Assurance for Federated Identity Management

Adrian Baldwin, Marco Casassa Mont, Yolanta Beres, Simon Shiu
HP Laboratories, Bristol

HPL-2008-25
March 20, 2008*

identity assurance, identity management, federation, privacy

Federated Identity Management is an emerging paradigm that is rightly getting a lot of standardization and research attention. One aspect that is not receiving enough attention is assurance. Given the challenges enterprises faced trying to demonstrate appropriate control of their internal and monolithic identity management systems, the problem of how to provide assurance to multiple stakeholders that controls, operations and technologies that cut across organisational boundaries, are appropriately mitigating risk, looks daunting.

The paper provides an exposition of the assurance process, how it applies to identity management and particularly to federated identity management. Our contribution is to show technology can be used to overcome many of trust, transparency and information reconciliation problems. Specifically we show how declarative assurance models can orchestrate and automate much of the assurance work, how certain enforcement technologies can radically improve identity assurance, and how an assurance framework can provide a basis for judging the assurance value of security technologies.

Internal Accession Date Only

Approved for External Publication

Submitted to Journal Computer Security (JCS)

Assurance for Federated Identity Management

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

Federated Identity Management is an emerging paradigm that is rightly getting a lot of standardization and research attention. One aspect that is not receiving enough attention is assurance. Given the challenges enterprises faced trying to demonstrate appropriate control of their internal and monolithic identity management systems, the problem of how to provide assurance to multiple stakeholders that controls, operations and technologies that cut across organisational boundaries, are appropriately mitigating risk, looks daunting.

The paper provides an exposition of the assurance process, how it applies to identity management and particularly to federated identity management. Our contribution is to show technology can be used to overcome many of trust, transparency and information reconciliation problems. Specifically we show how declarative assurance models can orchestrate and automate much of the assurance work, how certain enforcement technologies can radically improve identity assurance, and how an assurance framework can provide a basis for judging the assurance value of security technologies.

1 INTRODUCTION

Identity management is well supported by technology, there are standards for single sign on, authentication and authorization, 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][2] i.e. the process of ensuring that identity management is under appropriate control, difficult.

Even in the relatively mature realm of enterprise identity management, many organisations have been severely challenged by regulations such as the Sarbanes Oxley (SOX) which require them to demonstrate (i.e. provide assurance) that they have appropriate control over access to their IT systems. To be clear, demonstration here means able to provide broad and deep evidence (policy documents, events, and records) of the correct management of identities to (internal and external) auditors. The resultant audit reports help the executives, board and shareholders gain trust in the correct running of the systems and thereby the results. Rather than technology or architecture, the key problem is demonstrating that controls (typically processes such as termination of accounts, or segregation of duties) are being followed and that they actually are mitigating the identified risks.

The problem is more challenging in outsourced situations, which are the first step towards the service orientated world that federated identity systems support. For example, enterprises with subjective concerns have to trust generic providers; moreover, the service provider has the dilemma of having to share sensitive information in order to assure the customer and auditors. Historically data centres have obtained annual 3rd party certification (e.g. SAS 70 [3]) to demonstrate competence. However this is already too blunt an instrument for many situations, for example the financial service industry, through the BITS [4] shared assessment programme are creating templates for assurance better suited to their industry, and SOX [5] with more punitive consequences is forcing customers to demand more up to date (near real-time) assurance.

Another challenge enterprises are facing is how to leverage their identity management operations to allow partners and customers direct access to their applications. This is the beginning of federated identity management, and it introduces even more assurance problems. The main reason is that trust in critical parts of the identity management process is likely to cut across different organisations. To give a simple example, the decision to terminate an account needs to be taken by the business, but the action to do so is the responsibility of the service provider. This split of responsibility naturally makes it harder to collect and reconcile the necessary information.

Although the problem is mostly about the people and process challenges of assurance, this paper describes how we can use technology to address these challenges. First, we show the need for an appropriate assurance framework within which the whole lifecycle of identity management can be described. We then show how model based technology can be used to automate and improve the collection, analysis presentation and sharing of the required information.

Today's assurance processes mostly relate to IT management workflow (controls) checking, and the associated metrics and reconciliations required for this. However, we also believe that it is possible to change aspects of identity assurance, by relying on technology to enforce certain trust critical tasks. Moreover, this change will be necessary as assurance requirements become more complex e.g. through increasing privacy regulation. To address this we describe work done on rich policy enforcement that is particularly suited to assurance needs and in combination with the model based assurance, points the way to much more tractable ways of dealing with identity assurance.

Both the assurance process and the risks within identity management are not often discussed therefore section 2 gives some background to these processes and the audit methodology. Section 3 discusses the full lifecycle of managing identities, the most relevant risks that emerge, and how they are controlled. Section 4 expands this to risks in the context of federated identity management. Sections 5, 6, and 7 describe the need for and introduce a proposed framework for assurance. We then describe how model based technology which has already been shown to improve the efficiency and efficacy of enterprise identity assurance, can be used to address the specific challenges of assurance in federated identity management environments. Section 8 describes work in the area of privacy policy enforcement. Here we show how by "checking" these kinds of technologies we can simplify our assurance models and further improve the assurance situation. Section 9 discusses related work and section 10 summarises our conclusions.

2 Audit Assurance and Regulation

An enterprise must show that it is in control of risk to the business to government, regulators and partners who rely on their services. Here we address concepts around audit, assurance and regulation as background to the way we have approached the problem of identity assurance both within a simple enterprise scenario and as it is expanded to deal with federation.

Prior to defining what assurance is required, it is important that an enterprise understands the broad risk and threat landscape in which it operates along with its risk appetite. That is it should identify the critical services and information that it needs to protect along with the potential ways they may be disrupted or information changed or leaked. It also has to determine the risk appetite that will determine how willing it is to accept a risk rather than invest in mitigations.

Companies cannot look at risk in a complete vacuum, they operate within a regulatory framework that both ensures they mitigate certain risks and that risks are mitigated for customers, shareholders and the wider industry or public good. Current regulations include SOX, HIPAA, COPPA, BASEL II, and EU Data Protection Law; companies are revising their approach to risk and assurance due to the need to both meet regulation and demonstrate compliance.

Once risks have been explored, a framework for risk mitigation needs to be designed and implemented. Typically, the security office will define a set of security policies to ensure the uniform mitigation of risk across the company. Auditors and compliance offices describe mitigation strategies in terms of control objectives representing stages in processes[1] that should exist to prevent risk. In many case these follow standard frameworks such as ISO 27000 [6] for security policies and COBIT [7] for control objectives although the two frameworks can be mapped together and both map to ITIL [8] providing the overall framework for IT management. There may be additional external policies that companies must comply with due to the need to establish partnerships or trust relationships with third parties (federation); for example, a company must meet the PCI [9] guidelines if it is to be trusted to process credit cards.

When new enterprise IT systems are developed and deployed, they are designed so that they meet the necessary policies and that the procedures and processes used to manage them have controls which meet the control objectives. The resulting systems have the appropriate configurations and security checks and devices to ensure they meet the policy; for example, an access control system will be in place to control access to key resources and transactions. The correct management and maintenance

[1] When discussing IT audit we are generally referring to IT management processes but financial auditors may look for control steps within the finance processes such as accounts payable and receivable.

of the systems is ensured through process controls; for example limiting how changes are approved, who can change access rights and what changes can be made.

Auditors (internal and external) are chartered with providing assurance that risk is being managed appropriately. In general this means that auditors will look for evidence of an appropriate risk analysis, and will then proceed to check the appropriateness and effectiveness of the identified mitigations. The security office will also perform security reviews and audits to ensure the appropriate architecture and technical control mechanisms are in place. These roles have the responsibility of testing compliance to policies and control objectives, reporting both to the company board and through certifications to other interested parties thus supporting the necessary trust relationships.

This paper is concerned with assurance around the processes and systems that manage identity information and so in the next section we discuss the risks and necessary control objectives. The risks, assurance requirements and the trust relationships necessary for federated identity are addressed in section 4.

3 IDENTITY ASSURANCE

Identity assurance [1] is concerned with providing visibility into how risks associated with identity information are being managed. Identity information is never static but rather goes through various handling steps (e.g. acquisition, processing, disclosure, etc.), and for enterprises this requires having some type of identity management solution. Such identity management solutions [10] often 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 and authorization; (4) users' self-registration, provisioning and user account management; (5) single-sign-on and federation. They can also be combined to provide identity management services such as: identity lifecycle management; federated identity management; policy-driven access control; and privacy management.

3.1 Risks Associated with Identity Management Lifecycle

To perform risk analysis in an identity management context and define assurance requirements it is important to understand how the identity assets (e.g., user accounts, user profiles, user rights, etc.) are created, managed and used by an identity provider[2] and hence we start by looking at the identity information lifecycle.

[2] Rather than talking about an enterprise we will start using the term "identity provider" to refer to any 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.

In this paper we primarily focus on a “federated identity management” scenario (such as the one that can be enabled by federated identity management initiatives and technologies [10][11][12][13][14][15]) as it provides the most challenging and “interesting” aspects in terms of “identity assurance”, due to the key role played by trust. However, our observations and approach apply also to other scenarios such as single, stand-alone organizations/enterprises collecting identity information that need to assess their risks when handling identity information and report for compliance. We start by addressing this simple stand-alone case before expanding the analysis to the more general federated case.

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

Figure 1 identifies a number of operations within the identity management life cycle from the initial registration of a new identity through to the management of personal information associated with the identity and finishing with its disposal. A set of operations relating to how the identity assets are managed and used are also shown – it is these operations that must be properly controlled according to the policies defined by the identity provider. The degree of control depends also on the types of identity information and associated risks. Below are examples of some of the risks associated with individual steps in the lifecycle together with potential mitigating control objectives.

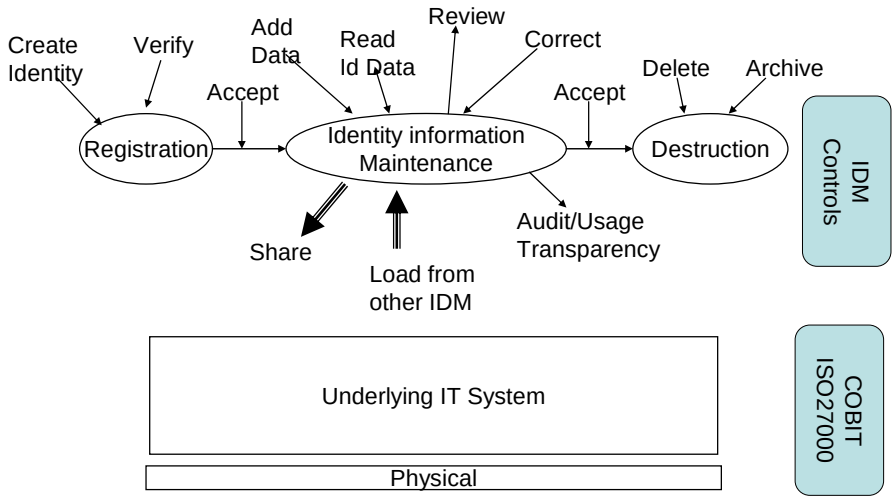

Figure 1. The Information Management process, operations and Controls

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.

Control Objective: Have a verify stage where the registration documentation is reviewed by a separate person.

Control Objective: Have appropriate IT controls over the identity database. In reality, there are a large number of IT controls that are necessary to help mitigate this risk from ensuring good access control on the database through ensuring good OS management for the host systems, good network management and physical security. These controls would typically be part of any enterprise assurance framework and in this paper we will generally refer to those risks and controls that specifically relate to the processes for managing identity.

Risk: Information associated with an individual is erroneous.

Control Objective: Ensure all additional initial information is fully reviewed.

Risk: The link between the individual and their identity is lost.

Control Objective: Ensure that the collection and management of authentication information is secure and in the case of biometric capture is carried out at well-controlled collection points by trained staff.

b. **Identity Information Maintenance**

Risk: Inaccurate information is recorded against an identity.

Control Objective: Limit those who can change and add identity information to those who have a need to perform the operation for their job. Ensure checks made on data added to the identity record are as strong as when the record is created – this may mean a review of data being added.

Control Objective: In cases where the user requests a change ensure that the user is appropriately authenticated (e.g. presence of id card, or change comes from an appropriate phone number).

Control Objective: Ensure those adding information into the identity record are recorded and can be held to account.

Control Objective: Have a review process by which data subjects can assess and correct inaccuracies in their information.

Risk: Identity information is accessible to the wrong people.

Control Objective: Ensure there is an access control system ensuring that only those with a need to access data can do so.

Control Objective: Ensure all those using identity information have appropriate security levels and management systems for the identify data.

Control Objective: Ensure different types of information accessible to different groups are clearly identified. For example, credit card details should be separate from addresses.

Control Objective: Where the need to access data is dependent on usage rules, ensure claims for usage are correct.

Control Objective: Ensure logging of data accesses and reviews of the logs so that those accessing data are kept accountable.

Risk: Data is retained for longer than contractually or legally allowed.

Control Objective: Ensure there is a data retention and deletion policy with regular reviews of the retained data.

Risk: Identity subjects are unable to access their own data and/or other authorised data.

Control Objective: Ensure the good management of credentials and biometrics used to associate identity subjects with their records. This should involve having password and password recovery policies as appropriate.

c. **Disposal**

Risk: Critical data is destroyed.

Control Objective: Ensure appropriate review of the reasons for disposal.

Control Objective: Limit those users with rights to delete identities; change aging policies and particularly any users with bulk deletion rights to a small number of appropriate administrators.

Risk: Leakage of deleted data.

Control Objective: Ensure all copies of data are removed including those shared with other parties.

The above examples are not intended to provide complete coverage but to illustrate the range of risks and mitigating controls that should be considered as part of the design process for an identity management solution, and that should be taken into account in defining an appropriate assurance strategy. The strength of a given control implementing a control objective will of course depend on the risk profile associated with the information. For example, the level of review on identity-related data should depend on the purpose for which the identity is used and the reliance that needs to be placed in the information.

Clearly, the identity management systems must also be run on IT systems that need to have appropriate controls for the level of risk associated with the use of the identities. For example, if we are reliant on the identities we should ensure the servers are run within a secure data centre with appropriate patching, network protection, security monitoring, physical protection, and so on.

4 Trust Relationships in Federated Identity Management

In the context of federated identity management, a federation of service providers revolves around an identity provider that enables users' interactions with these service providers. In such a scenario the identity provider collects users' identity information and mediates interactions and disclosure of this data. In more complex federations there might be not just one but multiple identity providers exchanging identity information. The trust relationships between these different stakeholders will often become intricate, and so we will examine them in more detail.

Figure 2 shows the different stakeholders each having different assurance needs and trust relationships. The simplest level of federation creates a *circle of trust* (CoT) between a number of service providers (e.g. within a supply chain, intra-governmental collaborations, consumer services or healthcare). An *identity provider* (IdP) within this circle will manage the majority of the identity information, but each service provider can keep additional information associated with identities.

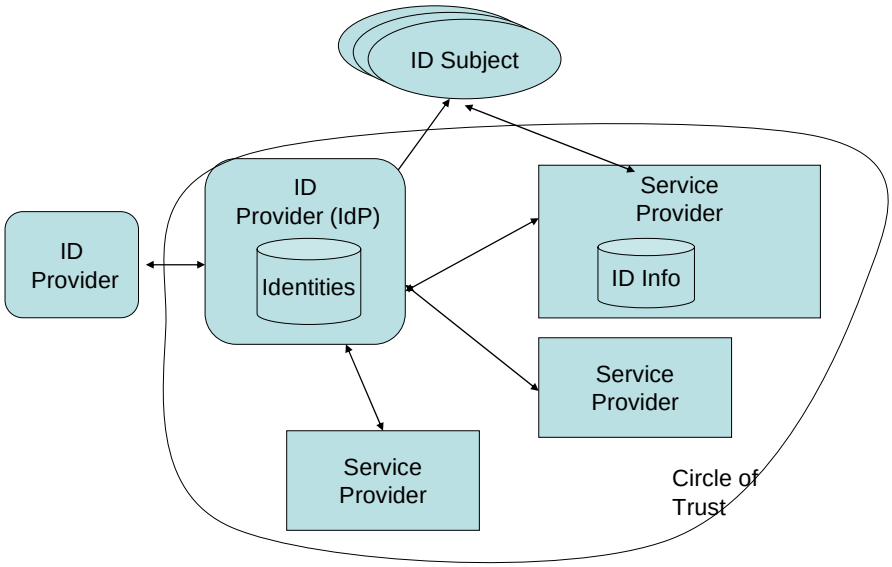

Figure 2. Stakeholders in Federated Identity Management

For example, in a healthcare scenario, a central authority within the circle of multiple health care providers is in charge of managing an electronic patient records (EPR) system identifying all patients. This central authority becomes an "identity provider" for sensitive information related to patients. Various healthcare providers communicate with it to obtain details about individuals or to update their details. In many cases, the EPR system holds a summary record for a patient, while more detailed treatment records are being retained locally with the individual healthcare providers.

Federation becomes more complex when dealing with interactions between different IdPs and circles of trust (e.g. inter-government or agency collaboration). Building on the healthcare example, we may have a situation where a patient is taken ill while travelling. The local hospital may now need to contact the remote EPR system as well as register the patient on a local EPR system. This interaction may be mediated by the local identity provider itself or may be a direct dynamic relationship between the hospital and the remote identity provider. In either case, it becomes harder to know if the

identity information to be disclosed to a party is going to be managed appropriately and so harder to trust identity providers from outside the circle. This can be even more complex when the cross boundary collaborations are dynamic and short-lived.

Each of the multiple stakeholders in the federated identity management scenarios will have different assumptions on what risks need to managing and so different requirements for what assurance they need from others. Although there is reliance that each participant is properly managing the provided identity information, the providers do not necessarily have good trust relationships.

Four critical trust relationships can be singled out in the federated identity management case: (1) an identity subject has to trust the identity provider (IdP) and its circle of trust; (2) the IdP needs to trust the service providers that they will correctly manage identity information; (3) the service providers have to trust that the IdP provides accurate identity information; (4) 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.

We believe that these trust relationships can be enhanced by having proper identity and IT assurance management in place. The identity provider has to ensure that risks similar to the ones described in previous section are properly managed and has to provide assurance showing the controls are in place. The service providers need to demonstrate that they have controls around usage of identity information, around verification process when new information is added, and around data retention. The IT systems of each provider must also be well run so that, for example, identity information they hold is not freely accessible via the internet but is held in a well-managed database behind a firewall.

In the healthcare scenario, the healthcare provider needs to trust the central identity providing authority on the quality of the information and its association with a given individual, as the accuracy of this information is critical in determining the treatment choices. The central authority in turn trusts each healthcare provider to keep information supplied confidential and more critically to ensure that all new information associated with the identity is accurate.

As we consider the relationships between multiple identity providers, there are even more trust dependencies. For example, the health record will have information submitted by a variety of professionals whose identity becomes a key aspect in trusting the information within the overall record. The submitter's identities would probably come from other identity providers – for example, doctors' identities and professional qualifications could come from a certification system run by the *british medical association* (BMA) who regulate doctors locally. Here the patient identity provider needs to trust that there are sufficient processes in place within the doctor registration process to ensure that all entries are people who have been identified as fit and proper to be doctors.

Below are some additional examples of the risks specific to federated identity management.

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.

Control Objective: Where there are relationships with other identity providers ensure there are contractual links concerning the standard of identities imported and the management of identities exported.

Risk: An identified person was inappropriately trusted with a task:

Control Objective: Check the validation of identity data coming from an identity provider.

The assurance between the identity provider's circle of trust provides a way to build trust in a consistent way. We believe that federated identity services will only work where each provider ensures that risks are appropriately managed and assurance about the implemented controls can be exchanged among all the stakeholders. In the next section, we describe requirements for an assurance framework that we believe can make this exchange possible.

5. REQUIREMENTS FOR ASSURANCE FRAMEWORK

In the previous sections we described how trust in identities and involved parties (handling identity information) can be enhanced by identity assurance and encouraging transparency over how identities are managed. There are limits to the degree of transparency that is acceptable and appropriate for service providers.

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. 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 assurance information – technology needs to support the mapping of the standards to the services controls. Tools are also needed to relate the results of automated or manual controls testing to the claimed standards. For full 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 SAS 70; however, such an approach is manual and costly. A SAS 70 [3] certification involves having an independent auditor certify that the controls reach a given standard and that the service provider follows the controls. This approach, however, is costly and fails to give the detailed insight into the strength of different controls that are required in many situations. Having a framework that supports automated audit testing ensures that an evidence trail has been created and it can be retained with appropriate integrity and confidentiality [16].

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 [17] can be used to create such an evidence trail allowing each Identity subject to verifiable access to the history of their identity. 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.

The control processes can themselves be complex and are often manual and hence error prone. Automation for enforcement of controls (with policy-based mechanisms and for checking how controls are operating) might be required to introduce further “peace of mind” in an assurance framework, as its presence (along with certified properties of these automated controls) could ease the burden of having to perform deep checks and test how controls are actually carried out in an “operational” environment..

In the following sections, we describe our solution for delivering a federated assurance framework, keeping into account these requirements and needs. In the next couple of sections, we describe how the assurance information is created by automating the testing of controls and how this information can then be shared among providers. Following this, we demonstrate how policy enforcement systems can be used to automate the controls; hence changing the risk landscape and simplifying the assurance models.

6. MODEL-BASED ASSURANCE APPROACH

The model-based assurance framework proposed in this paper for federated identity management leverages our previous, related work in the IT compliance and governance areas, mainly aiming at easing the burden of manual audits on enterprises IT systems as required for regulations such as SOX [5].

The aim was to capture the manual control testing methods and measurable aspects of the environment such as key risk indicators into a model-based approach, supporting model reuse across audits. Models are also used to automate gathering and analysis of control-related data from the IT environment. This model based assurance approach has been described in [18], [19] where we mostly concentrated on providing the necessary assurance for SOX.

After providing additional details about this work, section 7 shows how this can be leveraged to deal with identity assurance, in a federated identity management scenario.

6.1 Assurance Framework – main components

The main components of our assurance framework as shown in figure 3 include a model design tool a model repository, an analysis engine and reporting system, an audit data store, and instrumentation. The framework revolves around assurance models that capture and represent the enterprise control and risk mitigation architecture.

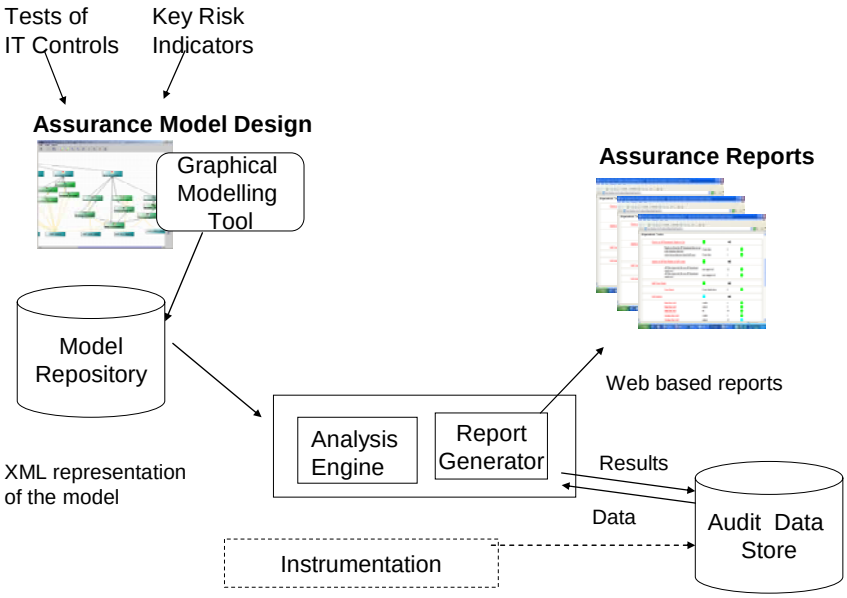

s

Figure 3. Model-based Assurance Framework

The model design tool is a graphical tool allowing auditors or other assurance staff to design models that describe how to check an individual asset, set of assets, and class of assets. For example, a model may describe the controls for a given critical application or the controls within a data centre for managing a set of servers or a model may describe controls that apply to all applications and this model can then be applied to every instance in this class. Overview models can be created to link together the assurance of these individual assets and provide the enterprise an overview onto aspects of assurance. All the models are placed within a model repository which when taken together represents the enterprise control architecture.

The models within the repository have a graph structure described later that is rendered into XML. This allows for a set of tools that support the assurance life cycle to be developed and driven from the descriptions within the model. The simplest of which is simply the rendering of the models into a human readable form (as web pages) hence making the model repository the single source for controls documentation.

The instrumentation system and associated audit data store are concerned with the collection and storage of evidence for checking and assessing the effectiveness of controls. Effective assurance requires that we can collect evidence from a wide variety of IT and management systems and assurance is only possible when agents or other plumbing for data collection is available.

The analysis engine uses the model to drive the analysis of the evidence within the audit store (as described in section 6.2). The results of this analysis are then placed back in the audit data store where they can be used for further analysis and by the report generator. The report generator (section 6.3) then renders the results of the analysis into a web-based report so that the results are easily browsed and understood.

The model design process itself proceeds in four steps:

1. Categorize IT Controls/Processes/Mechanisms needed for Assurance

2. Identify Measurable Aspects of these Controls
    a. Performance Indicators
    b. Correctness Tests
3. Build the Control Analysis Model
4. Use the model to monitor for changing conditions and to provide assurance reports

The starting point in this framework is to identify the key risks and corresponding controls that are important and what assurance information can support the evidence that these risks are mitigated. Following the health care example, we may decide to concentrate on assurance around the registration and the maintenance of a list of doctors. The key risk here is that a person might be included in the list without proper checks on his/her credentials as a doctor, and so there should be strong controls around the creation of identities and in particular around the checking and verification of evidence. More specifically a control might prescribe a list of documentary evidence that should be available and checked – for example information from a medical school and a passport to confirm the individual's identity.

Standards such as ITIL [8], ISO27000 [6] and COBIT [7] can help provide guidance as to the best practices for key operation processes and essential IT controls. A library of template controls based on these standards can be included in the assurance framework. During a deployment phase, a particular model need only be selected from the library and customised to meet the needs of a given organisation. This re-use of models can both ease the deployment process and help drive a gap analysis to find areas where controls are weak or non-standard.

The next step requires identifying what aspects of a control can be measured automatically. For example, to measure the 'doctor registration control' from the previous scenario we need to check that there are copies of evidence for all new entries on the list and that two persons had validated them. As a first simple test of this control, we may decide just to measure the existence of at least two validations by different people for each new entry. If possible, a further check could be made to validate the source of the evidence; e.g. that the person has studied at an approved medical school.

Once the measurable aspects of controls are identified, the model can then be designed for each of a set of the IT controls that will be measured and monitored for assurance. An example of such a model is given in the next section. Capturing the control testing and measurement aspects in a structured model means that it can be used for analysis, comparison as well as providing documentation that currently in organisations would typically exist across multiple spreadsheets and other documents.

Based on the model other components in the assurance framework analyse and report on assurance information at regular intervals. The intervals over which controls are checked depend on the type of controls e.g. varying based on factors such as frequency of occurrence. For example, controls managing who is authorised to access a given system are checked weekly or even monthly whereas checks that servers are patched and compliant with configuration policies may be run on a nightly basis. The analysis leads to reports concerning how well the controls are run for a given asset or set of assets (systems, applications or even business processes). Those reliant or responsible for the system can view the reports. Summary or benchmarking reports can also be generated to summarise results across multiple systems; for example, to provide a view on access management controls across all applications.

This model-based assurance framework has been piloted with considerable success [19] with audit saving considerable time on the audit tasks and also having more clarity on the results.

6.2. Representation of an Assurance Model

The assurance model (figure 3) in our approach is represented as an acyclic graph that defines the data flow upwards from the raw assurance data collected from the operational environment through a number of assurance tests that analyse sets of data – generating further sets of data and metrics as results – to metric judgement and status combination nodes. Each node within the graph has a number of incoming and out going connections as shown in figure 4 and the node describes how the data is analysed and where it fits within the controls framework.

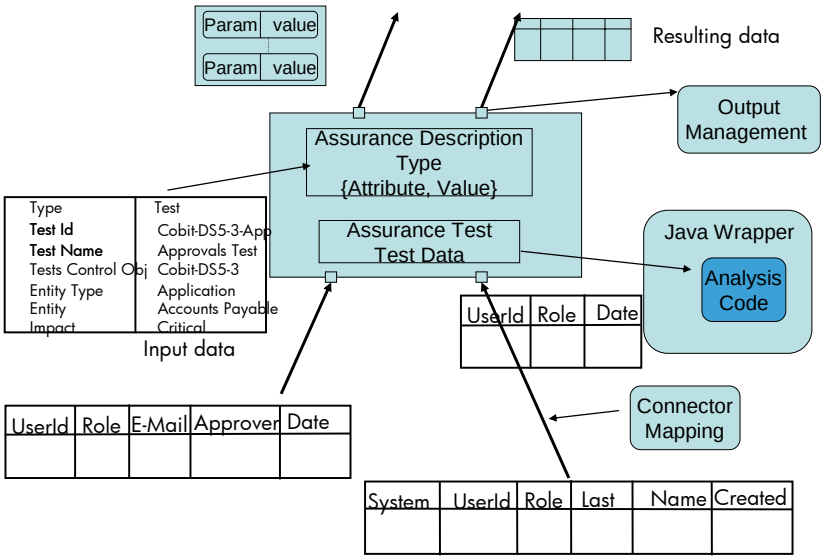

Figure 4. The Basic Node Structure within the Assurance Model

The raw assurance data collected from the operational environment are represented as nodes with no inputs. Nodes with no outputs form the overall result of a given assurance model. The data itself can take two forms, tables of data with a defined set of columns or a set of named parameters each with a given value. The parameter sets are used to represent both metrics along with status and trend values.

The assurance test[3] carried out within a node is one selected from a library of available tests and this defines the number and form of inputs and outputs. The definition of each test defines the form of input data expected and within the assurance model, each connector includes a mapping definition that maps and filters data from a source form in the required form. For example, a user table defining which users have roles within a number of ERP systems may be filtered using a system column so that only those entries for the system (asset) being tested are seen and then selected columns are made available to the test function (e.g. user id, role and account creation date). Other assurance tests may take inputs of a number of (defined or undefined) metrics or status values. Each test can have additional data specified within a model – for example acceptable values for metrics, the minimum or maximum period between two events.

[3] By assurance test we refer to some analysis carried out on data collected for assurance purposes which manually would be part of the audit testwork.

The test library contains a range of different tests from simple counts of incoming data matching a pattern, or finding groups within a data set, through tests that compare two sets of data (e.g. union, intersections and difference) or checking sequences of events (within given times). Other tests are more specific for the needs of audit, for example, finding role assignments that violate a segregation of duty matrix. Within the assurance framework we have models that generate status values or traffic lights that indicate how well controls are run and alert the viewer to issues. Here we have tests that define the relationship between the metrics and status values and for combining status values from different branches into a single status.

Along with the assurance test, the model includes a description and set of attribute value pairs describing the role of the node within the model and the asset to which the test is applied. These attribute value pairs provide useful descriptive information in generating a report. They also provide information used for creating overview reports – here we have an additional type with no direct inputs but where data is provided by searching over existing assurance test results where the test node has a set of attribute value pairs that match a given pattern. This, for example, allows results of all control tests relating to part of COBIT to be aggregated or for benchmarking of tests on a given set of assets.

Figure 4 shows a simple example of a test where the input data is a set of approval email data where users are approved for a particular role within an ERP system and a second set showing new users added to the enterprise ERP systems. This second set is filtered for the system of interest and just the user identity, role and date columns are retained. The test simply takes the difference between the two lists based on the user id and role columns (over a given time period) hence listing all unapproved users. Two outputs are shown the first would be two metrics the number of new users and the number of unapproved users; that is the size of second input set of data and the number of rows that are unmatched within the first set. The second output is a table of unapproved users – necessary evidence and useful for remedial action.

The graphical tool that has been developed as part of the framework for the construction of assurance models includes an extensive assurance test library. This allows for the creation of new assurance models and for the customisation of existing models. The tool also transforms the graphical representation of the models into an XML format for use by other components in the framework, such as the analysis engine and report generator.

6.3 Assurance Report Generation

In our approach, an assurance model describes how a given asset type, asset or set of assets is controlled and how to demonstrate that that control is mitigating risk. Once a model has been designed, it is registered against a number of assets, or sets of assets of the appropriate type. Data would need to be collected from these assets and related controls and then the model is used to drive the analysis by running each of the tests in turn. The analysis is run over a given time period (say the previous week) with the appropriate data being picked up from the audit database and with the results being written back. The analysis engine supports all the underlying data handing writing SQL queries to filter, retrieve and save the audit data.

The assurance system provides web based reports (see figure 5) that follow the structure of the model.

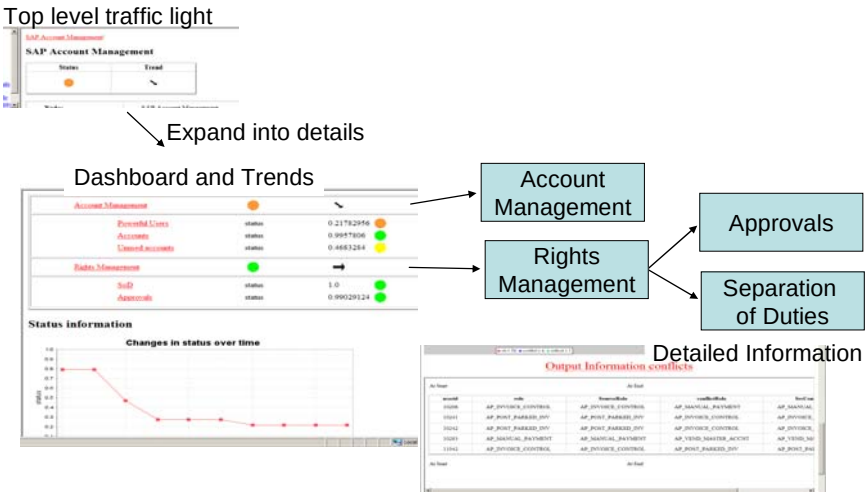

Figure 5. Structure of Assurance Reports

The first page is the top node (or nodes) where the outputs are unconnected. The user browsing the report will typically see a high-level traffic light for the overall report along with a trend showing how this traffic light varies for the asset over previous time windows. Along with the top level traffic light there is a more detailed view showing constituent status values or metrics each with a light. The traffic light status values from previous reports are also shown in a graph. The user can dig down through an area of concern, say rights management, by following the URL that effectively explores that branch of the model. This may lead to a report segment very like the top level report with an overall traffic light for that category showing more contributory areas. Eventually, navigating down the report will result in pages containing raw assurance data and metrics; for example, showing those people and roles that violate segregation of duty policies.

Overview and benchmarking reports have a very similar structure but rather than working from raw data they refer to other reports (some reports may be a mixture of the two). This can lead to reports showing how well a particular control or set of controls are run over a number of assets which provides a useful view for those responsible for some or all of the assets. Other views can follow categorizations given in standards such as COBIT or ISO27000 providing a useful overview for the director of audit, CISO or CIOs.

7 Developing Models for Identity Assurance

Our approach to identity assurance in federated identity management scenarios is based (and leverages) our previous work on model-based assurance. In a federated identity management context, both the identity provider and service providers need assurance models that can be used to measure how they handle identity lifecycle specific processes. In addition, they have to cover standard IT controls showing that the underlying IT systems are also well run.

The identity provider's assurance model needs to cover aspects of the identity management processes as identified in section 2. For example, under the identity creation step this would include high-level control objectives around the registration and verification process. The registration control objectives

would further decompose to include controls 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. Similarly, with the identity information maintenance and destruction steps we decompose the basic control objectives into a number of controls.

Following our healthcare example, during the identity maintenance stage the identity provider needs to demonstrate suitably strong controls for ensuring the correct update of doctor and patient information. For example, ensuring work locations are updated in a timely manner when changes are requested, as well as showing the presence of an incident management process to deal with the correction of erroneous details. Under each of these controls a number of detailed tests need to be modelled based on the measurable aspects of the controls. Such tests could be checking that all changes have a request, approval and an action and that they all happen within a given time period.

An example ordering and decomposition into control objectives is shown in figure 6, covering controls in identity creation, and identity maintenance. If a standard for identity assurance exists, it can be referenced by including attributes on the corresponding nodes in the model, both to identify the specific part of standard that control is measured against and also to refer to the level that is intended to be achieved.

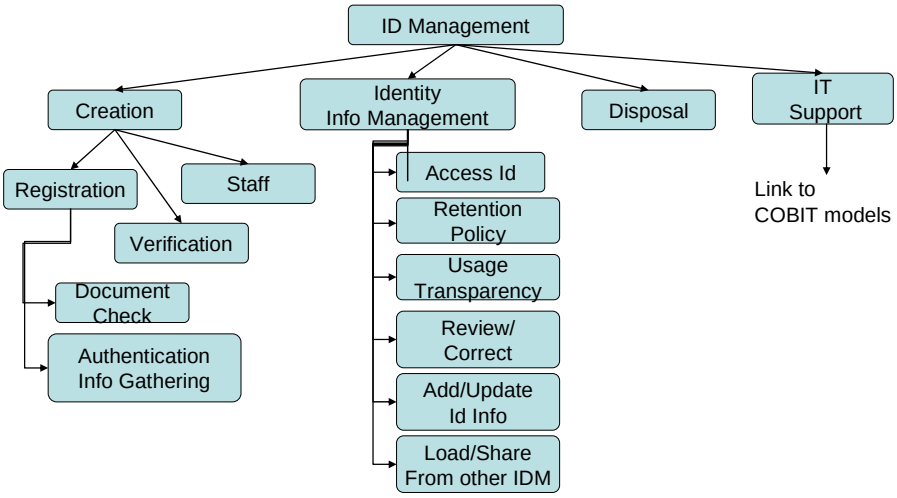

Figure 6. Organisation of Identity Assurance Model

The service provider will have a different set of control objectives that govern how identity information can be requested and used. These would include access management determining who can request information and for what purpose; data management controlling how data is managed within the business once received. Again, the assurance model would be organised as a hierarchical structure with tests of controls at the lowest layers that can be automated and used to show the strength of controls.

Other elements in the identity assurance models for both identity providers and service providers would include a set of key risk indicators (KRI) [20] that are indicative of the effectiveness of

controls. An example KRI would be a percentage of customers reviewing their identity data that lead to a correction. Such a number could be indicative as to how well the information is initially captured and maintained irrespectively of how the controls are operated.

In the healthcare example some appropriate key risk indicators would be:

- Number of requests to correct erroneous data - Where there is a well run registration and update process this number should be low and therefore, it is a good indicator of failures in the registration process, even if all controls appear to be running correctly.

- Numbers of exceptions - Many processes have associated exception processes but if their usage becomes common, this is an indicator that the overall process is badly designed.

- Number of new registrations - The number of new doctors entered every month may well be very static and predictable (low but lots in September as students graduate). Hence deviation from this pattern may indicated that unusual events are happening and hence that there are potential risks.

- Apparent Duplicate entries -- The quality of entries can be assessed, for example, looking for duplicate or similar names or multiple entries at a single address. The presence of such entries may cause concern.

- Average time to complete an entry/updates – Excessive times in the input of information may indicate risks in the process and that there is a window where information is inaccurate.

- Infrastructure KRIs – A number of KRIs could also be reported around how well infrastructure is run.

Figure 7 shows an example assurance model within our model design tool that shows various low-level control checks and metrics. The automated analysis and testing of these controls can now be performed by running our analysis engine, assuming that the required audit data has been gathered from the underlying identity management applications and IT systems. The analysis engine will run each of the tests and propagate the results up through threshold functions; and then combine 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 (non)compliance trends can be seen and differences highlighted.

Underlying the correct running of the identity management processes is the correct running of the supporting IT systems and so assurance reports should include aspects such as how the users running the system are managed. For example, checks that those with entry, change and verification roles have been approved and segregation of duty checks to assure that verification tasks are not performed by the same person that entered the data. Such controls can be simply tested, again, by comparing lists of events and finding mismatches. Other infrastructure controls would look at patching processes, anti-virus software, powerful accounts on the servers and databases, backup processes and so on.

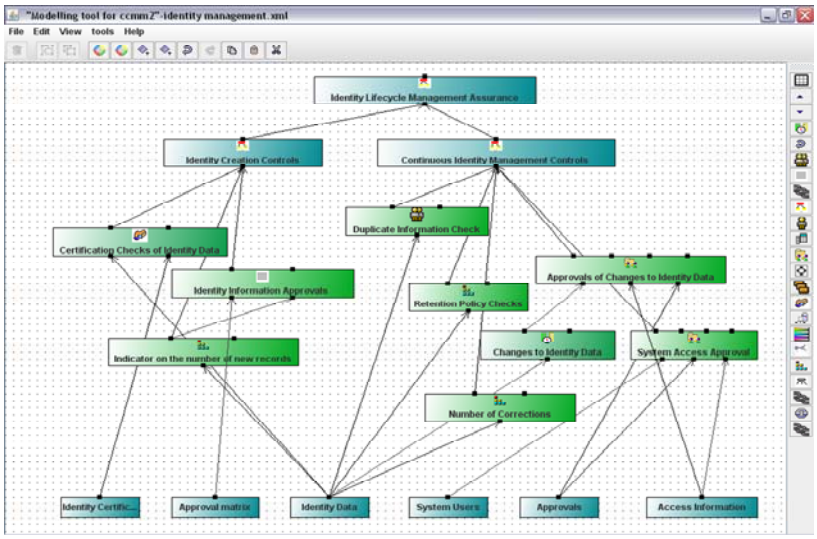

Figure 7. An Example of Identity Assurance Model Created Using our Model-based Graphical Tools

7.1 Mapping Identity Assurance Models to Federated Identity Management

Having a model based assurance system can help stakeholders in managing their own compliance and it 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 shared reports.

An identity provider should be willing to share with customers, identity subjects and partner identity providers the standards to which it adheres. Often the policies that an identity provider is aiming to meet will be defined within the relationships with other entities, for example, in the form of a privacy policy or in terms of a contract with service providers. Once captured in the assurance models an identity provider can share these policies in a form that can be used for comparisons, documentation, testing and reporting on the controls. Furthermore, they should be able to demonstrate how compliance is maintained over time by using the assurance model for automated testing. Sharing both the high-level assurance model and the corresponding report within the circle of trust should enhance the trust relationships amongst stakeholders.

This may lead to questions as to why other stakeholders should trust that the information provided correctly reflects reality – these models and reports could merely be assertions (i.e. not backed up by any data). We believe that trust in the assurance models and reports needs to be built by a combination of third party attestations and the evidence of undisputable audit trail kept by the provider.

One trust question is whether the identity management processes are sufficient to claim that the provider meets the control objectives to the specified level. Trust in the answer to such questions may be a matter of brand and public assertions, or it may be that third party reviews and certifications are required. Such reviews could be done either for the full set of processes around

identity management or just for a piece of it with certifications being contained as attributes within the assurance models and reports.

Another trust question is whether the provider actually operates to the claimed controls. We believe that this is where the assurance models and reports can be of significant use. The automated assurance reporting would contain all the evidence to support these claims. The only problem might be that the sharing of full results is inappropriate because of confidentiality and privacy issues. However, the framework allows a provider to assert that the data exists (and is archived), implying that if need arose there is the ability for a trusted third party auditor to validate the underlying data. This could be done on an occasional and sampled basis or in case of a dispute. Trusted audit solutions [17] can also be used to ensure the integrity of assurance evidence is demonstrable.

Following the above argument, it might be useful to mark parts of the identity assurance model and corresponding reports as 'public' to be shared with other stakeholders and parts as 'private' - as shown in figure 8. Assurance can be given to the stakeholders in the circle of trust and outside by sharing public parts and via audits and certifications on the private details. In more complex trust relationships, different levels of assurance information sharing with different stakeholders may be necessary. Here multiple overview models can be created on top of the main assurance model to provide these different views.

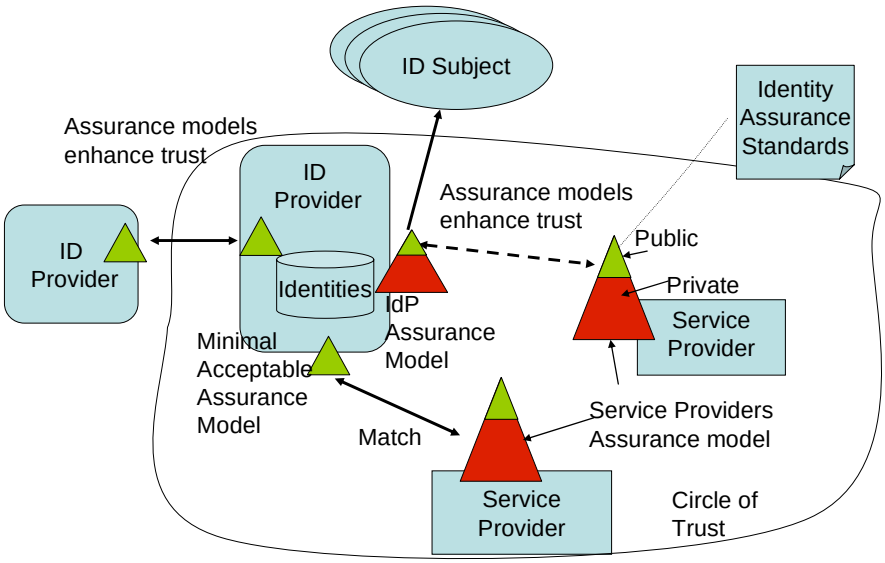

Figure 8. Using Assurance Models to Enhance Trust Between Identity Stakeholders

From the perspective of somebody reviewing assurance results given by the identity provider, this raises questions about their correct interpretation. We believe that a standard for identity assurance would be useful here specifying and communicating different controls and the level to which they mitigate risks. Someone relying on assurance models from different potential partners could build a minimum acceptable assurance model (MAAM) which expresses the lowest level of assurance they will accept (with reference to an identity assurance standard). Having two machine-readable models we can now do a simple comparison and check that the partner's assurance model is at least as strong as the MAAM and hence trust that they have sufficient controls in place. 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 the performance of one party against their assurance model could simply be done by looking at the top of the given assurance report. However, the MAAM can be used to not only to specify the acceptable level of control but also the way in which the receiving party wishes to view control information. In the same way that we create overview, reports we can map from the given assurance report to the MAAM model and hence present the information in a way that is meaningful for the receiving party. For example, this would ensure that an identity provider could see all assurance reports from their circle of trust in a standardised way and even benchmark performance across the members.

Let us examine how this can be applied in the healthcare example. The identity provider for the register of doctors (BMA) could make statements around the level of its controls – a level[4] 5 verification of entry qualifications; level 5 for change tracking, correction management and public access, with level 2 availability on all. Their assurance model contains this information encoded within the attribute, value pairs within nodes. On the other side, the electronic patient record (EPR) system offers a lower level of verification but with more controls around managing access and ensuring availability and resilience.

Before a new healthcare provider is added into the circle of trust of the EPR system, it needs to show the level of controls to which they operate. The ERP provider would have a MAAM defined, which specified the controls that they deem necessary and the associated levels of operation. For example, this may specify that it required controls of level 4 for processes involved in checking the trustworthiness of those entering new data into the system and level 5 controls for managing access to patient information. The healthcare provider needs to demonstrate that their assurance system meets these goals. This can be achieved by comparing their assurance model against the EPR providers MAAM. The healthcare provider may offer better levels but less would be unacceptable; hence, their membership would be rejected if they only operated level 4 controls on access management but would be accepted if they operated to level 5 or above.

The healthcare provider may pass on responsibility for managing controls to other providers; for example, the healthcare provider may rely on the BMA for the validation of those entering medical information. Here they may point to the BMA level 5 verification checks and hence satisfy the ERP as they effectively outsource or federate this part of the process.

7.2 Transitivity of Identity Assurance

When identity information is shared between among stakeholders from different circles of trust in the federated identity management scenario, it is necessary to gain assurance about the overall identity set. Consider the example of an identity provider bringing in or representing identities from other providers. Their internal assurance model and report now no longer accurately represent the assurance for all identities. Where they have used a MAAM to validate the other identity providers that they are getting data from this can be used to provide an accurate view over the minimal standard used for both the internally generated identities and those derived from other providers. This is the case since all the trust relationships have been based on the assessment of this model.

Following the example of checking the registration of doctors, a healthcare provider with assurance models in place can demonstrate how each identity service it uses conforms to MAAM. They could rely on the local BMA service to publish and mediate what is an appropriate level of conformance. Here the BMA would ensure that either doctors from other countries have similar levels of checks or the BMA does its own validation.

[4] Here level 5 represents a high degree of trust and control with level 1 being the bare minimum.

Whilst the minimal assurance models can support transitivity trust is not transitive. The service provider that relies on the identity information may wish to examine the assurance data from the provider of this identity information as part of its own business practices, with the result being that the identity information is simply rejected as the service provider is unwilling to trust a given identity provider (even though the evidence suggests otherwise).

7.3 Transparency of Identity Information Usage

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. Other approaches can help to achieve this. A secure usage log can be created and shared with each ID Subject using trusted audit techniques [17] allowing the user to verify the details of how their identity has been used. A sticky policy approach [21], [22] can be taken to ensure that a service provider cannot pass on identity data outside the circle of trust without appropriate audit or permission. Allowing each ID subject the opportunity to validate how information about them has been used helps ensure that those accessing identity information are accountable for the way the use it.

7.4 Combining All Identity Assurance Information

From the perspective of the circle of trust, 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 identity provider 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 also 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.

8 IDENTITY ENFORCEMENT POINTS

The model based assurance framework helps demonstrate that control is being maintained over identity information; but it does directly reduce the risk or ease the pain in running appropriate controls. Only stronger controls, better information management processes, and the control enforcement technologies can change how risk is mitigated, and hence reduce or ease the amount of assurance information that needs to be collected, analyzed and reported.

Within HP Labs, we have developed a number of automated identity based enforcement points to reduce operational costs and to reduce the likelihood of the human-based errors or of fraudulent use. In this section we describe some enforcement technologies that are driven by privacy policies, organisation guidelines and users preference. They help manage the lifecycle of identity information and enforce controls so that those tempted to override or break policy would have to hack or workaround the policy enforcement systems.

8.1 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* [23] 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.

8.2 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 [24] has been developed in the context of the PRIME project [25]: 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 being carried out in the context of PRIME [25], to ensure a modular and scalable approach to the control and enforcement of privacy obligations as well as to provide rich, audit logs about the OMS system.

8.3 Control Mechanisms and Assurance Frameworks

These examples demonstrate that certain controls can be automated and that as long as the automation systems are functioning correctly we should be able to trust that risks are mitigated. As well as specifying the levels necessary for controls the MAAM could mandate use of an (approved) automated system or recommend its usage if more complex assurance model comparisons are supported. However, the assurance model should include tests to ensure the automated controls are functioning correctly as well as KRIs that reflect its usage against possible exception processes.

9 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.

So far assurance and “identity assurance” has not been a primary concern and worry in the identity management community, despite all the issues and trust matters highlighted in this paper. There is usually an “oversimplifying” assumption that these aspects could be properly addressed by means of “legal terms” and contractual agreements among the involved parties. Despite this being a feasible way to address some of the involved issues, it is not the answer to all the trust and transparency concerns highlighted in this paper.

Standards groups operating in the Identity Management and Federated Identity Management space, such as the Liberty Alliance [11], OpenId [13], WS-*/WS-Federation initiatives [12], Microsoft CardSpace [14], Higgins [15], etc., are primarily focused on extending technologies for identity management and ensuring they inter-operate in a federated context. This is primarily achieved by defining suitable protocols and message formats (e.g. SAML) that can be used by all the involved members of an identity federation, to interact and exchange identity information.

There is a gap between these standardisation activities and the actual “identity assurance” requirements and needs. This paper discusses how the technology described here can shape and improve the way assurance can be done in federated identity environments.

Liberty Alliance has recently launched initiatives on an “Identity Assurance Framework” [26] and a related “Identity Assurance Expert Group (IAEG) [27]. Based on Liberty Alliance’s statement, IAEG is a “forum for identifying and resolving the market acceptance and commercial obstacles to broad deployment and adoption of identity assurance services. The first step has been development of a

global standard Framework (see below), which also defines support programs needed for validating trusted identity assurance service providers in a way that scales, empowers business processes and benefits individual users of identity assurance services. The Framework will be the basis upon which identity assurance providers and their services can be certified as compliant to common policies, business rules and baseline commercial terms; avoiding redundant compliance efforts and market confusion about the substance of identity assurance value delivered."

This is a promising declaration of intent, consistent with our suggested approach and ideas. However, it looks like in the short term the Liberty Alliance's Identity Assurance initiative is going to focus on specific aspects related to "certification" that consists of consolidating the Trust Framework of the EAP (Electronic Authentication Partnership), the Credential Assessment Framework of the US E-Authentication Federation, and other industry contributions. In our vision a more holistic and comprehensive approach should be taken.

In this paper we have shown how technology can be used to define and orchestrate the information collection and to automate analysis for assuring stakeholders that identity management risks are mitigated across a federated environment. We have also described how significant risks are changed using policy enforcement technologies suggesting that the assurance modelling approach provides criteria for judging the value of different identity management technologies. For example, it might show that from a risk perspective little is gained by using a certain kind of biometric system whereas the use of a good single-sign-on system and directory addresses many risks. Such analysis begins to address return on security investment, and the economics of security, which is a growing area of research, see [28].

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 [29] including provisions on the reliance on benchmarking and automated controls. The intention 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.

The authority for best practice assurance is ISACA [30], which serves as a very active professional body for auditors. This group focuses very strongly on sharing and improving best practices and for example, provides a good exposition on the risks auditors should look for in outsourcing, and an overview of best practice for federated identity management [31], [32]. Their work points to emerging challenges for efficient and effective assurance, which we have outlined and attempted to address.

Stolen et al [33] and Masacci [34] propose model based frameworks to support risk analysis, although this is primarily targeted at the requirements and solution design phases. In general, there is little work specifically addressing assurance and even less concerning a "holistic approach" to 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.

10 CONCLUSIONS

Security researchers and practitioners often, and quite rightly, point out the need to build security into IT products, services and solutions, rather than be left trying to bolt security on after the design has been completed. The large size of the assurance industry is evidence that although important it is not enough for enterprises simply to design good security into their solutions. In addition, they need to know, and be able to demonstrate that security controls are working effectively, the corollary being that assurance should be built into IT systems and processes rather than being considered as an afterthought. This paper covers these aspects for federated identity management. That is, it should be recognised that in addition to the technology, standardization and research looking at how to make federated identity management secure and effective, the assurance problem must also be considered.

The paper has outlined the top-down, risk driven nature of assurance, and particular risk and assurance issues that inevitably arise in federated identity contexts. We have shown the need for and proposed an appropriate assurance framework. We have further demonstrated that there is considerable scope for using technology to shape and define this framework.

More specifically the assurance modelling toolset shows how technology can be used to determine declaratively 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 simplify significantly the distributed controls and associated assurance.

REFERENCES

[1] IAAC, IAAC Position paper on “Identity Assurance (IdA): Towards a policy framework for electronic Identity”, available from http://www.iaac.org.uk, October 2005

[2] Dongwan Shin; Gail-Joon Ahn; Prasad Shenoy;, Ensuring information assurance in federated identity management Performance, Computing, and Communications, 2004 IEEE International Conference on 2004 Page(s):821 – 826

[3] The American Institute of Certified Public Accountants. Statement on Auditing Standards No. 70 (SAS 70) http://www.aicpa.org/download/members/div/auditstd/AU-00324.PDF

[4] BITS Financial Services Roundtable, "Financial Institution Shared Assessments Program", available at http://www.bitsinfo.org/FISAP/index.php

[5] 107th US Congress, Sarbanes Oxley Act http://frwebgate.access.gpo.gov/cgi-bin/getdoc.cgi?dbname=107_cong_bills&docid=f:h3763enr.tst.pdf

[6] ISO, ISO 27000 Series of Standards (Supersedes ISO17799) – http://www.27000.org, 2007

[7] ITGI, Control Objectives for Information and Related Technologies (COBIT), Fourth Edition, 2005

[8] V. Lloyd, Planning to implement service management (IT Infrastructure Library), The Stationery Office Books http://www.itil.co.uk/publications.htm, 2007

[9] Payment Card Industry (PCI) Data Security Standard https://www.pcisecuritystandards.org/pdfs/pci_dss_v1-1.pdf

[10] Casassa Mont, M., P. Bramhall, J. Pato, On Adaptive Identity Management: The next generation of Identity Management Technologies, HP Labs Technical Report, HPL-2003-149, 2003

[11] Liberty Alliance Project, The Liberty Alliance Specs, http://www.projectliberty.org/, 2007

[12] WS-Federation, Web Services Federation, http://www.ibm.com/developerworks/library/specification/ws-fed, 2007

[13] OpenId, OpenId Initiative, http://www.openid.net, 2008

[14] Microsoft, Micorosft CardSpace initiative, http://msdn2.microsoft.com/en-us/library/aa480189.aspx, 2006

[15] Higgins, Higgins Project, http://www.eclipse.org/higgins/, 2007

[16] N. Murison, A. Baldwin, Secure Distributed audit for shared customer environments, To be issued as Technical Report, 2006

[17] Baldwin, A., and S. Shiu, Enabling shard audit data. Int. Journal of Information Security Vol 4(4). Springer, 2005

[18] Baldwin, A., Y. Beres, and S. Shiu, Using Assurance Models to aid the risk and governance lifecycle. BT Technology Journal. Vol.25 No 1. January, 2007 (http://hpl.hp.co.uk/techreports/2007/HPL-2007-48.html)

[19] Baldwin, A., Y. Beres, and S. Shiu, Using Assurance Models in IT Audit Engagements. HP Labs Technical Report, nr. HPL-2006-148, October 2006. (Currently internal report but being made external)

[20] CCM, Continuous Control Monitoring: Enabling rapid response to control breakdowns, in research findings of Audit Director Roundtable, http://www.audit.executiveboard.com/ADR/, 2004

[21] Marco Casassa Mont, Siani Pearson, Pete Bramhall - Towards Accountable Management of Identity and Privacy: Sticky Policies and Enforceable Tracing Services - IEEE Workshop on Trust and Privacy in Digital Business 2003, TrustBus 2003 (DEXA 2003), 1-5 September, 2003, Prague, Czech Republic, 2003

[22] Marco Casassa Mont, Siani Pearson, Pete Bramhall - Towards Accountable Management of Privacy and Identity Management - 8th European Symposium on Research on Computer Security 2003, ESORICS 2003, 13-15 October, 2003, Gjovik, Norway, 2003

[23] M. Casassa Mont, R. Thyne, P. Bramhall, Privacy Enforcement with HP Select Access for Regulatory Compliance, HP Labs Technical Report, HPL-2005-10, 2005

[24] M. Casassa Mont, Dealing with Privacy Obligations in Enterprises, HP Labs Technical Report, HPL-2004-109, 2004

[25] PRIME Project, Privacy and Identity Management for Europe, https://www.prime-project.eu/, 2007

[26] Liberty Alliance, Liberty Identity Assurance Framework, http://www.projectliberty.org/liberty/content/download/3736/24651/file/liberty-identity-assurance-framework-v1.0.pdf, 2007

[27] Liberty Alliance, Identity Assurance Expert Group (IAEG), http://www.projectliberty.org/liberty/strategic_initiatives/identity_assurance, 2007

[28] Lawrence Gordon and Martin Loeb, "The economics of information security investment", ACM Transactions on Information and System Security v 5 no 4 (Nov 2002) pp 438-457

[29] 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

[30] Information Systems Audit and Control Association (ISACA), see http://www.isaca.org

[31] Nicholas Benvenuto and David Brand, "Outsourcing – A Risk Management Perspective", Information Systems Control Journal, Volume 5, 2005

[32] Leslie Pang, "A Manager's Guide to Identity Management and Federated Identity", Information Systems Control Journal, Volume 4, 2005

[33] Braber F den, Hogganvik I, Lund M, Stolen, K and Vraalsen F: "Model-based security analysis in seven steps – a guided tour to the CORAS method", BT Technology Journal 25, No 1, January 2007

[34] Giorgini P, Masacci F, Myloupos J and Zannone N: "Requirements Engineering meets Trust Management: Model, Methodology and Reasoning", in proceedings 2nd International Conference on Trust Management, LNCS 2995, 2004