Personalised Provenance Reasoning Models and Risk Assessment in Business Systems: A Case Study

Paul Townend, David Webster, Colin C. Venters, Vania Dimitrova, Karim Djemame, Lydia Lau, Jie Xu, Sarah Fores, Valentina Viduto
University of Leeds
Leeds, United Kingdom
e-mail: p.m.townend@leeds.ac.uk

Charlie Dibsdale,
Shyam Iyengar, Nick Taylor
Rolls Royce PLC
P.O. BOX 31
Derby, United Kingdom
e-mail: charlie.e.dibsdale@o-sys.com

Jim Austin[1], John McAvoy[2],
Stephen Hobson[2]
[1] University of York & Cybula Ltd.,
[2] Cybula Ltd.
Computer Science Building,
York, United Kingdom
e-mail: austin@cybula.com

Abstract— **As modern information systems become increasingly business- and safety- critical, it is extremely important to improve both the trust that a user places in a system and their understanding of the risks associated with making a decision. This paper presents the STRAPP framework, a generic framework that supports both of these goals through the use of personalised provenance reasoning engines and state-of-art risk assessment techniques. We present the high-level architecture of the framework, and describe the process of systematically modelling system provenance with the W3C PROV provenance data model. We discuss the business drivers behind the concept of personalizing provenance information, and describe the STRAPP approach to enabling this through a user-adaptive system style. We discuss using data provenance for risk management and treatment in order to evaluate risk levels, and discuss the use of CORAS to develop a risk reasoning engine representing core classes and relationships. Finally, we demonstrate the initial implementation of our personalised provenance system in the context of the Rolls-Royce Equipment Health Management, and discuss its operation, the lessons we have learnt through our research and implementation (both technical and in business), and our future plans for this project.**

Keywords—provenance, risk, trust, web services

I. INTRODUCTION

As real-world information systems increasingly require users to make business- and safety- critical decisions based upon data and analyses that arrive from many distributed and heterogeneous sources (both computer systems and human users), it is crucial for a user to be able to place trust in system outputs and to understand the risk of making decisions based upon these outputs.

Despite this, users are often unaware of the provenance of the data upon which they are asked to make a decision – where and who it has come from, what time it was created, based on what information, as a result of which decisions or 'value-added' operations, and so on. Such information is critical for allowing users to make informed decisions, and also for facilitating fault-diagnosis of system components (software, hardware, and human). Much research has been carried out on various approaches to capturing provenance at source (i.e. at the 'point of creation') e.g. [1,2], and tackling the challenges in representation for interoperability [3]. However, there is a lack of understanding as to how users can exploit provenance data

to enhance their work at the 'point of use'. A technology-driven approach to mining provenance could cause information overload or distrust by users as the information could be perceived as irrelevant or unreliable. By recording and presenting provenance information in a user-specific and context-aware form, and integrating a risk assessment concept which increases user awareness of possible risks, a significant improvement can be made to the trust that a user places on service data, as well as providing invaluable diagnostic information to system developers and maintainers. This enables users to make confident, informed, and timely decisions.

In order to address this problem, the STRAPP (trusted digital Spaces through Timely Reliable And Personalised Provenance) project has been established, funded by Rolls-Royce, Cybula Ltd, and the UK Technology Strategy Board. STRAPP has developed a generic framework to facilitate the creation of provenance-based, personalised trusted digital spaces for timely and confident decision making, for use in any multiple stakeholder domain in which critical decisions are made. Specifically, the project is targeting the OSys (a subsidiary of Rolls-Royce) Equipment Health Management (EHM) system, and a Medical Injury Index (MII) decision-support system developed by Cybula Ltd. This paper presents for the first time an initial implementation of the STRAPP system, based on the OSys EHM system, and addresses both the challenges faced in implementation, and also the lessons learned and future work required to improve the implementation further.

II. THE STRAPP PROJECT

The objective of the STRAPP project is to enable users to place increased trust on data shown by, and decisions made by, a system by allowing them to view the provenance of that data or decision, presented in a personalised manner (for example, based on their role; managers may need to view the provenance and risk of a decision at a different level than software engineers, etc.) Furthermore, the project aims to provide mechanisms to ensure users understand the risks associated with data and decision-making. STRAPP consists of three main internal components shown in figure 1 – the presentation service, personalization service and data management service. These are Web Services which interact with other internal components at their back end.
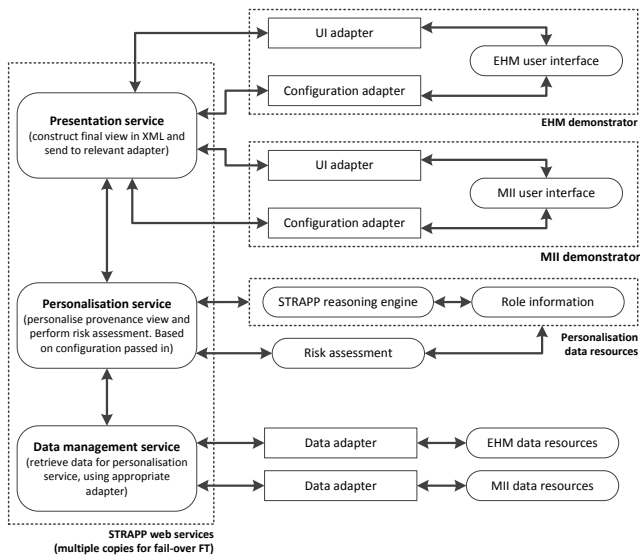
Figure 1: STRAPP Architecture

The STRAPP *presentation service* is responsible for formatting and displaying the final "view" of the STRAPP system (i.e. the information the system wishes to display to the user) and is also responsible for taking user input from an external component and passing that information (correctly formatted) to the appropriate STRAPP internal component. The "back end" of the presentation service consists of a "UI adapter" and a "Configuration adapter". These components are integrated into the external system. The STRAPP *personalisation service* is responsible for invoking the STRAPP provenance model and its reasoning engine, and the risk assessment components (discussed in sections 3 and 4). These components do not interface with any external components, and must retrieve any data they require through accessing the STRAPP data management web service. The difference between the personalisation service and the presentation service is that the presentation service may invoke the personalization service multiple times in the process of constructing a finished report. The STRAPP *data management service* is responsible for accessing external data resources on behalf of the personalisation service. It functions by taking a generic request, formatted in XML by the personalisation service and – through a bespoke data adapter – converting that generic request into a system specific request, and sending this to the target external component. On receiving data from a component, the data will be converted into a generic XML-based format and sent back to the personalisation service.

The major contribution of this work is the personalisation service, and the following two sections discuss the role of the provenance reasoning model and risk assessment components in greater detail.

III. PERSONALIZED PROVENANCE REASONING MODELS

In order to support the ability to make correct decisions, factors affecting the way the decision maker acts need to be considered. Presentation of provenance data in a way that ensures greater objectiveness in the decision making process is also required.

A. Modelling Provenance

Provenance information requires that the underlying system workflow of a target system be systematically modelled. Within STRAPP, we have named this workflow and associated provenance meta-data the *'Configuration Network'*. This Configuration Network will be unique for each system under observation and contains the linking between system personnel, processes and documents along with configuration management information as a connected directed graph.

The provenance of an end result of a target system can be derived, therefore, by following the complete path through the graph from the input data source to the end output. This provenance data will not just contain a list of entities from the workflow graph, but additionally will contain provenance specific meta-data such as:

- Versioning information about the software systems.

- Training data for software systems, for instance data used within event detection algorithms.

- Personnel associated with enacting system processes.

In order to model the configuration network and associated provenance meta-data we have built upon the W3C's PROV provenance data model standard [4]. This permits a formalization of the Configuration Network structure at the highest level. PROV defines the core concepts of **Entity**, **Activity**, and **Agent** to represent: documents; processes that act on Entities; and a human or software enactor of the Activity. The associations between these top-level concepts are illustrated in Figure 2. Furthermore, PROV permits meta-data to be attached to this basic structure; examples of this include activity start and completion times.

PROV is in a W3C Working Draft status as of September 2012. It must be stated at this point that this paper will not show the specific workflow created for the Rolls Royce Electronic Health Management decision support system due to commercial sensitivities; however, the Configuration Network will be discussed in general terms. STRAPP takes advantage of a provenance modelling approach to represent and reason about the Configuration Network and associated provenance meta-data. This modelling builds upon the W3C's de-facto ontological representation of PROV named PROV-O which is defined using the W3C's Web Ontology Language (OWL2). A strong motivation for adopting this approach rather than an *ad-*
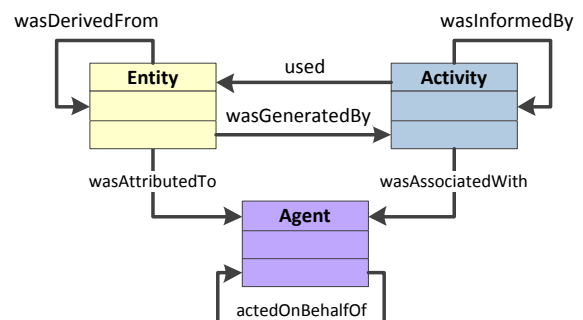

Figure 2: PROV Core Structures

hoc programmatic approach using SQL queries is that the modelling approach provides us with a standardized layer of normalized data to implement processing algorithms upon. This lends itself to the following advantages and abilities:

- Once the graph is created it can be parsed, modified, verified and serialized using standard and well tested tools.

- The ability to invoke reasoning engines (for instance, *Pellet* [5] or *FaCT++* [6]) to detect logical inconsistencies within the data and to detect missing elements. An example of this is the ability to detect if the provenance data is missing an agent associated with an activity.

- The ability to infer relationships between entities without data being explicitly present in the underlying database. This is supported by existing research in Description Logics. [5]

B. Personalizing Provenance Information to the User

Personalization is becoming one of the key drivers for successful deployment and user acceptance of complex systems. At the current time, many large companies have adopted personalization technologies as part of the services that they offer. For example, e-commerce sites like Amazon or eBay include customer-tailored recommendations to increase sales and improve user experiences, search engines use personalization features to filter-out irrelevant information, etc. In the case of a business system, personalization of information is needed in order to prioritize the information which is required for a user to complete a particular task according to the role that they are currently fulfilling. If the user's decision making activity is to be enhanced through the presentation of provenance information, then the personalization process needs to apply to the provenance information also. The personalization approach taken within the STRAPP project is of a *user-adaptive system* [7] style. Within this approach we define the following four components:

- **User Model** - a data structure (model) that contains explicit assumptions on the aspects of the user that are relevant to the adaptive behavior.

- **Context Model** - In STRAPP we will maintain a context model presenting information about the provenance structure (Configuration Network) and the user roles. In addition the context model will include a risk model to enable risk-aware personalization.

- **User model acquisition** - Within STRAPP, to start with we have considered role-based personalization, which considers a collective profile of the people within the same role. The motivation for this is so that user model acquisition will be simplified as much as possible.

- **User model application** - a module which includes algorithms for deciding what information to be shown to the user and how, considering the user (context) model. In STRAPP, this component will be the prime focus for personalization. The motivation for applying the context model is to decide what level of provenance detail to provide and how. This will also take into account the possible risk associated with the advice/notification/information provided by the system.

To provide a generic example for the use of personalization when presenting provenance in a decision support system, we can describe a generically applicable scenario; it would contain both an Analyst and a Support Engineer role. The requirements for provenance information between these two roles and the presentation of this information (e.g. the level of detail) would be quite different and thus require different information to be presented from the context model. A manager will want to see more information about the impact on business, in a summarized format on a single page. They will probably be ad-hoc users who do not want to learn how to navigate the system, they just want to get to data on their first click. On the other hand, a support engineer user would need a far more in depth look at detail, but is possibly still an ad-hoc user. The Analyst role would require the presentation of the chain of information (Entities and Activities) that led up to the end decision along with associated risk information. The Support Engineer's role would involve the task of determining why a false positive decision was made and to 'debug' the workflow. The Support Engineer would, therefore, require provenance information relating to the following:

- Configuration management and version control for datasets and training data associated with software agents which enact the activities within the workflow.

- Error and transaction logs associated with software agents.

- Verification and validation processes associated with software agents.

Whilst the Context Model would be expected to contain the graph representing the whole configuration network it is the role of the User Model Application to present a view of this model relevant to the user's role and current task. The underpinning assumption is that the way information is presented and the user experience have impact on the user's trust in making the decision support information 'actionable' from the specific context.

IV. RISK

Risk is often characterized by reference to potential events (combination of potential threats and vulnerabilities) and consequences (relates to potential impact), or a combination of these. In the context of STRAPP we define 'risk' as the 'likelihood' of an 'unwanted incident' and its 'consequence' where:

- Likelihood is the probability of something occurring.

- An unwanted incident is an event that directly or indirectly harms or reduces the value of an asset; an asset in this context can be something physical or conceptual to which a party assigns value to and for which the party desires to protect.

- The consequence is the impact of an unwanted incident on an asset in terms of the harm or a reduction in the value of the asset.

In order to assess the risk associated with making critical, high-value decisions based on evidence presented by a system, it is essential to know how the data was derived, processed and transformed. A traditional approach to understanding risk is through the process of risk assessment using a range of different methodologies and methods. For example, *Fault Tree Analysis* is a technique based on deductive logic where undesirable events are first defined and causal relationships of the failures leading to that event are then identified. In contrast, the approach adopted in STRAPP consists of using provenance for risk management and treatment; the output of risk assessment will then be used to evaluate the level of risk. Knowing the provenance of an object can help reduce or mitigate its associated risk as it provides evidence for data consumers to make informed decisions about its origins and transformations. This knowledge is essential to determine properties such as quality and integrity, and also for the user to assess the risk of using the data and the services that derived it. However, there may be different levels of risk requiring different levels of advice using different types of provenance.

A. Types of Risk

Within STRAPP, we have identified five types of risk: *technical origin risks* (e.g. sensors), *data-related* (e.g. integrity), *activity-related* (e.g. identify symptoms), *agent-related* (e.g. technician) and *risk of making a final decision*. In theory, objects that compose a provenance-aware systems expose their provenance and can be modelled using emerging W3C standards such as PROV (which provides the three objects of Entity, Activity, and Agent are discussed in section 3). The provenance of each PROV object can be used as the basis for calculating risk associated with each object. Considering technical origin risks, we are interested in the reliability of each technical element and its impact on validity of measurements performed. Data-related risk is associated with the data transfer from an aircraft to ground and extraction of this data. This type of risk looks at data reliability, likelihood of losing valuable data, likelihood of data corruption during transfer process and transformation of the raw data into context-aware information preserving its integrity and initial validity. Activity-related risks are the ones which relate to the data analysis process, which is based on specific requirements associated with particular activity types. Additionally, our findings reflect that risks associated with agents must be taken into account since agents can be humans and their knowledge used to determine whether diagnosis is valid may be crucial in this respect. The overall risk and impact of a notification can then be derived by following the complete provenance path (Shown in figure 4). With regard to risk of taking a final decision we consider an action or a combination of actions with respect to the trade-off of possible impacts, defined in terms of cost, and risks that may arise. The knowledge of a provenance chain and absolute values of risk with regard to these risk types increase an overall trust to the system and provides an ability to know, understand and manage risks.

B. Risk reasoning engine

The purpose of STRAPP is to enhance trust in decision-making processes based on data provenance and risk assessment. A simple approach to help achieve this is by creating a risk model based on the data provenance. However, this approach may prove to be inefficient in tackling issues such as large-scale evolutionary changes, temporal information, querying and inference, and user context. To address the limitations of a standard risk assessment and enable the above issues, a reasoning engine can be employed. Using CORAS (a security risk analysis method, which has a customized language for threat and risk modelling) we have developed a risk model to represent the main risk classes and relationships. The main classes of the model include:

- Assets
- Consequences
- Party
- Threats
- Threat Scenarios
- Treatment
- Treatment Scenarios
- Unwanted Incidents
- Vulnerabilities

The main relations of the risk model include:

- harms
- hasConsequence
- impacts
- initiates
- leadsTo
- occursWithDefined
- occursWithLikelihood
- occursWithUndefined
- occursWithUndefinedLikelihood
- treats

By building a reasoning engine based on the provenance and risk models, a platform is created where the data provenance and the risk model can efficiently communicate and combine to augment the decision support system, which can affect 'trust'.

V. IMPLEMENTATION

In order to assess the utility of the personalised provenance model described in section 3 in an industrial context, an initial implementation has been created that integrates STRAPP research with the Rolls-Royce "EHM" system. One of Rolls-Royce's business models in the aviation sector (specifically, aircraft engines) is to be a "TotalCare" service provider, taking responsibility for maintenance for a fixed price or guaranteeing availability of power or thrust over a given period of time. "TotalCare" relies on Equipment Health Management (EHM) systems, where sensor readings are taken from Rolls-Royce products and analyzed to monitor usage, life and detect early signs of failure. If a problem is found - which in real life is a rare occurrence - it is important that timely remedial action is started before the problem develops toward an unacceptable level. Deciding the appropriate remedial action always involves a human who has to rely on the information provided to support the decisions made. Critical to this process is provenance of the decision support information and a full understanding of the risk involved in making a decision.

The system design is shown in figure 3, and consists of three components. Firstly, a database containing anonymised Rolls-Royce EHM data was established, using the same schema and technology as the real system in order to allow for easy integration of STRAPP into the main EHM framework. Following this, a series of Web Services (containing the provenance reasoning model functionality) were developed and deployed on an Apache Tomcat server. This server was then replicated in order to improve availability through the use of an N-Copy fault-tolerance mechanism. All web service invocations are performed through this mechanism - should a service time-out or raise an exception, an alternate (identical) service is invoked on a different server; this process continues until either a response is obtained or all alternates have been tried. This provides basic "fail-over" fault-tolerance. Finally, the STRAPP EHM client itself was created; this is a standalone Java application, with a GUI that allows users to interact with and visualize the results from the Web Services, which in turn mine data from the EHM database. All system components are logically separate and can be executed on distributed nodes.

The STRAPP client implementation is focused on the aviation domain of EHM, which monitors aircraft engine data. This EHM system includes an "electronic diary", which supplements the main EHM processing application. The diary provides the ability to record notes and observations made by technicians and engineers about the possible problem events they observe. The EHM system produces a number of data trends and diagnostics where the system observes deviation from normal behavior. These deviations are initially inspected by the *technician* who marks up the significant events in the diary, before passing them on to the *engineer* who decides which diagnostic alerts to forward to Rolls-Royce and the airline customers, who then initiate remedial action. The EHM system along with the technician and engineer form a chain of triage for potential events, which helps ensure all genuine diagnosed events are suitably alerted, with a minimum of false positives. The EHM system, along with the electronic diary, form some of the sources for data mining initiated by the STRAPP services.
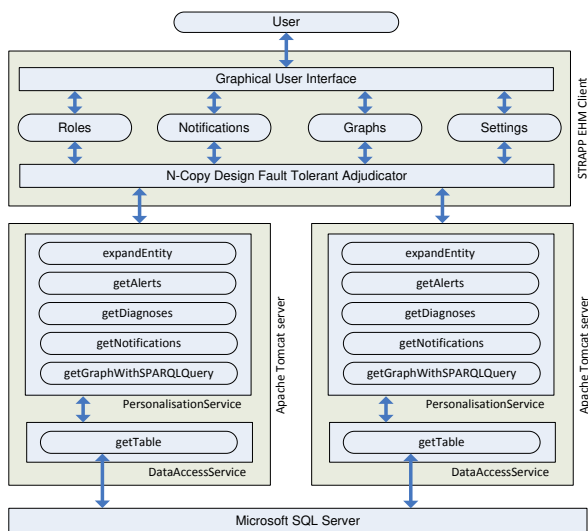

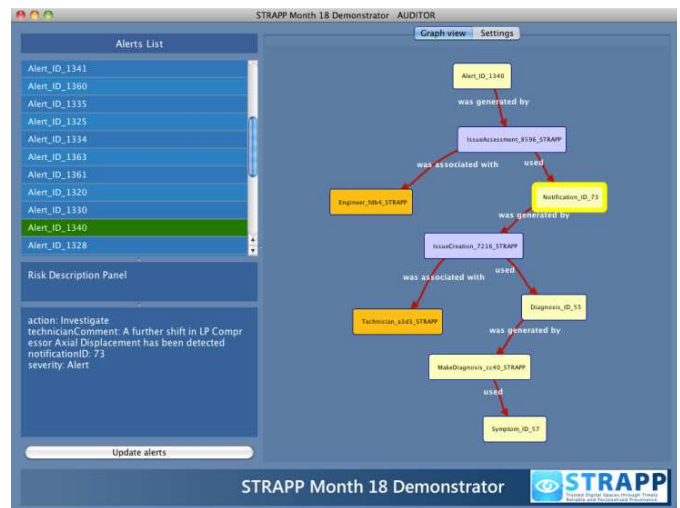**Figure 3: Main STRAPP system design**

Figure 4: STRAPP Client "Auditor" view

The client starts by allowing a user to log-in and select their role. In a production system, the user will already be assigned a role and the act of logging will assume the role authorizations and privileges. These roles can be one of *Technician*, *Engineer*, *Auditor*, or *Support Engineer*. A Technician's responsibility is to view data trends and a system diagnosis to identify if they require the attention of an engineer; if so, the technician resolves the diagnosis into a notification. An Engineer's responsibility is to view notifications to identify if maintenance actions need to be taken; if so, the notification will be resolved into an alert. The role of an Auditor is to look at each alert and verify the processes, data and documents that were used in its production. The auditing role is similar to that undertaken by an auditor in ISO 9000, using the provenance capability to ensure EHM processes and systems are operated and built to the design intent with inherent traceability. A Support Engineer is similar to an Auditor but assumes a more diagnostic role, determining why false positive alerts were made and 'debugging' the workflow.

Depending on a user's role in the system, they retrieve a list of notifications, alerts or diagnoses. The STRAPP framework supports each role by providing a personalised view of the provenance of each of these entities, and visualizing this information on screen in graph format. Through STRAPP technology, Technicians can quickly determine what diagnosis algorithm was used to identify each symptom that makes up a diagnosis; Engineers can view this information in addition to information about the technician who generated a notification; Auditors can view the entire chain of information that led up to the end decision to generate an alert; Support Engineers can view all the information that an Auditor can see plus additional provenance relating to the software systems that implemented the EHM processes (relating to configuration management and version control of datasets, error and transaction logs associated with software agents, verification and validation processes associated with software agents, etc.)

The client implementation initially provides a user with a limited view of the provenance chain appropriate for their role (in order to avoid information overload as discussed in section one); however, when a node in the displayed graph is double-

clicked, the STRAPP system streams additional data from the backend system's provenance database to add further provenance information about that node to the graph (which will itself be represented by further nodes). This allows a user to recursively mine the provenance database to a depth of their choosing. An example of this is shown in figure 4, which shows a screenshot of the STRAPP EHM Demonstrator system displaying an expanded provenance graph for the "Auditor" role. When a node is selected, any comments left in the provenance for that node are displayed in a pane on the left of the screen. Users can dynamically reposition graph nodes on screen to aid comprehension if so desired.

In the current system, a pane has been provided for the STRAPP Risk Assessment technology, but this has yet to be integrated into EHM. It is envisioned that each node in the provenance chain will have an associated risk score, which users can view to better understand the risks inherent in the decision they need to take, supported by an aggregated score of the inherent risks within the elements making up the provenance chain.

VI.   Lessons Learnt and Future Work

During the course of implementing the STRAPP framework and integrating it into Rolls-Royce EHM, a number of valuable lessons have been learned, both technical and social. The system has revealed a completely new use case within EHM, whereby provenance may be used for troubleshooting EHM Processes and helping to find where ICT infrastructure has failed. Additionally, the system has helped to reveal how the various reasoning engines can be used for inference, and how Rolls-Royce as a business can take advantage of this in risk, personalisation and provenance. Furthermore, the process of modelling workflow within the constraints of the PROV structure has exposed the requirement to provide additional configuration management data from the underlying system. In addition to this, the process of modelling the roles associated with the workflow has identified the need for privacy to be considered when presenting provenance information to a user, as it could be used to identify the performance of particular individuals. Indirectly, STRAPP has been a catalyst to review configuration management in Rolls-Royce development and in how the company run services; this has exposed opportunities for improvements to the benefit of the business. STRAPP can be regarded as a 'perspective' from which to view the current EHM service which is highlighting many opportunities for improving current practice.

In terms of future work, there is a great deal that can be done to further refine and – importantly – evaluate the effectiveness of the STRAPP system. We are currently organizing a series of user workshops to gain feedback from the Rolls-Royce EHM workforce in order to identify weaknesses and strengths of the system; this feedback will be used to inform a further round of development. We are also working to integrate our risk assessment technology into the demonstrator, with a goal of allowing the system to calculate a risk score for each node within a provenance chain, as well as the overall risk of the data and any subsequent decisions. Additional work will also be performed to assess the scalability of the scheme and refine its overall performance, and stress testing will be performed on a larger EHM dataset. Work will also continue on the application of STRAPP to the Cybula MII system in order to assess the effectiveness and utility of STRAPP in the health domain.

VII.   Conclusions

In this paper, we present the STRAPP framework, which seeks to combine the reasoning engine of a provenance model and a risk assessment model together with personalization to improve the trust that users can place in a business information system. We describe the high-level architecture of the framework, and describe the process of systematically modelling a system's "Configuration Network" with the W3C PROV provenance data model standard using a reasoning engine. We go on to discuss the business drivers behind the concept of personalizing provenance information, and discuss the STRAPP approach to enabling this through a user-adaptive system style. We then discuss the important of risk assessment, and the approach in STRAPP of using provenance for risk management and treatment in order to evaluate risk levels. We identify three types of risk, and discuss the use of CORAS to develop a risk model to represent main classes and relationships. By building data provenance and risk assessment models and using the reasoning engine over these models, we create a platform where data provenance and risk can efficiently communicate to augment decision support systems. Finally, we demonstrate the initial implementation of our personalised provenance reasoning system in the context of the Rolls-Royce Equipment Health Management system, and discuss its operation, the lessons we have learnt through our research and implementation (both technical and in business), and our future plans for this project.

References

[1] D. Barseghian et al., "Workflows and extensions to the Kepler scientific workflow system to support environmental sensor data access and analysis" in Ecological Informatics, pp. 42-50, Volume 5, 2010.

[2] D. Hull, K. Wolstencroft, R. Stevens, C. Goble, M. Pocock, P. Li, and T. Oinn, "Taverna: a tool for building and running workflows of services.," Nucleic Acids Research, vol. 34, iss. Web Server issue, pp. 729-732, 2006.

[3] L. Moreau et al, "The Open Provenance Model core specification (v1.1)", in Future Generation Computer Systems, 2010.

[4] K. Belhajjame et al., PROV-O: The PROV Ontology. W3C Working Draft 24 July 2012. http://www.w3.org/TR/prov-o/

[5] E. Sirin, B. Parsia, B. Grau, A. Kalyanpur, Y. Katz, "Pellet: A practical OWL-DL reasoner", in Software Engineering and the Semantic Web, Vol. 5, No. 2, pp. 51-53, June 2007

[6] D. Tsarkov, I. Horrocks, "FaCT++ Description Logic Reasoner: System Description" in Lecture Notes in Computer Science, Volume 4130, pp. 292-297, 2006

[7] A. Jameson, "Adaptive interfaces and agents" in A. Sears and J.A. Jacko (eds.) Human-computer Interaction Handbook, CRC Press, 2008