CHAPTER 7

Identity Management and Control for Clouds

In the last few chapters we covered the technologies, usage models, and capabilities that are required to enable trusted infrastructure in the cloud–one of the foundation pillars for trusted clouds. We looked at the concepts, solution architectures, and ISV components that establish and propagate platform trust, attestation, and boundary control, all of which are required to enable the trusted clouds. The other foundational pillar to enable them is identity management, and that is the focus on this chapter.

Identity management encompasses the management of individual identities and their authentication, authorization, roles, and privileges and permissions within or across system and enterprise boundaries, with the goal of increasing security and productivity while decreasing cost, downtime, and repetitive tasks. Identity management thus constitutes an essential capability for attaining trusted clouds. From a cloud security perspective, and given the distributed nature of the cloud, questions like, "How do I control passwords and access tokens in the cloud?" and "How do I federate identity in the cloud?" are very real and thorny ones for cloud providers and subscribers. In this chapter, we will provide a broad introduction to identity, survey the challenges and requirements for identity management systems, and describe a set of technologies from Intel and McAfee that address identity requirements.

The emerging cloud infrastructure connects remote parties worldwide through the use of large-scale networks and through a diverse and complex set of hardware and software technologies. Activities in various domains, such as e-commerce, entertainment, social networking, collaboration, and health care are increasingly being implemented by diverse sets of resources and services. These resources and services are engaged at various levels within those domains. The interactions between different parties at remote locations may be (and sometimes should be) based on the information that's needed to carry out specific transactions with little knowledge about each other beyond that.

To better support these activities and collaborations, it is essential there be an information technology infrastructure with a simple-to-use identity management system. We expect, for example, that personal preferences and profiles of individuals be readily available as a cloud service when shopping over the Internet or with the use of mobile devices. Extensive use of cloud services involving sensitive computation and storage should be done without the need for individuals to repeatedly enter user credentials. In this scenario, the technology for *digital identity management* (IdM) is fundamental in customizing the user experience, underpinning accountability in the transactions, and

complying with regulatory controls. For this technology to fully deploy its potential, it is crucial we investigate and understand the concept of digital identity. This in turn helps in developing solutions for the protection of digital identity in IdM systems, solutions that ensure such information is not misused and individuals' privacy is guaranteed. Moreover, several strong authentication techniques aimed at protecting digital identity from misuse and access control rely on multi-factor identity verification and strong identity factors.

Phillip Windley defines digital identity as "the data that uniquely describes a person or a thing and contains information about the subject's relationships."[1] We like this definition because it allows for practical ways to assert identity. Identity may simply be a collection of attributes that together disambiguates someone, or it may be a digital identifier with known unique properties.

Note that identity plays a role in many contexts, interactions, and transactions of everyday life. Examples of "contexts" include personal, social, work, government and e-commerce. The interpretation and view of the same identity information may vary based on other contextual information, thus increasing the complexity of the problem of managing such identities. Moreover, the policies, control, and management of the same identity information may differ based on:

- Identities owned and controlled by users or data subjects
- Identities controlled by third parties or cloud service providers but known to data subjects
- Identities controlled by third parties, such as credit rating agencies and unknown to data subjects

Analysis of the multi-dimensional aspects of the management of identity information and other related details regarding IdM components is important while assessing which identity solution best fits consumers' or business users' interaction with cloud services. In this chapter we focus on methodologies of IdM, and especially Intel technologies. We will not explore why users submit or share information in the various mentioned ways and for what purposes. That limitation notwithstanding, such legal, social, and behavioral contexts may be important when considering the management and use of identity information.

Identity Challenges

There are a number of obligatory considerations in the architecture of almost any identity system. These include issue identity, identity usage, identity modification, and identity revocation. Based on the simple identity credential lifecycle illustrated in Figure 7-1, we can identify some general shortcomings in current approaches to managing identity information.

[1]Phillip J. Windley, *Digital Identity* O'Reilly Media, 2005), 8–9.

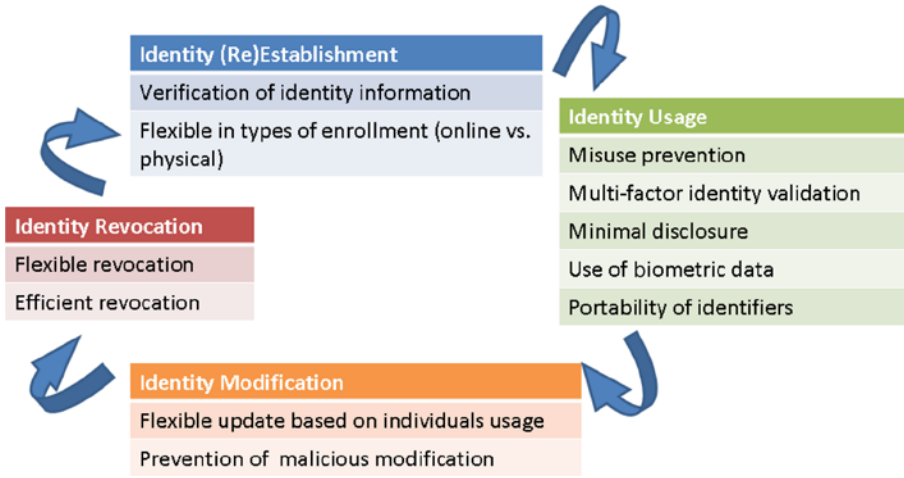

Figure 7-1. *Shortcomings of current federated IdM approaches in the credential lifecycle*

A limitation of current systems is that no information is provided about whether the strong and weak identifiers being enrolled and stored at the identity provider (IdPs) have been verified to be correct with respect to validity and ownership, as well as any indication of the strength of this verification. If an IdP has such information, then service providers are in a position to make a more accurate judgment concerning the trustworthiness of such identity information.

Furthermore, most IdM systems lack flexible enrollment mechanisms for the individuals who want to enroll in their systems. Enrollment can be in person at a physical location of an IdP or online. Current systems, however, do not provide alternative mechanisms for individuals to enroll. Moreover, the types of identity attributes that can be enrolled in most systems are restricted, based upon the nature of the IdP organization.

Identity Usages

A major drawback of current systems is that no specific techniques are provided to protect against the misuse of identity attributes stored at the IdPs and service providers. Even the notion of misuse is still being investigated and the solutions are in early stage of maturity. By "misuse" we refer to when dishonest individuals register fake attributes or impersonate other individuals of the federation, leading to the threat of identity theft.

To mitigate this threat, an upcoming trend is to require strong authentication. *Strong authentication* often refers to systems that require multiple factors, possibly issued by different sources, to identify users when they access certain services and applications. However, current approaches to strong authentication, such as those deployed by banks, enterprises, and government institutions, are neither flexible nor fine-grained. In many cases, strong authentication simply means requiring two forms of identity tokens—for example, password and biometric. Through prior knowledge of these token requirements, an adversary can steal and compromise that required identity information. Moreover, if the same tokens are repeatedly used for strong authentication at various service providers,

then the chances of these tokens being compromised increase. Yet, individuals should be able to choose any combination of identity attributes to perform strong authentication, provided the authentication policies defined by the verifying party are satisfied.

A recurrent issue in identity usage is the inability of some individuals to disclose minimal identity data about themselves to the service provider and IdPs, as per required to obtain the service requested. Digital identifiers have uniqueness properties that disambiguate someone or something within some domain of reference. For example, virtually every average-size company has two or more people with the same first and last names. Smaller companies have fewer name–space collisions; larger companies have more. To minimize the occurrence of these name–space collisions, identity management systems typically create unique digital identifiers. Interestingly, the identity management system could create a digital identifier that is globally unique, meaning that the identifier is not only unique within the company, but also may be unique at every other organization. This suggests that globally unique identifiers can be used to track and correlate activities between multiple organizations. Of course, such identifiers would be more than minimal, able to disambiguate individuals beyond what is required for the employer's use.

There are, likewise, several security and privacy concerns related to the extraneous identity information of the individuals stored at service providers and IdPs. Moreover, such data may be aggregated or used in a manner that could potentially violate the privacy requirements of those individuals.

Approaches need to be developed to address how biometric data can be used in an IdM system. Use of biometrics as an integral part of individual identity is gaining importance. At the same time, because of the nature of biometric data, it is not easy to use such data in a way similar to the traditional attributes. In theory, it should be possible to use biometric data together with other identity attributes to provide greater protection against identity attribute misuse. Biometric identifiers are designed to be globally unique. DNA biometrics are universally unique—it is believed that no human being has exactly the same DNA sequence as any other human who has ever lived or who will ever live.[2]

Another type of identity data becoming increasingly important in current systems is that related to individuals' histories of online activity. If this history can be verified and used for evaluating properties about an individual—for example, his or her reputation—then this data becomes part of that individual's identity. Consider a scenario in which an individual frequently buys books from an online store. This purchasing history can be encoded as an identity attribute of that individual, which in turn can be used to evaluate the person's reputation as a buyer. This history-based data needs to be better supported in current IdM systems. Companies like Amazon, Netflix, and Apple are using these types of attributes to classify customer buying habits and nature, in order to present a customized shopping experience.

Identity Modification

There are different approaches to take when it comes to finding mechanisms for the notification of changes in attributes. However, further investigation is required to develop flexible mechanisms for updating or modifying user-controlled enrolled identity

[2] *Encyclopedia of Espionage, Intelligence, and Security* Internet service. http://www.faqs.org/espionage/De-Eb/DNA-Sequences-Unique.html#b.

attributes. As the information is shared within the federation, updates performed on one system do not ensure consistency across the federation. Additionally, systems may fail to prevent malicious updates by attackers impersonating honest individuals.

Identity Revocation

Current federated IdM systems lack practical and effective revocation mechanisms. To enable consistency and maintain correctness of identity information, revocation should be feasible. Revocation feasibility for biometrics can be problematic, though. People can't simply change their fingerprints, irises, or DNA. Revocation in provider-centric systems, in which the IdP provides the required credential to the user each time, is relatively simple to achieve, however. A cryptographic digital identity can be mapped to a biometric identifier to create a credential with a manageable lifecycle. Such credentials are typically short term, and cannot be used without consulting the issuer again. If, however, the credentials are stored with the user, such as a long-term credential issued by the appropriate authority, then building a revocation system becomes more challenging and critical.

Identity Management System Requirements

In emerging paradigms of identity systems (such as user-centric identity) there are several distinct properties of the identity attributes that must be maintained. A key property is that of user control. While reasoning about the security and privacy properties of user control, we refer to the OECD countries. The OECD guidelines are widely accepted and they are the cornerstone of fair information practices and regulations designed to protect personal information around the world. In addition, Cameron's Laws of Identity are a recent set of prevalent guidelines regarding digital identity management.[3] They both aim at explaining the successes and failures of digital identity systems. In addition, design principles and rules to achieve several security and dependability properties are included. Figure 7-2 shows the properties of our taxonomy related to user control, illustrated as nodes. Taken together, these basic properties define what we mean by security and privacy in our solution.

[3] http://msdn.microsoft.com/en-us/library/ms996456.aspx.

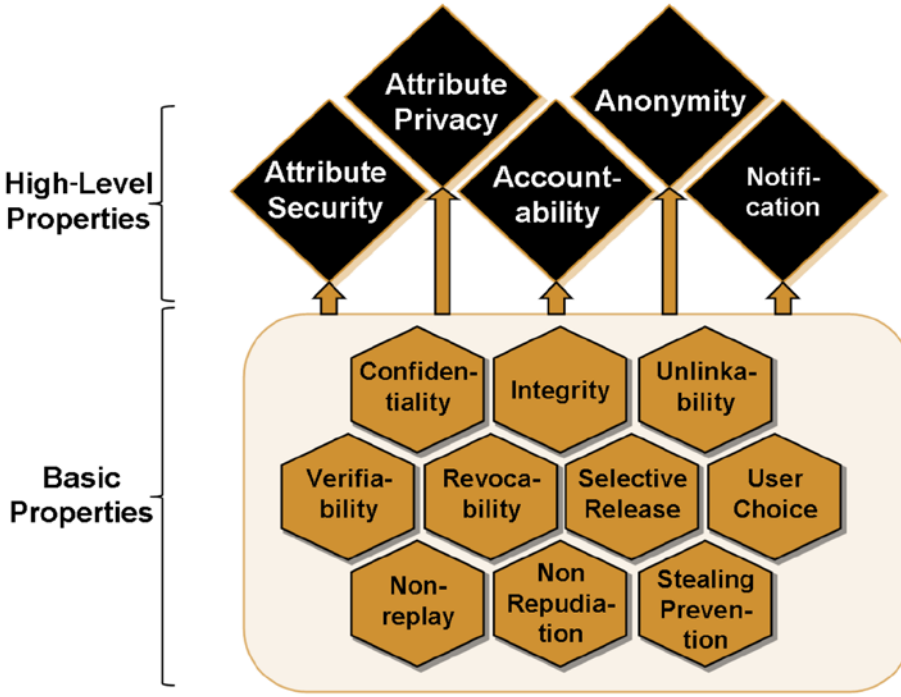

Figure 7-2. *Taxonomy of user control properties for identity attributes*

Basic User Control Properties

The basic properties related to the identity attributes either apply to the entire IdM system, to transactions in the system, or to the identity information and credentials of the entities involved. Although this classification is not exclusive, the semantics of the properties highlight which of the three they are relevant to. Table 7-1 briefly describes these properties.

Table 7-1. *Basic Properties Achieving Security and Privacy*

Property	Description
Confidentiality	This deals with the protection of information from unauthorized disclosure. This property applies to identity information and transactions in the system. Identity information should be accessible only by the intended recipients.
Revocation	Revocation of identity information is required to maintain the validity of information in the system. It should ensure that once invalid information is recognized, it is not used for identity verification purposes.
Integrity	This requires data not be altered in an unauthorized way.
Unlinkability	Ability to unlink two or more users or transactions so that an attacker, after having observed the transactions, cannot gain additional information by linking onto those transactions. Unlinkability prevents (illegitimate) merging of user profiles.
User Choice	The individual can choose among multiple IdPs and determine which attributes to release.
Verifiability	The individual can verify that the IdP provides the correct identity data about him or herself and according to that individual's intention. As such, an individual giving consent for what data is revealed, for what purpose, and to whom means that the individual's view of the transaction corresponds to the actual transaction and that the individual has agreed to execution of the transaction.
Non-replay	This prevents unauthorized parties from successfully using an individual's identity data to conduct new transactions. Non-replay is one prerequisite for obtaining the non-repudiation property.
Non-repudiation	The sending of non-repudiatable identity data cannot be denied by its sender and the ownership of the identity data cannot be denied.
Stealing Protection	This concerns the protection against unauthorized entities illegitimately retrieving an individual's data items. Stealing protection is required to achieve properties such as non-repudiation.
Selective release	Identity information can be released at a fine-granular level, as controlled by the individual. In this way, an individual can provide only the identity information that needs to be released for a service, without having to release additional information.

Key Requirements for an Identity Management Solution

The key requirements for an identity management system to ensure security and privacy of the identity data are as follows.

Accountability

Accountability refers to an ability to hold entities responsible for their actions in user transactions and for their use of identity information at the service provider and IdP. IdM systems have typically been focused on underpinning the accountability in business relationships and checking adherence to regulatory controls. In user-centric systems, the identity information of a user may be provided via the user's client. Therefore, it is required that, in addition to guaranteeing the integrity of the identity data, there should be accountability in providing such data. Accountability also becomes a significant issue if the user is to stay anonymous, as accountability and anonymity are, per se, contradictory properties. Nevertheless, conditional release of identity information can help in obtaining accountability in anonymous transactions. The eighth OECD accountability principle is devoted to understanding accountability, especially as it relates to privacy.

Notification

Notification Identity management (IdM):notification is desirable as a means for improving security by enhanced user control. Users should be able to receive and retrieve notifications regarding the usage of their credentials, so as to identify security breaches, and to estimate the extent of their compromised user identity information previously shared with external entities. It is desirable to allow users to collect data, whether under receive (push model) or retrieve (pull model) notifications regarding the usage of identity data. The sixth and seventh OECD principles of openness and individual participation can potentially be satisfied using comprehensive notification mechanisms.

■ **Note** Privacy legislation often requires notification of individuals impacted by release of privacy-sensitive *personally identifiable information* (PII). Identity credentials may be considered PII. Notification also helps individuals manage their privacy.

Anonymity

In transactions, *anonymity* deals with subjects remaining anonymous within an anonymity set—that is, with not being identifiable within some context or "set." In this context, something is more anonymous when it can be hidden in a population of otherwise indistinguishable members. A white sheep in a herd of white sheep is more

anonymous than a black sheep in that same herd of white sheep. Thus, anonymity is a specific notion related to data minimization, obtainable when the released attributes are not identifying the user.

Anonymity is supported by unlinkable transactions. Without such unlinkability, the anonymity set shrinks quickly when executing several transactions. *Pseudonymity*, or the use of pseudonyms as user identifiers, is a concept related to anonymity.[4] It plays a critical role in providing unlinkability and data minimization. There has been extensive work on the concept of pseudonymity, from both conceptual and implementation perspectives.

Note that conditional anonymity—that is, anonymity that holds only as long as a well-defined condition has not been fulfilled—can be based on conditional release of the identity information. In this way, the mechanisms providing for anonymity remain useful, as they are complemented by mechanisms for realizing accountability.

Data Minimization

Data minimization deals with minimal data release within a transaction. This can be achieved by having appropriate policy system support, by having unlinkable transactions, and by having a data release system that allows for selective and conditional release of identity information. This approach corresponds to the first OECD principle relating to collection limitation. It is also reflected in the European Data Protection Directive 95/46/EC and the national data protection laws within the European Union.

Attribute Security

The attribute *security property* reflects a comprehensive view of the security of a user's attributes. The main focus is on the correctness of attributes in the view of a service provider, meaning that the attributes belong to the person executing the transactions. This requires the attribute information to be integrity protected. Additionally, protections against stealing, and mechanisms to prevent sharing must be in place in order to stop another person from maliciously impersonating a user's identity. Furthermore, revocation of identity information must be feasible. Attributes in certain cases must be kept confidential with respect to parties other than the ones involved in the transaction.

Attribute Privacy

Attribute privacy refers to giving the user control over the attribute data. This is supported by system assurance and by allowing for user-chosen IdPs. Both those properties account for user-centric decisions regarding which IdP to trust. Anonymity and dependent properties very likely contribute to attribute privacy in that they help avoid the unnecessary release of (identifying) information. Data minimization also directly provides privacy.

[4] http://en.wikipedia.org/wiki/Pseudonymity.

An orthogonal property essential for reaching attribute privacy is support of privacy policy management, enforcement, and agreement. User control over attributed data helps the user remain anonymous outside the context or domain in which the identity is being used. Preventing disclosure of privacy sensitive information outside the context where it is needed is important; once this information is disclosed, it can't be reclaimed. Confidentiality measures ensure that privacy-sensitive attributes are not unintentionally disclosed to any party.

Identity Representations and Case Studies

There are various types of identity tokens used for device and user identification and for access control. Key examples are illustrated in Figure 7-3. From a security perspective, the prevalent method of conveying identity information that is certified by a trusted third party is through certificates.

Figure 7-3. *Types of identity tokens*

Based upon the representation of certified digital identity information, the resulting system may or may not satisfy one or more of the properties covered in the previous section. In the following, we discuss technical mechanisms that can be used to obtain an identity management system with the properties described in our taxonomy. We refer to three different core mechanisms. Note that what follows is not a complete survey of mechanisms but, rather, focuses on the more interesting properties relevant to the representation of certified digital identity information.

PKI Certificates

Standard certificates, like X.509, allow, in conjunction with a private signing key, a user to prove that attributes have been issued to him or her. The certificate contains attributes and a public key signed by the IdP (the issuer of the certificate). Note that in a typical IT enterprise, such certificates are used for managing users and client machines in order to establish secure channels between two enterprise entities, for provisioning, and for updating user machines or profiles.

To assert the attributes of a certificate to a relying party, the user engages in a challenge-response protocol with the relying party. This protocol requires the certificate to be sent to the relying party and a signature made with the private key. The step reveals all attributes of the certificate to the recipient. Technically, certificates are based on standard digital signature schemes such as RSA and are represented by standards like X.509,[5] which define the formats of the certificates.

Traditional certificate-based technologies allow for constructing systems in which a certificate is issued once and can be used repeatedly by users to reveal the attributes contained in the certificate. Thus, this technology allows for off-line IdPs. The tokens are generated by the user without involvement of the IdP, making this method flexible with respect to this aspect. This technology is, for example, used in multiple ID card proposals and public key infrastructure-based systems.

Security and Privacy Discussion

In the discussion of security requirements, note that the integrity of such schemes is accounted for by the user attributes being included in the certificate signed by the IdP, using standard signature schemes, and e being provided each time the attributes are asserted to a relying party. Confidentiality of attribute information is achieved by using encryption schemes in conjunction with public key infrastructure (PKI). Stealing prevention methods for standard certificate systems target protection of the master private key, as the certificates are made available to relying parties anyway. The following mechanisms can be used, also in a combined fashion:

- Binding all certificates to one master private key of the user and mandating appropriate protection of this key—for example, in a hardware token. As this requires the hardware token be used in each transaction, the portability of such tokens becomes important.
- Applying operating system mechanisms to prevent a user from sharing his or her key.
- Using multi-factor authentication makes it harder to share the token—for example, if it is derived from the biometrics of the user.

Finally, revocability can be achieved by the prominent mechanism of certificate revocation lists (CRLs) and associated protocols. This requires an additional protocol to be run in order to obtain the latest revocation list.

With respect to the privacy requirements, verifiability holds as a user can inspect the certificate and thus has control over the attribute information being revealed. Conditional release cannot be realized in the setting in which the protocols operate, as an IdP is not involved in the transactions. Technically, of course, protocols could be conceived that involve the IdP in a transaction to obtain the conditional release property, but by discussing this we would leave the basic paradigm of the system.

[5] http://en.wikipedia.org/wiki/X.509.

Selective disclosure is not possible in a setting that uses standard certificates, as these certificates always have to be revealed as a whole and no subset of their attributes can be revealed because of the properties of the employed standard signature schemes like RSA or DSA. Finally, unlinkability may not be achievable in this setting. This is because transactions done with multiple IdPs, or multiple transactions with one IdP, are linkable, as the same certificate bit string is being provided in every transaction.

Limitations

The main problem with using standard user-side certificates is the lack of overall privacy properties, and thus the strong trust assumptions that we have to make on the relying parties. Assuming stronger trust in a relying party may not be realistic, relying parties may benefit from gathering extraneous users identity information. The U.S. National Institute of Science and Technology (NIST) has defined comprehensive criteria for understanding and evaluating identity management systems.[6] Those criteria demonstrate how the principles of identity management may be applied when evaluating identity management systems for purchase or use.

Identity Federation

There are several enterprise identity usages that require nonemployee accounts, business-to-business (B2B) interactions, and interaction and use of data from multiple applications that may exist across different networks. *Identity federation* is a term used when organizations form trust relationships whereby identities or assertions of an identity can be shared by all applications within the federation. It is critical that business partners involved in a federation build a trust relationship with one another. This trust relationship, defined by business, technical, and legal agreements, describes the applications that are involved, the user profile information that is to be shared, and the responsibilities of all parties to manage and protect user identities.[7]

Several standardization initiatives for identity federation are being developed across the world. Among them, Kantara Initiative (http://kantarainitiative.org/) and WS-Federation (http://en.wikipedia.org/wiki/WS-Federation) are two significant efforts. These initiatives define an identity federation framework that allows assurance-levels mapping for various service providers. For example, the Kantara Identity Assurance Accreditation and Certification Program assesses applicants against its assessment criteria, including alignment with the NIST 800-63 Levels of Assurance (http://csrc.nist.gov/publications/nistpubs/800-63/SP800-63V1_0_2.pdf) and it grants successful candidates of the program the right to use the Kantara Initiative Mark, a symbol of trustworthy identity and credential management services at specified assurance levels. It also collaborates with Open Identity Exchange (OIX) and other related initiatives to allow an interoperable digital trust framework to promote adoption of a robust online trust ecosystem. Similarly, WS-Federation was created with goal

[6] csrc.**nist**.gov/publications/**nist**pubs/800-63-1/**SP-800-63**-1.pdf.
[7] http://assets1.csc.com/cybersecurity/downloads/FIM_White_Paper_Identity_Federation_Concepts.pdf.

of standardizing the way companies share principals and machine identities among disparate authentication and authorization systems that cross corporate boundaries. This translates to mechanisms and specifications that enable federation of identity attributes, authentication, and authorization information, but it does not include trust establishment and verification protocols.

The common objectives for federation proposals have been primarily to reduce the number of user-business interactions and exchanges of information such that critical private information is used only by appropriate parties. There is a need to ensure that user information is available to the SPs on demand, online and with low delay. Thus, user data is more up to date and consistent compared to the situation where each principal has to maintain its data in multiple places. Federations reduce costs and redundancy because the member organizations do not have to acquire, store, and maintain authorization information about all their partners' users. Also, both the federation initiatives try to preserve privacy, as only data required to use a service is transmitted to a business partner.

Single Sign-On

Single sign-on (SSO) improves security and usability. With SSO, user accounts and passwords are not reused across multiple sites or servers. SSO also improves usability by limiting the number of times the user must re-authenticate. Popular SSO systems include Kerberos, ActiveDirectory, SAML, and OpenID. The SSO systems work by converting the user authentication event into an access credential that is cryptographically protected. An access manager located at a remote server or within the same platform verifies the credential, rather than performing an authentication challenge with the user. The SSO credentials grant access for a period of time; that access is rescinded upon credential expiry. These systems make security and usability trade-offs that can be undesirable, however. If the credential timeout value is too long, malware can reuse the credential to prolong access that is otherwise unauthorized. If the timeout value is too short, the user must re-authenticate to continue the session.

An example of an SSO system is the McAfee Cloud SSO. It ensures that strong authentication is used for the customer's cloud-based software as a service (SaaS) applications, and helps allow SSO access to the cloud-based applications while complying with enterprise security policies. This solution is flexible and permits for an on-premises as well as SaaS-based solution, or both (hybrid model).

Intel Identity Technologies

Intel Corporation has developed several technologies useful for implementing identity management systems. Hardware support is often beneficial because it presents physical boundaries that can inhibit or prevent compromise of the identity management system by malware.

Hardware Support

Intel provides hardware support to enable hardened identity technologies on Intel platforms. Some basic underlying capabilities as of 2013 are as follows.

Virtualization Technology (VT)

Intel's Virtualization Technology (VT; see Figure 7-4) creates an additional layer of protection between physical memory and devices beneath the operating system.[8] Virtualization can be used as a security mechanism by isolating the operating system and applications from hardware using a small, and therefore well-understood software layer, that's also known as the hypervisor, ensures that hardware access follows some prescribed rules of behavior. The hypervisor implements a security policy designed to protect the integrity of information in memory, in peripheral devices, and in the CPU.

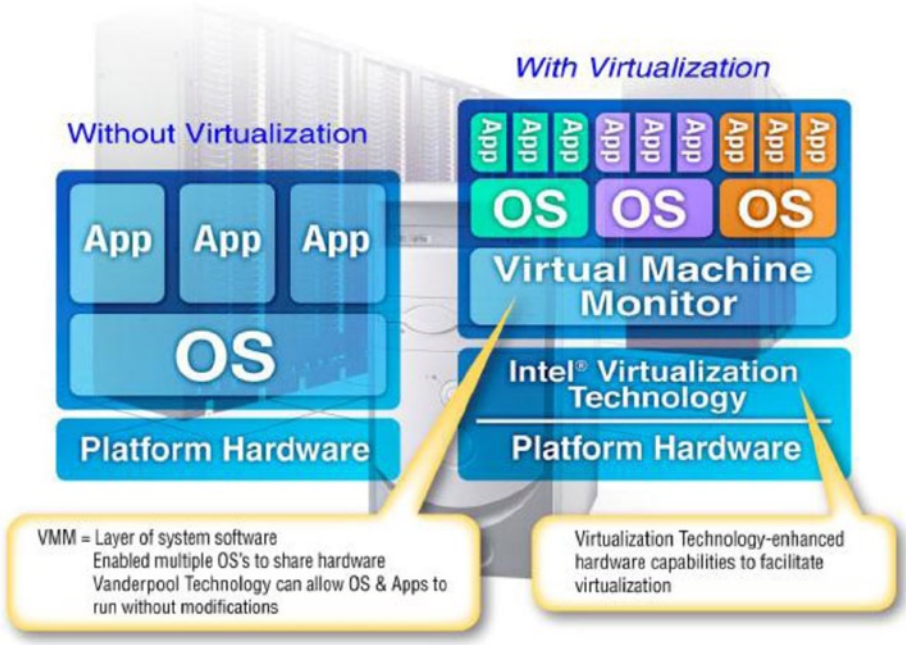

Figure 7-4. *Intel Virtualization Technology*

Intel Identity Protection Technology (IPT)

Intel's Identity Protection Technology (IPT; see Figure 7-5) consists of several credentialing and credential management capabilities for client platforms.[9] They are implemented in a security engine in hardware and offer an additional layer of security hardening and isolation from malware.

[8]For information about Intel Virtualization Technology, http://ark.intel.com/products/virtualizationtechnology.
[9]http://www.intel.com/content/www/us/en/architecture-and-technology/identity-protection/identity-protection-technology-general.html.

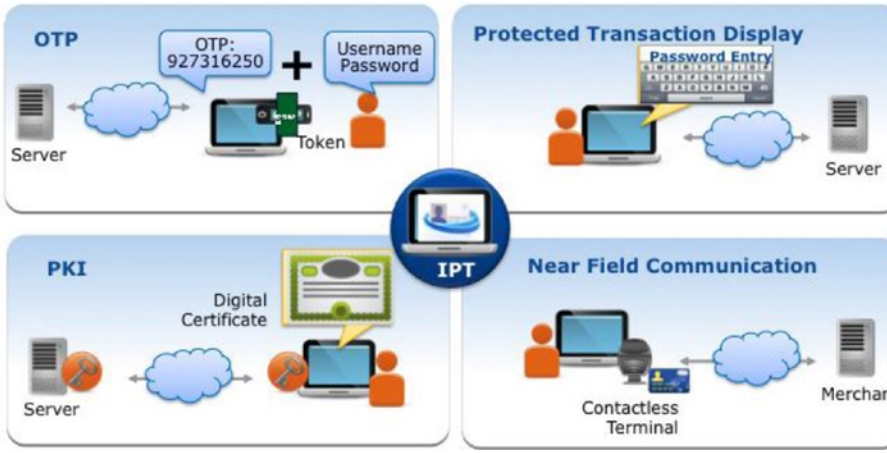

Figure 7-5. *Intel Identity Protection Technology*

- ***IPT-OTP.*** One-time passwords are single-use identifiers that cannot be anticipated or replayed by an attacker. Typically, the user and service provider agree to use a common "seed" from which a sequence of one-time passwords is generated. Keeping the seed secret is essential to security. IPT-OTP protects seeds in a hardware security engine.

- ***IPT-PKI.*** Public key infrastructure (PKI) is a set of hardware, software, people, policies, and procedures designed to create, manage, distribute, use, store, and revoke digital certificates.[10] Certificates are identity credentials that associate an asymmetric key[11] with an identifier. IPT-PKI is a cryptographic service provider that protects private asymmetric keys in a hardware security engine.

- ***IPT-PTD.*** In many cases, use of a credential requires user approval. Malware attacks that fake user approval may be a sufficient form of compromise to achieve the attacker's objective. IPT-PTD protects the output path between the hardware security engine and the graphics controller. Malware may not observe information displayed to a user. Protected output may be used to protect PIN input by rearranging a PIN pad display in a random order. When a user inputs the PIN using the randomized PIN pad, malware observing the mouse clicks cannot determine which (X,Y) coordinates map to which PIN digit. PINs are used by IPT-PKI and IPT-OTP to authorize use of a credential by a specific person.

[10]See Wikipedia, "Public Key Infrastructure." http://en.wikipedia.org/wiki/Public-key_infrastructure.
[11]See Wikipedia, "Public Key Cryptography." http://en.wikipedia.org/wiki/Public-key_cryptography.

- ***IPT-DeviceID.*** Use cases involving the computing platform when no user is present may require authentication. IPT-DeviceID associates a platform identifier with a credential. IPT-DeviceID protects the device credential in hardware.

Intel Security Engine

The security engine used to implement Intel's Identity Protection Technology has several capabilities that may be useful for enhanced privacy protections.

- ***Enhanced Privacy ID (EPID)***. The EPID is an asymmetric key provisioned at platform manufacturing time by Intel. It can be used to authenticate that an Intel platform security engine is performing a function securely. For example, the EPID key may be used to digitally sign the applet running on the security engine to prove its integrity and validity. EPID may also be used to prove an Intel security engine protects an IPT-PKI key. As the name suggests, EPID is privacy enhanced. This means the verifier can tell that the endpoint is an Intel security engine, but can't tell which one–even when the same platform returns a second time, the verifier can't correlate the second session with the first session.
- ***Sigma.*** The Intel security engine also implements a SIGn-and-MAc protocol (Sigma) based on a Diffie-Hellman key exchange that is signed using the EPID key. Sigma produces symmetric session keys for encryption and mac-ing of bulk data. Sigma allows a stream of data originating from the security engine to be transferred to a remote service provider. Sigma is useful for protecting logged event data, sensor input values, and configuring of policies.

The use of EPID and Sigma building blocks allows a client platform to interact securely without disclosing privacy sensitive information unnecessarily.

Intel's Manageability Engine (ME) implements security primitives for encryption, key exchange, and identity protection. It is integrated into Intel's chipsets. The ME (Figure 7-6) is isolated from the host operating environment and memory.

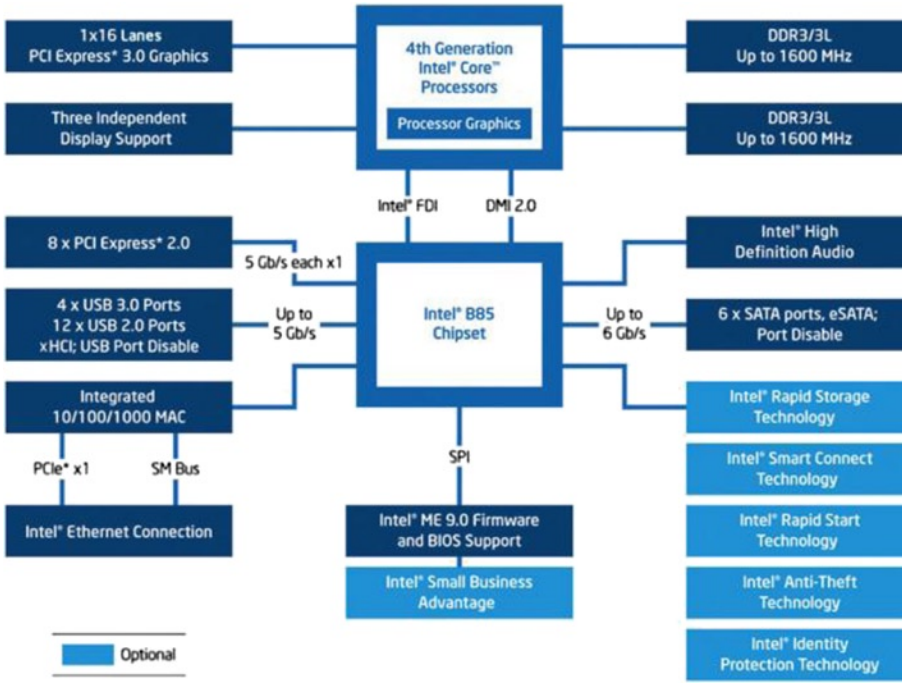

Figure 7-6. *Intel B85 chipset containing the Intel Manageability Engine*

Cloud Identity Solutions

Security services vendors such as McAfee provide a suite of security solutions for a wide range of enterprise and cloud-hosted services. Identity management is part of a comprehensive solution. Identity management services implement the credential lifecycle and ensure interoperability with a wide variety of services and applications.

The McAfee Cloud Ecosystem (see Figure 7-7) includes unified management, policy, reporting, and enterprise integration of pluggable security capabilities ranging from data loss prevention to web security. These capabilities are built upon an infrastructure that supports global threat intelligence monitoring and response, cloud-aware security, and enterprise-orchestrated policies. Such cloud-based security solutions offer dynamic protections that adjust as situational awareness changes. Cooperation among thousands of nodes participating in building a clearer picture of the threat landscape ensures that security incidents are processed and countermeasures are applied.

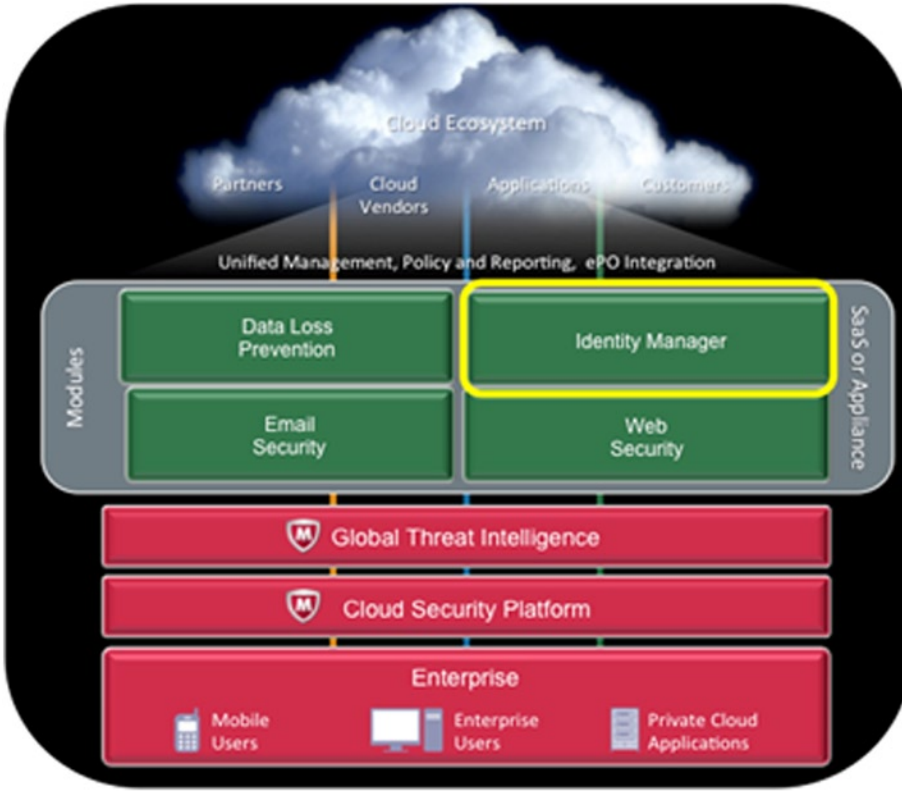

Figure 7-7. *McAfee cloud identity solution*

A cloud-based approach to security that includes identity management ensures that the known trusted users can be distinguished from the unknown and less trusted. Selection of a security services provider that implements such identity management comes with the implication that the provider is protecting the user's privacy in addition to ensuring computing security.

Summary

Identity management is an important component of a comprehensive cloud security infrastructure. This infrastructure must be rooted in sound identity management principles that not only ensure robust control of the identity credential lifecycle but also satisfies users' privacy desires. The identity management landscape is complicated by constant innovation and the evolution of authentication factor technology, identity credentials, and infrastructure investments. Complexity isn't necessarily good for security and privacy protection, but it appears to be an unavoidable reality. Taking the time to select a competent identity management provider can be an effective strategy for managing this complexity.

Computers that have deeply integrated identity protection technologies can be very effective in protecting user privacy and identity, while also delivering identity management solutions that interoperate with an already complex ecosystem of cloud services and that can promise continued support for an emerging Internet-of-things.

Identity management in the future holds many interesting challenges, especially when the Internet-of-things (IOT) is factored in.[12] The IOT promises Internet connectivity to a host of embedded systems, building automation control, smart appliances, and vehicles of all kinds. Technology advances make it practical to build wireless self-contained sensors that link directly to the Internet, feeding databases that analyze and infer new knowledge about the world. As people interact ever more widely with the world, sensors may be able to identify their unique properties using kinematics.[13] In an IOT world, devices will come equipped with device IDs to ensure they can be managed and controlled by authorized servers. They will have privacy-preserving capabilities that respect their user's right to privacy by filtering biometric data locally and translating it into digital credentials that more easily support credential lifecycle management.

In the next chapter, we focus on building and extending security, integrity, and confidentiality to applications and workloads that run in the cloud. As you would expect, the applications and workloads, which are typically packaged as virtual machines, anchor their integrity and trust in the chain of trust that is built with trusted compute pools and associated concepts and technologies that have been discussed in preceeding chapters.

[12]Intel adds Intelligence to Cloud for Internet of Things. http://iotinternetofthingsconference.com/2013/10/09/intel-adds-intelligence-to-cloud-for-internet-of-things/.
[13]See Wikipedia, "Gait Analysis Using Kinematics." http://en.wikipedia.org/wiki/Gait_analysis.