End to End Identity Management Architecture for the Enterprise

Chanaka Kannangara

Sun Microsystems

<chanaka.kannangara@sun.com>

ABSTRACT

Identity Management is fast becoming an important aspect in the organisation that the technologists and executives cannot afford to pay minimal attention anymore. Stringent, sometimes draconian regulatory requirements, ever increasing software licensing and operating costs and risks posed by identity theft are often seen as primary business requirements that drive the renewed interest of implementing comprehensive identity management framework within the organisations. Whilst the business drivers are clear and a compelling business case exists, it is common that CIOs , IT Managers and Architects find it difficult to determine and prioritise the implementation approach for an Identity Management framework. This is fuelled by the availability of numerous products and technologies which all play a part in Identity Management architecture and many interpretations from numerous vendors pushing mainly one aspect of Identity Management.

This paper looks at the end to end Identity Management architecture based on the fundamentals with a vendor independent view point. It focuses on the basics of attribute management, authentication and authorisation and how these fundamental Identity Management components are mapped to new and existing corporate information technology landscape. It also examines the often misunderstood and intangible benefits realisation aspects of Identity Management for the enterprise, its diverse user communities and partner organisations. The federation of identities are looked at as this brings about a new way of doing business with unprecedented efficiencies.

1. Identity Revolution

Network Identity has become an inseparable part of everyday life for most of us. In a world where more and more traditional transactions are now being done in the Internet or with some networked device, most of us have to become well aware of what our network identities are and be extremely careful about the proliferation of our identity information. As the owners of the attributes of our identity, it is up to us to protect if not jealously guard these attributes so that our identity information is not abused or misused by other parties.

If above is a concise summary of the challenges faced by owner of the identity, one can appreciate the amount of work required by the enterprise to gain and retain the goodwill and trust of the users so that they continue to engage with the enterprise using network identities. With more and more organisations seeing the tangible benefits of opening the information assets to the external stakeholders to streamline operations and increase efficiency identity management has become an essential plank underpinning today's enabled organisation.

In addition to this organisations are accountable for the way they store, process and handle the identity information by law. Government regulation and privacy acts have been in place for a while now, but with the emerging threats of terrorism and associated laws, stringent corporate governance and disclosure requirements, compliance is becoming a formidable task. Technologists and business executives are squarely seeing the challenge ahead. Are they prepared?

2. Basics

Is identity management about user single sign on, password self management or rapid user provisioning? Please don’t confuse the benefits of the identity management process with the essential building blocks of the identity management. As technologists, if we are to find a solution, it is paramount that we focus our efforts around getting the fundamentals right. The most important aspects; rather should I say the building blocks of Identity Management are;

- Attributes – information about a user or a device. Note that owner of the identity determines who get access to what attributes

- Authentication – the process of identifying a user or device based on the credentials/ attributes provided.
- Authorisation – the process of providing selective access to applications/entities based on the attributes of the authenticated user/ device.

Identity Management is complex as it spans from the clients tier, presentation tier, business logic and data layers of the IT landscape. Essentially, Identity Management is an enabler that provides a unique but important service to all these technology layers. If the right framework is in place the reusability is enormous. This is one service most of the existing and future applications use alike without the need for repudiation and reconciliation. Also the numerous non functional architectural aspects such as performance, high availability and security are also significant.

Identity management spans a heterogeneous application environments. Therefore it is important to pay attention to open interfaces, integratability and interoperability.

3. Getting the Momentum

Who should be the sponsor of the Identity management efforts within the enterprise? Is it Information Technology, Human Resources, Security or Operations? It is a vexed question as each of these business units have a stake in identity management. IT for its part in delivering the technology, HR for its role in inducing people to the organisation, Security in terms of verifying access control, and Operations from providing information to the relevant parties; they are all stakeholders of an identity management endeavour. The technologist find that most of his/her efforts to deal with technical issues are somewhat hampered by the challenges posed by the ownership and politics to get all business units on board to support the Identity Management efforts in the organisation. How does one go about securing support of all business units? It is imperative to get executive sponsorship for any Identity Management effort so that an enterprise wide cohesive approach can be adopted.

4. Technical Landscape

We discussed earlier the importance of attributes, authorisation and authentication. It is logical to start designing the identity framework by looking at the attribute management, thus looking at the various identity repositories within the organisation. Many applications have identity information embedded in the application itself. If self-service, self registration and servicing the external client base are important then it is preferable to build an enterprise directory by synchronising the common attributes of the identities to a central enterprise directory. It is a common misnomer that the enterprise directory has to be the source of truth for identities. It does not necessarily have to be. The authoritative source (the source of truth) could be the another system; in most organisations this is the Human Resource s system.

What are the technologies available to synchronise identity information? Traditionally this is the role fulfilled by the meta directories and synchronisation servers. However this functionality is now gradually being superseded by virtual directories. Unlike meta directories and synchronisation servers, virtual directories do not transport attributes between two physical repositories, but instead rely on virtual identity maps to track the relevant identity information. The information fetched in real-time from source repositories when requested.

The provisioning systems are architecturally similar to the synchronisation products but functionally are vastly different. Provisioning systems create, delete, update and manage user account life cycles, and also implement workflow functionality for approval process. They also play an important part in auditing and logging user account movements, which has become increasingly important in view of the executive accountability legislative requirements such as the Sarbane-Oxley Act. This is increasingly in the list of priorities of the CIO and as a result there is a renewed interest in the user provisioning and access control.

The connectivity to the target identity repositories plays an major role in selecting the appropriate technology for both provisioning and synchronisation. Many products in the market today require the installation of client side connectors to access the identity sources. This is potentially a security and management hurdle as some of these identity sources are outsourced or require special permissions from the security policy to install a connector in the same platform. The deployment effort in such cases could be prolonged increasing the cost of acquisition. On the other hand the agent-less systems do not require any connectors installed in the source systems, but rely on commonly available access methods to access the source repositories. Due to their non intrusiveness these agent-less systems are gaining popularity over the agent connector based systems. Table 1 shows a simple comparison of some non-functional aspects of an agent based and agent less systems available in the market today.

System characteristic	Agent Based	Agent Less
Manageability	Needs client side and server side monitoring.	Only server side monitoring required
Deployment Effort	Client connector require installations in client source repositories	No client side installations required, making deployment faster
Data Access	Rich proprietary APIs can be supported to extract non-standard information	Data access is mostly through standard interfaces.
High Availability	Client side failovers cannot be guaranteed	Server side failover determines the system availability metrics
Performance	Two step data transformations at server and the connector	Only one data transformation in server end producing higher performance levels.

Table 1: Comparison of some system characteristics of Agent based and Agent-less systems.

If an enterprise directory is considered for the identity management framework it is best to use a standards based directory such as Lightweight Directory Access Protocol (LDAP) based directory. Whilst similar to X.500 hierarchical structures, LDAP provides a less rigourous yet functionally rich directory schema which can scale from small enterprises' to huge ISP type organisations' needs. True LDAP directories are read-optimised for performance and the underlying implementation should support that. This is why a true LDAP directory is preferable over an RDBMS with an LDAP front end; they provide faster access rates, and with the right choice of products replication becomes much simpler.

Transactional authentication and authorisation play key parts in the overall identity management architecture of the enterprise. In this context there are two challenges: providing internal access to applications; and providing external access. There are numerous applications that require authenticated and authorised access within the enterprise and as shown in Figure 1, theses applications often rely on dedicated user attribute repositories to make access control decisions. One can argue that in an ideal situation it is preferable to make all these applications bind to a central directory for the attributes and policy details to determine the access decision. In the diverse heterogeneous technical landscape of the enterprise this is going to be extremely difficult if not impossible. Even if the most ambitious project works through all the schema anomalies and technology mismatches to make it happen, it will be very expensive to implement. Yet it would even more expensive to sustain such a solution when the applications and user stores are updated progressively.

Therefore the prudent approach for authentication and authorisation of the internal users is to allow the individual applications to do this using their own identity stores. Here the common attributes of these user stores can be abstracted to an enterprise directory so that services such as simplified sign-on, self administration, corporate whitepages, user provisioning and de-provisioning can be performed based on the information available in the central repository. If external user access to the backend applications and information assets are required, the enterprise directory and access control system can play an important role in facilitating this.

The main challenge associated with exposing information assets beyond the corporate security boundaries is that the user identity needs to be validated even before the applications are exposed to them. Traditionally external access has been provided using technologies such as Virtual Private Networking allowing direct access to the backend corporate network and applications that reside on it. This is not a very scalable or economical approach if the organisation intends to open up the information assets to stakeholders other than their employee base. Alternative technologies such as "VPN on demand" where access is provided via a secure enterprise portal are proven to be more economical and scalable. Role-based access plays a key part in providing access to a wider array of an organisation's stakeholders. With role-based access, an organisation can personalise the access based on the authenticated credentials and entitlements of the individual user, and expose the appropriate applications that are relevant to them.

The installation of a new identity management system will invariably have some disruption to users in their normal business activities. Therefore it is important to manage change, inform users and also make sure there enough incentives to adopt new changes. Single sign-on, self administration, personalised portals with virtual desktops, secure remote access for broad range of employees, are some of

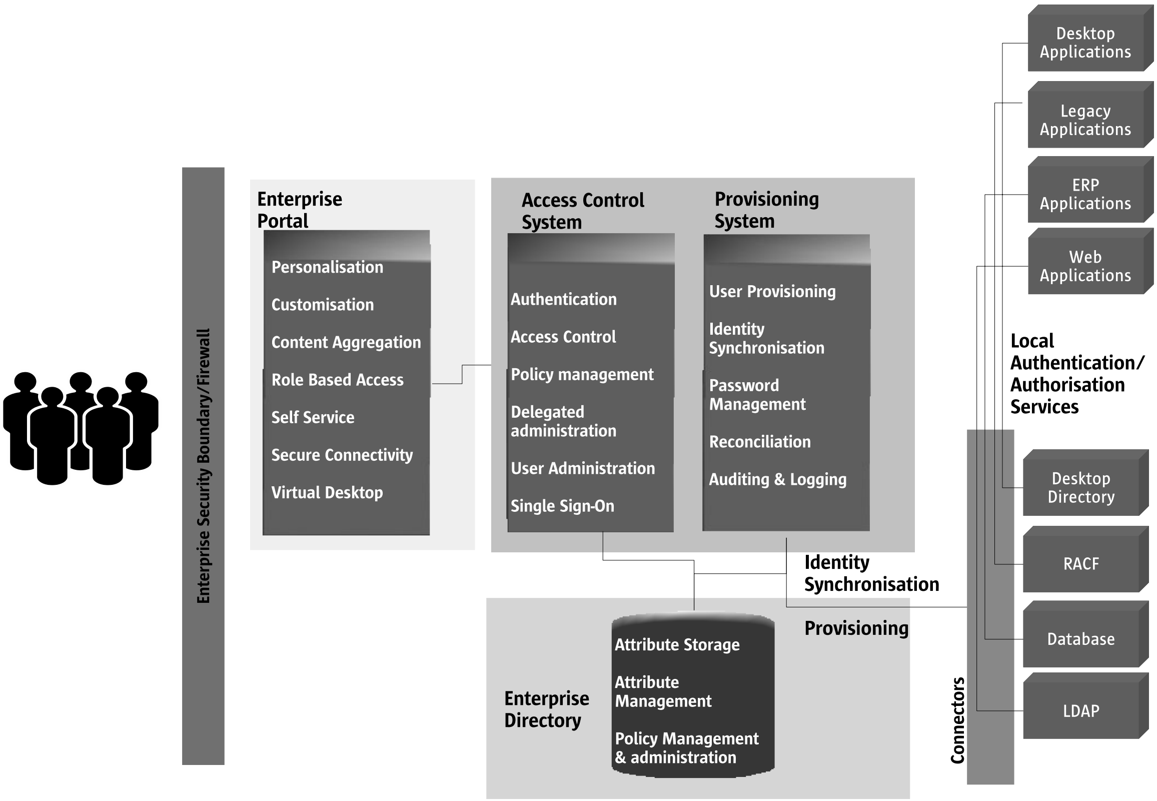

Figure 1: Architecural constructs that make up an identity management framework.

the notable capabilities that can be built upon an end to end identity management framework to realise substantial benefits for all user communities.

5. Interoperability

Identity federation and interoperability have been touted for many years but with not much success in implementations. When it comes to exchanging identity information technology alone cannot solve the numerous business challenges where trust is a key ingredient. Industry initiatives such as Liberty Alliance Project (http://www.projectliberty.org) has made significant strides in defining both technical and business frameworks to facilitate the exchange of identity information between organisations belonging to a "circle of trust". Some of the key benefits of identity federation for the users are cross-domain single sign-on, federated account linking and single sign-off, are covered by current Project Liberty specifications. The liberty technical specifications use Security Assertions Markup Language (SAML) to exchange identity assertions securely. Liberty specification-enabled identity management products are now available in the market providing unprecedented opportunities for organisations to tap into a wider customer base, generate efficiencies in the value chain and streamline business process across many partners and stakeholders.

Web services are becoming increasingly popular with organisations as a tool to drive efficiencies across an organisation's value chain. Web Services Interoperability Forum (WS-I) also has a security and trust model that articulates a framework for exchanging and federating identity information. Identity Management plays an integral part in any commercial web services framework. Both leading web services frameworks available today, J2EE and .Net, have considered identity management and made provision in the respective platforms. Java Authentication and Authorisation Services (JAAS) and .Net Forms provide sufficient hooks to incorporate identity management processes to web services deployed on these platforms.

6. Final Thoughts

Although it is not a new concept, the rapid changes in the current business environment makes Identity Management an important challenge facing the business today. The key to a successful implementation is to evaluate the current capabilities and define the overall architecture and then implement selective parts that generate significant user benefits with minimal disruption. It is important to get early runs on board.