Platform Capability Based Identity Management for Scalable and Secure Cloud Service Access

Abhilasha Bhargav-Spantzel

Intel Corporation

Email: abhilasha.bhargav-spantzel@intel.com

Steve W. Deutsch

Intel Corporation

Email: steve.w.deutsch@intel.com

Abstract—In the past identity management solutions evolved to solve the challenges with username/password based systems to provide a seamless single sign-on (SSO) experience for the user. With the advent of large scale cloud services, the existing SSO solutions for authentication using only username/password need to be revisited. We propose the use of platform capabilities and integrated credentials as a criteria for doing the authentication and authorization of the respective cloud service requesters. Cloud service requesters can be any type of device including PCs, TVs, laptops, phones, tablets and so on.

Based on the device type the capabilities can offer information that may be necessary and sometimes sufficient to provide access to a given service. More specifically, a user may not have to enroll to get certain types of cloud services because the platform capabilities and intrinsic certificates may be sufficient without user specific information or input. For example, if a device can provide secure geo specific information then services which are provided for devices in a certain geo can be qualified based on the provided geo information without any additional input. For services that are controlled for enrolled users, instead of establishing a username/password PKI certificates can be embedded on the device which is secured using the platform capabilities. This will allow secure yet seamless access to such cloud services. Such a model where user ID is not mandatory but definitely available per service requirements, allows for enhanced privacy without jeopardizing security. Additionally the flexibility of such a model may allow the scaled identity management policies as required for various types of cloud services.

Index Terms—Security, Privacy, Identity Management, Trusted Computing

I. INTRODUCTION

There are a number of cloud service providers (CSPs) that offer services and content that needs to adhere to a license agreement that corresponds to the technical aspects of protecting the data and ensuring privacy of its users. There may be additional requirements related confidentiality and security of the service itself [14], [15]. A license can be an explicit agreement provided by the content provider (e.g. Disney) or an implicit agreement corresponding the wishes or intentions of the user related to the use of the material. The license conditions can include restrictions on copying the material, viewing the material for a specified period of time, or restricting the material to a geographical location. The type of materials can include premium entertainment content, financial documents, personal health care information, and corporate confidential technical documents. The cloud service providers can range from entertainment studios, banks, health

care organizations and governments to major corporations world-wide.

With the advent of large scale malware attacks and advanced persistent threats [16] there is an increasing possibility that the user is using a system infected with such malware. The upcoming common consideration across all the services and providers is a *dependency on the behavior of the device hosting the user*. It is important for the service providers to be able to predict and enforce what will happen to the content they are providing once it is available to the user's device. An important element for the successful deployment of these services is to be able to ascertain the device's capabilities with some level of assurance before providing the service or access to the content.

Depending on the value of the content there may be a need for high assurance and non-repudiation of the device capabilities. The opportunity is to provide open platforms based on standards that have the capabilities to support the wide range of usages. The alternatives today are:

- 1) Deploy a specialized appliance [11]
- 2) Restrict usage to closed platforms
- 3) Build and maintain a customized platform
- 4) Risk business and governance on unknown platforms.

In this paper we investigate the requirements from both the user and device capability to access high assurance cloud services. As mentioned above to satisfy these license agreements we need to investigate user characteristics, profile, preferences, credentials etc. as well as the device characteristics, capabilities and credentials. We develop the notion of a license and provide an understanding of what are the fair usages of services and the data. We translate the license to a policy specification at the CSP and show how the combination of user and device capabilities can support that policy. As such we show how the device capabilities are a significant component of meeting the policy based on the service type. It helps achieve better non-repudiation because we have both the hardware and user assurance to provide a granular and holistic view of the client. Also we elaborate on how having assurance of the hardware integrity allows for a secure bootstrapping that can be used to ascertain additional information and confidence about the user himself. Interweaving the concept of device ID with the more traditional user ID based identity management systems requires understanding key service requirements, identity management lifecycle considerations, access control

policies and the resulting assurance based on the security and privacy properties.

Contributions. In this paper we elaborate on the key considerations to establish a platform capability based identity management system (IdM). Such an IdM would allow a flexible and fine granular access control based on both platform capabilities (denoted as device ID for brevity) and user credentials or user ID. We provide an analysis of service types and show where such a paradigm fits with the assurance requirements of the various services. We also provide a comparison of the identity lifecycle of device based IdM versus the user based IdM to understand the implications on the security and manageability of the respective identity information. We provide a policy representation of such an IdM system followed by a discussion of the privacy and security implications of the proposed solution. This helps in achieving the overall goal to ensure secure and privacy preserving access to high assurance cloud services.

The cloud environment makes certain important identity management requirements more challenging. In particular, identity federation that requires dynamic trust to be established with attribute based privilege management will be more consistent with assurance policies in the cloud. Another example is achieving mutual non-repudiation which requires both the platform assurance level as well as the user identity. Finally achieving confidentiality and privacy in cloud environment is critical to ensure there is minimal exposure of information and protecting the communication channels.

The rest of the paper is organized as follows. In Section II we provide a description of the types of cloud services and based on the value of the service a discussion on the identity assurance needs. Following that in Section III we provide a comparison of the lifecycle of the platform identity versus user identity. In Section IV we provide a logical representation of the policies that leverage device ID and user ID and key considerations of such policy specifications. We provide a workflow and example usage in Section V with an illustrative example tying together the various concepts presented in this paper. Finally in Section VI we provide a privacy and security discussion followed by the conclusion.

II. DEVICE AND USER IDENTITY ASSURANCE TYPES AND CLOUD SERVICES

By Identity assurance we mean the confidence we have on the non-repudiation of the identity claim. Based on the type of service, the user or device ID may or may not be required. Also, based on the value of the service being offered the assurance of identity can range from high to low. A break up of the possible combinations of the type of identity relevance and example services is provided in Figure 1. As such there are 4 possible models while considering Device and User Identity assurance - Type 1 where neither device nor user ID is not relevant; Type 2 where only the user ID provides user non-repudiation assurance; Type 3 where only the device ID provides device non-repudiation assurance; and Type 4 which

Non-repudiation Assurance				
Type	1	2	3	4
Device ID	X	X	✓	✓
User ID	X	✓	X	✓
Example Services	Free Services	E-commerce (Banking)	BlueRay (DRM)	ePharma (Healthcare)

Fig. 1. Device and User Identity Assurance Types

includes both user ID and device ID. In the following we detail the considerations for each type.

Overall we show how for high assurance services Type 4 is essential. More specifically, *independent of the strength of assurance of the user ID, there are requirements for protecting the service itself, ensuring user privacy, transaction and data security for high assurance services which in turn raise the bar for device capabilities.*

A. Type 1: No Assurance

Services which do not require ID assurance of the user or the device fall into this category. Examples include free services. Increasingly, services which are free in nature may also require some additional assurance because of potential of malware attacks from compromised platforms or need for accountability from the registered users.

B. Type 2: User Assurance Only

Most e-commerce based services that are available today are based on having a certain level of assurance of the enrolled user identity. This is typically in the form of usernames and password. Due to the proliferation of such username and passwords most of the identity management systems today focus on defining an easy to use Single-Sign-On solution to allow for an ease of use methodology to access multiple services. Such a model, however, does not ensure that the user ID is issued to the correct individual.

Increasingly, there is a need for having higher assurance for Type 2 services. Having a view of the platform capabilities and identity information helps understand the underlying platform where the user ID resides. Having better predictability of whether the user's system corresponds to a safe or a hostile environment enhances the security and assurance provided by the user ID itself. It also helps shift the responsibility of providing higher assurance from the user (e.g. by having increasingly difficult passwords to remember) to the hardware (e.g. by embedding hardware tokens [10], [13]). Figure 2 illustrates how on the same client platform we typically see different types of applications with varied degree of security assurance requirement. For Type 2 typically the user information is sufficient, however, understanding that there are potentially malicious applications residing on the same system a cloud service provider may require additional information regarding the platform to assess the security of the client's system.

Fig. 2. Different applications with Varied Assurance Requirements on the same Platform

C. Type 3: Device Assurance Only

In the world of Digital Rights Management (DRM) [12] one key methodology to ensure that secure content is delivered to the right system is by having certificates embedded on the platform by the manufacturer. This does not require user side provisioning or setup. For example the Blue-Ray is a DRM platform which ensures that Blue-Ray content from cloud services gets delivered to the right platform.

Such a model may become increasingly useful for other cloud services and usages such as geo location services supported by platform based information regarding the platform location, emergency alerts, m-commerce and so on. Creating the platform capability based identity management infrastructure would become important to avoid the problem that user ID proliferation had with the silo'd approach. Furthermore, the access control mechanisms would need to extend the policies to check for such capabilities to provide access to the platform based services. A discussion on these types of policies is provided in Section IV.

D. Type 4: Both User and Device Assurance

For high assurance services there is not only the need to ensure that the right user is making the service request but also that the user is on the right device. There are several examples such as the "Bring Your Own Device" [3] paradigm where we are seeing increasing threats and challenges because the device platform is not known and therefore the environment where the user is logging for can potentially be hostile. Moreover, in such environments it is hard for an infrastructure security service to protect or remediate the user data and the service itself.

In the process of accounting for user and device ID, the access to a service is guaranteed to provide non-repudiation for the user. This is an important property from the security and privacy perspective as discussed in Section VI. For example, the platform may be capable of providing an isolated

Fig. 3. Identity Management Lifecycle Comparison

environment which protects user ID and information related to a given secure service. Additionally, the assurance on the user identity claim itself is improved because the platform ID management lifecycle is more robust and regulated. Finally, the access control policies would need to be extended to capture the verification of the relevant platform capabilities under the right level of granularity.

An example of a high assurance service is ePharma (Please see Figure 2). Here we not only want to ensure the user ID is correct and captured by a secure hardware, but also that the information collected and exchanged during this transaction is not compromised by other (potentially hostile) applications running on that system. The privacy of the user would be compromised even if the data in transit was protected by secure network protocols. Furthermore the ePharma service would want to ensure that it can effectively mitigate and remediate a known (predictable behavior) platform in case of an attack. With the increasing number of malware attacks [16], this model will become more important to implement for all secure services.

III. IDENTITY MANAGEMENT LIFECYCLE COMPARISON

In this section we provide a comparison of the lifecycle of the platform identity versus user identity. In particular we show that platform capability based information denoted by platform identity has a low user touch and a well controlled methodology for issuance, usage, update and revocation. This contrasts with user identity based identity management which due to lack of control and scale is vulnerable to security attacks. Overall the key takeaway from this analysis is that platform identity has a higher degree of assurance in the entire identity management lifecycle and can add value from the security and usability perspective if used either by itself or in conjunction with the user identity. If used together with user identity, due to platform identities higher assurance, it can also strengthen the usage of user identity.

A. Identity Issuance

One key consideration during the issuance of identity is the verification of identity information itself. As such it is important to ensure that any identifier or certificate is issued to the right entity. There is high degree of concern for most user identities that are issued on the web [9] as they are often based on unverified arbitrary information (web forms). This leads to attacks such as identity theft [7]. The user may need to be physically present to ensure the identity is verified and issued

to the correct person. In contrast, the device ID and platform certificates are normally set by the manufacturer and cannot be easily forged or modified because of hardware protection. This leads to a higher degree of assurance, ease of provisioning and lesser risk of the compromise of the device ID or capabilities.

B. Identity Usage

Misuse prevention involves ensuring the identity information that is presented or claimed is done so by the actual owner of that information. There is a high degree of concern for the misuse of user ID information - Identity theft [7] is prevalent and hence there is a need for multi-factor authentication [5]. Moreover, in addition to the multiple factors there is an increasing need for *stronger* factors which are related to platform ID. Platform ID and identity protocols are normally driven and controlled by trusted computing groups and standards [10] and have a comparatively significantly less potential for misuse.

Another aspect related to identity usage is the minimal disclosure of information and ensuring that it is revealed on a need-to-know basis. Protection of User ID is important for privacy and compliance with regulations [8]. Most protocols involving sharing of user ID information do not ensure minimal disclosure. In contrast, platform capabilities are currently revealed in a controlled manner, although privacy implications need to be carefully considered [12].

Finally, while providing ease of use and ensuring usability by providing a seamless access is of high importance. For User ID there are SSO solutions to help manage several username/passwords and other user identity information (e.g. demographic, unique IDs etc.). However there is high degree of concern to ensure security while allowing better usability. In comparison, platform capability based attestation and service are inherently seamless and transparent (under the hood) - abstracting the complexity - and may not require user involvement.

C. Identity Modification

Ensuring flexible yet secure update of identity information is an important part of the identity lifecycle. There is a high degree of concern for a secure and flexible update of user ID. Such an update should be allowed only by authenticated users or correct authorities. In areas such as financial data or user's health data, lack of security for identity record modification may lead to dire consequences [7]. For platform ID however, there is less concern as there is a well defined processes in place to ensure platform capabilities are advertised and updated correctly.

D. Identity Revocation

The final stage of the identity lifecycle is to allow for flexible and efficient revocation of the ID information. There is high concern for user ID information and the inability to revoke it due to the untracked and mismanaged user ID information which may be widely distributed across several service providers. Revoked identity information can lead to identity theft [7] and service misuse. In comparison, platform

capability based attestation would fail if the certificates are revoked by a central entity.

IV. POLICIES OF PLATFORM CAPABILITY BASED ACCESS CONTROL

The access control policies would need to extend their definition to allow evaluation of the platform capabilities. As such capabilities influence the assurance level of the user ID and device ID claims, the policies need to be granular to express the access criteria and the assurance level of a given service. One example extension to a policy language for Platform Capability and service assurance is given below.

Access control policies define rules for accessing services related to authorization of the individuals and their service request. In order to reason about these policies - without having to deal with various specific syntaxes, we use a logical representation of privacy policies which is abstract and independent from any existing lower level language. Our aim is to explicitly represent device ID based concepts -such as device capabilities and assurance level- along with more traditional access control constraints, at a logical level. Please note that these concepts can be included in standard policy specification languages such as SAML [2], OATH [1] and others specified in various Cloud Standard specifications [6], [15], [14].

The components of the policy would be as follows -

- **User ID** $v = \langle v_1, v_2, v_3, \dots v_n \rangle$
Examples of v_i include username/password, biometric, demographic information etc.
- **Platform ID** $\pi = \langle \pi_1, \pi_2, \pi_3, \dots \pi_m \rangle$
Examples of π_i include platform capabilities such as TPM, trusted execution environment etc.
- **Platform enhanced User ID** $v_\pi = \langle v_{\pi 1}, v_{\pi 2}, v_{\pi 3}, \dots v_{\pi n} \rangle$
Examples of $v_{\pi i}$ include user certificate protected in HW secure storage [10], user one time password supported by HW based random number generator etc.

A combination of the above components would constitute of the single and composite policies as depicted below:

- **Single Policy** $\Psi_0 = \bigwedge (v_i \vee \pi_i \vee v_{\pi i} \vee \phi)$
- **Composite Policy** $\Psi_* = \bigvee (\Psi_{0i})$

There are 3 key properties that need to be satisfied by the above policies to achieve the necessary scale and expressiveness. They are *granularity* and *flexibility* and they are described as follows -

Granularity: Ability to express different low level user and device attributes and capabilities. Policy should have the ability to define multiple combinations of user ID and device ID attributes to express single and composite policies. It is important however that the policy should be simple to be evaluated efficiently (e.g. simple conjunction and disjunction of the IDs).

Flexibility: Allowing one or more composite policies to correspond to an assurance level for a given service. Multiple policy options should be provided to satisfy the security requirements

of the cloud service. This allows policy to be robust and support a wide range of platforms and Cloud service types.

Ordering: Composite policy options may be ordered so that the most robust requirements precede other alternatives. This would allow the user to satisfy the service requirements to the best of his ability. For example, the more secure your trusted computing base, the lesser is the dependence on the rest of the platform. Windows 8 for example ensures that the basic infrastructure e.g. the BIOS is signed so that the OS and applications residing on it can also be trusted.

Based on the above policy requirements the user and the client device need to respond with the appropriate information. We refer to this response as the platform quote which is defined as follows:

Platform Quote Q : Signed response from the client user and device to respond to the authentication challenge from the service provider. Q contains a combination of user attributes (or User ID) and device attributes or capabilities (or Device ID) required to satisfy a give policy Ψ .

V. WORKFLOW AND EXAMPLE USAGE

The Device Capability based IdM - ID usage workflow is provided using Algorithm 1.

An illustrative example usage of this workflow is given below. Consider a Cloud Service Provider responsible for e-prescribing medicines and creating and managing users medical records. Due to the nature of this service and compliance to regulations, such as HIPAA [8], this *E - Pharma* service requires high assurance identity verification of its customers (Refer to Section II Type 4). An example policy of *E - Pharma*:

$$\Psi_{E-pharma} = (v_1 \wedge \pi_1 \wedge \pi_2 \wedge \pi_3) \vee (\pi_2 \wedge \pi_3 \wedge v_{\pi 1})$$

where v_1 = username/password ; π_1 = isolated trusted execution environment ; π_2 = patched FW/OS; π_3 = HW based secure delete; $v_{\pi 1}$ = user cert in trusted storage Based on the above policy the *E - Pharma* service requires username/password provided in a isolated trusted environment or alternatively a user certificate stored securely in hardware. *E - Pharma* wants to ensure that the system is patched and has the ability to delete all residual data which may be generated when the user is accessing his medical records and prescriptions. When the user requests this service and receives the policy (step 3 of Algorithm 1) then as part of the the challenge response at step 4 the user can create this quote: $Q = [user\ cert\ signed\ by\ protected\ HW\ storage,\ signed\ proof\ of\ patched\ FW/OS,\ signed\ proof\ of\ enabled\ HW\ feature\ for\ secure\ delete]$

Once the *E - Pharma* verifies the elements in Q it can provide access to the medical records and prescriptions to the user.

VI. PRIVACY AND SECURITY DISCUSSION

When considering security it must include both privacy and confidentiality of the service, all the resources and data, and the user. This includes both adherence to the license for all

Algorithm 1 Device and User ID Usage and Verification

Require: Enrolled User ID list v and Device ID list π

Ensure: The the enrollment has been done securely and correctly (right ID is mapped to right entity)

```

{User Request}
1: User  $\rightarrow$  CSP : User requests for a high assurance service
2: User  $\leftrightarrow$  CSP : Secure connection is established to do the
   challenge-response
{Challenge-Response}
3: CSP  $\rightarrow$  User :  $[\Psi_{ServicePolicy}, timestamp, nounce]$ 
4: User : Create Platform Quote  $Q =$ 
    $[[v_i], \{\pi_j\}, timestamp, nounce]_{sign}$ 
5: User  $\rightarrow$  CSP : User sends  $Q$ 
{Quote Verification by CSP}
6: for each line  $i$  in  $\{v_i\}$  do
7:   CSP  $\leftrightarrow Verifier_{v_i}$  : Verify  $v_i$ 
8:   if verification == false then
9:      $Verifier_{v_i} \leftrightarrow$  CSP : Sends Verification Error and
       user ID reset suggestion
10:  end if
11: end for
12: for each line  $j$  in  $\{\pi_j\}$  do
13:   CSP  $\leftrightarrow Verifier_{\pi_j}$  : Verify  $\pi_j$ 
14:   if verification == false then
15:      $Verifier_{\pi_j} \leftrightarrow$  CSP : Sends Verification Error and
       Device ID remediation suggestion
16:  end if
17: end for {Service Provisioning}
18: if verification == true then
19:   User  $\leftrightarrow$  CSP : Secure connection used to provide
     service
20: else
21:   Reconciliation : User ID reset or Device ID remediation
22: end if

```

participating parties and any additional third parties that are either malicious or intent on violating the license agreement.

There are three key requirements that need to be satisfied to ensure the user's privacy -

- 1) Protect user's identity and credentials,
- 2) Secure the service transaction, and
- 3) Ensure that the residual data and information left by the transaction is protected.

The platform not only needs to have the capabilities to satisfy the above requirements but also must be able to attest to them in a confidential manner if required. The attestation is required so that the service can evaluate the client device to ensure it has the required level of assurance needed to qualify for the service.

Before the information is released it is important that there is a secure connection with the cloud service provider. As such, the communication with the service provider should not expose the user's identity if it is not needed for the transaction.

This is in line with the minimal data release policy [4]-ability to carry out the transaction based on the device capa-

bilities with minimal exposure of the user's ID or credentials. To ensure the security of the transaction the data should not be deleted, accessed or modified maliciously by any agent acting outside of either the service transaction or the license agreement (Please see Figure 2).

Additional security measures that need to be supported by the device capabilities include:

- 1) Protecting against DoS attacks from either the client or the service side
- 2) Protect against both user and service impersonation
- 3) Protect against theft of service and data or misuse of resources
- 4) Protect against malware and phishing attacks
- 5) Protect against replay, man in the middle and rough service providers

Therefore for high assurance services, resources, and data there can be a requirement for non-repudiation of both the device and user to achieve the above privacy and security requirements.

VII. CONCLUSION

In this paper we presented the concept of platform capability based IdM. Given the need for high assurance services in an increasingly hostile environment where these services and user data are attacked [16], [7] such an IdM would allow the device to provide a trusted base which can be used as a foundation to build identity verification based on user and device ID information. This model would be used to ensure the privacy and security of the user as he uses various cloud services. As part of future work it would be critical to address some of the challenges that come with such an approach including building the ecosystem support, policy and usability studies for such a paradigm shift.

VIII. CONCLUSION

The conclusion goes here.

ACKNOWLEDGMENT

The authors would like to thank...

REFERENCES

- [1] OpenID. <http://openid.net>.
- [2] Security assertion markup language specification set. <http://www.oasis-open.org/committees/security>.
- [3] Rafael Ballagas, Michael Rohs, Jennifer G. Sheridan, and Jan Borchers. Byod: Bring your own device. In *Proceedings of the Workshop on Ubiquitous Display Environments, Ubicomp*, 2004.
- [4] Abhilasha Bhargav-Spantzel, Jan Camenisch, Thomas Gross, and Dieter Sommer. User centricity: a taxonomy and open issues. In *DIM '06: Proceedings of the second ACM workshop on Digital identity management*, pages 1–10, New York, NY, USA, 2006. ACM Press.
- [5] Abhilasha Bhargav-Spantzel, Anna C. Squicciarini, and Elisa Bertino. Establishing and protecting digital identity in federation systems. *Journal of Computer Security*, 14(3):269–300, 2006.
- [6] William E. Burr, Donna F. Dodson, and W. Timothy Polk. Electronic authentication guideline: Recommendations of the national institute of standards and technology. In *NIST SP800-63*, page 64. NIST, 2006.
- [7] Mari Frank. Identity theft prevention and survival. <http://www.identitytheft.org/>.
- [8] <http://www.cms.hhs.gov/hipaa/>. The health insurance portability and accountability act of 1996.
- [9] <http://www.sourceid.org/resources/basics.html>. Sourceid: Open source federated identity management.
- [10] Steven L. Kinney. *Trusted Platform Module Basics: Using TPM in Embedded Systems (Embedded Technology)*. Newnes, 2006.
- [11] Paul Koster, Frank Kamperman, Peter Lenoir, and Koen Vrieling. Identity-Based DRM: Personal Entertainment Domain. In *T. Data Hiding and Multimedia Security*, volume 4300 of *Lecture Notes in Computer Science*, pages 104–122. Springer, 2006.
- [12] Radia Perlman, Charlie Kaufman, and Ray Perlner. Privacy-preserving DRM. In *IDTRUST '10: Proceedings of the 9th Symposium on Identity and Trust on the Internet*, pages 69–83, New York, NY, USA, April 2010. ACM Press.
- [13] Milan Petković and R. Koster. User-Attributed Rights in DRM. pages 75–89. 2006.
- [14] <https://cloudsecurityalliance.org/>. Cloud security alliance, 2008.
- [15] www.opendatacenteralliance.org/. Open data center alliance, 2011.
- [16] Rohit Varma. McAfee labs: combating aurora. Technical report, 2010. https://kc.mcafee.com/resources/sites/MCAFEE/content/live/CORP_KNOWLEDGEBASE/67000/KB67957/en_US/Combating%20Threats%20-%20Operation%20Aurora.pdf.