



Protecting Information

Steps for a Secure Data Future

A White Paper by:

Members of the Security Forum, a forum of The Open Group

January 2014

Protecting Information

Copyright © 2014, The Open Group

The Open Group hereby authorizes you to use this document for any purpose, PROVIDED THAT any copy of this document, or any part thereof, which you make shall retain all copyright and other proprietary notices contained herein.

This document may contain other proprietary notices and copyright information.

Nothing contained herein shall be construed as conferring by implication, estoppel, or otherwise any license or right under any patent or trademark of The Open Group or any third party. Except as expressly provided above, nothing contained herein shall be construed as conferring any license or right under any copyright of The Open Group.

Note that any product, process, or technology in this document may be the subject of other intellectual property rights reserved by The Open Group, and may not be licensed hereunder.

This document is provided "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, OR NON-INFRINGEMENT. Some jurisdictions do not allow the exclusion of implied warranties, so the above exclusion may not apply to you.

Any publication of The Open Group may include technical inaccuracies or typographical errors. Changes may be periodically made to these publications; these changes will be incorporated in new editions of these publications. The Open Group may make improvements and/or changes in the products and/or the programs described in these publications at any time without notice.

Should any viewer of this document respond with information including feedback data, such as questions, comments, suggestions, or the like regarding the content of this document, such information shall be deemed to be non-confidential and The Open Group shall have no obligation of any kind with respect to such information and shall be free to reproduce, use, disclose, and distribute the information to others without limitation. Further, The Open Group shall be free to use any ideas, concepts, know-how, or techniques contained in such information for any purpose whatsoever including but not limited to developing, manufacturing, and marketing products incorporating such information.

If you did not obtain this copy through The Open Group, it may not be the latest version. For your convenience, the latest version of this publication may be downloaded at www.opengroup.org/bookstore.

ArchiMate[®], DirecNet[®], Jericho Forum[®], Making Standards Work[®], OpenPegasus[®], The Open Group[®], TOGAF[®], and UNIX[®] are registered trademarks and Boundaryless Information Flow[™], Build with Integrity Buy with Confidence[™], Dependability Through Assuredness[™], FACE[™], Open Platform 3.0[™], Open Trusted Technology Provider[™], and The Open Group Certification Mark[™] are trademarks of The Open Group.

OASIS[™] and XACML[™] are trademarks of OASIS, the open standards consortium, the owner and developer of this specification.

CORBA[®] is a registered trademark of Object Management Group, Inc. in the United States and/or other countries.

Microsoft[®] is a registered trademark of Microsoft Corporation in the United States and/or other countries.

Protecting Information: Steps for a Secure Data Future

Document No.: W142

Published by The Open Group, January 2014.

Any comments relating to the material contained in this document may be submitted to:

The Open Group, 44 Montgomery St. #960, San Francisco, CA 94104, USA

or by email to:

ogspeccs@opengroup.org

Table of Contents

Executive Summary..... 4

Introduction..... 5

Requirements..... 9

Summary of Existing Protection Capabilities..... 13

Protection Issues with Current Technologies 15

Context 20

Challenges..... 26

A Data Protection Future..... 29

Strategies for Achievement 32

Summary 35

Appendix A: Notes on Terminology..... 36

Acknowledgements..... 38

References..... 39

About the Security Forum 40

About The Open Group..... 40



*Boundaryless Information Flow™
achieved through global interoperability
in a secure, reliable, and timely manner*

Executive Summary

This White Paper explains why information protection to meet today's and tomorrow's requirements needs to use stronger, more flexible protection mechanisms around the data itself.

It reviews the issues surrounding data protection today, describes properties that data protection mechanisms should include to meet current and future requirements, considers why current technologies don't deliver what is required, and proposes a set of steps for migrating to a data protection future that we expect will meet future needs, and strategies for achieving the stronger more flexible protection solutions we need.

In addition, business analytics are driving even larger volumes of data processing and data warehousing. The data provides the knowledge for companies to predict and react to business cycles and change. The data and knowledge it provides can be a company's most valued assets.

This White Paper supersedes the earlier Data Protection: Problem Statement and Requirements for Future Solutions (W12C) published in October 2012.

Introduction

Computers and the Internet have accelerated both the generation and amassing of large amounts of information or data. Data largely comes from two sources: the conversion of physical matter into electronic representations and the creation of new data that was either not useful or not possible in the pre-digital world. The result is a concentration of wealth in the form of electronic media.

Note that while information may be defined as the collection of or application of intelligence to data, in this White Paper the two words are used almost interchangeably as more general terms. A more precise set of classifications can be found in Appendix A: Notes on Terminology.

Data has increased in volume and value so quickly that there is an organizational lag in the perception of its value. This often results in a weak, but inaccurate, business case for deploying better data protection.

Enduring data protection requires two key changes to current approaches:

- Protection must move closer to the data itself; that is, away from reliance on network and platform security.
- Data protection must be automatically applied using policy-driven protection services.

This will not be accomplished by purchasing single standalone products, but rather through the implementation of cooperating components connected by vendor-neutral, standardized protocols.

While some changes can be evolutionary, new services and capabilities are required. Some examples are the externalization of access control decisions, the deployment of universal policy decision engines, and the automatic generation and management of metadata.

Information has value. Fred Smith, the CEO of FedEx, stated that the information about the packages they ship is more valuable than the package content. Admiral Grace Hopper at her retirement lecture said that one day information would be listed on balance sheets as corporate assets. But to leverage that value, it has to be accessible. Dan Geer notes that information is only valuable to the extent that it can be used.

E-business models require extensive, but controlled information sharing. The growth of network interconnectivity and subsequent demand for enterprise-to-enterprise information access means that information no longer sits in a closely guarded room but is widely distributed. While in general this accelerates business decisions and increases productivity, one risk of this (described in the book OVERConnected) is the movement from stable to volatile interconnections that can cascade out of control.

The value, and need for protection, can escalate rapidly. And the value also changes depending on the participants and their roles and responsibilities. Imagine the following email conversation:

Jim: *Ian, would you like to meet at Wimpy's for lunch?*

Ian: *Sure, I love their Quorn burgers. As long as we are meeting, can you bring the merger prospectus?*

Jim: *I'll bring a copy; in the meantime, here is a digital version. [File Attached]*

Current information security technologies and practices neither adequately secure information nor allow it to be shared when needed. They were designed for an earlier era and have not kept up with information as it has evolved from simple facts to large amounts of complex data and metadata. This usually results in some combination of two data protection approaches with their associated issues:

Protecting Information

- Implementation of global controls which lack the granularity needed to protect and share information appropriately
- Implementation of local controls based upon *ad hoc* attention-getting events rather than a comprehensive plan, which is usually too late, costly, and does not scale

As enterprises continue to integrate, existing information protection mechanisms are becoming too costly to scale. In the meantime the Internet has become a vector for advanced attacks designed to steal information or mine information without authorization. These attacks are growing in sophistication and volume and are overwhelming current information protection capabilities.

The collection of vast amounts of information in systems connected by the Internet present an ideal situation for a thief. Theft takes place remotely, resulting in little risk to the thief. Large amounts of data can be stolen rapidly – or rather copied – leaving the original in place thus avoiding detection. The advantages of data over traditional theft are illustrated below.



















		Theft Speed	Personal Risk to Thief	Chance Theft Will Be Detected	Chance of Keeping a Copy if Detected	Economic Value
 <p>Physical Objects</p>						
 <p>Physical Data Objects</p>						
 <p>Data Objects</p>						

Figure 1: A Thief's Risk Perspective

In spite of the accumulated data value, the ease of access by thieves, and the numerous successful reported cases of data theft, little has been done to evolve data protection beyond encryption systems. There is a growing industry built around Data Loss/Leakage Prevention (DLP) tools that is focused on the organization, tracking, and control of data, with one of the most difficult challenges being the protection of sensitive data and the identification, classification, and location of data in the control of an enterprise.

Data theft isn't always done with criminal intent. The collection of vast amounts of data and their placement on the Internet or even on closed networks that are accessible from the Internet provide opportunities for surveillance and espionage that could only have been dreamed of a decade ago. The only real difference between the two is whether or not a government is collecting data on its own citizens or those of another nation state. The technical capabilities are the same. The Wikileaks disclosures and the information released by Ed Snowden demonstrate the extensive use of data collection and aggregations.

Historically, data has been protected at different layers. Network-level protection has the ability to protect large segments of a corporate intranet, including the data that resides therein. Advantages include low cost, scalability, and manageability. Network-based protection tools include firewalls, routers that can encrypt

Protecting Information

traffic, and network-based access control systems. Reliance on network-based data security is falling out of favor for several reasons, including:

- Any breach tends to expose all of the data and resources – there is no granularity.
- Network-based security tends to be controlled by a single organization – so it does not work well between customers, partners, or suppliers.
- There is little granularity in the amount of protection – it’s either off or on.

Vendors that offer network-based security solutions have been adding some features to overcome these limitations, but are doing so by leveraging their existing technology that was designed to “fight the last war” and are rapidly losing ground.

At the next layer, data protection can be applied using operating system controls. Typically, these are file permission settings, Microsoft® Active Directory account settings, or global disk encryption. While the scope is more limited than that of network-based controls, it is still too broad. Complex applications – such as database management systems, CAD/CAM systems, etc. – may also provide application-wide data protection controls. While providing more local control, these systems are not granular enough to protect data, and when breaches do occur they tend to expose all data under control of the operating system or application. They are typically attacked by leveraging vulnerabilities in operating systems, and sometimes in the applications. We are also losing this race.

Tools for directly protecting data are currently limited. The most common tool is file encryption. This can be standalone using products such as PGP, or part of a larger application such as encrypting an email message using an email system. More recently, organizations have been using DRM (Digital Rights Management) products. These tools allow a user to decide what controls should be on the data and when to apply those controls. They provide more granularity; for example, with DRM you can send a message but prevent it from being printed or forwarded. Some of the flaws in current offerings include:

- The lack of standards – today it is difficult to use DRM to protect information that is shared between different organizations without requiring them to buy the same product, and even when they do, current DRM products often require connecting their internal identity management systems.
- Many encryption and DRM products are tied to your identity so an attacker can just steal the identity to obtain access.
- It’s still easier to attack the infrastructure than the cryptography on which the DRM is based. For example, a data thief would attack the DRM implementation by attacking the operating system platform that the identity management system¹ is running on and then use that access to attack the identity management system itself. Once it is compromised, the thief would just clone or subvert the account and identity information to become the target, thus inheriting the target’s rights to the data – bypassing the DRM.

Enterprises have typically used a method for protecting data called Discretionary Access Control (DAC). In a DAC system, users are given a set of tools and system properties which they use to protect information. Typically, written policies are guidelines issued to users, but the burden of selecting and applying the correct level of protection falls directly on the user. In contrast, in a Mandatory Access Control (MAC) system, access is controlled by systems that use machine-readable policies, information about the user, the data, the

¹ Adherence to Jericho Forum Identity Commandment #1 (see References) increases protection against but will not eliminate this form of attack.

Protecting Information

environment, and the desired action. The systems performing access control functions utilize all these types of information to render decisions.

Most current access control systems are of the DAC type. This is because the technology to build MAC systems has been limited and expensive, and standardized connecting protocols were lacking. Therefore, MAC has only been applied to the most critical information; e.g., national security controlled information. With the emergence of better technical solutions, standard protocols, and ubiquitous connectivity coupled with the growing value of corporate data, the time is right to begin a shift towards policy-based MAC.

Initially data wasn't protected at all. It was stored locally and, from a policy perspective, the Arpanet (the ancestor of the Internet) prohibited commerce until 1988 so there was little incentive for transmitting non-public data. From a technical perspective, there were many different, incompatible, network protocols. While there were some protocol translation gateways, they didn't work very well and there was little interconnectivity. Data was pretty much limited to the specific local networks of its owning organization.

Three events caused the development of early data protection mechanisms. First, the development of the Personal Computer (PC) meant that data was now scattered on small devices that were susceptible to theft. Second, the commercialization of the Internet allowed its use for financial transactions and the transmission of intellectual property. And third, the development of public key cryptography and the movement of cryptography in general to industry from governments allowed the development of technologies to protect data or rather provide enough confidence in its protection to enable the growth of e-commerce. This *ad hoc* data protection era, enforced by DAC mechanisms, is where we are today.

This is no longer adequate. The recent growth in the use of mobile devices (smart phone, tablets, etc.), high density portable storage (USB and small disk drives), and cloud computing services has greatly outpaced the ability to control the location and protection of information. Cryptographic and other protections are often incomplete or do not match protection requirements. It is the purpose of this White Paper to describe the strategies necessary to move to the next era – that of data-centric security.

While some specific data protection requirements are enumerated below, we are developing a new Open Group White Paper entitled “The Need for Data Commandments”:

- To explain why our industry needs to establish a set of high-quality Data Commandments
- To outline key considerations and complex issues involved in developing them
- To provide a draft set of data principles as a sound basis for undertaking further work to complete developing a set of Data Commandments

For further information on the availability/publication of this White Paper, contact The Open Group Security Forum at security-interest@opengroup.org or check [The Open Group Online Bookstore](#).

Requirements

Business Requirements

Technical solutions must be affordable, usable, and manageable, and must be aligned with and facilitate explicit business objectives. There should be a clear understanding of the threat they are mitigating. Moreover, they should improve user workflow, rather than hinder it. A user should be aware of the responsibility to protect information, but the data protection system should not constrain the user's task nor add significant task overhead. As some protection systems are already in place, replacement information protection services should be evolutionary where possible, allowing for smooth migration.

Enterprise culture – and end-user practices – must change so that employees, contractors, business partners, customers, suppliers, governments, and others who create and handle information realize the innate value of information and internalize safe handling practices. Data has increased in volume and value so quickly that there is an organizational lag in the perception of its value. Information protection policies must support these new business practices which make extensive use of data while relying on achievable implementation technology to protect it.

There needs to be a structured entitlement process at a business level. These entitlements – the decisions governing data access – are often based on a combination of informal processes, shifting priorities, policies (which often lag data use and threat changes), and other corporate cultural habits as described in the previous paragraph. A formal process for creating, documenting, and disseminating these entitlements is a necessary input to architecting and designing a set of capable information protection services. A standard data classification taxonomy should be associated with the entitlement process.

Security Requirements

Protection policies should be consistent throughout the lifecycle of the data. Typically, this includes creation, modification, aggregation, distribution, archive, and destruction. There should be consistent protection policies applied to the data at each of these stages, and those policies should only change when explicitly and authoritatively requested. The entity that owns or manages the data may make an explicit change to a given data object's protection profile, but this capability to provide lifecycle protection consistency should be the default. There should be mechanisms for changing the protection if the sensitivity of the data changes during the lifecycle, but there should not be gaps in the protection enforcement during these policy changes. Ideally the data should also be protected while in use, and while this is more difficult, techniques such as homomorphic encryption are being developed to do just that.


Protection policy should also be consistent across space as well as time. That is, the protection policies applied to the data should be the same regardless of its location or environment. This means that there is policy consistency regardless of whether the data is stored on a computer, stored in a database, attached to an email, in transit on a network, present on a smart phone, etc. Instead, the protection policy decisions and enforcement will depend upon context (policy attributes) of each particular situation in the lifecycle of the data. However, the owning entity should also have the ability to adjust protection policies as risk, threats, or legal jurisdiction change. For example, a new threat may require a replacement of the policy enforcement mechanisms and changing requirements in different export jurisdictions will often require protection policy changes.

While it would be typical for most organizations to control the policy for their data objects, there should be capability to negotiate or jointly manage policies between different organizations. This policy federation

Protecting Information

capability is necessary in today's environment where organizations engage in joint development with their suppliers, customers, and other business partners.

In order to effectively leverage data, these protections must be applied at the data layer itself. While protection at the lower layers (application, operating system, network, and physical storage) tends to be cheaper and more manageable as the controls get broader, this all-or-nothing approach results in significant collateral exposure and the protection mechanisms tend to stay with the environment and not the data. The result is either greater risk or less flexibility when data is used in an e-commerce environment. The scope of the protection should be specific to the information being protected. Ideally the protection will be carried with the data, making it consistent across any environment in which the data resides. This is not to say that data won't encounter additional protection provided by network, applications, and other environmental controls – just that it must not rely on them as their application will be inconsistent. Figure 2 is based on a conversation with Dan Hitchcock.²



See: Dan Hitchcock, *Evolution of Information Security Technologies*, 2005 at <http://movetheworld.wordpress.com>

Figure 2: The Data-Centric Access Control Future

This does not mean that other protection mechanisms should be discarded, just that their uses and priorities change. For example, the perimeter, even though not ideal for fine-grained access control, is still viable for coarse-grained protection of the systems that hold digital assets. Removing the necessity of fine-grained access control as a perimeter requirement allows perimeter design to be more focused on its primary role of protecting the availability of IT services. Similarly, while the primary role of platform security services is to provide a safe environment for the use of applications and their related data, their available data protection mechanisms are still viable as an additional protection layer.

Data protection systems – even if carried around by the data itself – need to operate in the real world which means being hosted by and interacting with other applications and operating system components. But that doesn't mean that they need to be trusted equally. Data owners should require strong authentication, utilizing trusted hardware and processes, from requesting subjects. The authentication trust roots should be independent of the applications and operating system. Using hardware-based trust roots for controlling the identity of entities accessing the data severely limits compromises that can come from infiltrating the

² For additional information, see his Internet site at: <http://movetheworld.wordpress.com>.

Protecting Information

applications and operating systems hosting the data protection systems. Ideally this would leverage a Trusted Platform Module (TPM), smart card, or other secure cryptographic device. While out of scope for this document, it should be remembered that these trusted authentication mechanisms require their own support, typically a functioning trusted PKI, and this is not trivial. This helps protect against attacks that compromise vulnerabilities present in large complex applications, operating systems, and platforms.

In order to enable e-business, data protection must be based on standards that are incorporated in diverse, competing products that members of the e-business community use. Enterprises that may use different products should be able to communicate data securely to their customers, partners, and suppliers. The use of proprietary protection mechanisms increases vendor lock-in, decreases interoperability, impedes e-commerce, and weakens security. The use of secure standard protocols allows enterprises with dissimilar IT infrastructures to share information securely.

While the properties of encryption worked well several decades ago, the modern business environment requires much more flexibility in the use of data. For example, businesses should be able to send a document to a supplier that can't be further distributed and that will self-destruct when the contract ends. In contrast, encryption is either on or off and once a document is decrypted there are no controls to further limit its use, duplication, or redistribution. Data owners should have tools at their disposal that allow policy-based encryption functions to be applied to data objects.

Where possible, there should be protection against data aggregation. This is difficult to enforce as data becomes more exposed the more it is used. Still there are some methods, such as the use of oracles, where data is used to generate answers to questions without revealing the data itself. As society and business becomes more data-centric, these and other techniques need to be developed and enhanced to help alleviate the growing risks.

These tools need to be administered securely. Ideally, there would be a separation of duties such that no one person can compromise the data protection systems. A typical two-person system would require one administrator to create roles and polices and a different administrator to enroll people in those roles. This is necessary to reduce insider threats as well as slowing down attackers from the outside. Administration should also be performed via trusted channels, and require multi-factor authentication to help prevent data loss (for example, losses may occur through the compromise of administration accounts). Finally, audit records need to be kept of all access as well as failed access attempts. The logs themselves need to be protected from unauthorized access. Unusual events should trigger real-time alerts.

Architectural Requirements

What follows is a description of the properties of an ideal data protection system. It is recognized that no such system exists today. It is also recognized that, in today's IT environment, many different systems need to work together to perform these functions. After describing the ideal system, this document will list current data protection services, near-term projects, and the roadmap for the next few years.

Data Types

Data protection systems should support both structured and unstructured data. Examples of structured data include drawings managed by computer-aided design systems (CATIA, DELMIA, ENOVIA, etc.), data held in Relational Database Management Systems (RDBMS), data managed by complex applications (email, ERP, SharePoint, WebEx, etc.), and data held in sophisticated storage management systems (EMC Documentum). Examples of unstructured data include office files (documents, presentation materials, spreadsheets, etc.), system logs, configuration files, generic text files, and other discrete digital data objects. Unstructured data files are commonly found on file shares, web servers, document repositories, and client platforms such as

Protecting Information

laptops and mobile devices. Data protection systems should also support streaming or network packet-based data such as voice, video, etc.

Data Usage

Data protection systems should be able to distinguish discrete activities or actions applied to the data and enforce permissions based on those actions. The system should also control the user's ability to create, modify, delete, read, infer from, distribute, convert formats, archive, manipulate, etc. Data protection systems must also perform fine-grained access control based on a rich set of attributes beyond the typical identity and action attributes. For example, environmental factors, such as location and time, should be considered.

Data User Types

While a detailed discussion of authentication is outside the scope of this strategy, data protection systems should support users who are fully authenticated, pseudonymous, or even anonymous. There should also be support for access by groups, roles, or other organizational structures.

Data Protection Attributes

While confidentiality is the most common data attribute addressed by data protection systems, other attributes such as integrity, location, access instances, and evidence of secure deletion should be supported.

Standardization

Standards are important for using data protection mechanisms with different organizations that have different products in their environment. They are also important to provide continued data access as technologies evolve and products – hardware and software – are replaced by newer ones. Standards should include the following:

- A standard container for encapsulating or protecting the information
- A standard programming interface for manipulating the protection around the data – essentially for unlocking it when valid access requests are received
- A standard protocol for communicating relevant data rights between data consumers and data owners
- A standard classification system for capturing the metadata necessary to process data access decisions including a simple extensible scheme and a standard system for labeling the metadata that can be read by Policy Decision Points (PDP) – while there are many different types of authorization architectures, the reference to PDPs does not advocate a specific architecture but refers to the requirement for the data protection system to make data access decisions

The leading candidate for both an authorization protocol and for making decisions based on metadata is the OASIS XACML. Another emerging candidate focused mostly on business-to-consumer transactions is the IETF OAUTH specification.

Summary of Existing Protection Capabilities

Just because we don't have the ideal set of tools to protect information, doesn't mean that it is completely unprotected. There are many solutions for protecting data, but they can be roughly categorized into three distinct layers or types which can be used independent of each other:

1. Protection controls may be directly applied to the data itself. The most common control in place today is some form of file encryption. Less common, but growing in popularity, are Digital Rights Management (DRM) services. These tools must evolve and become more sophisticated to meet business needs.
2. Protection may exist at the device, platform, application, or Operating System (OS) layer. These controls apply data protection in a local environment that is bounded by a specific piece of software or hardware. When the data leaves that environment, the protection is lost. Examples of this include whole disk encryption, access rights related to log-on credentials, PINs or passwords on mobile devices, and application or OS permissions.
3. Protection may also be done at the network or infrastructure layer at network "zone" or "perimeter" boundaries. As mentioned above, these controls are global, scalable, and relatively inexpensive. Protection at this layer is currently coarse-grained at best. Network protection mechanisms lose their advantages of scalability when they are configured to mediate access between specific individuals and discrete pieces of data; difficulties also emerge when trust in the network infrastructure is low (e.g., the data is crossing networks managed by others). For this reason, network protection mechanisms are usually configured to protect large amounts of data and resources. Once data crosses a protection perimeter – whether by intent or compromise – the data and other resources tend to be universally accessible.

There are several issues that limit these technologies when used for protecting data. While they are described in detail below, a summary list would include:

- Protection is too global and remote.
- Protection is neither granular nor interoperable enough.
- Protection is not integrated with centralized authorization services. While not all use-cases require centralized services, for those that do, these services are not generally available.
- Protection is often enforced by weak security services.

While each of these layers is capable of protecting data to some extent, the protection tends to meet more of the business and security requirements as it is moved closer to the data.

A rethinking of basic security indicates that the primary protection role of the network is availability of services or the protection of the enterprise from network events that would disrupt access. The primary role of devices, platforms, applications, etc. is to provide a safe place for data to be used or operated upon. While both networks and platforms can provide some data security, this is typically at a coarse-grained or bulk level and this should be secondary to the security functions just described.

This is illustrated in Table 1.

Protecting Information

Table 1: Appropriate Application of Security

	Network Security	Platform Security	Data Security
Primary Functions	Infrastructure availability and resilience	Operating environment integrity Application protection	Data confidentiality Data integrity
Secondary Functions	Data confidentiality	Data confidentiality	

Of course, this table – which focuses on technical security services – leaves out an important component of data security: the role of the users of these systems. Users have a direct impact on data security as they interact with data protection mechanisms. Although identity management is out of scope for this White Paper, users also have an indirect impact in data security via the confidence in their identity. A common path to gain unauthorized access to information is to subvert the identity of a person who already has legitimate access. See the Identity Commandments and supporting documents for more information on the role played by identity in access control.

Protection Issues with Current Technologies

Protection is too Global and Remote

Most data protection is enforced by network devices, operating systems, or applications. From an Open Group de-perimeterization viewpoint, the lack of cohesive control and interoperability is fundamentally flawed.

The primary goal of network devices, operating systems, applications, DBMS, etc. is not security, and their inclusion as part of the enforcement mechanism for protecting data potentially broadens the attack surface significantly. Not only does the part of these devices and services that deals with data protection have to be hardened against attack, but all of their functionality needs to be hardened as well. Often, bringing in one product requires adding another that it depends on, resulting in an expanded attack surface and an increased opportunity for vulnerabilities. Using the capability built into a DBMS or other data-centric application effectively “grandfathers” into the security model the weaknesses of the operating system on which they run. Adding a specialized security solution on top of this just adds to the complexity and opportunity for more vulnerabilities that can be exploited by the information thief.

The coverage of these various protection mechanisms is also often disparate, proprietary, and thus inconsistent. Imagine the situation where someone is using a document management system by connecting remotely to a server via a PC. The DMS will typically have some form of protection for data that it is storing and an access control mechanism. When it sends the data to the PC it will typically use some form of an encrypted tunnel to protect the data (and other un-encrypted traffic) while in transmission. The user might have a client on the PC that has some security, or more likely will rely on the PC OS to protect local copies of the data. Here we have three different protection models for the data: the client PC, the tunnel, and the DMS application. Data is vulnerable where it changes between the mechanisms – an especially egregious case being the common practice of decrypting tunnels at the perimeter, rather than the destination system (hopefully, but not usually, to check for viruses). It’s not just the transitions, but the difficulty of keeping the protections at the same level that make this fragmented security enforcement model dangerous.

Another issue with using global protection mechanisms is that they encompass too much data. Any breach of mechanisms that span large applications or protect perimeter information flows often results in large quantities of data being exposed.

Very global mechanisms, such as perimeter devices and operating systems, often have no context for adjusting protection to the level of sensitivity of the information. The result is protection that is too weak for sensitive information, or protection requiring excessive cost when all information is protected at the highest levels; thus negating the economic advantages of using risk management to allocate protection where it is appropriate.

Even specialized security systems, as they add functionality, require many more components as they add capability. If we take a look at a simple system involving the use of secret key algorithms for encryption, we generally have a couple of computers and their associated hardware, as shown below:

Protecting Information




Figure 3: Secret Key Encryption Trust Boundary

As is well known, secret key encryption has key distribution issues and is generally point-to-point, the result being lack of scalability and good protection for the secret keys involved. Public key cryptography using public-private key pairs offers more data protection services, scales better, and allows *ad hoc* communication. But it does this by adding more devices and services inside the trust boundary, making it easier to circumvent the encryption without actually breaking the algorithms themselves. A simple expansion to our model looks like this:



Figure 4: PKI Trust Boundary

PKI requires directories, certificate servers, registration authorities, and other components. Specialized data protection tools and services, such as DRM and DLP, add in even more devices and services:

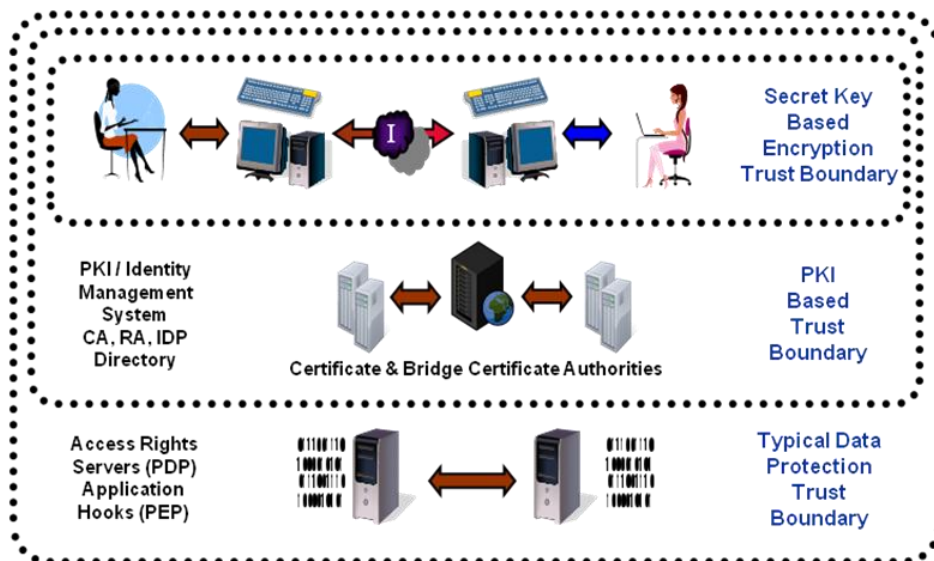


Figure 5: Typical DRM Trust Boundary

Protecting Information

The result is a system where the strong protection provided by the cryptography can be nullified by attacks against the other components.

Protection is Neither Granular nor Interoperable Enough

The only portable, interoperable technology that can be applied to data independent of its location or environment is some form of encryption. Encryption technologies and products have been around for some time and are fairly mature. A by-product of applying some encryption technologies is the ability to add a digital signature to attest to the provenance of the data.

However, encryption lacks persistence. It is either applied or not and the data being protected has no influence on its application. Once the data is decrypted, the content owner has no control over any future protection status, usage, or distribution of the information.

While encryption has served well and there continues to be many uses for using encryption, many of the new Internet-based uses for data do not match the protection model of encryption, for example:

- Global supply chains require that information be created and shared across multiple organizations, while also requiring restrictions on the distribution, timing of access, modification, and usage of that information.
- Privacy principles require that end users from different organizations or professions are limited in their ability to observe or modify different parts of the same data file.
- Governments also need to be able to apply different sensitivity settings to different parts of the same document.

To address some of these deficiencies, an emerging market has developed around protection technologies collected under the term Digital Rights Management (DRM). The original scope for DRM was to protect digital media from theft by copying and it was focused on fixed media forms, such as CDs and DVDs. The broader use of DRM technologies really requires a new term (some vendors are using the term e-DRM or Enterprise-DRM to distinguish corporate data protection capability from copy protection) as well as a more sophisticated open and standardized operation set. There are several proprietary products currently available from major IT and security vendors.

The major issue with these products is the lack of interoperability between them. At an organizational level, this limits their usefulness to the confines of private e-DRM eco-systems, such as a specific government, corporation, or other entity with the possible inclusion of other business partners who have selected the same product, and where Identity Federation can be achieved and supported by the e-DRM product. Generally, organizations have mature means (albeit often without cryptographic protection) of controlling data distribution and modification internally, but lack such controls for external use. The limited or nonexistent interoperability of e-DRM type products greatly limits their use in precisely the market where they would be most valuable.

Control is another issue, especially when data is being shared with or created and owned by multiple organizations. Many current information protection technologies are built on the assumption that only one entity will control access policies and enforcement mechanisms. Viable protection mechanisms need some form of negotiation between the entities involved that enables them each to adjust their risk or level of comfort appropriately. Typically these are encapsulated in a trust framework.

There are also implications on a personal level. The transition from traditional media to digital media for books, music, and video increases the need for additional protection of the digital content while at the same

Protecting Information

time conflicting with the consumer's need for portability among devices. Hardware (phones, e-book readers, digital players, etc.) is significantly more transitory and when the digital protection is tied to specific devices, firmware, file types, or operating systems, the efficacy and longevity of the content under protection is severely limited.

Rapid advances in technology, including big data analytics, are increasingly challenging the privacy of individuals and their personal data, and national privacy frameworks are often unable to evolve quickly enough to protect them, making individuals living through generational change – now moving to digital natives – more distrusting and discerning about their online activities and their personal data, especially when big data analytics are typically based on correlation, not causality. Data personalization, and how to protect it, is a further issue. Personalized data, by its nature, will differ from person to person at the principal's end, but will typically be aggregated at the organization end.

Protection is not Integrated with Centralized Authorization Services

While these services are not generally used there are some existing examples based on multi-level security solutions. These solutions are typically used in areas where high security is a requirement. The emerging MILS architecture from The Open Group aims to reduce the cost and complexity of these solutions.

As corporations automate their authentication and authorization systems, it becomes necessary to separate and centralize some of the components to provide flexibility and scalability. This requires applications and other programs that manage data to externalize their authentication and authorization requests. The authentication market space is – along with centralized reduced sign-on systems (typically using web-based interfaces) – being commonly deployed. Changes to the centralized system (such as replacing passwords with smart cards) automatically propagate to the protected applications.

Centralized authorization services are much less mature. In general, these require that the functions making the decision – often called a Policy Decision Point (PDP) – need to be separated from the enforcing functions – also called Policy Enforcement Points (PEPs). Typically, the PDP is a centralized application that ingests policy requirements and requests to operate on data and outputs a decision to the PEP. PEP functionality is either part of an application or deployed as a separate service in front of protected applications.

How does the PDP/PEP architecture relate to data protection? That data should be treated as any other resource – indeed the most valuable resource – in an environment, and protection of that data should also be driven by (at a minimum) the same conditions that drive other authorization decisions, and in a de-perimeterized world, an enhanced set of authorization decisions (or entitlement) based on not just user identity, but the relevant identity and attributes of all entities³ involved in the transaction chain. Current encryption and DRM-based technologies are PEPs or enforcement services, but they lack the connections to be driven from centralized PDPs or decision services and operate either as standalone services or as a feature of an application, such as email, document management, or a document office suite.

Protection is Enforced by Weak Security Services

It's a general principle that things that matter most should not be at the mercy of things that matter less. In the security world this can be restated such that weak things should not be used to protect stronger things. Typically, the biggest vulnerability to systems that use cryptography is to protect the keys with application or operating system components.

For enterprise use, current encryption and rights management products usually leverage centralized identity management systems which are based on centralized services and directories. These, in turn, are running on

³ The Jericho Forum Identity Commandments define five entity types: People, Devices, Organizations, Code, and Agents.

Protecting Information

servers that run operating systems. The attack surface has now been expanded from the cryptographic algorithms to a set of servers running applications. A data thief will bypass the cryptography and instead steal the identity of someone with access to the data by leveraging vulnerabilities in the identity management system or the platforms on which it resides.

The additional security systems, designed to protect information by trusting additional servers and applications, have actually increased the ability for a thief to steal the data.

Context

Before describing the components of a data protection system, it is important to provide some background on access control, the differences between access control and related security technologies, and the architectural options for deploying access controls. We should also remember that access control in general is composed of several different components (Account Provisioning, User Enrollment, Identification, Authentication, and Authorization, etc.) as well as supporting services (Directory, PKI Services, etc.). In order to keep this document simple, it is assumed that the user desiring data access has already been identified and authenticated.

At its simplest, access control is just the insertion of a guard between the user and the desired resource that mediates specific actions. Physical access control typically uses a variety of mechanisms for the guard including actual human guards, animals (dogs, chickens, etc.), isolation structures (fences, walls, gates, etc.), and their accompanying enforcement mechanisms (locks, guns, etc.). This is a crude form of access control as it often doesn't distinguish between identity verification (authentication) and permitted access (authorization). Possession of the appropriate key or identification is enough to pass the guard and, once past, all actions are permitted. To refine access control there may be additional guards or isolation structures (rooms with different keys, for example) providing a form of defense in depth. There may also be controls on what actions are permitted; for example, a user may be permitted to see some resource but not allowed take photographs of it.

The advent of electronics and IT has created significant new issues that require additional matching controls. These controls need to react at the speed and granularity of IT transactions.

The data protection equivalent to the above scenario would be a simple access control structure embedded in either the operating system or an application. The UNIX operating system offers an example of this by distinguishing between individuals and groups, and allowing discrete actions (read, write, execute, and navigate) on data or program execution. Typically a user makes a data access control decision by setting a group of permission bits associated with a file and the operating system enforces those permissions when access is requested.




Figure 6: A Simple Access Control Model

It turns out that, for a variety of reasons, it's better to separate the enforcement functions from the decision functions that the guard performs. Typically decisions are managed centrally and enforcement is done closer to the resource, allowing for better scalability and flexibility when designing systems. In the physical world, this is similar to a guard determining access by observing a credential that was issued by a separate, trusted authority. Separating these two roles also makes it harder to compromise the access control system.

Enforcement mechanisms are typically high performance and designed for volume, while decision functions are more intelligent but require much less bandwidth since they just respond to requests from the

Protecting Information

enforcement functions. This basic structure – logically separating enforcement from decision – was developed decades ago (ISO 7498-2:1989, ISO/IEC 10181-3, and XDSF) and remains the basis for nearly all authorization systems.

In a Discretionary Access Control (DAC) system, the decision function is often performed by a person, typically the person who either owns the information or who creates it. Whether it's an individual making a decision to sign or encrypt an email, or using a Risk Management System (RMS) to control a document's distribution, or setting the group permission on a UNIX directory, the person makes the decision and the operating system or application enforces that decision.

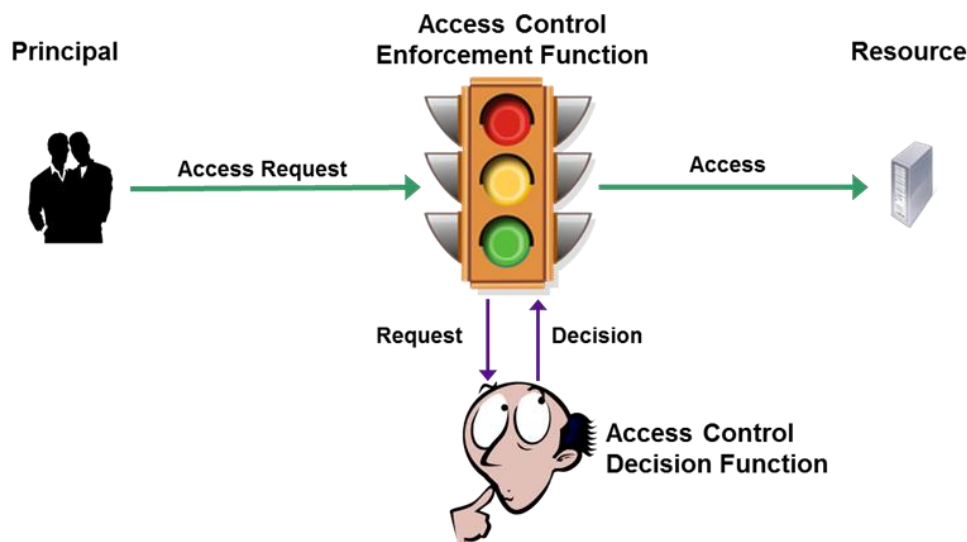


Figure 7: A Discretionary Access Control Model

For a variety of reasons, humans are not very good access control decision-makers. They are inconsistent (for example, the amount of data they request from users when making decisions varies enough to provide a lucrative social engineering market), they are relatively easy to compromise, and they make mistakes when entering data. Access control decisions must be made during a workflow that is unrelated to security and by individuals whose primary job function is not information security.

In a Mandatory Access Control (MAC) system, these decisions are made by the security system itself. In a MAC access control system, people create a set of machine-readable access control policies, and then the system makes individual access control decisions based on those policies. MAC-based access control systems have a lot of potential advantages:

- They take humans out of the day-to-day decision-making process, thereby improving accuracy and reliability. A MAC system will make the same decision given the same inputs while humans may vary, and if a human makes a mistake there is an opportunity for a breach.
- They scale both in complexity and volume.
- They are harder to compromise. The policy administrator and the user are two different roles. This means that collusion is required if an insider is going to violate access constraints.
- They react faster to policy changes or new attacks.
- They are more adaptable to making authorization decisions between enterprises.

Protecting Information

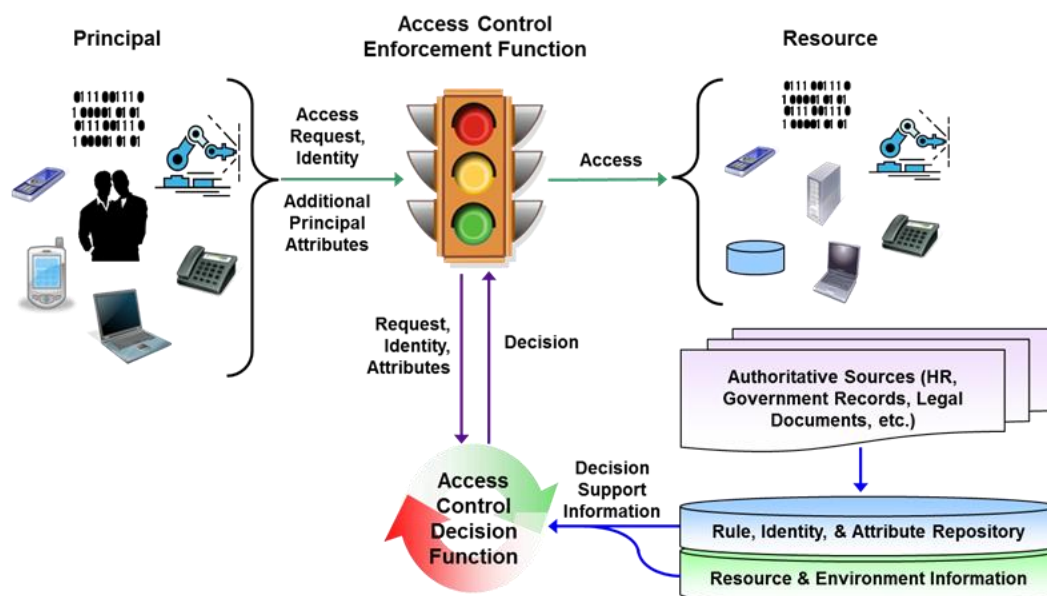


Figure 8: A Mandatory Access Control Model

MAC access control systems have a number of required components – some of which are informally present in DAC access control systems – but all of which need to be explicitly specified in a MAC system. MAC systems are currently rare as it's only been recently that technology has become available making affordable deployment feasible. These are usually made of separate components and capabilities that are embedded in products obtained from different vendors. To ensure interoperability between components from different vendors, these components must be connected using standard protocols. This is particularly important when access control systems are used between enterprises that use different IT vendors.

Here is a brief description of the major components of a MAC access control system:

- **Principal:** We have been using the word “user” to describe the entity needing access to the data. We should remember that much data access is also being done by applications, smart devices, etc. The word “principal” is a generally accepted industry term to describe any entity whose identity can be authenticated – and therefore can legitimately access resources. (Other industry synonyms include “actor” and “initiator.”) The access control system uses information about the principal when making access control decisions. This information, usually referred to as a set of attributes, is either retrieved from storage or presented with the request.
- **Resource:** While the focus of this White Paper is on data security, this is a general authorization model that can be applied to any resource including computers, industrial control systems, telephones, mobile devices, etc. as well as data. The access control system needs information, usually referred to as metadata, about the data (or other resources) in order to make the correct decision. Metadata can be embedded in the data itself or linked to the data and retrieved from storage via that link.
- **Enforcement Function:** This is sometimes called a Policy Enforcement Point (PEP). Enforcement mechanisms differ depending on the layer they operate at (firewalls at the network layer, permission bits at the OS layer, encryption at the data layer, etc.). This is another reason for separating access decisions from access enforcement. When we are looking at mechanisms for directly protecting data that satisfies the criteria listed above, the only practical solutions use some form of encryption to enforce confidentiality. The data is embedded in an encryption controlled software container that can understand

Protecting Information

the access control decisions. The enforcement function intercepts an access request from a user and then forwards that request to the decision function. If it gets a permit decision back, it then opens the gate and allows the user to access the resource.

- **Decision Function:** This is sometimes called a Policy Decision Point (PDP). The decision function is the brains of the access control system. When it gets a request from an enforcement function, it evaluates the request against a set of rules (or policy), metadata about the request object, user and environment, and information about the nature of the request (read, write, delete, etc.). It sends the resulting decision back to the enforcement function. A permission decision may also include an “obligation”. Obligations in an access control system consist of additional actions that are required; e.g., if X happens you must do Y. Obligations are often used to implement policy that is difficult to express in access control rules. The decision function needs to be able to consume a combination of stored metadata as well as metadata that might come in with the access request. Ideally, all resources would use a common access control service to make decisions at the device, network, application, operating system, and data layers. This provides consistency and blocks attacks that evade strong protections at one layer by leveraging weaker protections at other layers. Enforcement and request mechanisms will need access to the PDP. For enterprise-to-enterprise authorization capability, this means that the PDP will need to be accessible from outside the enterprise. For organizations that have common policies, cloud-based PDPs might be an option.
- **Administration [not shown on diagram]:** This is usually called a Policy Administration Point (PAP). While it is desired that a MAC access control system run automatically, at some point it needs configuration. Typically an information security organization would design and assemble the components as an enterprise service (much like a PKI or directory service) but the organizations that use the service would set the policy. The PAP is an administration function (typically a human-accessible GUI that talks to the decision function) that allows a person to enter a set of access policies and then translates those into machine language that the decision function understands. Large organizations often have enterprise-wide (export, privacy, intellectual property, etc.) as well as local (division, program, or project-specific) policies that need to be instantiated. An enterprise that is part of or has a large supply chain may also have policies that are set by customers, or rules based on organizational memberships.
- **Metadata:** As alluded to above, the decision function needs information to make the access control decisions. This information can be roughly grouped into four categories:
 - Information attributes about the resource – in this case metadata about the data
 - Information attributes about the requestor (the user or principal)
 - Information attributes about the desired action(s) on the resource
 - Relevant environmental information attributes (time of day, physical location, etc.)

While much of this data may be stored locally (typically accessible via an LDAP directory service), some of the metadata – especially that which is dynamic – will come in with decision requests from PEPs to PDPs. The decision service needs to be able to process both stored (static) and dynamic metadata. The differences are illustrated in Figure 9.

Protecting Information

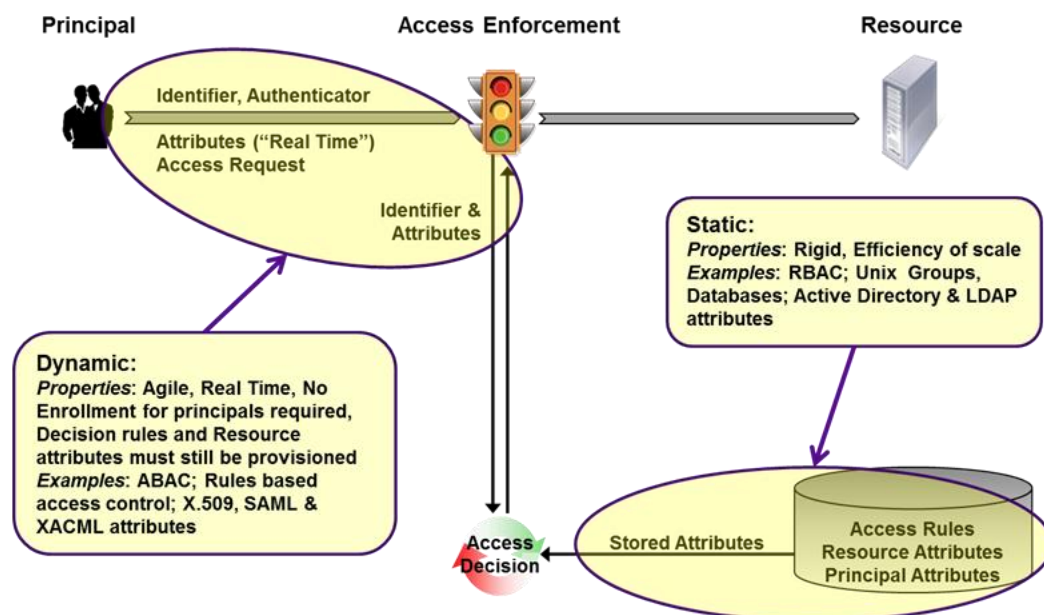


Figure 9: Static versus Dynamic Metadata

- Protocols and Standards:** A set of standard protocols are required to facilitate communication between the access control components. The standard protocol for initiating and evaluating access control requests (including obligation support) is the OASIS eXtensible Access Control Markup Language (XACML) protocol. The most common protocol for obtaining metadata information from stored data is the Lightweight Directory Access Protocol (LDAP). Other relevant protocols include OASIS Security Assertion Markup Language (SAML) for federated authentication and the Trusted Computing Group Interface for Metadata Access Points (IF-MAP) for network-layer access controls. Implementing standard protocols frees an enterprise from dependencies on specific products and vendors and facilitates enterprise-to-enterprise access control capability. Note that SAML and XACML are implemented as XML tags and messages that incorporate protocol-like request and response functions, but they still need a network protocol to transport the messages.

In summary, the right strategy is to implement an architecture that:

- Replaces DAC with MAC
- Separates the access decisions from the enforcement functions
- Connects the components with standard protocols which allows the transition to data-centric information security

It provides the structures that allow data protection to be applied to the data itself, independent of the environment. A mature access control ecosystem also allows for fine-grained data access control between enterprises, customers, and suppliers. This can help coordinate policies, protection profiles, audit information, etc. across a supply chain.

This architecture is really the confluence of two distinct, but necessary changes. First the shift to dynamic, policy-driven access control, and second the movement of data protection enforcement closer to the data itself. This can be illustrated with the following chart. The vertical axis shows a progression from the network layer to the application layer and then to the data itself. The horizontal axis shows a progression from either

Protecting Information

static or human-based access decisions to dynamic, policy-based access decisions. Note that, as was previously stated, although the system makes the individual access control decisions in the latter case, people still decide the policies that control those decisions. Policy-based access control decisions are just an automation or extension of business requirements. Some representative technologies and services are placed in their approximate positions on the chart as examples.

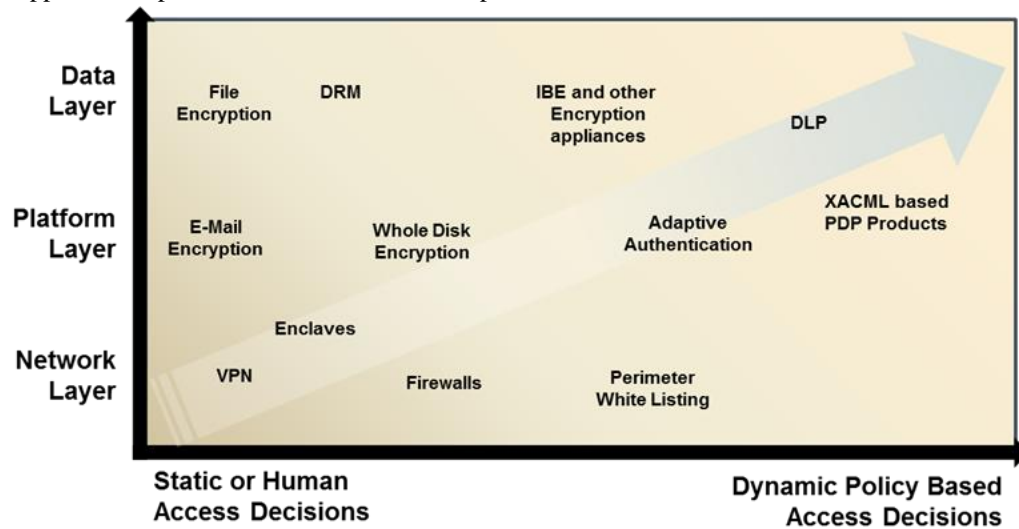


Figure 10: The Data Protection Shift

Challenges

We have made little progress in data-centric protection since cryptography went mainstream in the late 1990s. First, we will take a look at some of the general challenges or reasons for this lack of progress, and then we will look at each of the architectural components separately. Then we can examine strategies for moving forward.

Inertia

Most organizations are responding to data loss by incrementally patching or updating existing security mechanisms. However, these mechanisms were designed before the requirements imposed by extensive e-business appeared and in an era when threats were relatively simple. The trend towards state-sponsored espionage has awakened some organizations to the reality that data protection must improve, and there is now growing interest among some government agencies and more forward-thinking companies.

Vendor Recalcitrance

There are several reasons why IT and IT security vendors have been reluctant to pursue this path.

First, it is both easier and cheaper to continue to sell existing capability as long as possible, and when changes are required, it is also easier and cheaper to slowly evolve them. This has the added benefit of causing less disruption on the customer side. The IT industry is almost unique in that it went from theory to broad acceptance in less than a generation. As a result there are many practices and technologies that have become frozen in immature states. When this is coupled with the current business philosophy of short-term gains, the effort to move forward is just too great. Other examples of this thinking include the slow adoption of security technologies such as PKI, DNSSEC, secure routing protocols, and the revised network protocol, IPv6.

Monolithic solutions are also easier to sell and deploy. The ISO/IEC 10181-3 model requires a market for policy decision engines that separates them from the enforcement capability. It requires that devices and applications that enforce access control externalize their authorization decisions using standard protocols. It's difficult to launch a business making PDPs when there are no applications that will leverage those PDPs and it's difficult to make the business case for externalizing authorization decisions when not all of your customers have bought the PDPs to respond to those decision requests. Nevertheless, the market for vendors that make separate policy decision engines is growing. Among the first wave of vendors were Jericho Systems (no relation to the Jericho Forum), Securent (acquired by Cisco), and Identity Engines; more recent are Axiomatics and NextLabs.

Application Data Binding

In many large or complex applications, the data is so tied to a specific application that its uses, including access, are completely controlled by the application. Without cooperation from the application it is difficult to apply standards-based, interoperable protection to data. Ideally, both data and applications would exist separately in an operating system and the application would make an Application Program Interface (API) call or request to a common set of security services for authentication and authorization. While this is becoming more common, even some modern operating systems (e.g., Apple's IOS) have no concept of user-accessible files independent of their supporting applications.

This issue, similar to the general authorization issue above, is also largely dependent on vendors providing the required configuration options and APIs or protocols for the applications to call PDP/PEP base security

Protecting Information

services. Vendor reluctance to do so is understandable because they will need to continue to include internal security capability as long as they have customers who do not have centralized or enterprise-level authentication and authorization services.

Identity and Access Control Standards

The lack of good quality, pervasive identity standards and common access control mechanisms provides another challenge for protecting information. These components are required to implement the controls described above. And they need to be easy to use and work across different products in order to gain market acceptance.

Metadata

Metadata is required in order for automated access decisions to be made. Humans also use metadata, but are not taught to recognize it as such. Typically a person gets the clues necessary to make an access decision for physical documents from visual markings on the data, the way it is stored, or by asking someone. Many of these clues are also carried over into digital representations of data, often buried in file or directory names. A policy decision engine needs a deterministic way of finding this information and the solution is to attach this metadata either directly to the data itself or by a link to the file.

The most limiting factor is the creation and management of this metadata. Ideally it would be done without human intervention using information found in the data, information about the principal (which could be drawn from attributes stored in directories), and information about the environment (including programs or projects) which could be drawn from the operating system or applications. There should also be an override capability. One of the most prevalent but little known uses of metadata is the set of custom properties that Microsoft® Office inserts into office documents. Another is the EXchangable Image File Format (EXIF) data that digital cameras embed into photographs. While neither of these is commonly used for security purposes, they indicate that it is possible to automatically generate this data in both software and hardware. A further example is development of standard profiles of metadata for the domains of US export control and intellectual property control at the OASIS XACML committee.⁴

The Human PDP

There is a resistance to automating the policy decision function. While much of this is due to the aforementioned reasons, some of it is a cultural shift and a perception of losing control. Policy creation needs to be distinct from policy decision. Once that is understood, then it's easier to accept a human role as a policy. There is also the concept of the "Right to be Forgotten" in the form of expiration dates on personal information advocated by the Oxford Internet Institute's Viktor Mayer-Schonberger analogy of the world remembering forever, and applying the wrong data in the wrong circumstance.

Over-Connectedness

Being connected has made us more efficient, but there is now the risk of reacting so quickly that we don't give the thought we might have given to data protection, even ten years ago.⁵

⁴ For details, see <http://docs.oasis-open.org/xacml/3.0/ec-us/v1.0/cs02/xacml-3.0-ec-us-v1.0-cs02.html> and <http://docs.oasis-open.org/xacml/3.0/ipc/v1.0/cs02/xacml-3.0-ipc-v1.0-cs02-en.html>.

⁵ For details, see the book OVERconnected.

Protecting Information

Internet of Things (IoT)

The Internet of Things (IoT) is now entering the public consciousness – that vast amounts of data will be machine-to-machine. According to ABI Research more than 30 billion devices will be wirelessly connected to the Internet of Things (Internet of Everything) by 2020.⁶

Unauthorized Data Mining

There is increasing evidence of organizations, both commercial and government, collecting vast amounts of electronic data (both voice and Internet traffic) through programs such as PRISM.⁷

Anomaly Detection At Multiple Scales (ADAMS)

The Anomaly Detection at Multiple Scales (ADAMS) program creates, adapts, and applies technology to anomaly characterization and detection in massive data sets. Anomalies in data cue the collection of additional, actionable information in a wide variety of real-world contexts. The initial application domain is insider threat detection in which malevolent (or possibly inadvertent) actions by a trusted individual are detected against a background of everyday network activity.⁸

Data Contextualization

There is further evidence of data in data warehouses, despite being redacted and aggregated, still providing enough information to enable a person to reasonably guess the source of the data. There need to be appropriate data protection enforcement controls close to the data to protect against this – this may involve de-identification by the following measures:

- Reasonable measures to de-identify
- Company does not try to re-identify
- Entity that has done those two things prohibits downstream entities from re-identifying PI data

Or – as Paul Ohm suggested in his famous paper, to introduce a degree of difficulty in re-identification – so someone can't easily reverse-engineer the context, to determine the source.⁹

⁶ For details, see www.abiresearch.com/press/more-than-30-billion-devices-will-wirelessly-conneUnauthorized Data Mining.

⁷ For details, see www.smh.com.au/it-pro/security-it/australia-collecting-data-for-nsa-leaks-show-20131015-hv24k.html.

⁸ For details, see [www.darpa.mil/Our_Work/I2O/Programs/Anomaly_Detection_at_Multiple_Scales_\(ADAMS\).aspx](http://www.darpa.mil/Our_Work/I2O/Programs/Anomaly_Detection_at_Multiple_Scales_(ADAMS).aspx).

⁹ For details, see http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1450006.

A Data Protection Future

Data protection evolution can be broken into six stages. The first two stages essentially provide *ad hoc* data protection. The protection varies with the environment or location of the data and it is up to those entities that create and manipulate the data to apply protection.

The second two stages represent a shift towards applying Mandatory Access Control (MAC) to enforce data protection based on properties of the data, and access request types and entitlements granted to the user of the data. The discussion of DAC *versus* MAC in the previous sections sets the context for these two stages.

The final two phases add intelligence to the data giving it some measure of ability to participate in the access control decisions. Data enhanced in this way is often called smart data, but smart data can include other intelligence uses that are not necessarily security-related. The Open Group has published a separate paper on Smart Data. It defines smart data as: “data that remains appropriately protected when outside an entity’s direct locus of control (Jericho Forum Commandment #9)”, and notes that “it is not sufficient to only establish data as ‘smart’ – we also need mechanisms to verify whether that data remains “smart”, and also to enable smart data to promote awareness that it is “smart”. Put simply, it is data that is enabled to look after itself. In agency terms, smart data becomes its own agent and is able to make risk-based decisions about its own access.

While all data protection mechanisms expose the data to its environment when the data is accessed, the more sophisticated they are the more they need confirmation of the state of the environment before allowing access. *Ad hoc* data protection already relies on the environment and MAC-based policy access control can consider environment security as one of the attributes that the PDP takes into account when determining access, but smart data is more isolated and will require communication with a trusted entity to determine how much to trust the environment. These requirements and considerations are discussed in the Smart Data for Secure Business Collaboration White Paper.

These stages are evolutionary; that is, it is difficult to implement any given stage without building on the stages to the left of that stage. They are illustrated in Figure 11.

Protecting Information

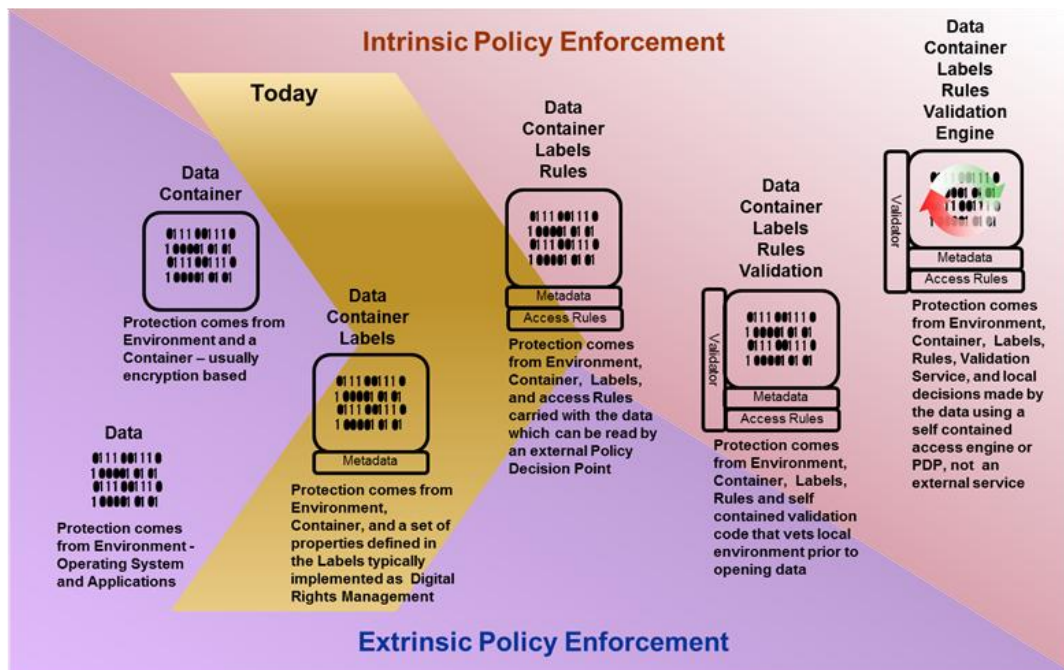


Figure 11: Data Protection Evolution

Data

Data itself is not directly protected – it is in the clear. Protection only comes from whatever the environment provides, typically from firewalls, encrypted tunneling protocols, etc. at the network layer or controls built into operating systems and applications. The protection varies depending on where the data is in an environment and what application is using it at any given time. Unfortunately this is true for the majority of data today.

Data + Container

The Container adds a protection boundary to the data. When protection is added to data, it is typically some form of encryption or encryption enhanced by digital rights technologies. This provides a protection layer that travels with the data. Encryption systems are often integrated with other non-security systems that are used to manage data.

Data + Container + Labels

The addition of metadata or labels, containing data attributes, gives an application that makes data access decisions some of the information necessary to grant specific accesses to the data. Rights management systems, although not standardized, represent the most common products of this type. Other, advanced data access control systems use a combination of labels and encryption to provide granular data access. Labels are necessary for any Mandatory Access Control (MAC) system.

Data + Container + Labels + Rules

A future data protection system might include the access rules as part of the metadata. Protected data would carry around its attributes and a set of access rules for interrogation by a local PDP. The rules might tell the

Protecting Information

PDP additional requirements for access, such as nationality, project relevance, physical location, etc. for the PDP to interpret. These rules would be used in conjunction with any that are already part of the PDP to enhance the decision process.

Data + Container + Labels + Rules + Validation

The data protection system could include, alongside the data, some code that would validate the safety and integrity of the environment before opening up or allowing access to the protected data. This could be done by validating a signature that originated with an application that had analyzed the environment, it could use a status stored by a machine health validation in a TPM, or other similar services.

Data + Container + Labels + Rules + Validation + Engine

A final stage would be for the data to carry a small PDP around with it. Ideally, this attached PDP would be small enough to not add too much overhead to the data and also to undergo formal testing for proof of correctness. It could be an add-on or plug-in where needed.

When access is desired, the PDP would first load into the host environment and then check the environment for safety before starting to process access requests. The PDP could work with other local PDPs to arrive at joint access decisions or independently if the integrity of the environment was suspect. Adding a small piece of executable code to the data will likely leverage virtualization technologies to create an enclave in a foreign host environment for the duration of the data access. OASIS XACML, already mentioned, is the only real contender for constructing a portable rules engine.

Strategies for Achievement

While the strategies below describe the components that are typically necessary for making complex authorization decisions used between enterprises, it is recognized that there is also a need for authorization models that protect data between individuals or individuals and organizations. Peer-to-peer and even business or government-to-consumer use-cases can often be supported by simpler authentication and authorization models. However, at a minimum these systems will also require a means of making data access decisions, protocols for communicating access requests and the resulting decisions, and the capability to enforce those decisions. The descriptions of PDPs and PEPs are also meant to include these capabilities in the more general sense. Even though these requirements are described separately from an architectural perspective, actual designs and the resulting solutions may combine them. While the protocols specified below are generally aimed at businesses, they can also provide a foundation for protecting data access at an individual level. In addition, newer protocols such as Open ID and OAUTH are specifically designed (but not limited to) protecting data in cloud and mobile environments. The intent, followed in this White Paper, of describing the more complex enterprise-to-enterprise requirements is that the simpler use-cases would also be supported as a subset; but there is no intent to limit other alternatives focused on these use-cases.

The first two strategies are general, looking at evangelizing and testing these new ways of protecting information. The four that follow are focused on specific components of the generalized data-centric security architecture.

Strategy 1: Industry Engagement

Convey a data-centric protection vision to users, customers, and suppliers. The benefits of data-centric security and the risks of maintaining the *status quo* need to be articulated. Vendors need to have a business case to develop the necessary tools. This can be accomplished by authoring and distributing position papers and communicating via print and Internet media. Presentation and talks at industry events or interviews are another channel to disseminate information. Communication and relationship with vendors and industry associations should also be pursued. This approach relies heavily on standards, which means that engagement with the appropriate standards bodies will be necessary.

Strategy 2: Proof-of-Concept

Implement services based on or evolvable to data-centric protection. The shift to data-centric security is a significant change to the typical IT environment and also changes procedures and process workflows to some extent. Given the reliance on IT services in most organizations, it will be necessary to take an evolutionary approach. Initially this means focusing on a narrow business application area with sensitive information. With maturity, the initial effort can be broadened. Significant value will be achieved when data-centric security can be extended externally across an organization's supply chain.

Strategy 3: Policy Decision Engines

Encourage the Commercial Off-The-Shelf (COTS) Policy Decision Point (PDP) market to develop products and deploy these PDPs internally. The policy decision capability (PDP) is the heart of data-centric security. In the past this decision capability has either been manual or typically buried in specific products. In order to have the necessary consistency, general-purpose PDPs need to be developed and the commercial market for these needs to be supported. Over the last ten years several PDP vendors have come and gone, but the market is slowly growing and there are several suppliers to choose from. Initial deployment could start with a

Protecting Information

commercial PDP vendor product that is used to govern enforcement services under the control of the organization. Initial experiences with these products will provide valuable information for improving them to better match enterprise security needs and business requirements.

Strategy 4: Policy Enforcement

Encourage Commercial Off-The-Shelf (COTS) Policy Enforcement Points (PEP) and applications to externalize access decisions. As has been explained, policy enforcement is typically embedded in COTS applications and COTS vendors will need to continue to support for customers that do not have separate PDPs. However, these vendors should also be encouraged to make available, either via open APIs or standard protocols, interfaces to the protection enforcement controls. This is typically referred to as “externalizing” the security access control decisions.

Note: Strategies 3 and 4 capture the dependencies of the PDP and PEP on each other and the realization that these will have to be an iterative process that requires the cooperation of vendors. PDP vendors will likely be in the security business, but most PEP capability is today included in products (database, CAD, and other structured data manipulation software) whose primary function is not security but rather the creation and management of information.

Strategy 5: Connecting Protocols

Support the development of and use internally LDAP, SAML, and XACML. This strategy has largely been achieved with LDAP (directory access) and SAML (authentication), but the benefits of using XACML (access control) are still largely untapped. More experience with these functioning in an integrated environment will aid in their evolution and relevance to solving real-world business problems.

Strategy 6: Automate Data Categorization

(And the generation and use of metadata for making policy-based data access decisions.) The generation and management of the metadata necessary to deploy MAC needs to be automated as much as possible. While some capability already exists in most applications that manage structured information, this needs to be expanded and tools need to be developed to manage metadata for unstructured information. Common labelling standards need to be developed and evolved to match business requirements. There is a relatively new labelling standard from OASIS that can be leveraged.

Strategy 7: Develop Smart Data and its Ecosystem

Moving from DAC to MAC and shifting the enforcement closer to the data is a complex process, but it’s just a beginning. As the concepts that govern the composition and definition of smart data develop, it’s important to evolve the corresponding data protection capabilities. This will include the movement of some of the access decision capability into the data itself. That capability will then interact with data-centric protection mechanisms such as encryption to make access control decisions independent of physical location. Sensor capabilities will need to be developed to ensure that the data is only accessible in trusted environments. The qualifications and ability to interrogate these environments and receive trusted attestations of safety will need to be developed. Likely, this will require leveraging independent third-party environment verification. Standards from the Trusted Computing Group, such as the Trusted Platform Module (TPM) and device health capabilities, are a good base to build on.

Protecting Information

Strategy 8: Transparency around Data Breaches when there is a Chance of Customer Harm

In order to maintain reputation in an increasingly digital world, organizations should be transparent when there is a chance of customer harm. This will involve development of data protection models and patterns, privacy maturity models, and voluntarily notifying customers if there is a data breach. The alternative is to be regulated and have to comply with mandatory data breach notification legislation.

Summary

The rapid growth in data volume, complexity of usage, and business criticality of information requires that new approaches need to be developed for protecting information. These approaches need to be risk-based and business-focused. They also need to be flexible and yet provide consistent protection commensurate with the data value and threat potential. This necessitates both the movement of protection enforcement closer to the data and a shift from user-applied to policy-driven protection models. This must also lay the foundation for the development of smart data.

Appendix A: Notes on Terminology

General Comment

As the audience for this document varies, I've tended to use common English descriptions for technical concepts and only introduce specific technical language where necessary.

Data *versus* Information

There exist several hierarchies for dealing with content. Russell Ackoff classifies mental content into five layers:

- Data, which can be thought of as a set of symbols or primitives
- Information, or data that has been organized and answers the who, what, where, and when questions
- Knowledge, which applies data and information to answer the how questions
- Understanding, which deals with the question of why
- Wisdom, which is the result of applying and evaluating understanding

Ackoff also makes the point that the first four categories deal with the past or what is known, but wisdom deals with the creation of the future. IT systems hold data or information. I have used the terms more generally than Ackoff in this White Paper to represent both facts and the analysis or IT applications of those facts.¹⁰

User, Subject, Principal, et al

ISO/IEC 10181-3 defines a principal as: “an entity whose identity can be authenticated.” And an entity can be a human, device, or piece of software. I have generally used the term “user” or “end user” when referring to a person desiring access, and the broader term “principal” when it was necessary to emphasize that access control must take non-human entities into account. The OASIS XACML standard uses the term “subject”. “Actor” and “initiator” are also used in some texts.

Policy Decisions

The document begins by describing the automation of policy decisions for protecting resources, including data. Gradually the term Policy Decision Point (PDP) is introduced. PDPs were introduced in the IETF Authentication, Authorization, and Accounting (AAA) Working Group documents and have generally replaced the older, but more accurate, ISO term Access Control Decision Function (ACDF). Policy decisions need to be made by a highly available service so the “point” in PDP is somewhat misleading; however, the term PDP is widely accepted today.

¹⁰ For Ackoff, see: www.systems-thinking.org/dikw/dikw.htm.

Protecting Information

Policy Enforcement

Similar to the treatment of the PDP, the Policy Enforcement Point (PEP) is introduced gradually, and likewise the term PEP has generally supplanted the ISO term Access Control Enforcement Function (ACEF). PEPs can be standalone boxes – a network firewall, for example, or security capability embedded in an application or operating system.

Acknowledgements

Lead author:

- Stephen T. Whitlock, Technical Fellow, Chief Strategist – Information Security, The Boeing Company, and founding Board Member of the Jericho Forum

The author wishes to acknowledge the significant contribution to the approach in this White Paper from:

- Adrian Seccombe, Research Associate, Leading Edge Forum
- John Sherwood, SABSA Institute, and Independent Consultant
- Paul Simmonds, Jericho Forum founding member
- Clark Thomborson, Jericho Forum member
- Shane Tully, Jericho Forum founding member
- Andrew Yeomans, Jericho Forum founding member
- Ian Dobson, Director Security Forum, The Open Group

References

- A Guide to Claims-Based Identity and Access Control, Microsoft Corporation, 2011.
- Broken Promises of Privacy: Responding to the Surprising Failure of Anonymization, Paul Ohm, UCLA Law Review, Vol. 57, p.1701 (2010), University of Colorado Law Legal Studies Research Paper No. 9-12 (August 2009); available at SSRN: <http://ssrn.com/abstract=1450006>.
- Computer Communications Security: Principles, Standard Protocols, and Techniques, Warwick Ford, 1993.
- Computer Security, John M. Carroll, 1995.
- CORBA Security Service Specification 1.0, Object Management Group, 1996.
- Distributed Security Framework (XDSF), Open Group Guide (G410), December 1994, published by The Open Group; refer to: www.opengroup.org/bookstore/catalog/g410.htm.
- Economics and Strategies of Data Security, Daniel E. Geer, 2010.
- Evolution of Information Security Technologies: Dan Hitchcock's Blog, posted January 16, 2008; available from: <http://movetheworld.wordpress.com>.
- Entitlement Management: The Next Security Wave, Linda Musthaler, Network World, March 12, 2007.
- ID Management: A Matter of Entitlement, Kelly Jackson Higgins, 2007.
- ISO 7498-2:1989: Information Processing Systems – Open Systems Interconnection – Basic Reference Model: Part 2: Security Architecture.
- ISO/IEC 10181-3:1996: Information Technology – Open Systems Interconnection – Security Frameworks for Open Systems – Part 3: Access Control Framework.
- Jericho Forum[®] Identity Commandments, Version 1.0, White Paper (W125), May 2011, published by The Open Group; refer to www.opengroup.org/bookstore/catalog/w125.htm.
- Jericho Forum[®] Paper: Collaboration Oriented Architectures (COA) – Secure Data: available from the Jericho Forum website at www.opengroup.org/jericho/publications.htm.
- MILS: Multiple Independent Levels of Security, MILS Presentation, Washington DC, April 2006 (available at www.opengroup.org/bookstore/catalog/r061.htm); and MILS Reference Material, Washington DC, April 2006 (available at www.opengroup.org/bookstore/catalog/r062.htm).
- OASIS eXtensible Access Control Markup Language (XACML); refer to: www.oasis-open.org.
- OVERConnected: The Promise and Threat of the Internet, William Davidow, 2012.
- Secure Computing: Threats and Safeguards, Rita C. Summers, 2000.
- Smart Data for Secure Business Collaboration, White Paper (W140), January 2014, published by The Open Group; refer to: www.opengroup.org/bookstore/catalog/w140.htm.

About the Security Forum

The Open Group Security Forum is a membership group of security experts from both the customer and supply sides of industry, government, and academia, who share a common aim to raise confidence levels in IT business operations. It was formed in 1991, shortly after the formation of X/Open, and continued through the merger of X/Open with the Open Software Foundation in 1996 as the Security Forum of The Open Group. Over that 22-year period, information technology has evolved rapidly through ever more powerful Personal Computing and Portable Devices to become a critical dependency to all business enterprises and also to a global population of online individuals intent on consumerization and social networking. This global explosion in computing power and usage continues to present huge security challenges over protecting IT systems from security breaches resulting in loss while also enabling high rates of availability over ever-increasing online processing speeds. The Security Forum has been and remains in the forefront of delivering effective security solutions, working alongside other security groups to deliver its open standards, guides, and other papers, all of which are available in The Open Group Online Bookstore (www.opengroup.org/bookstore/catalog/se.htm).

The Security Forum's current focus is on Security Management and Architecting for Security, including in vertical industry sectors. Over the past 10 years it has also been closely associated with the Jericho Forum on the challenges of de-perimeterization. The Jericho Forum sunset at the end of October 2013, and passed its legacy of publications and achievements to the Security Forum for safe-keeping and maintenance.

For more information, please visit: www.opengroup.org/getinvolved/forums/security.

About The Open Group

The Open Group is a global consortium that enables the achievement of business objectives through IT standards. With more than 400 member organizations, The Open Group has a diverse membership that spans all sectors of the IT community – customers, systems and solutions suppliers, tool vendors, integrators, and consultants, as well as academics and researchers – to:

- Capture, understand, and address current and emerging requirements, and establish policies and share best practices
- Facilitate interoperability, develop consensus, and evolve and integrate specifications and open source technologies
- Offer a comprehensive set of services to enhance the operational efficiency of consortia
- Operate the industry's premier certification service

Further information on The Open Group is available at www.opengroup.org.