Computer Standards & Interfaces 36 (2014) 689–703

Contents lists available at ScienceDirect

Computer Standards & Interfaces

journal homepage: www.elsevier.com/locate/csi

End-to-end policy based encryption techniques for multi-party data management

Michael Beiter [a], Marco Casassa Mont [b], Liqun Chen [b], Siani Pearson [b,*]

[a] Chief Technology Office, HP Printing and Personal Systems, Fort Collins, USA
[b] Cloud and Security Lab, HP Labs, Bristol, UK

ARTICLE INFO

Available online 31 December 2013

Keywords:
Cloud
Sticky policy
Policy enforcement
Privacy
Secret sharing

ABSTRACT

We describe a data management solution and associated key management approaches to provide accountability within service provision networks, in particular addressing privacy issues in cloud computing applications. Our solution involves machine readable policies that stick to data to define allowed usage and obligations as data travels across multiple parties. Service providers have fine-grained access to specific data based on agreed policies, enforced by interactions with independent third parties that check for policy compliance before releasing decryption keys required for data access. We describe alternative solutions based upon Public Key Infrastructure (PKI), Identity Based Encryption (IBE) and advanced secret sharing schemes.

1. Introduction

Lack of trust about privacy and security practice is at present a key inhibitor in moving to cloud models [1]. When sharing and storing information in the cloud, additional assurance is needed that appropriate security and privacy measures have been taken by Cloud Service Providers (CSPs). This problem is also present more generally as service provision chains become more global, complex and dynamic. Both business consumers and citizens are requiring more control over the usage and sharing of their personal and confidential information, as this is handled within potentially complex service provision chains.

Current commercial solutions primarily focus on traditional, back-end security controls (e.g. access control) on the data, once this data is stored on the service provider side. Privacy and data control aspects, such as fine-grained definition and enforcement of user's preferences and policies (on how to process data, to whom to disclose data, various obligations on data transformation, deletion, etc.), are rarely implemented unless in a coarse-grained way (e.g. via macro opt-in, opt-out options). These solutions usually do not scale across multiple control domains: users' preferences and policies on how to handle data are not necessarily propagated and enforced across a chain of data disclosures within multiple Cloud Service providers. This means that users have little end-to-end control about the destiny, usage and management of their data, once disclosed to a Cloud Service Provider.

To address this issue, we suggest a variety of solutions based on sticky policies, where policies and constraints are attached to data as it is transmitted and stored within the cloud. Sticky policies are strictly associated with users' data and drive access control decisions and enforcement of privacy and confidentiality.

Our solutions ensure that data disclosed within cloud services is used, accessed, processed, stored and shared, etc. based upon agreed (potentially fine-grained) policies and constraints and degrees of assurance are provided by independent (trusted) third parties about compliance to these policies. Mechanisms using data encryption, driven by policies, can be used to ensure degrees of (fine-grained) data protection; Trusted third parties (called Trust Authorities (TAs)) can be used to provide compliance checking, enforcement and audit capabilities. Our solutions provide a practical solution to enhancing user control and providing accountability within the cloud, removing business barriers in the sense that organizations might be willing to move more of their sensitive operations to the cloud model.

We believe that approaches based on cryptography are suitable to make significant progress towards providing the required level of control and accountability on personal and confidential data. This paper illustrates how this can be achieved by focusing on three solutions, one is general using ordinary Public Key Infrastructure (PKI), a second uses Identity Based Encryption (IBE) [2] and a third is based upon secret sharing [3]. In the PKI-based approach, it is assumed that all the stakeholders have certified public/private key pairs from trusted Certificate Authorities (CAs). In this context, these CAs can play the role of Trust Authorities. Policies are bound to data by encrypting the data under a symmetric key that a sender and receiver conditionally share based on fulfilment of policies, and sticking the data to the policing using public-key enveloping techniques. If IBE techniques are used instead for this binding then it means that a third party needs to check certain properties (as specified within the sticky policy associated with data) at the time of decryption, before an IBE decryption key for that data is released. These IBE techniques are conceptually equivalent to the PKI once, however they leverage a different cryptographic schema. An alternative approach consists in leveraging secret sharing techniques: in this case, the parties involved in the data management solution are enrolled in several, cascaded secret sharing schemes. By recreating the shared

* Corresponding author. Tel.: +44 117 3162558.
E-mail addresses: Michael.Beiter@hp.com (M. Beiter), Marco.Casassa-Mont@hp.com (M. Casassa Mont), Liqun.Chen@hp.com (L. Chen), Siani.Pearson@hp.com (S. Pearson).

0920-5489/$ – see front matter © 2013 Published by Elsevier B.V.
http://dx.doi.org/10.1016/j.csi.2013.12.004

secrets, the parties can compute encryption keys required to access managed data. Instead of needing to a priori define all trusted authorities that will supervise access to the managed data and to manually enable each such Trusted Authority for each asset, the customer only has to provide a share of the secret sharing scheme. The resulting approach has more manageable computation, storage, and transmission bandwidth requirements as compared to prior solutions, and yet can still provide fine-grained control over access and usage of customer data. In general, the most appropriate solution will vary according to the context and trust models involved.

2. Cloud scenario

This paper focuses on a Cloud Scenario consisting of multiple Service Providers, end users and enterprises. In this scenario, both end-users and employees within enterprises make active use of services in the Cloud, as shown in Fig. 1.

Cloud service providers can use services provided by other providers in the cloud, in order to supply the required capabilities. For example, a storage service provider might use third-party data back-up services and information lifecycle management capabilities and part of their offering.

In this context, personal information, confidential data, etc. can flow from one service provider to another one, due to a chain of service interactions and dependencies. For example, a user might disclose personal data to a CSP, during a business interaction and/or the provision of a service. The CSP might then need to interact with other service providers in the cloud, in order to provide the desired service. This might require sharing some of the personal data.

We consider situations, such as within health service provision, access to applications and services in the cloud (storage, computing, etc.), and so on, as shown in Fig. 1, where a customer (that might be a citizen, employee or an enterprise) indeed needs to reveal personal and even sensitive information in order to receive a service, but wishes to control the way in which that information is used.

In this paper we describe our approach to provide this control capability as well as degrees of assurance. We want to:

- enable users to express their (privacy and security) preferences and policies when disclosing their personal data;
- provide mechanisms to protect data whilst it is shared across parties;
- provide mechanisms to increase the level of accountability.

3. Our general approach

We define a system and mechanisms to enable the protection of data to be shared by a user (or service) with service providers, based on agreed policies and privacy preferences. The user can be actively involved in the selection of multiple, interchangeable services called Trust Authorities (TAs), that will track and audit for the fulfilment of these policies.

Our solutions use sticky policies associated with data to dictate how to handle data at the receiver side. Our schemas involve three types of parties: Cloud Service Providers (CSP), which store and process the user's data, Trusted Authorities (TA), which audit that the CSPs handle the user's data according to the sticky policies defined by the user, and the users themselves, who own the assets and define access restrictions in sticky policies.

We aim to enable the users to define policies which are preferences or conditions about how that information should be treated. The policy governs the use of associated data, and may specify for example the following:

- The purposes of using data (e.g. for research, transaction processing, etc.).
- That data may only be used within a given set of platforms (with certain security characteristics), a given network or a subset of the enterprise
- Other obligations and prohibitions (allowed third parties, people or processes; blacklists; notification of disclosure; deletion of data after a certain time)
- A list of trusted TAs (potentially the result of a negotiation process).

The policy may be represented in any convenient format.

The basic mechanisms of the proposed sticky policy solutions, also shown in Fig. 2, are as follows:

- In order to be able to more easily interpret and enforce end user policies, instead of offering free expression of policies from end users, their preferences and policies are defined within a framework imposed by organizations. There are different ways of achieving this: one mechanism is that SPs offer a 'smart notice' containing the list of supported (macro) policies and TAs, where these policies relate to access control and obligation behaviours supported by the organization, and the end user can choose from these [4].
- A user (customer) — interacting with a SP — can select the granularity of how policies apply to items or specific subsets of personal data to be

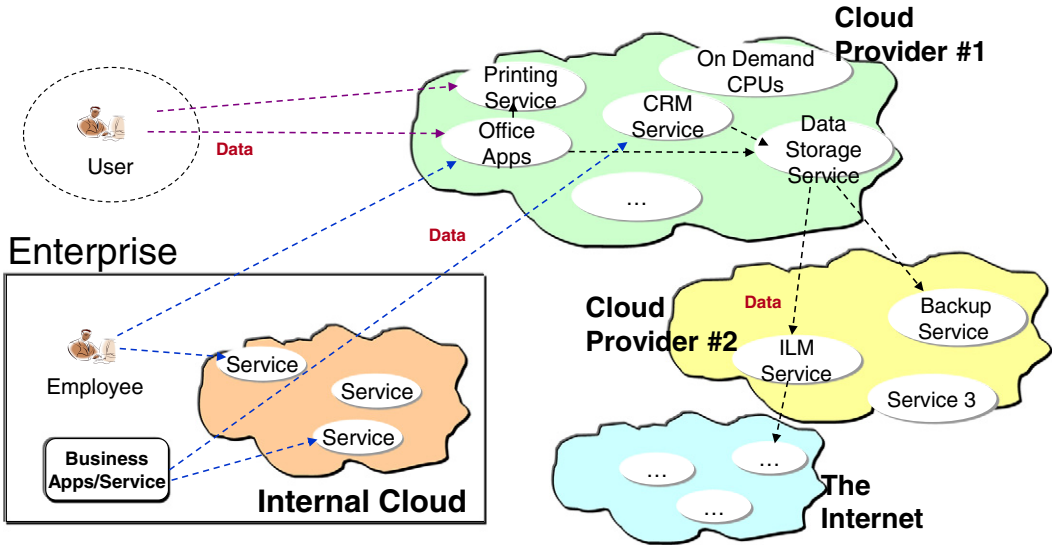

Fig. 1. Example data management scenario.