



On Identity Assurance in the Presence of Federated Identity Management Systems

Adrian Baldwin, Marco Casassa Mont, Yolanta Beres, Simon Shiu
Trusted Systems Laboratory
HP Laboratories Bristol
HPL-2007-47R1
January 29, 2008*

identity, assurance,
identity assurance,
trust, federation

In this paper we address the appropriate management of risk in federated identity management systems by presenting an identity assurance framework and supporting technologies. We start by discussing the risk mitigation framework that should be part of any identity assurance solution. We then demonstrate how our model based assurance technologies can be used to report success of an identity assurance programme. We discuss how this approach can be used to gain trust within a federated identity management solution both by communicating the nature of the assurance framework and that risks are successfully being mitigated. Finally, we show the importance of automation of controls in easing operational costs; providing improved audit information and changing the risk mitigation landscape.

* Internal Accession Date Only

Approved for External Publication

Published in the Proceedings of DIM'07, November 2, 2007, Fairfax, Virginia, USA
© Copyright 2008 Hewlett-Packard Development Company, L.P.

On Identity Assurance in the Presence of Federated Identity Management Systems

Adrian Baldwin, Marco Casassa Mont, Yolanta Beres, Simon Shiu

HP Laboratories, Bristol, United Kingdom

adrian.baldwin, marco.casassa-mont, yolanta.beres, simon.shiu@hp.com

ABSTRACT

In this paper we address the appropriate management of risk in federated identity management systems by presenting an identity assurance framework and supporting technologies. We start by discussing the risk mitigation framework that should be part of any identity assurance solution. We then demonstrate how our model based assurance technologies can be used to report success of an identity assurance programme. We discuss how this approach can be used to gain trust within a federated identity management solution both by communicating the nature of the assurance framework and that risks are successfully being mitigated. Finally, we show the importance of automation of controls in easing operational costs (and we describe related approaches developed at HP Labs and PRIME project); providing improved audit information and changing the risk mitigation landscape.

Categories and Subject Descriptors

K.6.4 [System Management]: Management audit, quality assurance.

General Terms

Management, Measurement, Security, Legal Aspects, Verification.

Keywords

Identity, risk, assurance, models, control, audit.

1. INTRODUCTION

In many senses identity management is a mature discipline within enterprises. There are standard technologies for single sign on, directories and for group or role based access control. However, many aspects remain procedural and reliant on people doing the right things. This makes identity assurance [1], i.e. the process of ensuring that identity management is under appropriate control, difficult.

As industries move more to outsourcing of IT, business processes, and ultimately to federated services the reliance on process and people becomes more problematic. In this context, standards for federated identity management are being worked on, but these focus on extending the reach and interoperability of identity technologies. Such approaches seem unlikely to address questions such as: how a business can convince an auditor that they have sufficient control and visibility of the people and processes being applied by service providers a few steps away and outside of their control. Identity assurance is all about ensuring that these processes are well controlled and therefore risk is mitigated.

The thesis and contribution of this paper is to show that technology can be used to directly support and improve the process of federated identity assurance. The technologies used bring together previous work on policy enforcement, that was done as part of PRIME project [14] and model based assurance [2]. We believe these technologies address identity management in very different ways than traditional focus for Identity Management, which often improve or address point areas. Without addressing the identity assurance issues it is unlikely that federated identity systems will be adopted for many enterprise tasks.

The problem of federated identity assurance is not often discussed, so the next section spells out the problem and why it will remain important to address. Section 3 provides an overview of a technical framework in support of identity assurance with Section 4 shows how model based assurance can be applied to this problem. Section 5 addresses issues of automation to improve risk mitigation and how this simplifies the assurance problems. Section 6 discusses the wider significance of the research, how it relates to standards for identity management, and the audit industry.

2. IDENTITY ASSURANCE

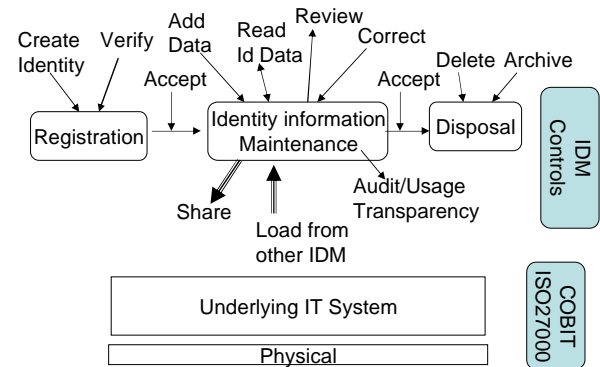
Identity assurance is concerned with the proper management of risks associated with identity management. The term “identity management” [5] is currently associated to technologies and solutions, mainly deployed within enterprises, to deal with the storage, processing, disclosure and disposal of users’ identities, their profiles and related sensitive information. They provide the following core functionalities: (1) storage, indexing and retrieval of identity information. Related technologies include databases, LDAP repositories, meta-directories, virtual directories, etc; (2) identity and credential certification; (3) authentication,

authorization and audit; (4) users' self-registration, provisioning and user account management; (5) single-sign-on and federation. These technologies can be combined to provide identity management services such as: identity lifecycle management; federated identity management; policy-driven access control; privacy management.

As anticipated in the introduction, processes define how identity information has to be managed; identity management technologies ease the burden of dealing with them, by automating some of the related operational aspects. However, it is of paramount importance to ensure that these processes are well controlled and therefore risk is controlled – hence the need for identity assurance.

Prior to defining an identity assurance framework, a risk analysis needs to be carried out identifying the identity assets (e.g. user accounts, user profiles, user rights, etc.) and the impact if there is a loss of confidentiality, availability or integrity along with threats that could lead to such losses. From an understanding of risks an enterprise can make decisions about the control objectives (strategies for mitigating risks) they need and ultimately design the controls that need to operate to achieve these objectives. Typically controls will be additional stages in management processes designed to mitigate risks (e.g. an approval step) although they may be technological mechanisms. Performing such a risk assessment from a clean sheet is complex and so many will rely on best practices such as COBIT [6], ITIL[7] and ISO 27000 [8] for general IT and IT security; here we seek to explore identity specific issues.

In understanding identity assurance it is important to understand how the identity assets within an enterprise are created, managed and used and hence we start this discussion by looking at the identity information lifecycle. Figure 1 shows such a life cycle from the initial registration of a new identity through the management of personal information associated with the identity and finishing with its disposal.



A set of operations concerning how the identity assets are managed and used are also shown – it is these operations that must be properly controlled according to the identity provider’s policies to mitigate risk. Here, we use the term “identity provider” to refer to an entity that collects identity information (of customers, employees, etc.), process it and potentially discloses it to other parties. The role of “identity providers” is central to contexts involving federated identity management, simplifying users interactions with service providers (e.g. via single-sign-on) mechanisms [9].

Underlying the identity management solution there are a raft of IT systems where risks must also be managed to gain assurance about the identity processes. Equally there must be physical safeguards ensuring a safe and secure IT operating environment. Companies have experience in operating to frameworks such as COBIT, ITIL and ISO 27000 and automation tools [2] help in documenting, monitoring and communicating compliance to these frameworks.

Whilst trust in an identity provider will be coloured by their ability to run their IT systems, it is the management of the processes around identity management that are the most critical aspect in building trust in an identity provider. As with all information assurance, controls are often process driven rather than technology controls – although technology can help in the automation of the operations, monitoring and sharing of the controls.

Fig 1 identifies a number of operations within the identity management life cycle and here we will briefly examine some of the factors that need to be considered with in an identity assurance framework. The degree to which any controls operate will of course depend on the types of identity information and associated risks. Here we list some of the risks with some sample control objectives:

a. Create Identity

Risk: An identity is created that doesn’t correspond to the physical or virtual being that it is intending to represent.

Control Objective: The registration process should ensure that enough documentary evidence of sufficient quality has been provided.

Risk: Checking process fails or is bypassed.

Control Objective: Ensure that those operating the registration processes are fit and proper for the task and ensure that they have had adequate training.

Figure 1. The information management process, operations, and controls

be a single IdP for all identities and hence there will be a need for IdPs to share information.

Within the CoT there is a reliance that each participant is doing the right thing and properly managing the identity data with which they are provided or that they help to create/augment. Despite this mutual reliance the service providers do not necessarily have trust relationships. There are four critical trust relationships: an ID Subject has to trust the IdP and its circle of trust; the IdP needs to trust that the service providers will correctly manage identity information; the service providers need to trust that the IdP provides good accurate identity information; the service providers also need to trust that the identity provider ensures that all members of the circle of trust (i.e. other service providers) behave properly.

These trust relationships can be enhanced by ensuring that the correct identity and IT assurance frameworks are in place. The identity provider will have an assurance framework to manage risks very much like that described at the start of this section. They need to ensure that the service providers have controls mitigating risks around ensuring that the identity information is correctly used; when additional information is associated with an identity that this has been verified; and that data is not retained longer than necessary. They also need to know that the service providers IT systems are well run and that, for example, identity information they hold is not accessible to the internet but is held in a well managed database behind a firewall.

The service provider needs to understand the IdPs' assurance framework to ensure that the controls are sufficient to mitigate risks around their use of identities. For example, if the IdP has a simple self registration process this is clearly not suitable for a service provider relying on the identities for financial transactions.

The service provider should also look at the assurance framework that controls the way others become members of the circle of trust and around usage transparency and incident tracking within the CoT. Such additional elements in the assurance framework are necessary due to risks associated with federation. Here we would have a set of additional risks sample control objectives

Risk: Service providers misused identity information

Control Objective: Check and regularly review each service providers controls over the use of identity information.

Control Objective: Ensure that each service provider has well run IT systems and applications within the boundaries receiving Identity information.

Risk: No accountability for handling of identity information

Control Objective: Ensure that there is a logging system showing when a service provider gets information about a given Identity and when they destroyed the information.

Risk: Failures in the identity controls or identities are not recognised.

Control Objective: Ensure that there is an incident management system where problems with identity information can be reported and logged.

Risk: Enforcement of controls is not possible.

Control Objective: Ensure there is a contractual relationship behind the circle of trust.

Having such an assurance framework that mitigates risks associated with membership of the IdPs circle of trust help build trust in the consistency of the overall federated identity assurance framework. Assurance issues become complex where the circle of trust becomes a dynamic domain.

From an identity perspective the ID subjects trust relationship is with the identity provider; that is not to say they need no trust in the service provider but that trust is about service delivery. The identity subjects therefore need to be assured that the identity provider is properly managing their identity and the CoT members who may gain access. Identity subjects will probably have far less specific concerns with trust and assurance being gained via the ability to review and correct their information and through usage transparency allowing them to see how their identity has been shared.

A much more complex trust relationship exists outside of the bounds to the CoT. For example, identities may be shared between IdPs or IdPs may act as a conduit through which a service provider can find out about a customer who has been verified by an alternative IdP. Where the IdP mediate federated identities it is up to them to ensure that each party has appropriate assurance controls.

Where the identity provider is passing on identity information from other providers they need to mitigate risks around the different potential meanings of the information. That is they need to ensure that the processes in registering, validating identities and managing associated information are strong enough for the reliance that their customers place on them.

Much of this discussion has assumed different parties can see into the internal controls operated by other entities. Clearly this level of transparency is impractical and limiting the sharing of assurance data is addressed in subsequent sections.

3. IDENTITY ASSURANCE FRAMEWORK

From the identity assurance section it is clear that trust in identities and associated information could be very much enhanced by having an identity assurance framework encouraging transparency over how identities are managed. However there are limits to the degree of transparency that is acceptable and appropriate for service providers. The control processes can themselves be complex; often manual and hence error prone. Hence within an identity assurance framework we need automation support for both checking controls are correctly operated and to simplify controls with policy based technologies.

Clearly there are issues with the free form sharing and assessment of assurance information. Companies will not share details of their internal processes and should not share detailed audit samples showing that they are run correctly. In trying to gain trust in a provider a detailed assessment may be too expensive a process for the required level of trust. This implies that there needs to be common standards and ontology for the sharing of identity process information – technology needs to support the mapping of the standards to the services controls. Tools also need to relate the results of automated or manual controls testing to the claimed standards. To fully enable the

sharing of identity assurance information, standard frameworks need to be underpinned by appropriate legal and regulatory frameworks.

It is, of course not sufficient to assert compliance to a given identity assurance standard there needs to be an evidence trail that is auditable and producible in the case of a dispute. Currently within outsourcing contracts such trust leads to clauses requiring the ability to audit or audit certifications such as SAS70; however, such an approach is manual and costly. Having a framework that supports automated audit testing ensures an evidence trail has been created and it can be retained with appropriate integrity and confidentiality [10].

As well as the macro level assurance produced by such an assurance framework there needs to be a usage transparency service creating details of how each identity is used. Secure audit technologies [11] can be used to create such an evidence trail that is accessible and verifiable by each ID subject. Linking this to the macro identity assurance systems for each identity provider and service provider in a federated identity system ensures there is a complete assurance picture for each individual.

This paper aims to demonstrate how automation technologies can be used in delivering such a federated assurance framework. The next section describes how the assurance information and performance can be tested and shared. The following section provides examples to demonstrate how policy enforcement systems can be used to automate the controls; hence changing the risk landscape and simplifying the assurance models.

4. MODEL-BASED IDENTITY ASSURANCE

This section introduces and discusses our work to enable an identity assurance framework. It starts with a brief overview of a model based assurance system allowing the automated testing of controls and the creation of risk views on this information. This solution provides the basis for our framework on the management and sharing of identity assurance information described in 4.2 onwards.

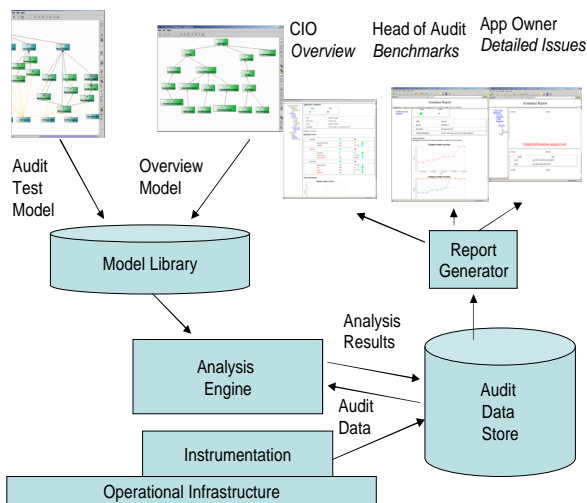


Figure 3. The model based assurance system

4.1 Model Based Assurance

A model based assurance management framework has been developed and piloted with considerable success [2]. It allows an enterprise control framework to be captured in a series of models. These models range from detailed audit test models to those producing a risk overview. The models both capture the control environment and enable automated risk reporting solutions. As well as providing for automated testing the controls models capture the enterprise assurance environment. For many companies these exist in a series of documents or spreadsheets. Having captured the assurance framework in a structured way along with automated (or manual) test results we can now manipulate and share this data in different ways.

The model based assurance system (Figure 3) consists of tools to support the creation of assurance models or customisation of models from a standard model template library. There is a data collection system to load an audit database with information relating to controls. An analysis engine takes the assurance models and applies them to information in the audit database. A report generator then shows the analysis results.

The models range from detailed models capturing tests that would currently be performed manually by an auditor to overview models. The overview models contain a specification of the control framework detailing the relationships between risk, control objectives and controls on different parts of the enterprise IT architecture. These models can refer to the results of detailed test models and other overview models and hence can be used to create different view for different audiences. At the lowest level the audit automation model encodes tests carried out by auditors (using a test library); for example, to look for system users who are no longer employees, or to find *segregation of duty* (SoD) issues. This derives detailed results listing users violating policies so that the application owners can make immediate corrections. The higher levels of the report include traffic light indicators based on comparing certain metrics (e.g. number of SoD issues) against a threshold function to produce a status for a given control area. Examples of such models can be found in [2], where they have been applied to support audits of account management processes.

Benchmarking across systems is performed at the level of these control areas based on the status values. Here results that are lodged against a given set of systems (e.g. financial applications) can be compared (even as the details of the tests vary) and graphed over time (based on having standard control naming). This would allow those with oversight functions such as the director of audit or the manager of a group of applications see where they should apply their efforts.

The highest level models lead to a high level dashboard showing how the enterprise is performing in mitigating different risks. An overall traffic light is given in the report along with ones for each of the risks identified in the model. Such models can provide a quick overview for the CIO or CISO.

4.2 Application to Identity Assurance

IdPs and service providers each need assurance models relating to how they handle identity specific processes. These models should also refer to standard IT controls showing that the underlying IT systems are also under control. The Identity provider's assurance model covers aspects of the identity management processes as identified in section 2 with figure 4 showing a potential organization for grouping controls within the assurance model.

Taking a few examples under the identity creation area there should be high level control objectives around the registration and verification process. The registration control objectives would further decompose to include objectives around how identities are checked and how and what authentication information is captured (e.g. passwords, biometrics). Other control objectives under the creation area would include controls on the staff running the processes both to check that they are adequately trained and that they are fit and proper people for the task. Each of the control objectives within identity information management would again be further decomposed into sub objectives.

Each of the control objectives within the identity provider model is identified and described by a number of attributes. To support federation this should include references to control objectives identified within an identity assurance standard and where the standard contains options of the level of checking or control it should refer to the level that is intended to be achieved. These are abstract concepts within the model that serve to document the aims of the controls on identity processes and serve as an organisational structure to report performance against the controls. Under each of these control objectives there would be a number of detailed tests that are modelled based on the detailed controls implementing the control objectives.

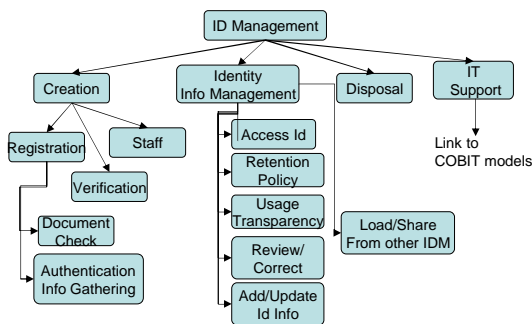


Figure 4. Organisation of an identity assurance model

The service provider will have a different set of control objectives that control how identity information can be requested and used. This will include access management around who can request information and for what purpose; data management controlling how data is managed within the business once received. Again there will be a high-level organisation of the assurance information in a hierarchy with tests of controls at the lowest layers that can be automated and used to show compliance with policy.

Other elements in the identity assurance models would include a number of key risk indicators (kri) [12] that are indicative of well run controls. For example, a kri for an IdP could be to look at the percentage of customers reviewing their data that lead to a correction. Such a figure can be indicative as to how well the information is initially captured and maintained; even if the controls are correctly operated.

Once the models are populated with low-level control checks which are plumbed to the identity management systems automated testing of the controls can be performed. The analysis engine within the audit framework will run each of the tests and propagate the results up; through threshold functions; and then combining status results to produce a report following the hierarchy of the model with red, yellow and green lights showing compliance to each (or groups) of control objective(s). Reports are generated at regular intervals so that compliance trends can be seen and differences highlighted.

4.3 Mapping Identity Assurance Models

Having a model based assurance system can help each stakeholder in managing their own compliance and also forms the basis of a federated identity assurance solution. The model itself describes the control objectives to which an organisation is trying to comply; the detailed tests underneath provide much more information as to how they are being achieved. The reports generated against these objectives show how well they are being achieved. Hence the sharing of assurance reports would meet many of the goals around federated identity assurance; however, as discussed earlier there are practical issues with the ease of assessment and confidentiality of such reports.

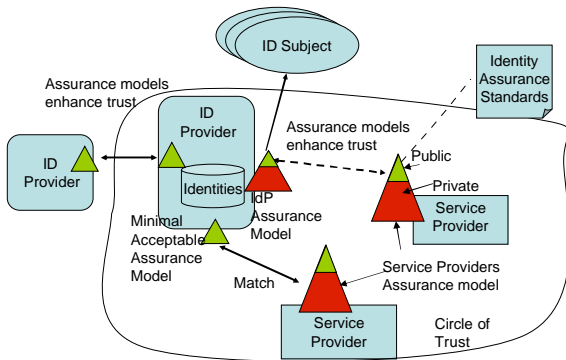
As an IdP I should be willing to share with customers, identity subjects and partner IdPs the standards to which I adhere. These standards are set out in an organised and readable manner in the high levels of the assurance model. Further I should be able to proudly state that I am compliant to this assurance model showing I've maintained compliance over time. Following this argument the high level assurance model and corresponding reports should be shared within my CoT.

This may lead to questions as to how my partners can trust this information correctly reflects reality – these models and reports could merely be my assertions. This is where trust in the assurance model and reports needs to be built by a combination of third parties and the existence of evidence. The first trust question is: are my processes sufficient to claim that I meet the control objectives to the specified level? Such questions may be a matter of trust based on my brand and public assertions or it may be that third party reviews and certifications are required. This could be done either as a whole system or piecemeal with certifications being contained as attributes within the assurance model.

The second trust question is whether I am operating to the claimed controls. The automated assurance reporting contains all the evidence to support these claims but sharing such data is inappropriate. In using such a framework I am asserting that the data exists (and is archived) hence there is the ability for a trusted third party auditor to validate the data. This could be done on an occasional and sampled basis or where there is a dispute the data exists to show control has been maintained. Trusted audit solutions [10] can then be used to ensure the

integrity of assurance evidence is demonstrable. The loss of this data could be taken as an admission of failure.

Following the above argument we can mark parts of the identity assurance model as 'public' to be shared with partners and parts as private - as shown in figure 5. Assurance can be given to partners by sharing public parts and via audits and certifications on the details. There may be a need for more complex trust relationships involving different levels of sharing with different partners. In this case alternative overview models can be created for these different customers.



From the other perspective, receiving and reviewing the public assurance information could be quite a task particularly if there are not clear standards. If we assume that there are identity assurance standards as someone relying on identity information we could build a minimum acceptable assurance model (MAAM). Having two machine readable models we can now do a simple comparison and validate that the assurance model being validated is at least as strong as the MAAM. This could be a simple binary decision or a more complex comparison could lead to a trust indicator whose value is based the presence of optional controls (perhaps weighted by the importance of such controls).

Judging performance against the identity assurance controls could be done simply by looking at the top of the given assurance report. However, in designing a given MAAM an alternative report could be produced based on the looking at the compliance with the individual goals in each report. One way of further gaining assurance is to look at the key risk indicators and some standard *kris* should be defined and also included in reports allowing further customisation of the judgement as to whether a partner is meeting their assurance goals.

4.4 Transitivity of Identity Assurance

When identity information is shared between federated providers it is necessary to gain assurance about the overall identity set. Consider the example of an IdP bringing in or representing identities from other providers. The assurance model and report now no longer accurately represent the assurance for all identities. Instead the MAAM represents this and this can be used to provide an accurate view over both the internally and federated identities that are being managed. This is the case since all the trust relationships have been based on the assessment of this model.

4.5 Usage Transparency

Assurance at the level of individual identities can be partially achieved by gaining assurance over the way the overall set of identities is managed. However, this is a broad brush assurance and does not help show an individual that their identity has been used properly. A secure usage log can be created and shared with each ID Subject using trusted audit techniques [11] allowing the user to verify the details of how their identity has been used. Giving each ID subject the opportunity to validate this also helps ensure those accessing identities are accountable for the way the use them in a way that even carefully designed controls could not.

4.6 Overview of Identity Assurance Information

From the perspective of the CoT, the assurance model and performance reports represent an overview as to how well risks with identity management are being mitigated. As identities are passed across the IdP domains there is no clear authority for overall identity assurance information. One option would be to have a regulatory authority ensuring the overall compliance to the identity assurance standards.

From an individual's perspective the usage transparency log could form an overall assurance record of everyone who has touched their identity. For this to be the case not only does every interaction need to be logged, but the message format must include a reference to the public elements of the assurance model of each party. This creates an overview of all the assurance information for each individual.

5. IDENTITY ENFORCEMENT POINTS

The model based assurance framework helps demonstrate that control is being maintained over identity information; but it does not in itself help reduce risk or ease the pain in running appropriate controls. Here other automation technologies can change the risk landscape and hence reduce or ease the amount of assurance information that needs to be collected, analyze and reported. Within HP Labs we have developed a number of automated identity based enforcement points that reduce operational costs and reduce likelihood the human-based errors or possibility of fraudulent use.

In particular these enforcement points are driven by privacy policies, organisation guidelines and users preference and help manage the lifecycle of identity information. The mechanisms enforce these privacy policies so that those tempted to override or break policy would have to hack or workaround the policy

Figure 5. Using assurance models to enhance trust between identity stakeholders

enforcement systems.

5.1 Privacy-aware Access Control

Privacy-aware access control is required to ensure that identity information is only accessed upon satisfaction of predefined policies and users' preferences. This is particularly important to preserve privacy. Traditional access control solutions (that involve users, their roles, protected resources and access rights) are necessary but not sufficient in the enforcement of privacy constraints over identity information. These solutions need to be

extended to keep into account the purpose for which data has been collected, consent given by data subjects and other conditions.

This work focused on research and development of a *privacy-aware access control system* [3] that enforces privacy policies (defined by privacy administrators and based on data subjects' privacy preferences) on personal data stored in heterogeneous enterprise data repositories. In this system, privacy policies explicitly define the purposes for which personal data can be accessed, how to keep data subjects' consent and which actions need to be fulfilled at the access time (filtering-out data, blocking access, logging, etc.). Our solution provides the following key functionalities: it allows (1) administrators to graphically author policies involving both privacy and access control aspects; (2) fine-grained modelling of personal data (stored in relational databases, LDAP directories, etc.) subject to privacy policies; (3) deployment of policies and decision-making process based on them; (4) enforcement of these policies at the data access time; (5) logging and auditing capabilities.

At run-time, our solution transparently intercepts attempts made by applications and services to access personal data stored in various repositories. This is achieved via *Data Enforcers* – i.e., privacy-aware *Policy Enforcement Points* (PEPs). Multiple Data Enforcers can be used, one for each type of data repository. A Data Enforcer component extracts relevant information from queries (e.g. requestor's credentials and any metadata) and asks the *Policy Decision Point* (PDP) to make a decision based on relevant privacy policies. This decision could allow a data requestor to have partial access to data subject to the satisfaction of associated constraints. Decisions made by these PDPs, related enforcements made by PEPs and the overall contexts are logged and can be further analysed by the assurance system for compliance checking and to report privacy violations in a wider context of identity assurance.

The audit capability provides fine-grained log information usable by the model-based assurance system to provide identity assurance reports. Having such systems in place ensures that controls around the usage of personal data are enforced as a standard component of the software system and hence need not be checked in detail. Thus the risk of misuse is very much reduced and this can be reflected in the pruning of the identity assurance model.

5.2 Privacy-aware Information Lifecycle Management

As well as controlling access to identity information according to privacy policies it is important that identity information is managed throughout its lifecycle. Hence policy systems are needed to manage privacy obligations, such as duties and expectations on data deletion, data retention, data transformation, etc. For example, data might need to be deleted after a predefined period of time, independently from access policy. Traditionally these life-cycle management tasks would be carried out by manual review –obligation management technologies automating these tasks again ensure that risks around identity lifecycle management are reduced.

Our obligation management model has been developed in the context of the PRIME project [14]: it includes an Obligation

Management System (OMS) that explicitly manage privacy obligations on personal data, providing the following functionalities: (1) explicit representation of obligations as *reaction rules*; (2) scheduling of obligations; (3) enforcement of obligations; (4) monitoring of enforced obligations. Obligations are automatically derived from privacy preferences (e.g. requests for deletion or notifications) expressed by ID subjects/or administrators. These obligations are scheduled by the OMS system based on relevant events. If triggered by these events, OMS enforces privacy obligations, for example by deleting data, sending notifications or triggering workflows. Enforced obligations are monitored for a predefined period of time for compliance reasons. A fully working prototype of this system has been developed demonstrating the feasibility of such automated identity lifecycle management.

Obligations are associated to identity information either within an enterprise or disclosed to third parties: their enforcement has an impact on the overall identity assurance. The automation of obligation management processes further simplifies the definition and need for controls in an identity assurance model. Instead of checking policies are correctly enforced on each piece of identity data we need only check that the obligation system is functioning as expected with the correct policies.

Further R&D work is currently carried out in the context of PRIME, to ensure a modular and scalable approach to the control and enforcement of privacy obligations as well as provide rich, audit logs about the OMS system.

6. DISCUSSION AND RELATED WORK

Assurance requirements and processes are defined by regulators, auditors and groups such as the Public Company Accounting Oversight Board (PCAOB). These groups focus on improving the assurance process almost independently of the technology involved. At the same time standards groups for Identity Management such as the Liberty Alliance [9], are focused on extending technologies for identity management and ensuring they inter-operate in a federated context. There is a gap between these activities, and this section discusses how the technology described here can shape and improve the way assurance can be done in federated environments.

We have shown how technology can be used to define and orchestrate the information collection and analysis needed to assure stakeholders that identity management risks are mitigated across a federated environment and how significance of risks are changed using policy enforcement technologies. This suggests the assurance modelling provides criteria for judging the value of different identity management technologies. For example, it might show that little is gained from a risk perspective by using a certain kind of biometric system whereas the use of a good single-sign-on system and directory reduces many risks.

Often current system designs include a security review but take little account of the overall operational environment; it is rare for auditors to be consulted up front. This results in systems where it is hard or expensive to gain adequate assurance although the costs associated with SOX are starting to drive changes to this approach. From this perspective there has been a lot of interest in automated controls testing and the PCAOB has recently released

draft guidance [12] including provisions on the reliance of benchmarking and automated controls. The intension of these guidelines is to have a more principled approach to designing auditable systems. This debate, centred on financial reporting, is concerned with the tradeoffs between benchmarking vs. designing controls in a risk based way that supports audit.

There is little work specifically addressing federated identity assurance. The IAAC group have run workshops on the topic and produced a paper [1] that describes the problem, with slightly more emphasis on individuals and citizens. The paper suggests a framework is needed that takes account of the numerous stakeholders, and that IAAC will be active in leading the community. In many ways this paper can be read as a contribution to that agenda showing that technology can and should play a role in the resulting framework.

7. CONCLUSIONS

Federated identity management is a complex area with a lot of technology, standardization and research effort. This paper has shown that identity assurance puts a different, important and often overlooked perspective on the problem. It was shown that a framework for assurance will be necessary, further it has been shown that there is large scope for using technology to shape and define this framework.

More specifically the assurance modelling toolset shows how technology can be used to declaratively determine what information needs to be shared, including allowing service providers to determine and control the level of granularity that should be shared, and to automate the sharing and analysis. The combination with privacy policy enforcement shows both that technology can be used to significantly simplify the distributed controls and associated assurance.

8. REFERENCES

- [1] IACC, IAAC Position paper on "Identity Assurance (IdA): Towards a policy framework for electronic Identity", available from <http://www.iaac.org.uk>, October 2005
- [2] A. Baldwin, Y. Beres, and S. Shiu, Using Assurance Models to aid the risk and governance lifecycle. BT Technology Journal. Vol.25 No 1. January, 2007
- [3] M. Casassa Mont, R. Thyne, P. Bramhall, Privacy Enforcement with HP Select Access for Regulatory Compliance, HP Labs Technical Report, HPL-2005-10, 2005
- [4] M. Casassa Mont, Dealing with Privacy Obligations in Enterprises, HP Labs Technical Report, HPL-2004-109, 2004
- [5] M. Casassa Mont, P. Bramhall, J. Pato, On Adaptive Identity Management: The next generation of Identity Management Technologies, HP Labs Technical Report, HPL-2003-149, 2003
- [6] ITGI, Control Objectives for Information and Related Technologies (COBIT), Fourth Edition, 2005
- [7] V. Lloyd, Planning to implement service management (IT Infrastructure Library), The Stationery Office Books <http://www.itil.co.uk/publications.htm>, 2007
- [8] ISO, ISO 27000 Series of Standards (Supersedes ISO17799) – <http://www.27000.org>, 2007
- [9] Liberty Alliance Project, The Liberty Alliance Specs, <http://www.projectliberty.org/>, 2007

- [10] N. Murison, A. Baldwin, Secure Distributed audit for shared customer environments, To be issued as Technical Report, 2006
- [11] A. Baldwin, S. Shiu, Enabling shard audit data. Int. Journal of Information Security 4. Springer, 2005
- [12] 12.CCM, Continuous Control Monitoring: Enabling rapid response to control breakdowns, in research findings of Audit Director Roundtable, <http://www.audit.executiveboard.com/ADR/>, 2004
- [13] Proposed Auditing Standard, An audit of internal control over financial reporting that is integrated with an audit of financial statements – and related proposals”, PCAOB Release No. 2006-007 PCAOB Rulemaking Docket Matter No. 021 Available from PCAOB website <http://www.pcaobus.org/>, 2006
- [14] PRIME Project, Privacy and Identity Management for Europe, <https://www.prime-project.eu/>, 2007