviewed
V2.0	21/12/2010	Custodix	Version with Healthcare Integration Trial completed
V3.0	23/12/2010	Nottingham Kenteq Custodix	Final version with reviewers comments addressed