To Trust or Not to Trust? Developing Trusted Digital Spaces through Timely Reliable and Personalized Provenance

Colin C. Venters, Jim Austin, Charlie E. Dibsdale, Vania Dimitrova, Karim Djemame, Martyn Fletcher, Sarah Fores, Stephen Hobson, Lydia Lau, John McAvoy, Alison Marshall, Paul Townend, Nick Taylor, Valentina Viduto, David E. Webster and Jie Xu

Abstract—Organizations are increasingly dependent on data stored and processed by distributed, heterogeneous services to make critical, high-value decisions. However, these service-oriented computing environments are dynamic in nature and are becoming ever more complex systems of systems. In such evolving and dynamic eco-system infrastructures, knowing how data was derived is of significant importance in determining its validity and reliability. To address this, a number of advocates and theorists postulate that provenance is critical to building trust in data and the services that generated it as it provides evidence for data consumers to judge the integrity of the results. This paper presents a summary of the STRAPP (trusted digital Spaces through Timely Reliable And Personalised Provenance) project, which is designing and engineering mechanisms to achieve a holistic solution to a number of real-world service-based decision-support systems.

Index Terms—Provenance, risk, sense-making, service-oriented computing, trust

INTRODUCTION

Increasingly real-world information systems require users to make critical, high-value decisions based upon data and analysis that have been derived from distributed and heterogeneous sources and services. As a result, it is critical for a user to be able to place trust in system outputs and to understand the risk of making decisions based upon these outputs. Nevertheless, users are often unaware of the provenance of the data upon which they are asked to make informed decision. In addition, the emergence of service-oriented computing as the dominant computing paradigm, which relies heavily on the dissemination, exchange and reuse of data sets, has exacerbated the need for a mechanism to engender trust in the data utilized in and between the services. Tsai et. al., [1] suggest that due to the dynamic nature of service-oriented systems, it is critical to consider not only the security and integrity of the data but also its trustworthiness. The decision to trust is based on evidence to believe or to be confident in someone or something [2]. To address this, it is suggested that provenance is critical to building trust in the data and the services that generated it [3].

This paper presents a summary of the STRAPP (trusted digital Spaces through Timely Reliable And Personalised Provenance) project, which is designing and engineering mechanisms to achieve such a holistic solution as well as applying and evaluating the developed mechanisms to a number of real-world service-based decision-support systems in the aerospace engineering and healthcare domains.

- *Colin C. Venters is with the University of Huddersfield. E-mail: c.c.venters@hud.ac.uk*
- *Jim Austin is with the University of York. E-mail: austin@cs.york.ac.uk*
- *Charlie E. Dibsdale is with O-Sys. E-mail:charlie.e.dibsdale@o-sys.com*
- *Vania Dimitrova is with the University of Leeds. E-mail: v.g.dimitrova@leeds.ac.uk*
- *Karim Djemame is with the University of Leeds. E-mail: k.djemame@leeds.ac.uk*
- *Martyn Fletcher is with Cybula Ltd. E-mail: martyn@cybula.com*
- *Sarah Fores is with the University of Leeds. E-mail: s.fores@leeds.ac.uk*
- *Stephen Hobson is with Cybula Ltd. E-mail: stephen@cybula.com*
- *Lydia Lau is with the University of Leeds. E-mail: l.m.s.lau@leeds.ac.uk*
- *John McAvoy is with Cybula Ltd. E-mail: mcavoy@cybula.com*
- *Alison Marshall is with the University of Cumbria. E-mail: alison.marshall@cumbria.ac.uk*
- *Paul Townend is with the University of Leeds. E-mail: p.m.townend@leeds.ac.uk*
- *Nick Taylor is with O-Sys. E-mail: nick.taylor@o-sys.com*
- *Valentina Viduto is with the University of Huddersfield. E-mail: v.viduto@hud.ac.uk*
- *David E. Webster is with the University of Leeds. E-mail: d.e.webster@leeds.ac.uk*
- *Jie Xu is with the University of Leeds. E-mail: j.xu@leeds.ac.uk*

1 THE STRAPP PROJECT

The STRAPP project has been established, funded by Rolls-Royce, Cybula Ltd, and the UK Technology Strategy Board to facilitate the assessment of provenance-based, personalised trusted digital spaces where timely and critical decisions should be made. The objective of STRAPP is to enable users to place increased trust on data shown by, and decisions made by a system and by allowing them to view the provenance of that data or decision, presented in a personalised manner. For example, 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 visualization mechanisms to ensure users understand trust and the risks associated with data and decision-making. These mechanisms are integrated to both the Equipment Health Management system developed by O-Sys, a subsidiary company of Rolls-Royce PLC, that provides customers primarily in the aerospace, marine and energy sectors with the ability to diagnose and predict equipment faults, and to the Brain Injury Index system developed by Cybula Ltd that assists researchers and practitioners in the healthcare industry, with a focus on neuroscience.

STRAPP consists of three main internal components: the presentation service, personalization service and data management service. The presentation service is responsible for the input and output of the system as well as passing that data to the appropriate STRAPP internal component. The personalisation service is responsible for invoking the provenance model and its reasoning engine, and the risk assessment components. The data management service is responsible for accessing external data resources on behalf of the personalisation service. These services are all implemented as web services, which interact with other internal components at their back end. A major contribution of this work is the personalisation service, and the following two sub-sections outline the role of the provenance-reasoning model and risk assessment components.

1.1 Personalized Provenance Reasoning Models

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'; 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. Our provenance modelling builds upon the W3C PROV-O standard [4], which itself is an instance of the PROV modelling standard [5], encoded using the W3C's Web Ontology Language (OWL2). The provenance model is represented in RDF against the PROV-O ontology. The provenance of an end result of a target system can be derived 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.

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. The personalization approach taken within the STRAPP project is of a user-adaptive system [6] style. Within this approach we define the following four components: user model; context model; user model acquisition; and user model application. The approach to personalised provenance adopted in the STRAPP project is discussed in greater detail in Townend et. al., [7].

1.2 Risk Reasoning Engine

Risk is often characterized by reference to potential events and consequences, 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; 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. Within STRAPP, five types of risk have been identified: technical origin risks e.g. sensors; data-related e.g. integrity; activity-related e.g. identify symptoms; agent-related e.g. technician; and the risk of making a final decision.

In order to assess the risk associated with making critical, high-value business decisions based on evidence presented by a system, it is essential to know how the data was derived, processed and transformed. In theory, objects that compose a provenance-aware system expose their provenance and can be modelled using emerging W3C standards such as PROV. The provenance of each PROV object can be used as the basis for calculating risk associated with each object. A quantitative risk assessment approach is applied within STRAPP to estimate the level of risk possessed by the provenance data recorded within the PROV-O data model; thus an identification of the elements of risk within the provenance chain becomes important. The general STRAPP risk assessment model, divides the assessment process into the following stages: vulnerability identification; threat identification; current control analysis and effectiveness; event analysis; quantative risk analysis; and decision support. For more detailed discussion on the approach to risk assessment adopted in the STRAPP project detail see Viduto et. al., [8].

2 Summary & Conclusions

In this paper, we present a summary of the STRAPP project, 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. Increasing trust is a multifaceted problem. There is a need to understand how and why data consumers trust or mistrust data or the services that generated that data in order to build robust models, which incorporate a human dimension. Emerging evidence suggests that understanding trust from an end-users perspective is essential and that they should play a pivotal role in the validation process to assess whether any of these approaches actually improve trust. The decision to trust is a complex interplay between the physical evidence associated with a data object or service and an individual's confidence and belief based on their subjective perception of the truth. However, trust is not a binary function of trusting or mistrusting but rather there are degrees of trust, which may be fluid and influenced by the degree of associated risk. This raises questions concerning the concept of uncertainty, how this can be represented, measured and mitigated. It is argued that provenance can increase trust in heterogeneous data and services. Provenance is a documented, historical representation of the origins, processes', and transformation of data, which provides a qualitative dimension. However, current definitions suggest that this should be accompanied by quantitative metric. 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'. The most recent findings from our experience of developing STRAPP are summarized as follows: the relationship between risk assessment and provenance; the need for a layered architecture; the need to limit the processing of provenance in a large-scale system; optimizing the trade-off between transactional granularity and system performance; persistence and usefulness of provenance data sources; and persistent risk assessment results. Future work will focus on refinement and evolution of the effectiveness of the STRAPP system. A series of user workshop have been organised to evaluate the effectives of the system, which will then be used to inform a further round of development.

Acknowledgments

The STRAPP project (Trusted Digital Spaces through Timely Reliable and Personalised Provenance) is funded by the UK Technology Strategy Board (grant reference 1926-19253), Rolls-Royce plc, Osys Ltd, Cybula Ltd, and the UK Engineering and Physical Sciences Research Council Knowledge Secondment Scheme. Their support is gratefully acknowledged.

References

[1] W-T. Tsai, et al., "Data provenance in SOA: Security, reliability, and integrity." Service Oriented Computing and Applications, 1(4): p. 223-247, 2007.

[2] Y. Yamamoto, "A morality based on trust: Some reflections on Japanese morality." Philosophy East and West, 40(4): p. 451–469, 1990.

[3] L. Moreau, et. al., "The provenance of electronic data." Communications of the ACM, 51(4): p. 52-58, 2008.

[4] T. Lebo et al, "PROV-O: The PROV Ontology", http://www.w3.org/TR/prov-o/

[5] Y. Gill et al, "PROV Model Primer", http://www.w3.org/TR/prov-prime/

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

[7] P. Townend et. al,. "Personalised provenance reasoning models and risk assessment in business systems: A case study." Proceedings of the 7th IEEE International Symposium on Service Oriented System Engineering, March 2013, San Francisco Bay, USA

[8] V. Viduto. Et. al., "Trust and risk relationship analysis on a workflow basis: A use case." ICIMP'2014: Proceedings of the 9th International Conference on Internet Monitoring and Protection, July 2014, Paris, France, 2014.