Digital Identity in Perspective

John Tolbert

For over a decade, many IT security professionals have been focused on the subject of digital identity. Most software applications rely on digital identity as a means for granting access. Software applications, like the people who create them, are often flawed, insecure, and vulnerable. Knowledgeable network and system architects liturgically repeat that protocols such as HTTP are stateless, while knowledgeable security architects lament this fact. Statelessness presents somewhat difficult problems to overcome when designing secure applications. Browsers, web applications, and the back-end systems to which they connect exemplify these conditions. Data and network security often seems to be an afterthought. Nevertheless, the systems do work, despite their inherent problems. The Web and the Internet are now pillars upon which global commerce and much else rests. How did we get here, specifically considering digital identity?

In the beginning (of IT), there was the computer. Programmers saw that not all users should have been created equal. Therefore, in order to assign different types and levels of access, administrators created accounts for users. Realizing that it is not good for accounts to be alone, the technicians coded groups into existence. Administrators employed these constructs to grant different permissions to different users, to each according to his/her business requirement; thus access control lists (ACLs) were born. [Depending on one's background, permissions seem to be synonymous with privileges and entitlements.]

We have made great progress in solving previously discovered problems in digital identity, for example: user directories, web access management, and identity federation. These developments have improved the user experience, while decreasing system administration costs. Most organizations now have a serviceable identity infrastructure.

Within the realm of digital identity, an archipelago of islands from different sectors has emerged from the sea of standards. Education, enterprise, consumer, government, defense, health care, finance, and other sectors utilize many of the same products and protocols to accomplish their disparate missions. Industry-specific solutions evolve and morph to meet the needs of their verticals. However, these solutions are generally all based on common standards and protocols used by everyone. For instance, LDAP is ubiquitous, and SAML and profiles of SAML are used by many different types of businesses and governments. To solve the remaining and yet-to-be identified problems in digital identity, we must improve existing standards. We should seek harmonization and convergence between identity schemes and protocols. In the areas where disagreements exist regarding protocol utilization, organizations should attempt to define use cases and frameworks for mapping the appropriate use of technology associated with digital

identity. For instance, in the business-to-consumer (B2C) realm, SAML is often said to be “too heavy”. However, SAML is used in business-to-business (B2B) scenarios millions of times per day for high value transactions. Conversely, while enterprises avoid the use of OpenID for high value transactions, they do employ it in other solutions.

During the decade of “focus on identity”, great strides have been made in strong authentication and credential management. Two key business drivers in this field include the ability to create and broker assertions of strong authentication and the reduction of the number of credentials to manage. The long-heralded benefits of identity federation are appearing, though by quiet materialization between identity providers and service providers, rather than by a noisy revolution. The search for a universal identity provider (or providers) has thus far been fruitless. Many countries are offering PKI-based digital identifier cards to their citizens for authentication to web-based government services. These government-issued IDs are also being used as authenticators for transactions with private sector service providers. In the United States however, no such single identity provider has appeared, nor is one likely to appear in the near future. Thought leaders have predicted that various types of non-governmental enterprises might step into this role, but so far this has not happened.

Digital identity is not an end-in-itself. Recall that accounts exist to distinguish users from one another, and serve as containers of attributes that can be evaluated in access control or other policy enforcement decisions. The next iteration of identity evolution produced role-based access control (RBAC), whereby accounts were assigned to roles (also expressible as attributes), ostensibly to simplify the assignment of permissions. However, organizations who deployed RBAC often realized that the number of roles needed to satisfy requirements exceeded the number of users within said organization.

The current access control paradigm is attribute-based access control (ABAC), which is transitioning to policy-based access control (PBAC). ABAC could be viewed as a subset of PBAC. Perhaps for the sake of simplicity we may combine them at this stage: Policy- and attribute-based access control (PABAC). The OASIS eXtensible Access Control Markup Language (XACML) standard functions as both a standard policy language and protocol for transmitting both policies and access control decision requests/decisions between systems and domains.

Identity is a single component of an access control decision. Other factors must be examined in authorization decisions, such as “what is the requested resource?”, “what action is requested?”, and “what environmental conditions are stipulated by policy?”. We must consider the relative position of identity in applications such as access control, privacy, and personalization. While identity is certainly an important foundational element, it is not the sole driver of access control decisions.

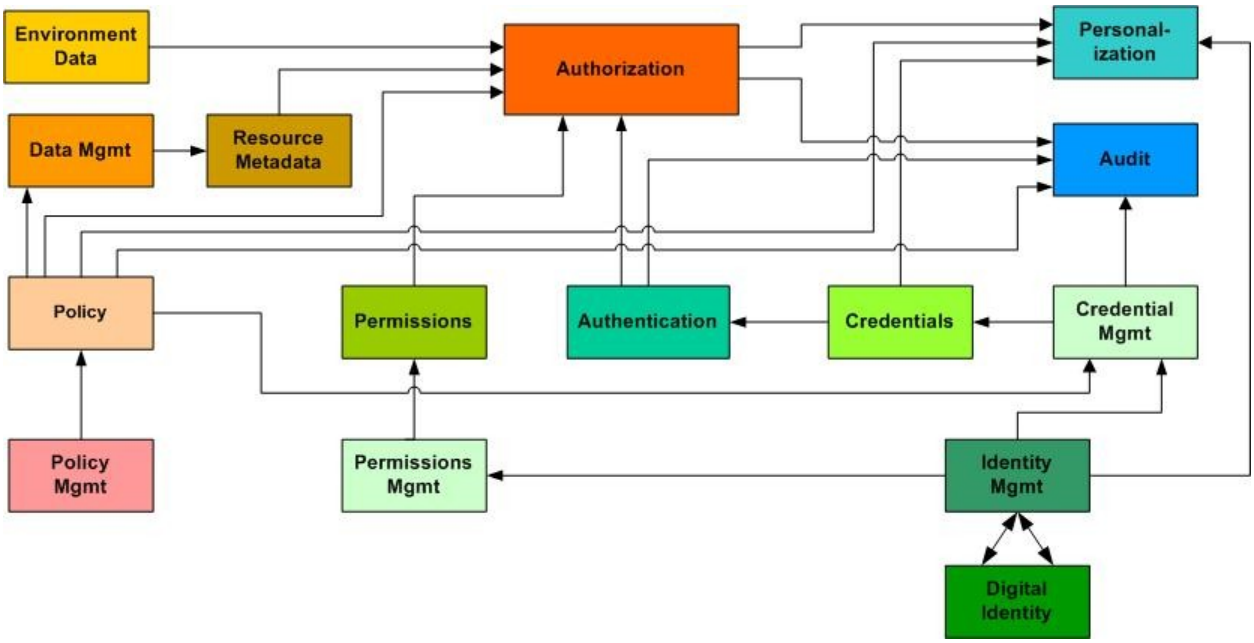

Identity can be viewed as an amorphous amalgamation of attributes about a given subject. Identities exist somewhere, but are nebulous until invoked in an authentication event. Given that "identity solutions" are often packaged with authentication and access management systems, the terms identity, authentication, and authenticators are often conflated in the minds of security product managers and architects. *The purpose of an authenticator is to prove that you know a secret, possess a token, or have certain biometrics, and then binds you to some permissions*. An authenticator doesn't define a user.

Considering the collective drive to introduce fine-grained authorization controls, the concept of digital identity is expanded to include device and application identifiers. In this environment, identity is broader than a simple mapping to human identities. Stricter policies, reflecting the current state of technology mandate that access control and privacy decisions must also compute these data points. Certain types of resource access should and can be limited to defined devices in specific locations. For examples, access to export-controlled information can be restricted by physical location; access to licensed intellectual property can be restricted to certain authenticated computers; and access to sensitive personal information can be prohibited from mobile devices. Additionally, device identity is useful for customization and personalization, as exemplified by usage of HTTP user-agents to notify web servers when content is being requested from mobile devices, allowing mobile-friendly content to be sent to the requestor rather than full content.

Recently publicized attacks such as APTs (Advanced Persistent Threats) show that digital identities become practically irrelevant as security controls if unauthorized users gain complete control over endpoint computer systems (including browsers) and network infrastructure. The goals of APT are to employ advanced techniques to

compromise and control remote machines, then exfiltrate data to attackers' systems. Once compromised, the malware on a machine can utilize authentication and identity services to gain access to everything the authorized user would normally be permitted. One of the most popular exfiltration methods is usage of standard browser protocols such as HTTPS.

Exploitation of computing resources by outside agents remains a serious concern. The mechanisms utilized have improved, and it is significantly easier for adversaries to achieve results remotely without risking physical assets. Likewise, the insider threat is also a top concern. The expanded capacity of storage devices and their concomitant reduction in size have allowed for removal of large amounts of data. In the case of insider threat, permissions associated with digital identity are abused because the actor has a valid digital identity (or identities) within systems, and has been granted access to resources. The relatively low cost of entry for the perpetrator and potential loss of intellectual property via APT and insider threat represents an existential threat to businesses and organizations.

It is logical to surmise that more effort should be directed at solving access control and privacy concerns rather than maintaining a continued focus on identity. We must find innovative ways to protect data within all levels of electronic sytems. Enhanced tools must be developed which will:

1. Inventory data
2. Categorize data
3. Tag/mark data with meaningful metadata
4. Cryptographically bind metadata to data objects
5. Allow administrators to grant permissions to individual data elements, based on user/device/application identities as well as resource metadata, actions, and environmental factors.

Within the web browser, opportunities for improving the identity portion of authorization transactions exist. Users still have many accounts and passwords. Account and password management is very problematic: users must remember a myriad of usernames and passwords for dozens of systems. However, people have multiple roles in life and may well want them to be separated in terms of digital identity also, so simply consolidating digital identities is not likely to be the answer. Passwords are inherently a very weak authentication mechanism. This situation is exacerbated by the fact that the accounts and passwords are transmitted regularly over the web. Further, accounts and passwords can be easily phished or key logged. Thus, the account/password is an extremely easy credential to acquire by a nefarious party, and the theft does not become apparent until harm is already done and discovered. The use of stronger, federated authentication mechanisms has

been slow to be adopted across enterprises, but the trend must be accelerated. In order to improve digital identity in web browsers, we must expand the utilization of strong authentication and identity proofing mechanisms which can be leveraged by browsers and web applications.

Perhaps the best solution would be to add better cryptographic functionality to web browsers, such that browsers could generate, utilize, and prove possession of a cryptographic key without requiring transmission of that key across the internet. This would be useful for human user identity with browser interactions. Given the nearly ubiquitous distribution of Trusted Platform Modules (TPMs) in computing hardware today, having a programmatic method for accessing the TPM would create a much more secure means of authenticating and asserting device identity, as well as other parameters (such as software patch levels, device firewall status, and anti-virus software signature file levels). Though TPMs won't prevent all attacks on users and their devices (DLL injections of keyloggers after bootup, for example), the goal in this space should ultimately be to provide degrees of *platform assurance*, based on TPM, machine identity, and machine health.

Collaboration with international standards organizations that have produced valuable and widely-used specifications in the realm of digital identity would benefit us greatly. To address known deficiencies, we should work with the appropriate extant bodies:

- We must load-balance our efforts within IT security to adequately address the wider notions of policy enforcement, recognizing that identity in the browser is a part, rather than the whole. The work of the Trusted Computing Group should be evaluated and leveraged as much as possible in the realm of platform assurance.

- Extensions and profiles of current specifications should be pursued before launching entirely new efforts: for example, consider OASIS SAML profiles.

- Interoperability of identity concepts, products, and protocols, across multiple platforms and even across multiple sectors should always be a primary motivation, and synergizing with the Kantara Initiative would be a productive means of accomplishing these tasks.

Paraphrasing a popular political slogan of the 18th century, "standards that specify least specify best" might be a good maxim in our endeavor to improve digital identity in the browser. The standards that govern identity management today, such as LDAP and SAML, though well-defined and feature-rich, tend to be implemented in a minimalist manner. Standards which specify too much, or are overly cumbersome to understand, deploy, and use fall by the wayside. Let us heed the lessons of x.500, WS-Federation, and SPML when venturing forth on Identity in the Browser.