An Oracle White Paper
November 2010

Human Capable Information Security

Disclaimer

The following is intended to outline our general product direction. It is intended for information purposes only, and may not be incorporated into any contract. It is not a commitment to deliver any material, code, or functionality, and should not be relied upon in making purchasing decisions. The development, release, and timing of any features or functionality described for Oracle's products remains at the sole discretion of Oracle.

Introduction

If we assume that the digital information explosion will continue at an exponential rate, and that there is a fundamental need to improve information security whilst reducing its cost per transaction in both enterprise and cloud computing models, then we can agree that a declarative, heuristic security approach is the only feasible mechanism for the application of information security in the medium to long term.

Declarative, heuristic security is a concept with which we're intimately familiar as human beings – it is the way in which we place value levels on the security or privacy of information which we hold as individuals and then make real-time, context sensitive, risk-based decisions to decide upon our willingness to share that information.

In the human world, we take advantage of a lifetime of sensory input to help us effectively manage relationships with strangers, friends and family alike; all within the unwritten rules of our community behaviour. In the digital world, organisations are currently far less able to replicate these essential trust models but, through usage of identity-centric technology, we can help mimic dynamic relationship management and solidify identity assurance and usage.

So, what if we could rebuild the human elegance, security and infinite adaptability of decisioning in the digital world? Does this help solve the problems outlined above as well as mitigating complex security problems such as the insider threat?

If so, then we might one day see the realisation of the following statement:

> *"The goal of information security is to replicate human levels of automated authorisation capability in the digital environment; including heuristics, risk awareness and contextual decision making."*

On the basis of the above, this document considers the following three topics:

1) Positioning an information security strategy in a declarative context
2) The implications of pursuing a humanistic approach to information security
3) What technology components are needed to fulfil our goal?

.

Positioning an Information Security Strategy in a Declarative Context

Information Security is largely built around the Authentication, Authorisation and Accounting/Audit (AAA) model i.e. Who are you? What can you do? What have you done?

- The "Who are you?" question is already well-governed with a range of identity assurance tools and **Authentication** techniques such as Oracle Adaptive Access Manager or other two-factor Authentication attempting to ensure that the right digital identity has been adopted by the right human
- Various solutions specialise in providing **Accounting** / **Audit** functions and monitoring the environment for signs of out-of-policy security breaches. However, by definition these are detective, after the fact security controls
    - Security Information and Event Management (SIEM) solutions
    - Data Loss Prevention (DLP) tools
    - Audit Log sniffers
    - Pattern analysis technology
    - Application-specific auditing solutions

This leaves **Authorisation.** In today's IT environments, the decision logic that provides for Authorisation decisioning upon request is typically programmatically encoded in each application. i.e. the application takes Authorisation decisions in isolation, based upon its own business logic and view of internally held users/privileges:

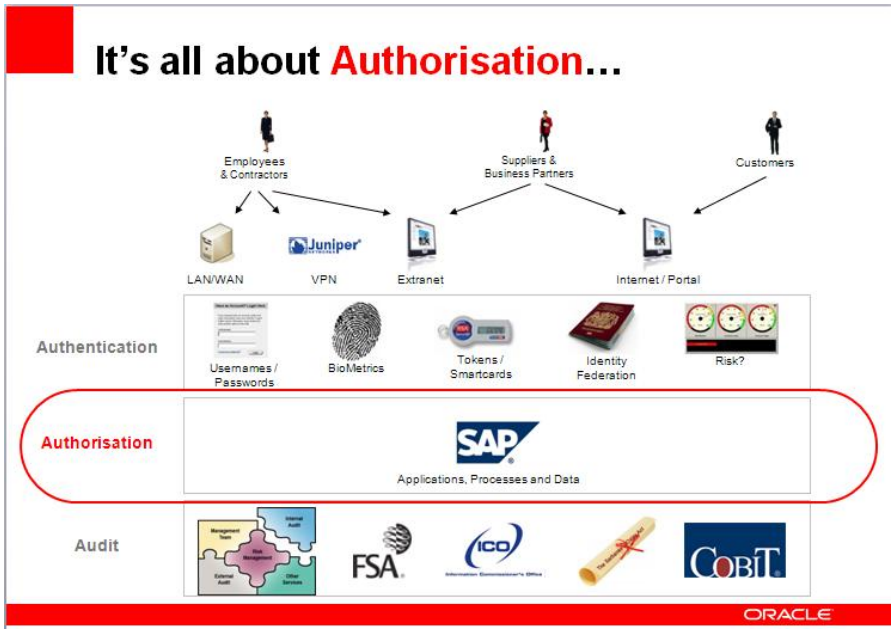

Figure 1 – Programmatically embedded Authorisation as part of the AAA Model

Unfortunately, this has generally led to a silo'd access control environment where individual applications not only enforce their own Authorisation logic but also have application-specific requirements for Authentication and Audit:

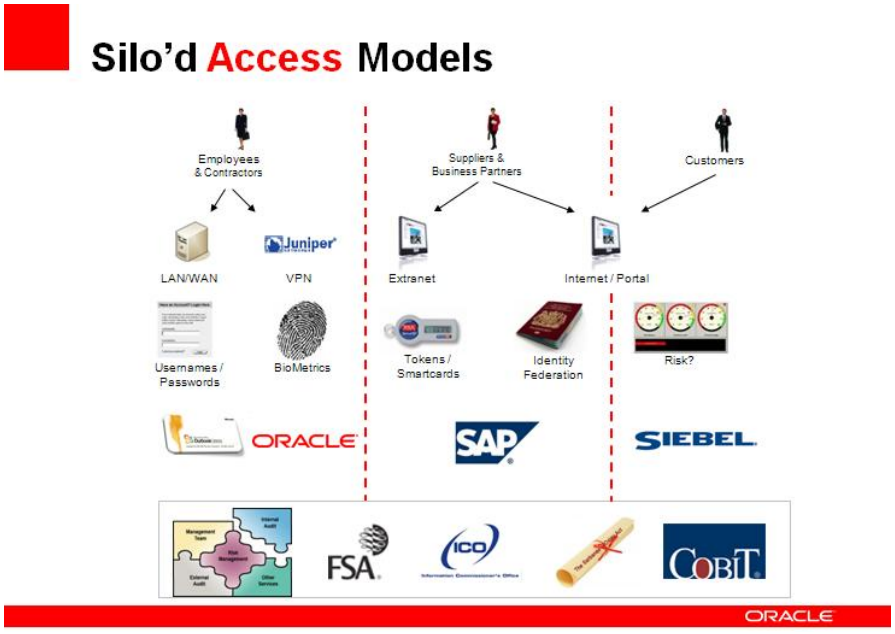

Figure 2 – Silo'd Access Models

When extrapolated across hundreds or thousands of applications, this leads to serious complexity in both identity lifecycle management (i.e. ensuring the right users are in the right user tables with the right privileges at any given time) and the application of consistent access control.

Much of the traditional major-vendor Information Security portfolio is dedicated to helping customers mitigate the symptoms which may be traced back to programmatically embedded logic, such as:

- User Provisioning – ensures that the right users / privileges exist in the right places across the environment at the right time
- Virtual Directory – takes silo'd directories, often serving specific sets of applications, and presents them as a consolidated view
- Enterprise Role Management – enable organisations to understand who *should* have what privileges and provide a mechanism to attest to who *does* have what
- Web Access Management and Federation – provides single sign on and coarse grained Authorisation to discrete web assets which would otherwise maintain their own AAA approaches

In the model depicted in Figure 2, above, updates to corporate Information Security (authorisation of access) policies have to be propagated and managed across multiple targets. In addition, new applications have their programmatic security logic re-coded from scratch each time which is expensive and creates margin for error and risk.

However, many Oracle customers are seeing benefits in no longer coding unique security models into their new applications. Rather, the target application, process or data can *declare* what it requires in terms of security or authorisation. This centralised model of policy enforcement determines the user's entitlement and their request for access is then granted or denied in real time. Companies such as Oracle recognise these issues and have developed declarative security technologies, such as Oracle Entitlements Server, to provide centralised (notwithstanding the ability to disperse PIP/PDP/PEPs via the OES Security Modules) Access policy evaluation and enforcement. This approach is known as *declarative* security:

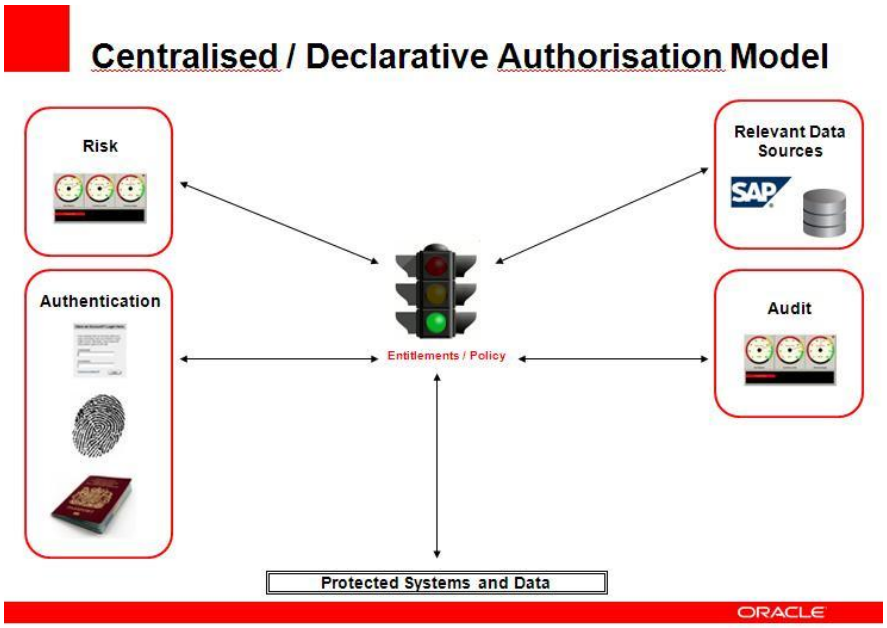

Figure 3 – Centralised and Declarative Authorisation Model

In Figure 3, the programmatic authorisation logic from the protected systems and data has been abstracted and delegated to a centralised Entitlements service which, upon demand, evaluates the appropriate access policy in the context of other information and returns a grant/deny decision (and any associated obligation) to the requesting target. In this way, the organisation achieves centralised management of access policy across their environment with all of the associated risk and cost reduction that this implies.

It should be noted here that traditional IA&M techniques (such as Authentication and Identity lifecycle management) play more of a supporting role rather than taking centre stage as has traditionally been the case. i.e. the Authorisation process generally relies on a range of identity-centric data inputs in order to process its policy decision.

The centralisation of Authorisation also enables more advanced IA&M techniques, such as risk-based Authorisation and centralised Audit, to be delivered more easily. Risk-based authorisation is fundamental to the extensibility of a declarative architecture and Oracle is focusing heavily in this area in order to help prepare organisations for the next iterations of security technology.

It is, however, important that a set of best practice guidelines are developed to assist the understanding as to which Authorisation functions should be abstracted and centralised rather than left programmatically encoded within the target application. It may not be beneficial or achievable for example, certainly in the short term of a declarative strategy, to fully disengage the business Authorisation logic specific to that application. Rather, it is the organisation's global security policies, which can be replicated in Entitlements Policy, updated and called as required, which should be the first set of Authorisations to be declaratively called in this way.

Oracle is currently working with many organisations to help them begin their journey. How do our customers achieve a predominantly declarative environment? Where do they start? What are the pre-requisites and likely costs?

The Implications of a Humanistic Approach to Information Security

So far, so good – most of the centralised, declarative authorisation strategy described above is well-supported by Oracle technology and we can already help organisations to set out on the right path to delivery of this architectural goal.

However, even with a declarative approach to Authorisation, we merely achieve a more elegant (and hence more manageable) mechanism for access policy enforcement. This doesn't address three specific issues surrounding Information Security enforcement that the human is considerably better equipped to deal with:

Digital Information Security Challenge	Human Capability
Well maintained access policies only keep "the bad guys out" and do not mitigate the insider threat whereby a previously trusted individual makes the conscious decision to maliciously compromise the information access that he or she has.	The most basic survival instincts of the human brain ensure that we continually re-evaluate our most trusted relationships with people and surroundings to detect even small anomalies in behavioural change.
Most organisations do not regularly review and re-classify their information assets as their sensitivity and importance changes over time – thus making access policy enforcement (programmatic or declarative) difficult to execute appropriately.	The human decision to authorise access to the information that we hold continually revises and uses our perception of the sensitivity of that information as the primary evaluation criteria.
Finally, access policies themselves require constant updating to reflect the organisation's ever-changing attitude to risk and the context and content that are being governed.	Context-aware and adaptable to an almost infinite number of parameters, human decision making is quick, elegant and surprisingly accurate.

These three additional challenges are reflective of the fact that there is already too much to do to keep one's information security house in order using manual intervention and management processes. But what if we could achieve even partial automation of these? What management overhead do we save and what information security benefits could we achieve?

Background Thinking

Human systems cope well with real-time authorisation decisioning, even given a wide range of variables; known and unknown. The process involved in decision making can simply be described thus:

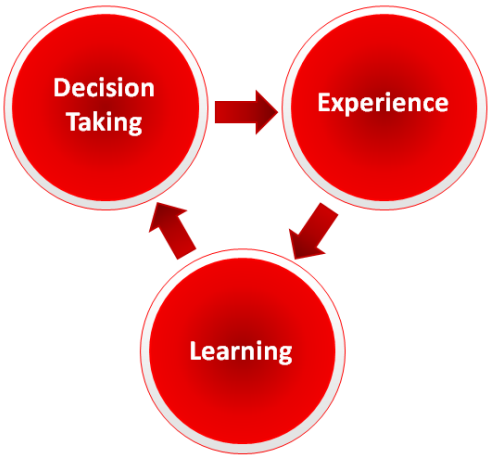

Figure 4 – Simplified decision making feedback loop

In Figure 4, each component is not necessarily as simple as shown:

- *Decision Taking* – human policy execution generally requires additional information as part of its processing and is therefore often a contextual, iterative, indefinite process
- *Experience* – incorporates memories derived from the five senses or, where direct memories do not exist, is capable of substituting analogous experiences
- *Learning* – applied subjectively and in the context of all experience; not just the experiences related to the specific decision process in hand

Different levels of importance also flex this model. For example, it doesn't take many cycles of the feedback loop for a child to learn that a kettle can be hot and therefore touching it should be avoided. However, a keen salesman might call a prospect fifteen times before finally deciding not to pick up the phone again.

When thinking about human-capable security and its analogies to the digital world, it can be useful to reflect on the terminology with which we're intimately familiar as humans and make the mental leap to the digital equivalents:

Human Behaviour or Concept	Digital Equivalent	Example Supporting Oracle Technologies
Introduction	Registration Enrolment Federation	Oracle Access Manager Oracle Identity Manager Oracle Identity Federation
Recognition	Authentication	Oracle Access Manager Oracle Adaptive Access Manager Oracle Enterprise Single Sign On Oracle Identity Federation
Referral or *Collaboration*	Federation	Oracle Identity Federation Oracle Access Manager
Recollection or *Memory*	User Stores and Directories Audit	Oracle Directory Services Oracle Identity Manager Oracle Identity Analytics Oracle Audit Vault
Risk	Contextual, real-time management of digital security	Oracle Adaptive Access Manager Oracle Mantas
Relationship	Identity Lifecycle Management	Oracle Identity Manager Oracle Identity Analytics
Suspicion or *Obedience*	Digital behavioural anomaly detection	Oracle Adaptive Access Manager Oracle Mantas
Authority	Authorisation Entitlements Roles Segregation of Duty	Oracle Access Manager Oracle Entitlements Server Oracle Identity Analytics Oracle GRC Manager Oracle Database Vault Oracle Label Security
Learning	Heuristics	Oracle Adaptive Access Manager Oracle Mantas Oracle Sigma Dynamics
Custom or *Tradition*	Authorisation Entitlements Roles	Oracle Access Manager Oracle Entitlements Server Oracle Identity Analytics
Confidentiality	Information Rights Management Encryption	Oracle Information Rights Management Oracle Advanced Security

Trust	Identity and Access Management	Oracle Identity and Access Management portfolio
Citizenship	Compliance Standards and Interoperability	Oracle Identity Analytics Oracle's support for industry compliance standards

By extrapolation from the humanistic model of decision making, we can draw an analogous digital feedback loop as relevant to Information Security, described thus:

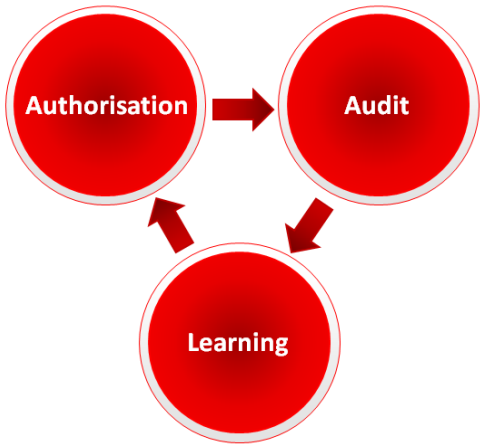

Figure 5 – Information Security Decision Making

As noted previously, we can broadly substitute existing technologies such as Oracle Entitlements Server for the Authorisation step and a range of data storage and reporting tools for Audit.

However, the Learning component is a less well developed area in the digital information security technology market and is instead traditionally filled by a plethora of ill-managed, expensive and error-prone human processes.

The model described above is a simplified schematic of our Human Capable Information Security goal and the components of this decision making process are discussed further below.

What Technology Components Are Required To Fulfil Our Goal?

AAA (Authentication, Authorisation and Audit)

As noted earlier in this document, traditional AAA technologies are readily available; although they are typically still applied in a programmatic architecture.

In the Authorisation space, there are a number of leading technologies capable of fulfilling the requisite roles in our preferred, declarative and centralised Authorisation architecture. Oracle Entitlements Server, for example, allows organisations to create centralised security policy and have it referenced on demand by applications and infrastructure which is delegating some or all of its Authorisation function.

Additionally, Oracle Platform Security Services allows enterprise developers to adhere to standards of declarative platform security and build applications which can more easily leverage centralised Authorisation standards moving forwards.

Given that "what you know about who you know" is key to supporting the Authorisation decision, Oracle has also invested heavily in delivering a comprehensive and integrated identity lifecycle and identity assurance services platform. Technologies such as Oracle Identity Manager and Oracle Identity Analytics securely and efficiently control the lifecycle of our digital users whilst Oracle Adaptive Access Manager and Oracle Advanced Security help enforce the levels of identity assurance that we wish to be attached to Authorisation of our key information assets.

In the Audit space, we seek to leverage the transactional information that can be captured from a user's interaction with our digital environment. SIEM technologies exemplify the desirable, cross-platform data collection engines that harvest security events and correlate these to produce valuable insights into user and system behaviours. Oracle has a plethora of platform or discipline-focused Audit repositories; from dedicated technologies such as Oracle Audit Vault and Oracle Identity Analytics, to the logs available from other Enterprise and/or security-centric tools such as Oracle Applications, Oracle Information Rights Management and Oracle Access Manager.

Data Classification Management

In order to help meet the challenge of applying appropriate Authorisation to our data assets over their lifecycle, a more complete understanding of those assets is required.

Many organisations are currently undertaking a data security classification project with the expected output of a better understanding of what information their organisation holds, where it holds it and what safeguards should be enforced to protect it.

However, the technology space in this area is yet to mature and organisations continue to find it difficult to discover, categorise and maintain the vastness of their information stores. Yet, this data regularisation process is an essential foundation on which to base the automation of processes in a human capable authorisation model.

Automated Learning

To properly replicate human behaviour, as described in the model above, it is necessary to look in the direction of Artificial Intelligence; specifically the AI area known as Artificial Neural Networks (ANN).

An ANN is a program that can recognise patterns in a given data set and subsequently develop a model for that data. It resembles the human brain in two aspects:

1) Knowledge is acquired by the ANN through a learning process (trial and error)

2) Inter-neuron connection strengths known as synaptic weights are used to store this knowledge

An ANN is a set of processing elements and connections with adjustable strengths:

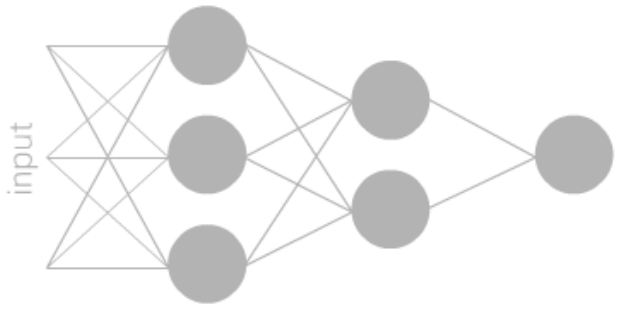

Figure 6 – A Neural Network Schematic

To train the network, we must present it with data and wait for the network to produce an output which is then compared with the desired output. The network weights are then adjusted to compensate for the error and the process is repeated until the ANN consistently reproduces the required behaviour.

When applied to Information Security, Artificial Neural Networks are a significant enabler toward a more automated approach to Information Security delivery and can help us to meet the three key remaining challenges of information security described above:

Identifying the insider threat.

- By automating the monitoring of identity-related information, we should be able to not only better identify when a breach has occurred but also raise the suspicion levels ahead of that breach and attempt its prevention. It must be noted here that a true ANN approach to learning normal and abnormal behaviour is significantly different from most of the current, "heuristic" technologies available in this space today.
- In order to train the ANN towards this goal, we would require access to significant levels of clean, historic, user-centric data such as that available in Oracle Identity Analytics, Oracle Audit Vault and generic SIEM solutions.
- It would also be highly beneficial to enrich this organisational-centric data with other user-centric information such as social networking input, Linked-In details, HR information, previous psychometric results etc. However, care must be taken to ensure that the organisation's appetite for the privacy of its employees is consistent with its delivery of this type of solution.

Figure 7 – ANN Step 1 - Application of better detective controls by applying ANN capability to identity-centric audit

Automating the modification of classifications in meta-data.

- If and when an organisation has undertaken a data classification exercise, has enabled the data sources with meta-data related to those classifications and is perhaps applying associated Authorisation policy as the routine for data access, then automatically updating those classifications during the data lifecycle is highly desirable.
- If the ANN is able to replicate the human behaviour it has learnt, then classifications can be auto-flexed to better represent the true picture of the organisation's risk position on the data it holds at any given time.
- Clearly, iterative steps towards a fully automated data classification environment are advisable; particularly when reducing classification sensitivity over time.

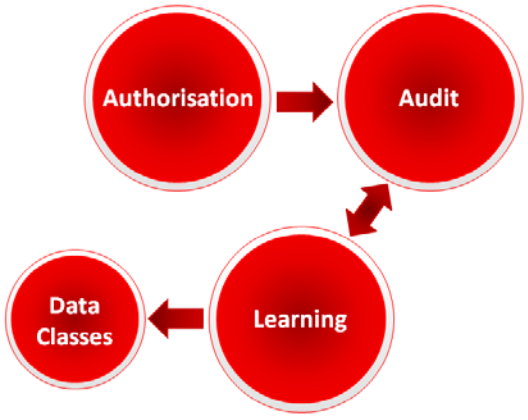

Figure 8 – ANN Step 2 – Auto-updating of information asset classification during the data lifecycle

Automatically updating the access policy for the business.

- If the Authorisation policies are the InfoSec equivalent of the human decision making node in Figure 5 above, then the ultimate aim of ANN learning is to reach a point where we can automate the updating of these.
- As an example, if the ANN learns that, every time a project team member requests a certain document related to that project, access is granted, then it is reasonable to suggest that the associated Authorisation policy may be gradually amended to reflect a similar level of repeat access.
    - Note that this may infer that the access policy was incorrectly set in the first instance. As such, we may see the initial application of ANN technology in this area as an automated method of alerting to potentially incorrect access policies; rather than remediating the policy automatically.

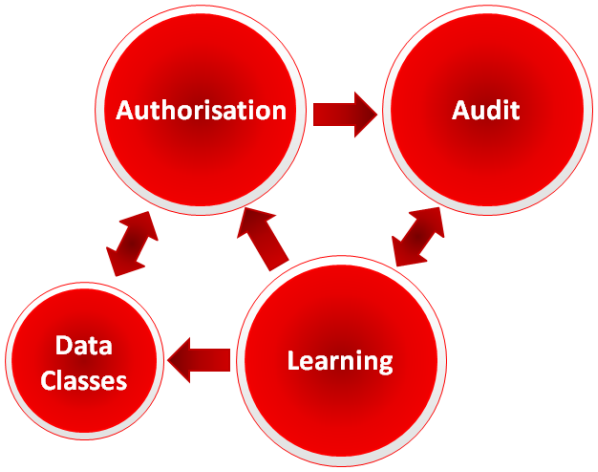

Figure 9 – ANN Step 3 – Auto-updating of Authorisation policy based on learning of human equivalent actions over time

As noted above, organisations would be likely to require a gradual period to develop a comfort level with releasing some of these tasks to an automated process.

Further information on the usage of Artificial Intelligence can be found at the Association for the Advancement of Artificial Intelligence (AAAI) website.

There are already moves afoot within the technology industry to use AI technology in identity-centric use-cases. ANN engines are already in use, for example, as a means of providing predictive, enriched Human Resources analysis and the underpinning technology has been ratified, adopted and proven by a number of leading global companies. Further research into this type of technology and its abilities to be integrated seamlessly with our Authorisation environment is critical to the development of successful, long-term information security technology solutions.

Conclusion

Solving information security on a macro scale is the biggest challenge facing the digital world today. Without evolving the existing models in search of a solution to this problem, our creativity in digital expansion will continue to be hampered by real-world events and concerns.

There is much to be learned about future-state digital information security architectures by studying how humans cope with much the same set of problems. By shifting our architectural focus towards declarative, risk-aware and contextual systems, we can better provide the foundation on which automated learning, the final piece of the puzzle, can be built and leveraged.

White Paper: Human Capable Information Security
November 2010
Author: Duncan Shores

Oracle Corporation
World Headquarters
500 Oracle Parkway
Redwood Shores, CA 94065
U.S.A.

Worldwide Inquiries:
Phone: +1.650.506.7000
Fax: +1.650.506.7200

oracle.com

Hardware and Software, Engineered to Work Together