

Unified Model for Data Security – A Position Paper*

Raja Naeem Akram and Ryan K. L. Ko

*Cyber Security Lab, Department of Computer Science, University of Waikato
Hamilton, New Zealand.*

{rnakram, ryan}@waikato.ac.nz

Abstract—One of the most crucial components of modern Information Technology (IT) systems is data. It can be argued that the majority of IT systems are built to collect, store, modify, communicate and use data, enabling different data stakeholders to access and use it to achieve different business objectives. The confidentiality, integrity, availability, auditability, privacy, and quality of the data is of paramount concern for end-users ranging from ordinary consumers to multi-national companies. Over the course of time, different frameworks have been proposed and deployed to provide data security. Many of these previous paradigms were specific to particular domains such as military or media content providers, while in other cases they were generic to different verticals within an industry. There is a much needed push for a holistic approach to data security instead of the current bespoke approaches. The age of the Internet has witnessed an increased ease of sharing data with or without authorisation. These scenarios have created new challenges for traditional data security. In this paper, we study the evolution of data security from the perspective of past proposed frameworks, and present a novel Unified Model for Data Security (UMDS). The discussed UMDS reduces the friction from several cross-domain challenges, and has the functionality to possibly provide comprehensive data security to data owners and privileged users.

Keywords-Data Security, Data Provenance, DRM, Access Control, Information Sharing, Cryptography.

I. INTRODUCTION

A crucial element of our modern inter-connected digital life is *data*. It can manifest as holiday snaps, bank statements, health records, next generation product design or military tactics. The importance of data for ordinary users through to multi-national corporations and governments cannot be over-emphasized. Data protection was a critical concern even before the advent of the Internet. Realisation of the importance of securing data led to the development of the science of cryptography and secure storage (i.e. safes and locks). In the age of the Internet and increased connectivity, the ease of sharing data with or without authorisation has created new challenges for data security. These challenges have also seen the redesign of data security from basic constructs of confidentiality and access control to domains of privacy, availability, auditability, quality of data and traceability.

*This work is a position paper that discusses notion of Data Security from various aspects. We also present a new model for data security termed as “Unified Model” with its objectives and requirements. The paper does not present an implementation of the proposed “Unified Model”.

A. Growing Ineffectiveness of Traditional IT Security

The proliferation of the IT industry in different industries presented new challenges associated with the protection of data. Some early systems were designed to cater for particular sets of problems associated with specific industries; i.e. Bell-LaPadula Model [1] and Biba Model [2], while other mechanisms like Role Based Access Control (RBAC) [3] and Access Control List (ACL) [4] were proposed as general-purpose data (and services) security models. With increased inter-connectivity and large scale data collection and sharing between organisations, many of the traditional data security paradigms solely based on *ex ante* (up-front) controls are no longer effective [5]. An example is copyright control and the need to accommodate fair use. Modern data security paradigms not only require a gate-keeper approach, where each entity has to authenticate before accessing the data, but should also satisfy additional properties such as traceability and tamper-evidence, discussed later in this paper. Another example is the increasing ineffectiveness of IT perimeter defence (firewalls, hardening, etc), as witnessed by the recent increase in security breaches of high profile institutions.

New thinking and a new approach to data security is evidently needed. In this paper, we will evaluate the evolution and current state-of-the-art of data security paradigms, discussing how these paradigms are implemented and what application areas they are suited to, along with their possible shortcomings.

The contribution of this paper is to set an agenda for the exploration of new and innovative data security paradigms. The paper proposes a set of properties that a data security mechanism should provide for future inter-connected cyber systems.

B. Structure of the Paper

In section II, we begin with a brief discussion of the problems associated with existing data security paradigms. We then examine the area of data, data-users and ownership/control, which provides the definition that we subsequently use in this paper. To elaborate on the current state of data security, we discuss traditional access control mechanisms in section III. This discussion leads to the idea of Digital Rights Management (DRM) in section IV. Subsequently, in section V the paper diverges into the concept of data provenance and how it relates to data security. Sections III to V may seem to be a random collection of different data security mechanisms: in this paper

we succinctly discuss them to high-light work done in these well-established but somewhat isolated fields that originally focused on solving the issues of a niche market but which have now inspired the Unified Model for Data Security (UMDS). In section VI, we articulate the rationale for rethinking data security not as an isolated but as a comprehensive unified mechanism. Finally, in section VII we summarise the findings of this paper and present the agenda for future research on UMDS paradigms.

II. THE CURRENT STATE OF DATA, USER, AND OWNERSHIP

While IT innovation has progressed by leaps and bounds, IT security techniques are still lagging behind. In this section, we present the rationale behind the importance of data security and why a rethinking exercise is necessary.

Subsequently, we discuss the concept of data users and its various sub-categories. The section concludes with a brief description of data control and how it is defined from the perspectives of different stakeholders.

A. Data Security: The Problem

1) *Defining Data*: Data is defined as “factual information (as measurements or statistics) used as a basis for reasoning, discussion, or calculation” by the online Merriam Webster dictionary¹. The same dictionary defines data from a computer science perspective as “information in numerical form that can be digitally transmitted or processed”.

The nature of data in the digital space is different from the general meaning of the word: data is any collection of numerical (e.g. binary) values that a computer stores, processes, and communicates. This can include software and associated information generated and/or consumed by software, hardware, and users. Therefore, any assemblage in binary form, whether as part of a software package or associated information, can be categorised as data. If we follow this definition, then anything in a digital environment (i.e. computers, mobile phones, networks, etc.) is data. This definition is too vague and encompasses a very broad range of devices, which is not useful for our discussion. Therefore, we have tried to narrow down the definition of data in the context of this paper:

“Data is a collection of numerical values (i.e. binary values) that make a unique identifiable set, representing a passive entity. Collection of data requires an active entity (i.e. software, and hardware modules) to collect, process, and communicate it.”

Programmes such as Microsoft Word, Excel, Acrobat PDF and information stored in a database are categorised as data. In each of the given examples, the collection has a unique identifiable representation that notifies the software of the nature of the collection, whether the collection is a PDF or a Microsoft Word file. An executable file (e.g. *exe*) can also be categorised as a data file until it starts execution, at which stage it takes the role of an active entity that manipulates passive entities (i.e. data files). Therefore, in this paper, data security

means the security of passive entities and not the process executor’s (i.e. software and hardware) security.

2) *The Obsolescence of Traditional Data Security Approaches*: Most computer systems, including software and hardware, are built to service data in one form or another. Security of data has been recognised since ancient times, which led to the development of cryptography and locked cabinets. With the advent of computers and especially the internet, readily available access to data stored at remote locations mandated the development of security mechanisms. Early implementation translated the traditional data security controls (i.e. locked cabinets and cryptography) to the digital world. This attempt brought in access control mechanisms (e.g. password-based access) and secure channel communications (rather than insecure network connection). Passwords replaced keys to the locked resources which in the majority of cases were related to data or some related services (i.e. email).

Soon it was realised that simply translating the physical world’s security measures had limiting effects on digital space. The overarching desire to share data and make it readily collectable and available to a wide range of services can be construed as a major cause of the limited effectiveness of traditional data security mechanisms. For auditing purposes, specialised mechanisms were designed so that system administrators could track who had accessed particular resources. However, with the increased complexity of computer systems and a vast array of data structures (files and databases) that a user can use in a variety of ways, log recording became cumbersome and for system administrators it became a hassle to audit such logs. The log-based data protection mechanisms were useful when data remained within well-defined boundaries.

However, well-defined boundaries are difficult to ascribe in a modern IT infrastructure. Modern IT infrastructures are inter-connected with third party IT systems, employees use mobile and ubiquitous technologies that might be under the control of an IT infrastructure administrator. Therefore, logs are well suited to environments where data remains in a clearly defined and centrally controlled environment; however, with fluid boundaries and continuous flow of data across the boundaries, such mechanisms have their limitations. In recent years, data provenance has been put forward as a possible way to track data across heterogeneous independently managed environments. We will return to data provenance later in this section V.

The Internet, and in recent years the advent of cloud computing, has enabled malicious insiders (software including worms and viruses, and disaffected employees) to ex-filtrate sensitive data [6]. The difficulty in preventing information leakage is exacerbated by cloud computing and over-reliance on traditional protection mechanisms [7]. Furthermore, control of data, its storage (e.g. location) and transformations is difficult to (independently) guarantee to be secure and traceable (logs and provenance) – unless you trust your cloud provider.

The data is the most crucial element of any computer system. Whether it belongs to an individual user (i.e. consumers) or a large-scale organisation, the importance of data cannot be undermined [8]. The paramount importance of data

¹Website: www.merriam-webster.com/dictionary/data

security is understood, and substantial work has been put into achieving it. There are a number of different proposals that do solve the problem in a confined view; however, from a broad view of IT systems they have limited effectiveness (e.g. access control against and information leakage by trusted software or malicious insiders). We discuss a select few of these data security proposals in sections III to V, illustrating their application and effectiveness.

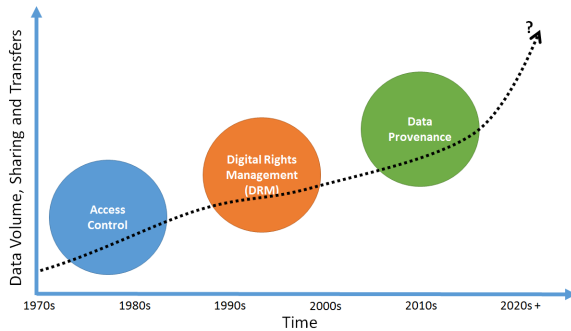


Figure 1. Evolution of Data Security and Quality Mechanism

With increases in the volume of data being generated/consumed and the proliferation of data-sharing technologies, different data security and quality mechanisms were proposed. In Figure 1, we group many of these mechanisms into three broad categories: access control, DRM, and provenance. These groups of mechanisms are also evolving internally and providing their stated goals. However, these technologies in their current state might not be the solution to future demands of data sharing with strong security, privacy and auditability of data items.

The problem in state-of-the-art data security is that there is no holistic approach that looks at data security from a comprehensive system view, rather than as a problem of access control, cryptography, operating system or applications. An effective cryptography will not be effective if there is insecure key management, operating system vulnerabilities or applications using the decrypted information and leaking those to malicious entities. We are of the opinion that a holistic approach to data security will express it as a system problem – not just a cryptography or access control issue.

Before we begin the in-depth discussion of various data security mechanisms, subsequent sections will discuss the concepts of data ownership, use and control. Definitions of these concepts are necessary in order to follow the discussion in later sections of the paper.

B. Users, Data-Ownership and Control

The concept of data ownership is one of the most contentious issues related to modern digital life. When a user fills in an online form, and provides her details, who owns the information that is part of the provided details? In most cases, this information is most probably under the control of the organisation that is collecting the information. Do they own the data, even when it belongs to the users? What role do “terms and conditions” play in such collection of data?

A recent example is the Instagram² that made it acceptable to publicly use customers’ uploaded images for commercial purposes [9]. The language used in the terms and conditions clearly gives Instagram the right to use customers’ data for commercial purposes, but after a strong reaction from their users they had to back down and pledged to remove the condition from their terms and conditions [10].

As an example, Google’s privacy statement clearly points out that contents uploaded by their users remains the users’ intellectual property. However, in the same privacy statement they also list the following:

“When you upload or otherwise submit content to our Services, you give Google (and those we work with) a worldwide license to use, host, store, reproduce, modify, create derivative works (such as those resulting from translations, adaptations or other changes we make so that your content works better with our Services), communicate, publish, publicly perform, publicly display and distribute such content.”

So in simple words, users retain the intellectual property rights to any work they upload on Google, but Google has the right to snoop into the upload contents. Therefore, partial ownership is still with the users, but they don’t have complete control of what their data is used for. Similar examples can be quoted from other online service providers.

Current data security techniques can offer limited protection to a user, including access control and cryptography. Most of the issues stem from the basic concept of data ownership and control, along with the different roles a user performs. Data ownership is defined as “the legal rights and complete control over a single piece or set of data elements³”. Legal rights are difficult to manage and contest in a court of law by individual consumers, and the definition is closer to the organisational ownership of their data, not an individual’s ownership of her data, although legal rights and strong accountability are proposed as a possible solution to the lack of effectiveness of traditional data security measures [5]. From a data security point of view, data ownership is the ability to control the access, modification, and transmission of data along with the ability to track and audit the data and associated processes.

The data owner, whether it is an individual user or an organisation, has complete control over how, where, and by whom their data can be accessed. In addition, the data owner should have the ability to track the data and the processes performed on it. A data owner also has the ability to delegate the administration of its data to third parties, which act as data custodians (i.e. Google in case of Google Docs). In the current context, the data owner intrinsically transfers their ownership rights (e.g. functionalities to control the data) to the data custodians. The data owner has to trust the data custodian and has no “technical” means of controlling any aspect of their data beyond what is sanctioned by the data custodian.

In the context of the data users, they are individuals or organisations that utilise the data (after explicit or implicit permission from either the data owner or custodian). For example, users may visit a news site to read some articles.

²Instagram: It is a popular photo-sharing app, owned by Facebook. Web link instagram.com

³Web link: <http://www.techopedia.com/definition/29059/data-ownership>

The owner/custodian of the news article has given permission to access the information to the respective users. However, they do not have any functionality to track the historic derivation and possible future processes performed on the data.

Therefore, a data owner has complete control over a single piece or set of data items. The control includes but is not limited to the i) confidentiality, ii) integrity, iii) availability, iv) access & access provisioning, v) communication, and vi) tracking. The data custodian manages the data on behalf of the data owner, but the provisions made with respect to the control of data should be determined by the data owner. For example, the data owner might delegate access and communication control to the data custodian; as in a newspaper article written by an author. The author has delegated the ability to allow access to her article and communication to the newspaper (organisation). A user is an individual or an organisation that accesses or handles the content of the article (data items) without any particular privileges in relation to it.

Although it represents an over-generalisation of the different roles played by various actors in the IT infrastructure, we have restricted them to distinct categories to make it less complicated to understand the subsequent discussion. In subsequent sections, we will initially look at the access control mechanism, followed by DRMs, information sharing and provenance.

III. DATA SECURITY THROUGH ACCESS CONTROL

Access control-related security mechanisms existed before the advent of computers. The fundamental role of an access control is to act as a gatekeeper, which will check individuals arriving at the gate to ensure they have the required clearance to pass through it. Such a mechanism has been successful in the real world, so it is logical that early researchers in computer science implemented similar mechanisms to protect data and services. In this section we will look at different access control models and why they are still evolving to mitigate the issues associated with future challenges.

A. Access Control Models

To understand different access control models, understanding of the concepts “subjects” and “objects” is crucial [11]. A subject is an entity that has the ability to manipulate and use other entities in a computer system. Examples of subjects can be human users, and processes. In comparison, an object is an entity manipulated and/or used by subjects. Examples of objects can be files, printers and even other subjects.

Therefore, an access control mechanism has to check whether a subject has the right to access an object. The policy that stipulates whether a subject can access an object, and under what conditions, is referred as “access policy”. Access policy stores the information regarding the privileges assigned to individual subjects in relation to objects. Traditionally, general categories of access control mechanisms were discretionary, mandatory, and roles-based [12]. In discretionary access control, policy on an object is defined by the object’s owner: the entity that has created the object. In mandatory access control, the policy is defined at the system level and

usually by the system administrator. Finally, in role-based access control, subjects are assigned roles, such as students and cashiers. The access policy is then assigned to individual roles rather than individual subjects. Role-based access control is a different model to mandatory and discretionary systems [3], even though there might seem to be an apparent similarity to mandatory access models (i.e. replacing subjects for roles).

Early access control models were evolved from implementations in the military domain, which has clear and hard delineations of individual subjects and objects and their relationships, which in most cases do not change dynamically. The adoption of computer technology in civilian organisations that have a more fluid and dynamic relationship between subjects and objects required an enhancement of the existing models. Role-based access control is an example of one such enhancement.

Increases in complexity and changes to organisational requirements, technical capabilities, structures and relationships with other organisations required new models to address such changes. To keep up with the changes, more complex and fine-grained access control models were proposed, such as attribute-based access control, policy-based access control, and risk-adaptive access control.

Attribute-based access control makes decisions based on the attributes associated with requesters (subjects), environment (IT system), and resources (objects) [13]. An important functionality of this model is to enable a requester to gain access without any prior registration with the system or resource. This enables an organisation to implement an access control model that accommodates unanticipated requesters.

Different organisations have different requirements in terms of implementing risk management, accountability and compliance with relevant laws and regulations. Examples include the Health Insurance Portability and Accountability Act (HIPAA) for health care [14] and Sarbanes-Oxley (SOX) for corporations [15]. Policy-based access control is a new model that aims to address the requirement of having a mature and secure access control model based on abstract policy and governance requirements [16]. In most implementations, policy-based access control models are implemented as an extension to attribute-based access control models [17].

Modern organisations are dynamic and changes to their operational environment are constant. Such changes are initiated by legal requirements, economic and financial realities and risk factors. This dynamic nature extends to their security and access control mechanisms. Access control models do not adequately address this dynamic and constantly evolving nature of today’s cyber security. The risk-adaptive access control model aims to take into account real-time, adaptable, and risk-aware mechanisms [18]. Risk-adaptive access control models are implemented over the traditional access control models with the right to change access policies dynamically [19]. In addition, such models require Trusted Platform Modules (TPMs) [20], automated behavioural analysis [21], and machine learning algorithms [22] which are still either not widely adopted or not integrated with access control models [17]. Furthermore, risk-adaptive access control requires a well-defined and unambiguous mechanism that generates and

communicates the environmental conditions on which the risk-adaptive access control model makes the decisions. A well implemented risk-adaptive access control model will require data confidence regarding the environmental conditions, which can be provided by data provenance [23] discussed briefly in section V.

B. Wrapping Up on the Discussion

Except for the risk-adaptive access control model, all other models are based on static policies that require external interference. This means that with fast-changing threat landscapes, the reaction time to avoid any data breaches is very short. To be effective in such a rapidly changing environment, the access control mechanism has to have dynamic adaptability. Furthermore, access control is point-of-entrance security. Once a malicious entity has passed this point, it is difficult to restrict its actions within the privileges of the entity. The data might be breached in a manner that is totally permitted by the privileges assigned to the entity; for example, covert channels [24] or unintended consequences of ambiguous policies. Access control mechanisms are an important element of data security, although their effectiveness is still debatable. However, any data security mechanism has to implement access control in one way or another – it is difficult to implement a comprehensive data security without access control.

IV. DATA SECURITY THROUGH DIGITAL RIGHTS MANAGEMENT (DRM)

In our inter-connected world, software and data are transacted as commodities. A software company might want to sell its products online, but does not want its customers to make illegal copies of their software. Similarly, the media industry might want to sell movies and songs to consumers with an assurance that their intellectual property will be protected even when it is no longer under their direct control.

Digital rights management systems have been proposed for the above-mentioned requirements [25]. DRMs represent the set of technologies deployed by hardware manufacturers⁴, and software/data copyright holders to control their contents (products) after they are sold. In most cases such devices, software, and data are in the possession of their consumers (or their devices). Consumers can have malicious intentions and may want to reproduce the contents (illegally [27]).

In the next section, we will briefly discuss DRM architecture and related challenges.

A. Digital Rights Management in a Nutshell

A DRM system has essential components including, i) contents coding (language), ii) content identification, iii) packaging (cryptography), iv) distribution, v) digital rights assertion and usage, vi) tracking and monitoring [28]. The DRM system requires all the participants to implement the listed functionality.

⁴To avoid hardware counterfeiting, manufacturers might deploy Physical Unclonable Functions (PUF) [26] that can be categorised as a DRM technology.

We consider a generic DRM system. A content generator will code its digital content as required by the content distributor. The content generator is the intellectual property owner and the content distributor (provider) is the entity that distributes the contents to individual users. For example, iTunes⁵ is a content distributor while the media companies that make their contents available to consumers via iTunes are content generators.

Content providers will transform the digital content and associated rights using Digital Rights Expression Languages (RELs). Examples are Markup Language (XrML) [29], Open Digital Rights Language (ODRL) [30] and MPEG-21 [31]. The RELs can be considered as a way to convey the access control information to the consumer's content reader [32]. The digital content is communicated securely (using secure channel protocols) to individual consumers. At the consumer end, a DRM enforcement entity will manage the download content and associated policy.

One of the most crucial items, prone to a wide array of threats in the DRM ecosystem, is the rights enforcement and management entity. To strengthen it, a four-layer security framework was proposed that was based on content protection, rights enforcement, rights management and trust management [33]. Rights enforcement and management are similar to access control models, and many of the models discussed in the previous section are implemented as part of the DRM ecosystem. In addition to implementing traditional access models, the DRM system led to the development of the persistent access control model [34].

Content usage generates content tracking information: the DRM enforcement entity on the client side can record (and may transmit) the tracking information to the content provider. The tracking information is to interoperate with the rights management to ensure that the consumer abides by the lease agreement governing the digital contents. The lease agreement might have limits on the number of accesses, or time restrictions⁶.

When implementing DRM systems (only) in software, supported by general purpose operating systems, it is a challenge to provide a strong notion of remote attestation and seal (cryptographic) techniques [35]. Traditional protection mechanisms implemented as part of the operating system (e.g. access control) cannot protect decrypted contents and provide assurance that the DRM policy has been securely enforced [36].

A secure, reliable and trusted execution environment with a remote attestation mechanism is necessary for a secure DRM system. The inclusion of the TPM provided a logical choice to provide such an environment for the DRM system [37].

DRM provides a cross-platform and/or infrastructure assurance that the data (or software) will be used as sanctioned in the associated policy. Traditional access control mechanisms can only enforce the access policy with regard to the data items that are directly under its control, whereas DRM tries to

⁵iTunes: It is a digital media management, distribution and usage application developed by Apple Inc. Web link: <http://www.apple.com/itunes/>

⁶BBC iPlayer has a time limit on downloaded contents that at maximum can be 28 days.

provide the same level of security but on a remote device that is not under the control of the content generator (data owner).

B. DRM Technologies: The Verdict

In modern cyber life, a huge amount of personal data is uploaded to different companies (e.g. Facebook and Google). A possible solution (in a limited sense) can come from the DRM-based mechanisms that enable a data owner to control her data [38]. As DRM might be getting a new lease of life as a privacy-preserving technology, it is slowly falling from relevance in its traditional settings⁷.

Whatever the final outcome, the basic design principle of remotely managing data access even under a malicious host/user is necessary for any future data security strategy. DRM might not be implemented in the same way as today, but the lessons learned from DRM research will definitely help any future proposals.

V. DATA SECURITY THROUGH DATA PROVENANCE

In the previous sections, we have looked at mechanisms that protect data from unauthorised access or usage. In this section, we will look at data provenance that can help in data auditing and tracking of changes that happen to the data.

A. Traditional Provenance Perspectives

As the amount of data produced increases exponentially, there is an ever-growing interest in understanding how a particular piece of data was created and what manipulations the given data has gone through in the past.

Data provenance can be defined as a snapshot of all the transformations a data item has gone through during the process that created the data item. In other words, as defined by [23, 41, 42] provenance is the meta-data of the derivation history of data. Data provenance is an important component in many data-intense studies and/or industries like eScience [43] and healthcare [44]. In such environments, provenance ensures the quality of data and repeatability of results.

Early provenance systems were mainly concerned with the data quality of a database, to ensure that no error crept into large and lengthy calculations. Even if such errors appeared, a provenance system allowed a query mechanism to search the source of such an error, so other possible data items that in their provenance record included the particular node which was found to be the culprit could also be adjusted to avoid the proliferation of errors produced. A substantial body of work has been conducted in the domain of data provenance related to databases [45].

B. Emerging Provenance Challenges

In recent years, the advent of cloud computing and document security across distributed systems have given a new dimension to data provenance design and requirements [41,

46]–[48]. Furthermore, provenance is collected at different levels in a system and between systems [41]: some mechanism collects the provenance at the application layer⁸. Examples of provenance mechanism collection at the system layer are HP's TrustCloud (Flogger) [46, 47], S2Logger [49], DataPROVE [41] and PASS [42]. There is also an increasing call for a data-centric view over the traditional system-centric view for cloud computing [50]. Zhang et al. [41], after considering the traditional provenance [51] and emerging provenance criteria, stipulated the following core requirements and properties of cloud data provenance systems that address common challenges listed, and make provenance in the cloud truly useful:

- 1) Coordination between storage and computing facilities: Provenance, like typical digital data in a cloud, is generated and processed by computer facilities, and maintained on storage facilities.
- 2) Interface/API that allows customers to record provenance of their objects. If data originating from non-provenanced sources is to be stored in the cloud, a solution to incomplete provenance is to allow verified/signed manual input of provenance for the data, with the help of a well-designed interface and/or API.
- 3) Security elements that need to be provided for reliable provenance:
 - Integrity: The assurance that provenance is not forged or tampered. An extensive range of papers have been presented on how to secure the provenance record [52, 53]; however, it was difficult to find any provenance-related work that presents a mechanism to provide complete data security.
 - Availability and Auditability: An auditor can check the integrity and the correctness of provenance information, though how to prohibit or detect suspicious user annotation and false provenance fabricated by malware is still an open question.
 - Confidentiality: Provenance may contain sensitive information about the data it describes, or it may be sensitive information by itself. Encryption methods and access control policies for provenance are a necessity to prevent information leakage from provenance. It is difficult to ensure confidentiality when inside intruders such as privileged administrators and cloud service providers are involved.
- 4) Provenance data consistency: Provenance information must be consistent with the data it describes. Inconsistency in provenance and its data can mislead both customers and service providers.
- 5) Atomicity: Provenance must be recorded atomically with the data it describes. Atomicity pertains to provenance storage, whereas consistency pertains to provenance and data retrieval. Atomicity and consistency together assure that provenance accurately and completely describes the data, i.e., provenance data-coupling.
- 6) Causal ordering: A provenance system must ensure that an object's ancestors and their provenance are persistent

⁷Apple and EA decided that DRM is not a viable technology for content distribution [39, 40]. Apple went for watermarking, so illegal copies can be tracked back to the users who bought them and then they might be held accountable.

⁸most if not all of the database provenance mechanisms collect provenance in this layer

before making the object itself persistent, which prevents dangling pointers from appearing in the provenance recorded as a directed acyclic graph (DAG).

- 7) Data independent persistence, also referred to as long-term persistence: A provenance system retains an object's provenance even after the object is removed. Although an object is removed, its provenance must still be present in the provenance DAG as some other objects' ancestor; deleting the object's provenance will make the DAG disconnected. An object's provenance can be removed only if it has no descendants.
- 8) Efficient query: The primary use of provenance data is for users to check the lineage properties of a corresponding object of interest, through external queries. Considering the graph structure of provenance and the large size of the cloud and the objects stored in it, efficiency of querying affects directly the value of provenance.

We consider that provenance can become a post authentication data security mechanism, where discovery and trace of a data breach can easily be discovered using the system provenance. However, papers on the topic of research issues, and survey papers on data provenance do not elaborate on any such proposal [54]. Data provenance, if implemented as a light-weight mechanism at system level, can provide an excellent auditing tool, which in our opinion is an important component of a holistic approach to data security. Furthermore, data provenance records might help risk-adaptive access control mechanisms to make better informed decisions.

In the applicability of provenance in data security, implementers should consider the elements that are not required to be stored – even their history can lead to security and/or privacy violations. For example, storing the internal state (especially the seed values) of a pseudorandom number generator [55] as discussed by [42].

VI. UNIFIED MODEL FOR DATA SECURITY

In subsequent sections, we briefly discuss a few of the data security and quality mechanisms that we think adequately represent the gaps between different fields. With the gaps addressed, we aim to work towards addressing data security according to modern day challenges.

A. Objectives of the Unified Model

The main aim of the unified model is to bridge this gap and create a data security model that is flexible enough to accommodate complex policies and scalable to be deployed from embedded environments like cyber physical systems to the cloud computing environment.

The unified model takes into account the three main components of a data security (see Figure 1):

- 1) ex ante (gate-keeping),
- 2) policy enforcement, and
- 3) audit/trackability.

The unified model is not the overlapping part of the three components. It is the complete integration of all three of them. We propose that the unified model can be implemented




Figure 2. Overview of Unified Model

as a standalone component, like a virtual machine on a device, and all data accesses are handled by the virtual data machine. It can also be implemented at the same layer as the kernel, between system services and hardware, but it separates data-related routines from general routines (existing kernel) and includes them in a secure and trusted data kernel. Both methodologies, whether the unified model is implemented as a virtual (data) machine or data kernel, should have the support of tamper-resistant hardware (e.g. TPM).

The unified model can be understood as a proposal that merges and enhances existing mechanisms proposed for different data security (sub-set) requirements. In subsequent sections, we will describe a list of possibility functionalities that are considered to be crucial for the unified model.

B. Agenda for a Unified Model

This section presents the list of requirements for the unified model.

- 1) Confidentiality: Access to data is only granted to authorised and trustworthy entities/processes after explicit verification of identity.
- 2) Integrity: Modification to data can only be permitted if carried out by authorised and trustworthy entities/processes. Any unauthorised modification should be detectable.
- 3) Availability: Unauthorised entities/processes should not have the ability to monopolise the resources (i.e. data).
- 4) Data Confidence: A data consumer (i.e. entities and/or processes) can gain assurance about the historic transformation that the respective data has gone through.
- 5) Trust: Data consumers should have a mechanism to evaluate and validate the trustworthiness of data, including the trust evaluation of individual entities, and processes that have handled the data along with individual operations in the past.
- 6) Auditability: Each access and transformation performed on data is recorded as an integral part of the data. Such a record will facilitate the data audit to establish confidence and trust in the data. This feature can also be used by system auditors to validate the derivation and security aspects of the data.
- 7) Accountability and Appropriate Use: Effective auditability could enable the accountability of malign entities/pro-

cesses that handled the data. Such entities/processes might be authorised to access data but their purview did not include the actions they performed.

- 8) Forensic Evidence Preservation: The system has the resilience to protect the forensic evidence, enabling a secure and tamper-evident data provenance that can assist investigators to analyse possible data breaches.
- 9) Privacy (and Accessibility): Privacy of data should be protected and any policy defined by the data owner should be enforced in a secure manner. In addition to privacy, certain data items require accessibility in certain limited situations, like healthcare data in an emergency. Therefore, the policy should be flexible enough to cover such rare events.
- 10) Dynamic Policy Changes: To accommodate the changing environment and diversity of host devices (system) that will handle the data, the policy should be dynamically modifiable to accommodate any local environment-specific changes.
- 11) Risk Averse and Automatically Adaptable: The unified model should be able to detect the changing landscape of risk and adapt the security and privacy policies to adequately represent the data owner's requirements with changing threat vectors.
- 12) Data Ownership: Data control should be completely given to the data owner. The unified model should ensure that no other authority, whether an authorised entity or a malicious entity, gains any data control privileges unless explicitly expressed by the data owner.
- 13) Revocation: The data control privileges and data access rights given to data custodians or data users should be revocable if required by the data owner.
- 14) Inescapable: Access to the data items should not be circumventable around the unified data model components (figure 2)
- 15) Feasible: The implementation of the unified model is technically viable and lightweight, so its inclusion does not adversely affect the overall performance of the host environment.
- 16) Scalable: The model is scalable from resource constraint embedded devices to cloud computing environments.
- 17) Flexible: The policy descriptor should be unambiguously expressive enough so that it can represent data security policies in a control-rich manner, enabling a wide variety of data owners to define the policy.

The above list is not an exhaustive one; however, it is a fundamental requirement stipulated as a point of reference for the development of the Unified Model. We do not propose the model to be a panacea for all information security problems, however, this model might be able to streamline the issues facing data security in rapidly changing environments and react to it, possibly in real-time. Taking the Figure 2 and the requirements listed in this section, we can make an informed guess that a possible unified model can be constructed by integrating risk-adaptive access control mechanism as ex ante, DRM style policy management and enforcement, and data provenance to provide secure data audit and trackability. All

these mechanisms will feed into each other, providing valuable information to each of the components and enabling adequate decisions to be made.

Using the unified model, an individual consumer or an organisation can manage, control and track (audit) their data in the cyber world. For example, social media consumers can upload their data to the remote servers (may be in cloud) but still able to control the dissemination of their data along with track it (and associated actions performed on it). The data repository on the social media servers has to negotiate the rights (access, communication, usage and tacking policies) with the consumer. In addition, the data repository has to provide adequate security assurances to the consumer (unified model's implementation on the consumer device) that it will use the data in accordance with stated policy. The security assurances has to be independently verifiable and trust in social media company or their infrastructure should not be implicit. The unified model can be considered as a first step towards the empowerment of data owners (end-user or an organisation) to control the access, security, dissemination, usage and trackability of their respective data.

VII. CONCLUSION & FUTURE RESEARCH DIRECTIONS

This paper presented the security challenges associated with a crucial component of modern day IT infrastructure: data. A data breach might damage the image of an organisation, or if not, drive it out of business. We discussed the definition of data, and different entities that interact with it in different capacities. Before beginning a generic overview of the proposed unified model, we elaborated on the three prominent domains that deal with the data in one form or another. These domains included access control, DRM and data provenance. The description of these domains provided the initial foundation of the unified model. The unified model covers different data security, privacy, and quality mechanisms to provide a streamlined approach to data protection. The main goal of the unified model is to be flexible and scalable. We also provided a list of potential requirements that possible future work on the unified model has to take into account.

Future research directions will look into the security, privacy and usability challenges to the unified model, followed by the requirement of a trusted execution environment and how such an environment might be provisioned in systems ranging from embedded to cloud. Further work will involve the integration of the three main components of the unified model and how they can establish a beneficial symbiotic relationship to enhance current data security needs.

REFERENCES

- [1] D. E. Bell and L. J. LaPadula, "Secure Computer Systems: Mathematical Foundations," MITRE Corporation, Tech. Rep. 2547, March 1973.
- [2] K. J. Biba, "Integrity Considerations for Secure Computer Systems," MITRE Corp., Tech. Rep. ESD-TR-76-372, 04 1977.
- [3] D. Ferraiolo and R. Kuhn, "Role-based access control," in *In 15th NIST-NCSC National Computer Security Conference*, 1992, pp. 554–563.
- [4] J. H. Saltzer, "Protection and the control of information sharing in multics," *Commun. ACM*, vol. 17, no. 7, pp. 388–402, July 1974.
- [5] D. J. Weitzner, H. Abelson, T. Berners-Lee, J. Feigenbaum, J. A. Hendler, and G. J. Sussman, "Information Accountability," *Commun. ACM*, vol. 51, no. 6, pp. 82–87, 2008.

- [6] F. Rocha and M. Correia, "Lucy in the Sky Without Diamonds: Stealing Confidential Data in the Cloud," in *41st International Conference on Dependable Systems and Networks Workshops*, ser. DSNW '11. Washington, DC, USA: IEEE CS, 2011, pp. 129–134.
- [7] M. Van Dijk and A. Juels, "On the Impossibility of Cryptography Alone for Privacy-preserving Cloud Computing," in *Proceedings of the 5th USENIX Conference on Hot Topics in Security*, ser. HotSec'10. Berkeley, CA, USA: USENIX Association, 2010, pp. 1–8.
- [8] L. Kaufman, "Data Security in the World of Cloud Computing," *Security Privacy, IEEE*, vol. 7, no. 4, pp. 61–64, 2009.
- [9] D. McCullagh. (2012, December) Instagram says it now has the right to sell your photos.
- [10] D. McCullagh and D. Tam. (2012, December) Instagram apologizes to users: We won't sell you photos. Online. CNet.
- [11] M. A. Harrison, W. L. Ruzzo, and J. D. Ullman, "Protection in Operating Systems," *Commun. ACM*, vol. 19, pp. 461–471, Aug. 1976.
- [12] P. Samarati and S. D. C. d. Vimercati, "Access Control: Policies, Models, and Mechanisms," in *Revised Versions of Lectures Given During the IFIP WG 1.7 International School on Foundations of Security Analysis and Design on Foundations of Security Analysis and Design: Tutorial Lectures*, ser. FOSAD '00. London, UK, UK: Springer-Verlag, 2001, pp. 137–196.
- [13] V. C. Hu, D. Ferraiolo, R. Kuhn, A. Schnitzer, K. Sandlin, R. Miller, and K. Scarfone, "Guide to Attribute Based Access Control (ABAC) Definition and Considerations," *NIST SP*, vol. 800, p. 162, 2014.
- [14] Centers for Medicare & Medicaid Services, "The Health Insurance Portability and Accountability Act of 1996 (HIPAA)," USA, 1996.
- [15] United States Code, "Sarbanes-Oxley Act of 2002, PL 107-204, 116 Stat 745," Codified in Sections 11, 15, 18, 28, and 29 USC, July 2002.
- [16] L. Zhi, W. Jing, C. Xiao-su, and J. Lian-Xing, "Research on Policy-based Access Control Model," in *Networks Security, Wireless Communications and Trusted Computing, 2009. NSWCTC '09. International Conference on*, vol. 2, 2009, pp. 164–167.
- [17] National Institute of Standards and Technology, "A Survey of Access Control Models," Online, NIST, August 2009.
- [18] R. W. McGraw, "Risk-Adaptable Access Control (RAAdAC)," Online, NIST, Tech. Rep., August 2009.
- [19] S. Kandala, R. Sandhu, and V. Bhamidipati, "An Attribute Based Framework for Risk-Adaptive Access Control Models," in *Availability, Reliability and Security (ARES), 2011 Sixth International Conference on*, 2011, pp. 236–241.
- [20] "Trusted Computing Group, TCG Specification Architecture Overview," The Trusted Computing Group (TCG), Beaverton, Oregon, USA, revision 1.4, August 2007.
- [21] G. Owen, "Automated forensic extraction of encryption keys using behavioral analysis," in *Cyber Security, Cyber Warfare and Digital Forensic (CyberSec), 2012 International Conference on*, 2012.
- [22] T. M. Mitchell, *Machine Learning*. McGraw-Hill, NY, USA, 1997.
- [23] P. Buneman and Susan, "Data Provenance - the foundation of data quality," September 2010.
- [24] S. A. Ahmadzadeh and G. Agnew, "Covert channels in multiple access protocols," in *ACM SIGCOMM Computer Communication Review*, vol. 41, no. 4. ACM, 2011, pp. 404–405.
- [25] Q. Liu, R. Safavi-Naini, and N. P. Sheppard, "Digital Rights Management for Content Distribution," in *Proceedings of the Australasian Information Security Workshop Conference on ACSW Frontiers 2003 - Volume 21*, ser. ACSW Frontiers '03. Darlinghurst, Australia, Australian Computer Society, Inc., 2003, pp. 49–58.
- [26] R. S. Pappu, "Physical One-way Functions," Ph.D. dissertation, Massachusetts Institute of Technology, March 2001.
- [27] E. H. Freeman, "The Digital Millennium Copyright Act," *Information Systems Security*, vol. 11, no. 4, pp. 4–8, 2002.
- [28] C.-C. Wu, C.-C. Lin, and C.-C. Chang, "Digital rights management for multimedia content over 3G mobile networks," *Expert Syst. Appl.*, vol. 37, no. 10, pp. 6787–6797, 2010.
- [29] J. Y. Halpern and V. Weissman, "A Formal Foundation for XrML," *J. ACM*, vol. 55, no. 1, pp. 4:1–4:42, Feb. 2008.
- [30] R. Pucella and V. Weissman, "A Formal Foundation for ODRL," *CoRR*, vol. abs/cs/0601085, 2006.
- [31] *ISO/IEC 21000-5: Information Technology-Multimedia Framework Part 5: Rights Expression Language*, 2004.
- [32] X. Wang, "Design principles and issues of rights expression languages for digital rights management," *Proc. SPIE*, vol. 5960, 2005.
- [33] E. Diehl, "A Four-layer Model for Security of Digital Rights Management," in *Proceedings of the 8th ACM Workshop on Digital Rights Management*, ser. DRM '08. NY, USA: ACM, 2008.
- [34] J. Park and R. Sandhu, "The UCONABC Usage Control Model," *ACM Trans. Inf. Syst. Secur.*, vol. 7, no. 1, pp. 128–174, Feb. 2004.
- [35] J. F. Reid and W. J. Caelli, "DRM, Trusted Computing and Operating System Architecture," in *Proceedings of the 2005 Australasian Workshop on Grid Computing and e-Research - Volume 44*, ser. ACSW Frontiers '05. Darlinghurst, Australia, ACS., 2005, pp. 127–136.
- [36] "Secure Data Management in Trusted Computing," in *Cryptographic Hardware and Embedded Systems - CHES 2005*, ser. Lecture Notes in Computer Science, J. Rao and B. Sunar, Eds. Springer Berlin Heidelberg, 2005, vol. 3659, pp. 324–338.
- [37] A. Yu, D. Feng, and R. Liu, "TBDRM: A TPM-Based Secure DRM Architecture," in *CSE (2)*. IEEE Computer Society, 2009, pp. 671–677.
- [38] T. Sinonite. (2013, October) Microsoft Thinks DRM can Solve Privacy Problem. Online. Technology Review.
- [39] E. Hardy. (2009, January) Apple Removes DRM Restrictions from iTunes. Online. Brighthad.
- [40] J. Brightman. (2013, March) EA: "DRM is a failed dead-end strategy". Online. Games Industry International.
- [41] O. Q. Zhang, M. Kirchberg, R. K. L. Ko, and B. S. Lee, "How to track your data: The case for cloud computing provenance," in *Cloud Computing Technology and Science (CloudCom), 2011 IEEE Third International Conference on*. IEEE, 2011, pp. 446–453.
- [42] K.-K. Muniswamy-Reddy, D. A. Holland, U. Braun, and M. Seltzer, "Provenance-aware Storage Systems," in *Proceedings of the Annual Conference on USENIX '06 Annual Technical Conference*, ser. ATEC '06. Berkeley, CA, USA: USENIX Association, 2006, pp. 4–4.
- [43] R. Bose and J. Frew, "Lineage Retrieval for Scientific Data Processing: A Survey," *ACM Comput. Surv.*, vol. 37, no. 1, pp. 1–28, Mar. 2005.
- [44] "EU PROVENANCE Project: An Open Provenance Architecture for Distributed Applications," in *Agent Technology and e-Health, 2008*.
- [45] T. Cadenhead, M. Kantarcioglu, and B. Thuraisingham, "A framework for policies over provenance," in *3rd USENIX Workshop on the Theory and Practice of Provenance*, P. Buneman and J. Freire, Eds. Crete, Greece: USENIX, June 2011.
- [46] R. K. L. Ko, P. Jagadpramana, M. Kirchberg, Q. Liang, M. Mowbray, S. Pearson, and B. S. Lee, "Trustcloud - a framework for accountability and trust in cloud computing," in *IEEE 2nd Cloud Forum for Practitioners (IEEE ICFP 2011)*. DC, USA: IEEE CS, July 2011.
- [47] R. K. L. Ko, P. Jagadpramana, and B. S. Lee, "Flogger: A file-centric logger for monitoring file access and transfers within cloud computing environments," in *3rd IEEE International Workshop on Security in e-Science and e-Research (IEEE ISSR 2011), in conjunction with IEEE TrustCom 2011*. Changsha, China: IEEE, 2011.
- [48] R. Lu, X. Lin, X. Liang, and X. S. Shen, "Secure Provenance: The Essential of Bread and Butter of Data Forensics in Cloud Computing," in *Proceedings of the 5th ACM Symposium on Information, Computer and Communications Security*, ser. ASIACCS '10. New York, NY, USA: ACM, 2010, pp. 282–292.
- [49] C. H. Suen, R. K. L. Ko, Y. S. Tan, P. Jagadpramana, and B. S. Lee, "S2logger: End-to-end data tracking mechanism for cloud data provenance," in *Trust, Security and Privacy in Computing and Communications (TrustCom), 2013 12th IEEE International Conference on*, 2013, pp. 594–602.
- [50] R. K. L. Ko, M. Kirchberg, and B. S. Lee, "From system-centric to data-centric logging-accountability, trust & security in cloud computing," in *Defense Science Research Conference and Expo (DSR), 2011*. IEEE, 2011, pp. 1–4.
- [51] K.-K. Muniswamy-Reddy, P. Macko, and M. Seltzer, "Provenance for the Cloud," in *Proceedings of the 8th USENIX Conference on File and Storage Technologies*, ser. FAST'10. CA, USA: USENIX, 2010.
- [52] Z. Bao and S. B. Davidson, "A fine-grained workflow model with provenance-aware security views," in *3rd USENIX Workshop on the Theory and Practice of Provenance*, P. Buneman and J. Freire, Eds. Crete, Greece: USENIX, June 2011.
- [53] N. Swamy, B. J. Corcoran, and M. Hicks, "Fable: A language for enforcing user-defined security policies," in *IEEE Symposium on Security and Privacy*, 2008, pp. 369–383.
- [54] Y. S. Tan, R. K. L. Ko, and G. Holmes, "Security and Data Accountability in Distributed Systems: A Provenance Survey," in *the 15th IEEE International Conference on High Performance Computing and Communications (IEEE HPCC13)*. Zhang JiaJie, China: IEEE CS, 2013.
- [55] R. N. Akram, K. Markantonakis, and K. Mayes, "Pseudorandom Number Generation in Smart Cards: An Implementation, Performance and Randomness Analysis," in *5th International Conference on New Technologies, Mobility and Security (NTMS)*, A. Mana and M. Klonowski, Eds. Istanbul, Turkey: IEEE CS, May 2012.