Using Modeling and Simulation for Policy Decision Support in Identity Management

Adrian Baldwin, Marco Casassa Mont, Simon Shiu

HP Laboratories
HPL-2009-56

Keyword(s):

Identity Analytics, IAM, Policy Decision Support, Modelling, Simulation, Identity Management, Provisioning

Abstract:

The process of making IT (security) policy decisions, within organizations, is complex: it involves reaching consensus between a set of stakeholders (key decision makers, e.g. CISOs/CIOs, domain experts, etc.) who might have different views, opinions and biased perceptions of how policies need to be shaped. This involves multiple negotiations and interactions between stakeholders. This suggests two roles for policy decision support tools and methods: firstly to help an individual stakeholder test, refine their understanding of the situation and, secondly, to support the formation of consensus by helping stakeholders to share their assumptions and conclusions. We argue that an approach based on modeling and simulation can help with both these aspects, moreover we show that it is possible to integrate the assumptions made so that they can be directly contrasted and discussed. We consider, as a significant example, an Identity and Access Management (IAM) scenario: we focus on the provisioning process of user accounts on enterprise applications and services, a key IAM feature that has an impact on security, compliance and business outcomes. Whilst security and compliance experts might worry that ineffective policies for provisioning could fuel security and legal threats, business experts might be against policies that dictate overly strong or bureaucratic processes as they could have a negative impact on productivity. We explore the associated policy decision making process from these different perspectives and show how our systems modeling approach can provide consistent or comparable data, explanations, "what-if" predictions and analysis at different levels of abstractions. We discuss the implications that this has on the actual IT (security) policy decision making process.

External Posting Date: March 21, 2009 [Fulltext]   Approved for External Publication
Internal Posting Date: March 21, 2009 [Fulltext]

To be published in IEEE Symposium Policy 2009, London, UK, 20-22 July 2009

Using Modelling and Simulation for Policy Decision Support in Identity Management

Adrian Baldwin, Marco Casassa Mont, Simon Shiu
Systems Security Laboratory
Hewlett-Packard Labs
Bristol, United Kingdom
adrian.baldwin@hp.com, marco.casassa-mont@hp.com, simon.shiu@hp.com

Abstract*—The process of making IT (security) policy decisions, within organizations, is complex: it involves reaching consensus between a set of stakeholders (key decision makers, e.g. CISOs/CIOs, domain experts, etc.) who might have different views, opinions and biased perceptions of how policies need to be shaped. This involves multiple negotiations and interactions between stakeholders. This suggests two roles for policy decision support tools and methods: firstly to help an individual stakeholder test and refine their understanding of the situation and, secondly, to support the formation of consensus by helping stakeholders to share their assumptions and conclusions. We argue that an approach based on modeling and simulation can help with both these aspects, moreover we show that it is possible to integrate the assumptions made so that they can be directly contrasted and discussed. We consider, as a significant example, an Identity and Access Management (IAM) scenario: we focus on the provisioning process of user accounts on enterprise applications and services, a key IAM feature that has an impact on security, compliance and business outcomes. Whilst security and compliance experts might worry that ineffective policies for provisioning could fuel security and legal threats, business experts might be against policies that dictate overly strong or bureaucratic processes as they could have a negative impact on productivity. We explore the associated policy decision making process from these different perspectives and show how our systems modeling approach can provide consistent or comparable data, explanations, "what-if" predictions and analysis at different levels of abstractions. We discuss the implications that this has on the actual IT (security) policy decision making process.

Keywords: Identity Management, Identity Analytics, Policy Decision Support, IAM, User Account Provisioning, Security Analytics

I. Introduction

The process of defining IT (Security) policies within organizations is complex. Key decision makers (e.g. CIOs and CISOs) make the final policy decisions, but these are reached through a consensus-building process, involving stakeholders and experts from security, business, financial, legal and HR. It is a considerable challenge to help this diverse group bring their skills and perspectives to the discussion, whilst limiting conflicts and misunderstandings.

The main contribution of this paper is to show how modeling and simulation can support the policy decision making process by allowing stakeholders to convey consistent explanations and predictions to different audiences, at the right levels of abstraction.

We illustrate this by means of an Identity and Access Management (IAM) case study. IAM is important for protecting and securing the organizations' resources, enabling the right people to access legitimate resources for the right purposes. It is a rich area in terms of the policies that could be defined. In this context, IAM is also a business enabler and has a direct impact on business applications and services. At the very core, IAM solutions [22] provide provisioning, enforcement and auditing capabilities. In short IAM policy decisions have a direct impact in terms of people behaviors, costs, productivity, losses and availability. We focus on a core IAM capability, the *provisioning process of user accounts* to enterprise applications and services. The relevant policies might, for example, dictate levels of automation to be achieved by enterprise provisioning processes, acceptable accuracy levels, required approval and configuration times and the number of authorization requests that are necessary, depending on the context and types of resources to be accessed and protected. Relevant questions are: what are the consequences of setting particular policy decisions? Which people have relevant knowledge or concerns? How do we capture and use their inputs?

The remaining part of this paper is structured as follows: Section 2 expands on our analysis of the policy decision making process, specifically in an IAM context. Section 3 provides further details about enterprise identity management and the provisioning process. Section 4 illustrates how modeling and simulation approaches can effectively help to support the policy decision process. Section 5 describes, in more details, our approach and methodology along with an overview of the specific model we have built for the provisioning process, related simulations and the types of results and analysis that can be provided to the stakeholders. Finally, Sections 6, 7 and 8 discuss related work, next steps and conclusions.

II. On the Policy Decision Making Process

Triggers for changing or analyzing security policy can come for a number of reasons: a large number of policy "exception" requests are usually a good sign that something is wrong. It can also be any of the stakeholders (i.e. decision makers and domain experts) feeling that the inherent trade-offs are inappropriate, for example IT operations may feel the burden/resources required to maintain a particular policy is too large, or conversely a security officer may feel the threat environment has changed and so a tighter policy is warranted. In these cases either the policy can be changed, or investments and resources can be re-aligned to more efficiently meet the policy.

There are numerous challenges to helping the stakeholders, with relevant concerns and subject matter expertise, to express and share their knowledge. Fig. 1 shows the decision making framework that this structured sharing must support, i.e. allowing the stakeholders to reach one of these forms of conclusion.

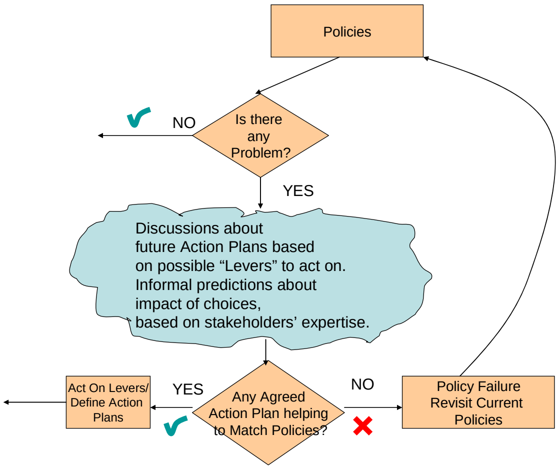
Policies
NO
Is there any Problem?
YES
Discussions about future Action Plans based on possible “Levers” to act on. Informal predictions about impact of choices, based on stakeholders’ expertise.
Act On Levers/ Define Action Plans
YES
Any Agreed Action Plan helping to Match Policies?
NO
Policy Failure Revisit Current Policies

Figure 1. Basic Policy Decision Process

The main theme of this paper is to explore and illustrate how systems modeling [16,17] can provide this support, see Fig. 2.

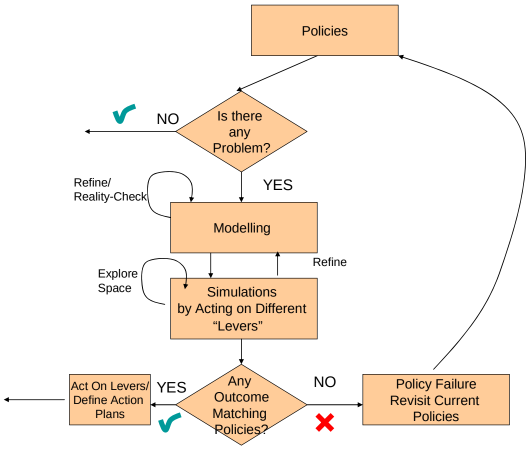
Policies
NO
Is there any Problem?
YES
Refine/ Reality-Check
Modelling
Refine
Explore Space
Simulations by Acting on Different “Levers”
Act On Levers/ Define Action Plans
YES
Any Outcome Matching Policies?
NO
Policy Failure Revisit Current Policies

Figure 2. Policy Decision Making Support

Specifically, we show how a combination of executable process models, probability theory and Monte-Carlo style experimentation (based on simulations) can be used to help stakeholders explore their own intuitions and assumptions, share these with others in a coherent and consistent way and jointly investigate the consequences of investments and policy changes. The next section provides some background about IAM, necessary for our case study.

III. Enterprise Identity Management

Identity and Access Management (IAM) solutions for enterprises [22] include functionalities such as authentication, single-sign-on (SSO), authorization, auditing, compliance and assurance management, provisioning, data storage, link to legacy systems and data consolidation. These functionalities are, in general, used for user account and access control management, federated identity management and privacy management. A more detailed description of various components and related capabilities is available [22].

For the purpose of this paper, we focus on *user account provisioning solutions*. These solutions are used by enterprises to deal with the lifecycle management of user identities and accounts on protected resources, including the enrolment, customization, modification and removal of user accounts associated with users, employees and customers along with setting rights, permissions and access control information. Getting the right *provisioning* in place is as important as getting the right *enforcement* (authentication, authorization and access control) in place. A wrong or poor provisioning process could give more than necessary rights to users or prevent them from accessing legitimate resources. This is an IAM area that is still in evolution, as are the related processes of defining enterprise users’ roles and access control permissions.

At the very core, user account provisioning solutions aims at ensuring that valuable resources (such as business applications and services) are protected against unauthorized accesses. Provisioning processes keep into account changes in the workforce (i.e. people joining, leaving, changing their roles) and organizational changes (re-organisations, large lay-offs, mergers, etc.).

Provisioning of user accounts (and access control permissions) in enterprises usually involves two phases: **(a)** ***approval phase***: the creation, modification or removal of user accounts (associated to a user, for a specific application/service) need to be authorized by one or more people that have managerial responsibilities (e.g. line managers or supervisors); **(b)** ***deployment and configuration phase***: in case of a successful approval, this phase consists in carrying out configuration activities, to actually create, modify or remove a user account on a system/application/service, along with related user rights.

Depending on the kind of adopted provisioning solution, there might be different degrees of automation, ranging from *ad-hoc, manual processes* to *fully automated and centralized processes*. The former might rely on human interactions and system administrators. The latter might involve the execution of workflows and automated configuration scripts. These phases could have degrees of failures or different implementations, depending on cultural attitudes and working environments. A typical set of IAM provisioning policies might be expressed as:

- ***P1***: Employees’ user accounts should be provisioned within an organization in max 3 days

- ***P2***: No user account must be provisioned without management approval
- ***P3***: All user accounts to be provisioned (added, modified, changed) on core business applications and services require 2 levels of approval
- ***P4***: Users accounts of people leaving a company must be removed within 2 days the departure date
- ***P5***: The accuracy of the provisioning process (in terms of correctly configured user accounts on protected resources) should never be less than 99%

The CIO, CISO or maybe a risk manager (decision maker) would be responsible for defining these policies and their appropriateness. However, policy analysis and decisions will require the input and consent ("buy in") of several stakeholders, including: **security experts**, that understand the vulnerability of the provisioning process and can articulate the technical consequences**; business experts and application/service owners**, that understand the criticality of appropriate access to business objectives, and to some extent the business burden the policies create; **compliance experts**, that are driven by the need to be compliant to internal guidelines, laws and legislation (such as SOX), being able to pass auditing sessions, etc.; **IT Operation experts**, that have an understanding of how the IT infrastructure runs along with the involved performance, service delivery aspects and costs.

IV. Policy Decision Support for Provisioning Management

The policy decision support challenge for IAM provisioning is how to allow the different stakeholders to convey their knowledge and concerns. To focus this discussion, we assume a situation where there is some centralized automation provisioning for enterprise applications, but that many applications still maintain "ad-hoc" manual provisioning processes (e.g. carried out by local system administrators). The security/compliance manager (domain expert) feels intuitively that more applications should adopt the automated process because she believes it will improve risk and compliance issues. Formally, the *security manager* will be challenged to produce a business case (perhaps a cost-benefit analysis) for the investment, informally there will be a lot of negotiation involving application owners and IT operations (other domain experts). Specifically, the *application owners* will be concerned about disruption to user (aka business) productivity and the IT operations team about the costs and burden that any changes require.

We argue that modeling and simulation can support the overall decision making process. Our aim is to produce a model of the IAM provisioning systems (and related processes) deployed in the organization that will show how to help these stakeholders express and explore their subjective concerns. A useful first step is to identify the different *metrics* that these stakeholders will be interested in:

- **Security/Compliance Officer**
  - o **Access Accuracy**: the number of correctly configured user accounts, against the overall number of created accounts, including badly configured accounts and hanging accounts;
  - o **Approval Accuracy**: the number of approved provisioning activities, against the overall provisioning activities, including the unauthorized ones.
- **Application Owner (Business)**
  - o **Productivity Cost**: these are the costs, in terms of loss of productivity (for employees), due to delays during the approval and configuration/deployment phases of the provisioning process.
- **IT Operations (IT Budget Holder)**
  - o **IAM Provisioning Cost**: this is the cost of deploying (IAM) automated provisioning solutions, for a specified timeframe (involved license fee, fixed and variable costs);
  - o **Provisioning Effort**: this is the actual number of provisioning "transactions" carried out by the organization, in a specific timeframe, giving an idea of the effort and involved workload.

With these metrics in mind we can build an executable process model of the provisioning systems. A high-level schematic of this model is shown in Fig. 3.

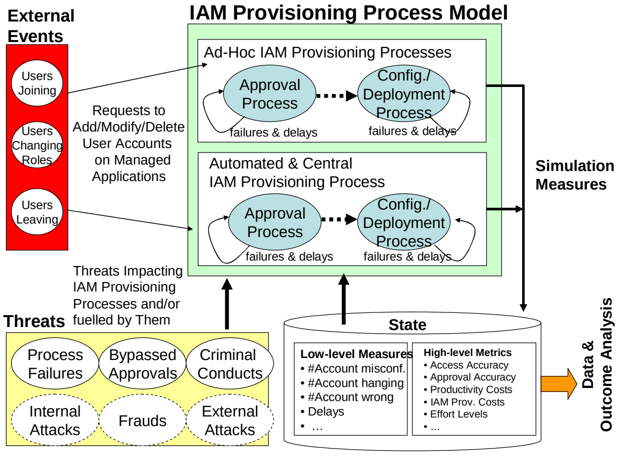

Figure 3. High-level Provisioning Model

More details about the model are provided in Section 5, but roughly we (mathematically) model the actual approval and deployment processes. As they execute they affect the model state, which reflect the metrics we are interested in. These processes are triggered by external events (e.g. employees joining or leaving the organization or changing their role, hence requiring their user accounts to be updated) which we represent stochastically. A simulation, based on the model, proceeds by sampling relevant probability distributions which determine when the external events cause the execution of provisioning processes. By repeating this simulation many times (i.e. in the style of Monte Carlo analysis), we start to build a picture of how different assumptions (e.g. about how processes execute, how often they are triggered or fail) can affect the measures and metrics we are interested in. The threat processes can be folded into this analysis to explore specific failure or attack situations.

Low-level measures (used to calculate the metrics mentioned above) are tracked by the model and calculated during simulations, including: number of correctly configured and mis-configured user accounts; number of hanging accounts (people that left); overall approval time (delays) for provisioning requests; overall configuration/deployment time (delays); number of lost approval and deployments/configuration requests; number of bypassed approval processes. For example, Fig. 4 shows the probability density functions (pdf functions) of some of these measures, as determined by simulations of our model, over the period of time of a year.

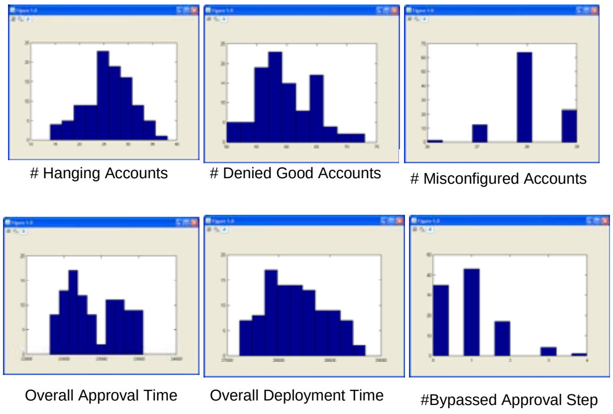

Figure 4. Experimental Results: Pdf of Low-level Measures

The different stakeholders are well placed to compare these fine-grained results with their tacit knowledge, and in some cases with empirical data. A typical next step for an interested stakeholder is to understand and challenge how these results are being derived, e.g. posing the questions “what is it in the assumptions that leads to these results?”, and “do I agree with them?”.

In addition to supporting this exploration it is important that the model provides a meaningful aggregated view so that all the stakeholders can coherently discuss their inputs. The aggregated view should also be meaningful to the key decision maker(s). The graph in Fig. 5 illustrates an example of how this may be done, by means of the high-level *metrics*, derived from low-level *measures*.

The *cost and accuracy metrics* (shown in Fig. 5) may vary depending on the number of provisioning work items, and so the view shows the results for the assumed (modeled) effort level. Section 5 provides additional details for the approach used to produce this normalized view and how to calculate these *metrics*. The key point though is that the assumptions about how this normalization is done are transparent, and potentially subject to discussions.

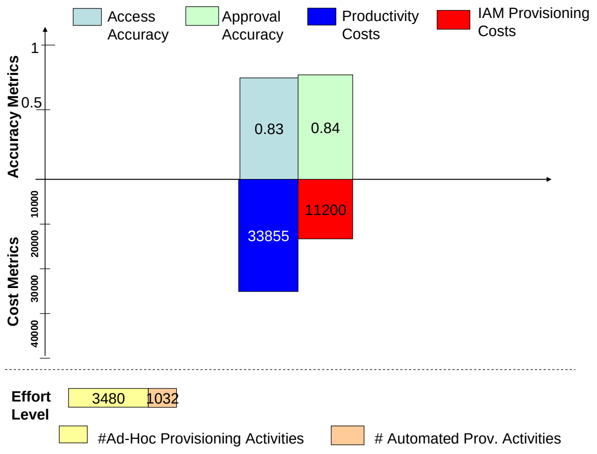

Figure 5. Experimental Results – High-level Metrics

In our case study, we consider the case where the enterprise has 5 core business applications and 100 non-core, lower-priority applications. In the current state, only 2 core applications and 10 non-core applications are currently provisioned with automated and centralized IAM processes. Again, Fig. 4 and Fig. 5 show the *measures* and *metrics* that represent the implications of current enterprise investments in IAM provisioning processes (simulated over a year timeframe). These figures indicate lack of policy compliance (see policy examples in Section 3). For example, policy *P5* is violated as access accuracy is far smaller that 99%.

In an attempt to be compliant, the stakeholders might want to explore the impact of introducing more IAM provisioning automation for protected resources (core and non-core applications/services), by running them under centralized, common processes rather than on an ad-hoc basis. This is one of the *“levers”* a decision maker can act on to change the current situation (another option is to change the policy). Hence, the stakeholders might want to investigate the implications of automating additional applications, in a year timeframe, by considering different automation cases, as shown in Fig. 6.

Simulations of the model can be carried out for each case of interest and the results can be compared. The outcomes, in terms of high-level *metrics*, are shown in Fig. 7.

Fig. 7 shows that accuracy measures are increasing by investing more in automation of IAM provisioning processes. Similarly, productivity costs decrease but IAM provisioning costs increase. This shows that, for certain values of the “lever” (e.g. case 4 - full provisioning automation) the corresponding IAM investment costs are too high, compared to the productivity costs. Further analysis of which applications require more provisioning or different assumptions about future workload might change this analysis. The point is that these metrics can be used to qualitatively and quantitatively show the impact of policy choices. Similarly, results show that Policy *P2* (see Section 3) will never be met (approval accuracy always less than 1); hence policy *P*2 might need to be changed.

Experiments	Core Business Applications (5 Apps)	Non Core Business Applications (100 Apps)
CASE #1 – Provisioning CURRENT SITUATION	automation: 2 Apps ad-hoc: 3 Apps	automation: 10 Apps ad-hoc : 90 Apps
CASE #2	automation: 3 Apps ad-hoc : 2 Apps	automation : 40 Apps ad-hoc : 60 Apps
CASE #3	automation: 4 Apps ad-hoc : 1 Apps	automation : 70 Apps ad-hoc : 30 Apps
CASE #4	automation: 5 Apps ad-hoc : 0 Apps	automation: 100 Apps ad-hoc: 0 Apps

Figure 6. Experiments - “What-if” Cases

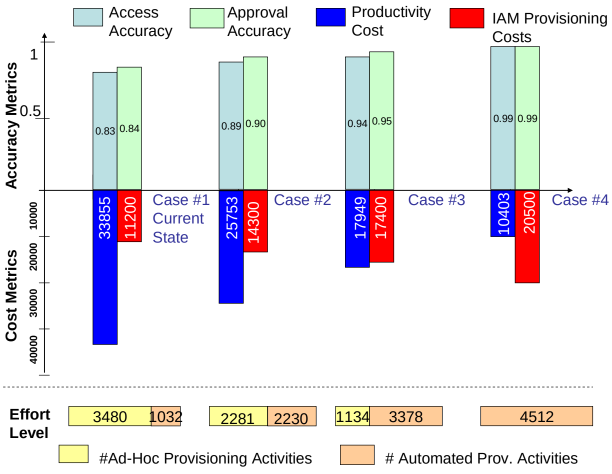

Figure 7. Experiments – Prediction Outcomes for Different “What-if” Cases

V. Our Modeling Approach

Our modeling approach relies on mathematical models and related simulations. The use of mathematical models in engineering has a long and distinguished record of success ranging over mechanical, civil, environmental and electrical/electronic engineering areas. The mathematical methods used in these fields are mainly concerned with continuous phenomena and typically use techniques from calculus such as differential equations. For modeling security and identity management operations the appropriate mathematical methods are more discrete, being drawn from algebra, logic, theoretical computer science and probability theory. In order to apply these methods, we require a conceptual analysis of the relevant aspects of the systems of interest.

In the IAM provisioning case study, we specifically model the difference between ad-hoc and centralized IAM provisioning and explore the impact of choices on existing policies and/or to shape new policies. We seek to illustrate this through the impact on the *measures* and *metrics*, introduced in Section 4.

Our model, discussed in details in [24], explicitly focuses on the representation of IAM provisioning processes, by considering the various steps involved in the approval and deployment/configuration phases. When a user joins, leaves or changes their role (external events), based on their role, a relevant set of applications - that need to be provisioned/de-provisioned - is identified (by means of probability distributions). For each affected application, either centrally managed or ad-hoc managed, the relevant IAM provisioning/de-provisioning steps are modeled along with various measures.

Fig. 8 and Fig. 9 provide additional details about the aspects represented in our model, including the external events (user joining, leaving and changing roles), involved applications and triggered enterprise IAM provisioning activities and processes.

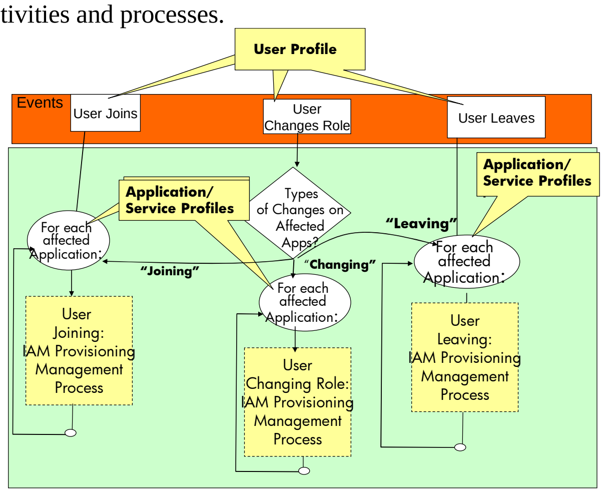
User Profile
Events
User Joins
User Changes Role
User Leaves
Application/ Service Profiles
Application/ Service Profiles
Types of Changes on Affected Apps?
For each affected Application:
“Joining”
“Changing”
“Leaving”
For each affected Application:
For each affected Application:
User Joining: IAM Provisioning Management Process
User Changing Role: IAM Provisioning Management Process
User Leaving: IAM Provisioning Management Process

Figure 8. Discrete-event Probabilistic Model – IAM Provisioning Processes

Specifically each type of provisioning activity, involving a user and one or more applications/services, is explicitly modeled as a “process”, see Fig. 8. Fig. 9 provides the details of the modeled “provisioning workflow” for ‘Users Joining’ the organization: this includes approval and deployment phases, delays and failures (including bypassing the system) along with the points where measurements are taken. Similar workflows are built to model “User Leaving” or “User Changing” roles. Each modeled workflow defines probabilities for how the measures (and related high-level metrics) are affected by each execution of the provisioning process: these will vary depending on whether the application has adopted automated or ad-hoc provisioning.

The model has been built in Demos2k [17,18,19], a specialist language and discrete-event probabilistic simulation tool, which allows such processes to be expressed and executed. Demos2k implements a modelling framework based on the mathematical foundations of a synchronous calculus of resources and processes, together with an associated modal logic [24]. The mathematical framework behind the Demos2K programming language [17] revolves around three key concepts: **resources,** capturing the essentially static components of the system; **processes,** capturing the dynamic components of the system; **environment** within which a system functions.

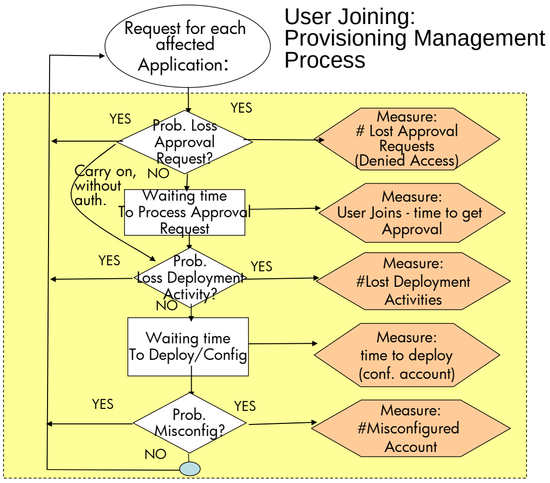

Figure 9. Schematic of the Executable Process Model for IAM Provisioning - New Users Joining an Organisation

The IAM provisioning processes, described in Fig. 8 and Fig. 9, are dynamic entities and have been represented by means of Demos2k processes, A full copy of the implemented model (about 1500 lines of code), based on the conceptual model shown in Fig. 3 and the workflows shown in Fig. 8 and Fig. 9, is available [24].

External events, such as the arrival of a new user, are modeled stochastically, i.e. with appropriate probability distributions. Intuitively, the more IAM provisioning processes are centralized, automated and managed under common policies the more their behaviours are similar, as opposed to ad-hoc processes. However, the more centralization and automation is introduced, the higher the impact of IAM costs (license fees) and faults. We test and explore this trade-off using a Monte Carlo style simulation which can be run with parameterized assumptions about which applications have automated or ad-hoc provisioning. This allows us to build a picture of how different choices will lead to different outcomes.

An instance of a simulation specifies the number of core and non-core applications and the number of applications having automated and ad-hoc provisioning. Fig. 6 shows various assumptions we made in terms of applications and automation levels. Within the model, there is a range of parameters determining the probability distributions for how often the different processes are triggered (typically varying means on negative exponentials), and probability distributions for which applications are affected by the different user centric processes. To illustrate the way the difference between ad-hoc and IAM automated provisioning processes are handled, for each of the category, there are parameters for:

- Waiting times for Approval Request and Deployment/Configuration of User Accounts: modeled as normal distributions;
- Probabilities of Loss of Approval Request and Deployment/Configuration Requests: modeled as Bernoulli tests;
- Probability of Bypassing the Approval Process: calculated dynamically as a Bernoulli test where the probability of the event is: 1–1/(1+ num_approval_failures). The more failure happens in the approval process, the higher is the probability this test succeeds. This might be particularly true in case of centralised IAM provisioning processes.

The details about the definitions of these parameters and probability distributions are provided in a related, extended HP Labs Technical Report [24], along with a full copy of the implemented model.

It is important to notice that some of the probability distributions mentioned above have been tuned, within our model, based on empirical values provided by customers and HP business groups. They can be modified to reflect the reality of specific provisioning processes. The current model has been kept simple: it can be further refined and extended, depending on the level of details needed or available.

As anticipated in Section 4, the model keeps track of a range of cumulative *measures* including: number of approval requests; number of lost approval requests; number of bypassed approval processes; approval time; deployment time; number of misconfigured user accounts; number of denied, legitimate user accounts; number of wrong user accounts (that should not exist - hanging user accounts). The model uses these measures to derive the *high-level metrics*, see Fig. 10.

Experiments have been carried out by running simulations (by executing 100 times the same model), over a predefined period of time (e.g. 1 year) for each experimental case described in Fig, 6. These simulations produce, as an outcome, statistically significant low-level *measures* and derived high-level *metrics*. This information (about 1 MB of data for each experimental case) has been processed, analysed and the outcomes (in terms of measures and metrics values) have been displayed, as shown in Fig. 4, 5 and 7.

This model can be run by different stakeholders (decision makers and domain experts) to directly carry out “what-if” experiments, by acting on available “levers” and changing model parameters. Stakeholders can focus on low-level measures or high-level metrics, depending on the desired level of abstraction they work at, compare results across multiple “what-if experiments” and, if required, delve into the details (e.g. up to the level of the probability density functions of output measures/metrics). This enables stakeholders to improve their understanding of the overall aspects involved in a specific scenario, map predicted outcomes to current policies and compare against their intuitions; it provides them with additional evidence to back their opinions and positions.

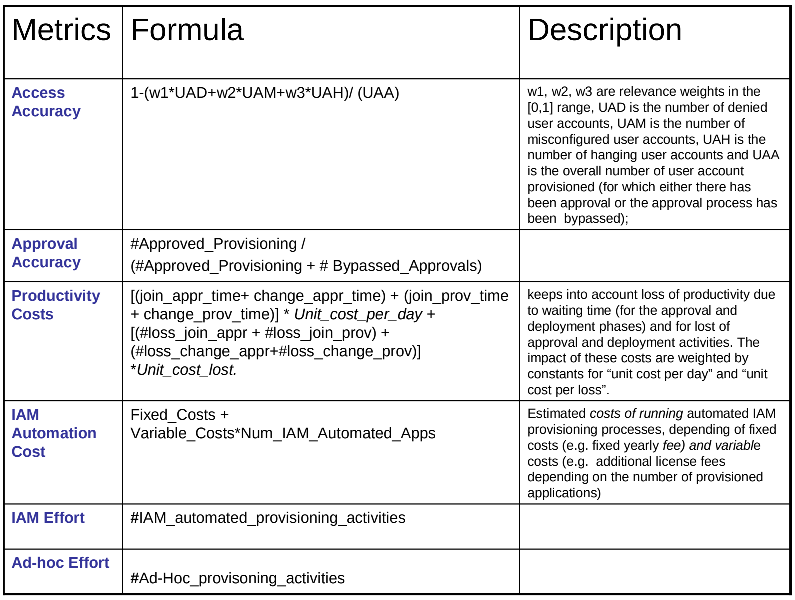

Metrics	Formula	Description
Access Accuracy	1-(w1*UAD+w2*UAM+w3*UAH)/ (UAA)	w1, w2, w3 are relevance weights in the [0,1] range, UAD is the number of denied user accounts, UAM is the number of misconfigured user accounts, UAH is the number of hanging user accounts and UAA is the overall number of user account provisioned (for which either there has been approval or the approval process has been bypassed);
Approval Accuracy	#Approved_Provisioning / (#Approved_Provisioning + # Bypassed_Approvals)	
Productivity Costs	[(join_appr_time+ change_appr_time) + (join_prov_time + change_prov_time)] * *Unit_cost_per_day* + [(#loss_join_appr + #loss_join_prov) + (#loss_change_appr+#loss_change_prov)] **Unit_cost_lost.*	keeps into account loss of productivity due to waiting time (for the approval and deployment phases) and for lost of approval and deployment activities. The impact of these costs are weighted by constants for “unit cost per day” and “unit cost per loss”.
IAM Automation Cost	Fixed_Costs + Variable_Costs*Num_IAM_Automated_Apps	Estimated *costs of running* automated IAM provisioning processes, depending of fixed costs (e.g. fixed yearly *fee) and variabl*e costs (e.g. additional license fees depending on the number of provisioned applications)
IAM Effort	#IAM_automated_provisioning_activities	
Ad-hoc Effort	#Ad-Hoc_provisoning_activities	

Figure 10. Modelling – Definition of Metrics

VI. Related Work

There is a lot of literature on how to use mathematical modeling to affect policy decisions, see papers in the Management Science Journal [27] as well as papers in specific areas such as hydrology, land usage and environmental contexts [1,2,3] or social science [4]. In contrast this work is focused on security and the challenge of helping multiple stakeholders gain consensus and shared understanding.

The area of policy decision support for security, privacy and identity management has not yet been widely explored. A case for using modeling and simulation in information security is made in [5]. Paper [23] explores risk metrics for identity management but it uses a traditional bottom-up risk management approach, based on the assessment of auditing metrics.

Modelling and simulation have been used in specific contexts of identity management and privacy, to explore the impact of technical choices on policies, such as password policies [6,7], identity phishing [8] and security polices for network access control [9]. This is important related work. However, it does not describe how to effectively provide support to different stakeholders in the policy decision making process and focuses just on a few aspects of identity management.

Our work aims at exploring and advancing the state of the art in this space, for a wide range of IAM aspects. This R&D work is part of the HP Labs Security and Identity Analytics project [10,11]. We are not aware of current research or commercial solutions that aim at modelling and simulating the overall complexity of identity management and related policy decision making process. Standards such as ISO 27001 [12], CoBit [13], ITIL [14] describe best practices and methodologies respectively in terms of information security management, IT governance and service management. Decision makers still need to understand, interpret and instantiate them in their specific operational environments. We can use these standards as drivers and references but our work adds the value of grounding the reasoning to specific environments, related policies and the underlying IT infrastructures (possibly along with human and social behaviours).

Our work is complementary to studies on policy refinement and deployment. These studies (e.g. [25]) primarily focus on how to refine policies, once they have been agreed, in order to enforce them. We focus on the policy decision making process and how to support it.

We leverage the work done by HP Labs in the Open Analytics project [15,16], that we consider as a reference. Specifically, we use Demos2k [17,18,19] as the reference tool for our modelling and simulation activities. Finally, an important aspect of our work is the studies in the space of economics and social science. We aim to leverage work done in [20] to build mathematical models that realistically reflect users’ behaviours and the associated impact.

VII. Discussion and Future Work

We have implemented a fully working model [24] of an IAM provisioning management process along with measures, metrics and analysis of outcomes of relevance to different stakeholders. It has been (internally) tested to support the policy decision making process in the IAM provisioning space. This model can be extended in various directions. More detailed descriptions of IAM provisioning processes can be introduced (if information is available) along with a representation of user behaviours (e.g. [11]), to explore, for example, their impact during the approval and deployment phases, on regional and cultural basis. The enforcement side of IAM (e.g. authentication, authorization, etc.) can also be factored in to explore investments trade-offs, based on (policy) choices and various assumptions made by stakeholders. Initial work in this space is described in [21]. Further areas to be investigated include the modeling of the impact of security threats on IAM processes (and in particular for provisioning processes), involved risks and how to support related policy decision making processes.

Our future R&D work includes exploring additional IAM areas (where support could be provided for policy decision making), including: enterprise single-sign-on, authorization and authentication, auditing, IAM outsourcing, IAM-as-a-Service and implications of IAM in cloud computing and Web 2.0 scenarios.

Ultimately, the goal is to create a model library, covering key, relevant IT aspects and policy concerns in the IAM area that can be systematically leveraged by decision makers and domain experts. To achieve this, we are looking for opportunities to engage with HP customers (and other parties) in technology trials, to further validate our approach (to support the policy decision making process) against their current approaches, refine our models and methodology.

VIII. Conclusions

This paper describes current challenges in making effective policy decision within organisations, both in terms of how to form good opinions and then dealing with painful politics and the process of reaching consensus. We illustrated how modeling and simulation methods help to address these aspects, providing objective and relevant analysis for all the involved stakeholders at appropriate levels of abstractions. We focused an IAM provisioning scenario, where relevant (and conflicting) policies might apply. We illustrated how the outcomes of our modeling and simulation activities, based on "what-if" analysis, can explain and predict the impact of specific (policy) choices, from different viewpoints. This is work in progress. We will engage in customer trials to further tune our approach and models. Part of this work will be carried out in the context of the HP Labs' Identity and Security Analytics project [10,11].

References

[1] N. Becu, A. Neef, P. Schreinemachers and C. Sankapitux, "Participatory computer simulation to support collective decision making: Potential and limits of stakeholder involvement", ScienceDirect, Elsevier, 2007

[2] P.W. Adams and A.B. Hairston, "Using Scientific Input in Policy and Decision Making", Oregon State University, 1995

[3] H.H. Khoo, T.A. Spedding, L. Tobin and D. Taplin, "Integrated Simulation and Modelling Approach to Decision Making and Environmental Protection", Kluwer Academic Publisher, 2001

[4] C. Kennedy and G. Theodoropoulos, "Towards Intelligent Data-Driven Simulation for Policy Decision Support in the Social Sciences", School of Computer Science, University of Birmingham, UK, 2005

[5] J.H. Saunders, "The Case for Modeling and Simulation of Information Security", GSEC National Defense University, http://www.johnsaunders.com/papers/securitysimulation.htm, 2001

[6] R. Shay, A. Bhargav-Spantzel and B. Bertino, "Password policy simulation and analysis", DIM 2007, 2007

[7] A. Adams and M.A. Sasse, "Users are not the enemies", Communications of the ACM, 1999

[8] T. Moore and R. Clayton, "The Consequence of Non-Cooperation in the Fight Against Phishing", 3rd APWG eCrime Res. Summit, 2008

[9] J.Y. Koh, M. Yi, T. Cho and H. Kim, "Knowledge-Based Modeling and Simulation of Network Access Control Mechanisms Representing Security Policies", Springer, Information and Communications Security LNCS book, 2002

[10] Security Analytics, HP Labs Project, Systems Security Laboratory (SSL), http://www.hpl.hp.com/research/systems_security.html, 2008

[11] M. Casassa Mont, A. Baldwin and S. Shiu, "On Identity Analytics: Setting the Context", HPL TR, HPL-2008-84, 2008

[12] ISO, "ISO 27001", Information Security Management, , 2005

[13] ISACA, "Cobit, IT Governance", http://www.isaca.org/, 2008

[14] ITIL, "ITIL IT Infrastructure Library for Service Management", http://www.itil-officialsite.com/home/home.asp, 2008

[15] D. Pym, R. Taylor, C. Tofts, M. Yearworth, B. Monahan anf F. Gittler, "Systems and services sciences: a rationale and a research agenda (Open Analytics Project)", HPL-2006-112, 2006

[16] R. Taylor and C. Tofts, "Model Based Services Discovery and Management," PICMET 2008, 2008

[17] Demos2k, "Demos 2k modelling and simulation environment", http://www.demos2k.org/, 2000

[18] G. Birtwistle, "Demos, discrete event modelling on Simula", Macmillian, 1979

[19] B. Monahan, "DXM - The Demos eXperiments Manager", HP Labs Technical Report, 2008

[20] Trust Economics, "Trust Economics Project", UK DTI grant P0007, Trust Economics Project, 2008

[21] A. Baldwin, M. Casassa Mont, B. Monahan, D. Pym and S. Shiu:, "System Modelling to Support Economic Analysis of Security Investments: A case Study in Identity and Access Management, unpublished, 2009

[22] M. Casassa Mont, P. Bramhall and J. Pato, "On Adaptive Identity Management: The Next Generation of Identity Management Technologies", HPL-2003-149, 2003

[23] G. Peterson, "Introduction to Identity Management Risk Metrics", IEEE Security & Privacy, 2006

[24] M. Casassa Mont, A. Baldwin and S. Shiu, "Identity Analytics – User Provisioning Case Study: Using Modelling and Simulation for Policy Decision Support", HPL TR, HPL-2009-57, 2009

[25] M. Sloman, N. Dulay, B. Nuseibeh, "SecPol: Specification and Analysis of Security Policy for Distributed Systems", 1997

[26] D. Pym, B. Monahan, "A Structural and Stochastic Modelling Philosophy for Systems Integrity". HPL TR, HPL-2006-35, Feb 2006

[27] Management Science, Management Science Journal, http://mansci.journal.informs.org/, 2009