

Document Provenance in the Cloud: Constraints and Challenges

Mohamed Amin Sakka^{1,2}, Bruno Defude¹, and Jorge Tellez²

¹ Novapost, Novapost R&D, 13, Boulevard de Rochechouart 75009 Paris-France
{amin.sakka,jorge.tellez}@novapost.fr

² TELECOM& Management SudParis, CNRS UMR Samovar, 9, Rue Charles
Fourrier 91011 Evry cedex-France
{mohamed_amin.sakka,bruno.defude}@it-sudparis.eu

Abstract. The amounts of digital information are growing in size and complexity. With the emergence of distributed services over internet and the booming of electronic exchanges, the need to identify information origins and its lifecycle history becomes essential. Essential because it's the only factor ensuring information integrity and probative value. That's why in different areas like government, commerce, medicine and science, tracking data origins is essential and can serve for informational, quality, forensics, regulatory compliance, rights protection and intellectual property purposes. Managing information provenance is a complex task and it has been extensively treated in databases, file system and scientific workflows. However, provenance in the cloud is a more challenging task due to specific problems related to the cloud added to the traditional ones.

1 Introduction

With the advent of cloud computing and the emergence of web 2.0 technologies, the traditional computing paradigm is shifting new challenges. This new paradigm brings issues related to the trustworthiness of data as well as computations performed in third-party clouds. Certainly, the digital form of electronic data has many advantages of being easily accessible, accurate and more useful. However, the mobile and modifiable nature of digital entities requires specific metadata and techniques ensuring information probative value. More specifically, in a cloud-computing scenario, most of the data will reside in data clouds, while the applications will run in the cloud. In this case, users require additional insurances regarding the confidentiality, privacy and integrity of information. Required metadata is provenance and can be defined as information about the origin, context or history of the data. Depending on the domain, provenance is useful for many purposes and its systems can support different types of usage. According to studies on information provenance [9,5,7], it can serve for audit trail and justification, regulatory compliance and forensics, replication recipes, attribution and copyright, as well as data quality and informational purposes. Despite it seems obvious, information provenance issues introduces hard challenges. Provenance is not obvious because it overtakes widely recording the whole

data history [14]. Provenance is listed as a hard problem in some information science recent surveys like the UK Computing Society's Grand Challenges in Computing [10] and the InfoSec Council's Hard Problems List. Provenance is challenging because it's difficult to collect and is often incomplete, it comes from untrusted and unreliable sources, it's heterogeneous and non-portable and it requires confidentiality and privacy insurance. In this paper, we present information provenance challenges in the cloud and how it can serve to enhance information integrity. Especially, we focus on services for electronic documents and more precisely storage and archival services.

Traditionally organizations and companies buy storage systems for each of its sites or acquire equipment for its data center locations. Generally, this represents enormous costs especially when we take into account the cost of managing, supporting, and maintaining these systems. For these reasons, outsourcing storage and archival to external clouds is a good alternative to reduce costs. It permits also to discard the hard task of locally managing and maintaining storage systems according to legal compliance and security recommendations [11]. However, this leverages many issues about probative value and implies an immediate need to information provenance and lifecycle.

This paper is organized as follows. Starting from a case study, section 2 illustrates provenance needs and constraints for documentary clouds. The third part presents provenance as a cloud challenge. Section 4 presents provenance management approaches, some related works and their lacks. In the last section we conclude and present our future works.

2 Introducing Provenance Constraints

2.1 Context: A Bank Record Case Study

We are going to present a case study illustrating information probative value and provenance challenges for electronic documents. This case study comes from Novapost (www.novapost.fr) which is a French company specialized on providing collect, distribution and archiving services for electronic documents (especially human resources documents like pay slips, or personal confidential documents like bank records). Novapost aims to consolidate the existing services by providing dematerialization and long term archiving solutions allowing to improve archive management (formats migration, information integrity, document lifecycle management, interoperability and reversibility issues), ensure compliance with regulations relative to the probative value of electronic documents, ensure continuity and reliability in the management of information lifecycle and its provenance and guarantee seamless integration with customers through SaaS (Software as a Service) solution provided in partnership with IBM and based on innovative technology and highly secure infrastructure.

The case study scenario is the following: a customer of Novapost services (a bank) sends its customers bank reports to Novapost in AFP¹ format. Novapost

¹ AFP is a document format originally defined by IBM to drive its printers and support the typical form printing on laser printers.

documents service (provided as a cloud service) applies a specific processing depending on the received document type (splitting, templating, copying ...), then generates individual bank reports in PDF/A [22] format for archiving, interacts with external trust authority (signature service composed by a signature server and a timestamp server), communicates with IBM cloud for legal archiving to archive them in IBM's datacenters. The archived files are accessible to bank customers via web portal or a widget plugged on the bank website.

2.2 Provenance Constraints

As we have mentioned before, the heart of Novaposts solutions is document lifecycle management and documentary workflows requiring high probative value of the processed documents. The purpose of the R&D department in Novapost is to provide solutions to address the problems confronting dematerialization actors. We address the issues of ensuring electronic documents probative value, their lifecycle traceability as well as guaranteeing the readability and the sustainability of their processing tracks for long term.

In France, the probative value of electronic documents is subject to compliance with the conditions requested by the Article 1316-1 of the Civil Code mandating to identify the emitter persons/entities and to provide an integrity insurance through the whole document lifecycle. To ensure documents probative value and to perpetuate the informational capital of electronic archives, we need to keep a plausible trustworthy history of their provenance. This explains the strong requirements for traceability mandated by record management standards [15,17,19].

If we consider the standard lifecycle of a document, it includes steps of generation, additional processing, transfer and archiving. This means that document's technical integrity could change while its legal integrity is the same. At present, traceability management solutions for informational flows, especially documents are based on an accumulation of log files, coming from heterogeneous systems and multiple remote legal entities. Document provenance information's are distilled in multiple files that we need to cross to retrace its history. Add to that the sustainability of the responsible entities for the conservation of these trace elements is also problematic. So, provenance in the cloud is considered as a transverse problem, at the crossroad between different types of constraints. According to our study, provenance constraints can be categorized into legal, business and technical constraints (cf. Fig 1):

- Legal constraints: they are relative to the regulatory compliance about conservation duration of personal information as well as compliance with legislation governing information privacy like CNIL recommendation in France (www.cnil.fr) or the CDT in USA (www.cdt.org). Add to that the recent international law and the principle of territoriality which is not always in harmony with the distributed nature of the clouds over countries. For our

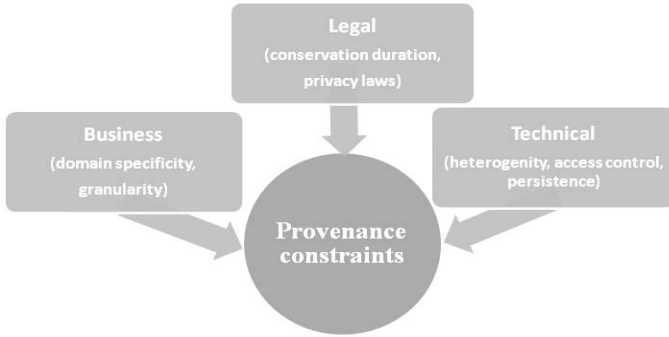


Fig. 1. Provenance constraints

case study for example, track elements are distributed across multiple systems and managed by multiple legal entities. We don't know what is the legal entity responsible for keeping these tracks neither where should they be kept and for how long.

- Business constraints: first of all, the word track is not generic, it depends on the business. For example, a track element in the bank system is different from a track element in the archival system. Also, the granularity of provenance depends on the context and the specificity of each cloud. Business constraints involves the definition of accessibility rules and profiles to provenance since confidentiality requirements imposes hiding the participation of some elements or some actors. These constraints are relevant to the definition of the suitable provenance collection and management policy to alleviate semantic interoperability issues.
- Technical constraints: they are relevant to provenance collection and management. It involves interoperability at the syntactic level between provenance records generated by heterogeneous systems, security issues including provenance integrity, access control and privacy. Performance, reliability and scalability constraints are also important because we are in a cloud context mandating from a provenance service to be scalable and available. Add to that trust and reliability about track elements provided by external systems/services during their creation or their transfer. In our case study, the identification and the crossing of log files about the processing of a document (known that document identifiers are not the same across different systems) is problematic and may rise probative value loss. Even if we can cross these elements, the heterogeneity of formats and data structures affects their readability. The other issue is guaranteeing confidentiality and privacy: an administrator of the archival system should not be allowed to access the contents of customer's bank records. Inversely, a customer should not access to the details about the systems that generates its bank record.

3 Challenges for Provenance in the Cloud

As illustrated by the case study, provenance becomes more and more important and challenging. The need for a provenance technology provided as an end-to-end service to guarantee information integrity and probative value is accentuated. Principally, the challenges can be divided into two categories. The first category includes known provenance challenges. It involves object identification and coupling with provenance, provenance reliability and confidentiality. It contains also the level of genericity of such a service. The second category contains new challenges which occur only in a cloud context and that are essentially relevant to performance, availability and scalability. Clouds are characterized by their ability to scale dynamically on demand. They can scale up and down according to the workload. So, how can we define a provenance service in the cloud and what are its characteristics? Would it be possible to use existent cloud computing frameworks like Hadoop (<http://hadoop.apache.org>) or more high level frameworks like Chukwa (<http://hadoop.apache.org/chukwa>) or Pig (<http://hadoop.apache.org/pig>) to meet high scalability and fault tolerance characteristics?

3.1 Identification and Mapping Digital Objects with Their Provenance

The first question is how to identify objects flowing between different clouds? The problem is that every cloud has its own policy to identify objects and that these policies are generally confidential. The challenge is that the technical integrity of a numeric object can change whereas its legal integrity remains the same. For these reason, traditional identification techniques based on hash computing are not suitable and cannot provide a unique object identifier. The second question is how to keep the link between a digital object and its provenance information along the different lifecycle phases and how to ensure provenance persistence? Persistence is the ability to provide an object provenance even if the object was removed. This data independent persistence is mandatory because in many cases where the ancestor object is removed, the provenance graph will be disconnected and the descendents objects provenance will be undetermined. So how provenance-data mapping should be managed? Would it be possible to consider provenance as an intrinsic metadata of each object or it will be better to consider the document and its provenance separately and just keep a descriptor referencing the object and its provenance.

3.2 Trust and Reliability

In the cloud, information crosses the traditional boundaries and passes through untrusted clouds. Generally, clouds are untrusted since the guarantee provided regarding the origin and the transformations are minimal, unclear or unreliable. So provenance can be forged or tampered. Having a tempered or low quality provenance can be worse than having no provenance at all. To ensure trustworthy provenance, we need to provide the mandated security elements [15,19]:

- Integrity: the assurance that provenance data is not tampered. Any forgery or tampering in provenance records will be detected.
- Availability: an auditor can check the integrity and the correctness of provenance information.
- Confidentiality: provenance information should not be public, only authorized authorities can access to it. This means that systems must have the ability to restrict views of data and its provenance according to different confidentiality levels.
- Atomicity: at storage time, both provenance and data should be stored or neither should be stored.
- Consistency: at retrieval time, data returned should be consistent with provenance.

3.3 Heterogeneity and Granularity

As illustrated by the use case, provenance is generated and managed by different entities according to different policies. It is often heterogeneous and not interoperable. Depending on the context, provenance needs and capabilities differ from one cloud to another and can't be exactly the same. The problem is that there is no portable and standard plausible one format fits all provenance that can be wired into general-purpose systems. Provenance should be tracked at different levels of granularity by the use of a flexible granularity approach. Such an approach should be based on a standardized provenance format that can be extended to a coarse-grained format or reduced to a fine-grained format. In this scope, a standard format for provenance called Open Provenance Model (OPM) [16,23] was proposed to the scientific community. This model standardizes provenance syntax and proposes a vocabulary for provenance. However, it presents security lacks because traditional access control techniques don't apply to provenance graphs and that the adaptation of existing access control models to provenance is insufficient [3]. That's why we think that the current OPM model is not sufficient for covering provenance in the cloud needs, especially security and privacy needs.

3.4 Cloud Constraints: Extensibility, Availability and Scalability

The cloud is designed to be scalable and available on demand. Existing provenance solutions in other environment (workflow, database...) dont consider the availability or scalability in their design. In addition, cloud is not extensible and couldn't be modified or extended. Furthermore, the availability of the provenance service has to match the high ability of the cloud. If it's not the case, the overall availability is reduced to the limited availability of the provenance service. Regarding scalability, the use of a database can make provenance queriable but not scalable. This limitation is due to updates synchronization between clients. These updates can cause a distributed service lock introducing a distributed deadlock or a scalability bottleneck due to a single lock. In this case, the need is to communicate with a cloud database because the use of a parallel

database is hard to maintain and is in contradiction with clouds model. Storing the provenance in a separate service introduces the issue of coordinating updates between object store (documents) service and the database service .

4 Provenance Approaches and Related Works

4.1 Provenance Approaches

To address the aforementioned challenges, we present in this section approaches for provenance management. We have identified three types of approaches for provenance collection and management:

1. Centralized approach (cf. Fig 2): it consists on the definition of a centralized policy for collecting and managing provenance from end-to-end. In this case, a unique trusted authority provides a provenance service. This approach permits to alleviate data structure heterogeneity by using a unique standardized provenance format. However, we have to imagine that this authority can access external clouds through specific API. This authority should provide strong authentication techniques permitting to its clients to trust this centralized service and to communicate without revealing critical information.
2. Hybrid approach (with import/export feature): this approach allows the import and the export of provenance between a centralized provenance service and external clouds (cf. Fig 3). The import feature allows retrieving the requested information without an additional management effort because it's supposed that the imported information is directly exploitable. In the same way, the centralized authority can directly export exploitable provenance to external services requesting it. Within this approach, the provenance authority should define a management policy to alleviate syntactic and semantic interoperability issues.
3. Federated approach (cf. Fig 4): in this approach provenance management responsibilities are delegated to each participant cloud which should handle it autonomously. Provenance is distilled through multiple files that we need to

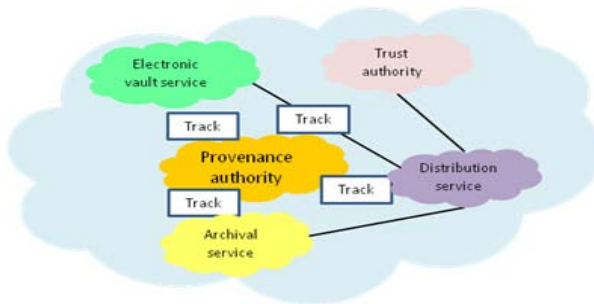


Fig. 2. Managing provenance within a centralized approach




Fig. 3. Managing provenance within a hybrid approach




Fig. 4. Managing provenance within a federated approach

cross to produce the provenance chain. Within this approach, limits are heterogeneity between provenance formats coming from different heterogeneous entities implying semantic and syntactic interoperability issues. To have a probative value, it is necessary that provenance elements can be folded and be consistent at both ends of a transaction. Since one of the parties don't meet these conditions, the whole provenance chain becomes suspicious. More importantly, the accuracy of information implies a need for corroborating provenance elements between two clouds. This limits evidence character to transactions where only two players are concerned.

After the identification of the different constraints and the presentation of the approaches, we have achieved the following classification according to the management approach (cf. Table 1). This table illustrates that the challenge level about trust and reliability is high for the three approaches because provenance access control is always challenging since traditional access control techniques are not relevant for provenance DAG [3] and that provenance tempering is possible even within centralized approach. In the same way, managing cloud constraints has always the level (H) in the three approaches.

Table 1. Classification of provenance constraints

Constraint/ Approach	Identification Mapping	Trust / reliability	Heterogeneity / granularity	Cloud specific constraints
Centralized	L	H	L	H
Hybrid	H	H	M	H
Federated	H	H	H	H

Challenge level: - High: H
- Medium: M
- Low: L

4.2 Related Works

Previous researches on provenance tackles principally the problem from one of two sides: database and files management or workflow and grid environments. Works on database and files provenance [2,4,21] focused on data- driven business processes and relies on the ability to trace information flows through SQL manipulations or by increasing file headers and tagging them [1,20]. On the other side, workflow provenance research [8,9] tackles scientific processing (like scientific workflows in collaborative environments for biologists, earth scientists...) and deals with less transparent but predefined process executions. Many other works on provenance addresses specific problems like provenance in software development known as versioning systems which manages distributed intra-files provenance but are not able to manage inter-files provenance. We also mention recent work on securing provenance from tempering based on cryptographic mechanisms [12].

The particularity of the provenance needs in our work is the superposition of a set of traditional provenance constraints with cloud specific constraints. Most of these works assumes the ability to alter the underlying system components, and dont propose a native provenance infrastructure addressing the problem as an end-to end property corresponding to all the aforementioned constraints.

5 Conclusion

In this paper we have presented a pragmatic case study illustrating the need for provenance to ensure electronic documents probative value. It allowed us to illustrate and analyze the challenges and to categorize the inherent constraints motivating our research on information provenance.

Provenance is becoming more and more important as information is shared across networks and exceeds internal boundaries. Actually, the most used medium for information sharing, the web does not provide this feature. But in the cloud, which is not yet a mature technology, provenance issues can be addressed from now.

Provenance in the cloud is an open research problem and our future work will be focused on provenance deployment as a service in the cloud. We aim to propose a centralized provenance service model for the cloud. This model should

be generic and extensible. It should provide an abstraction of application domain and accept proprietary extensions to cover the specifics of heterogeneous business processes. In this model, we aim to address security, privacy and access control issues. Regarding security issues, we want to introduce semantics to enhance provenance confidentiality. Regarding privacy, we wish to develop techniques assuring information usefulness in respect with privacy needs without any leakage. This requires the definition of new access-control techniques and data access schemas for provenance graphs. In the future, we think that the challenge for provenance systems is to introduce remembrance which is a kind of an RFID for all types of digital objects. This concept was introduced in [13], it consists in augmenting all data objects with persistent memory and to consider that provenance is an intrinsic property of data objects. Remembrance can be very interesting in cloud computing scenarios and can immediately show that a piece of data is originated from an untrusted object or application. It exceeds traditional versioning and snapshots systems and tries to give a complete end-to-end history as an atomic property of any digital object. To be useful and exploitable, remembrance should be coupled with strong privacy and access control techniques.

References

1. Agrawal, P., Benjelloun, O., Sarma, A.D., Hayworth, C., Shubha, U., Nabar, C.U., Sugihara, T., Widom, J.: ULDBs: Databases with Uncertainty and Lineage. In: *Trio: A System for Data, Uncertainty, and Lineage*. VLDB 2006, pp. 1151–1154 (2006)
2. Bhagwat, D., Chiticariu, L., Tan, W.C., Vijayvargiya, G.: An Annotation Management System for Relational Databases. In: VLDB, pp. 900–911 (2004)
3. Braun, U., Shinnar, A., Seltzer, M.: Securing provenance. In: *Third USENIX Workshop on Hot Topics in Security (HotSec)* (July 2008)
4. Buneman, P., Khanna, S., Tan, W.C.: Why and Where: A Characterization of Data Provenance. In: Van den Bussche, J., Vianu, V. (eds.) *ICDT 2001*. LNCS, vol. 1973, pp. 316–330. Springer, Heidelberg (2000)
5. Cameron, G.: Provenance and Pragmatics. In: *Workshop on Data Provenance and Annotation* (2003)
6. Cheney, J., Chong, S., Foster, N., Seltzer, M., Vansummeren, S.: Provenance: A Future History. In: *International Conference on Object Oriented Programming, Systems, Languages and Applications*, pp. 957–964 (2009)
7. Da Silva, P.P., McGuinness, D.L., McCool, R.: Knowledge Provenance Infrastructure. *IEEE Data Engineering Bulletin* 26, 26–32 (2003)
8. Davidson, S., Cohen-Boulakia, S., Eyal, A., Ludascher, B., McPhillips, T., Bowers, S., Freire, J.: Provenance in Scientific Workflow Systems. *IEEE Data Engineering Bulletin* 32, 44–50 (2007)
9. Goble, C.: Position Statement: Musings on Provenance, Workow and Semantic Web Annotations for Bioinformatics. In: *Workshop on Data Derivation and Provenance* (2002)
10. Grand challenges in computing research conference, UK Computing Society (2008), <http://www.ukcrc.org.uk/press/news/challenge08/gccr08final.cfm>

11. Hasan, R., Yurcik, W., Myagmar, S.: The Evolution of Storage Service Providers: Techniques and Challenges to Outsourcing Storage. In: Proceedings of the 2005 ACM workshop on Storage Security and Survivability (2005)
12. Hassan, R., Sion, R., Winslett, M.: Preventing History Forgery with Secure Provenance. *ACM Transactions on Storage* (2009)
13. Hassan, R., Sion, R., Winslett, M.: Remembrance: The Unbearable Sentience of Being Digital. In: Fourth Biennial Conference on Innovative Data Systems Research (2009)
14. INFOSEC Research Council (IRC) Hard problem list. Technical report (November 2005), http://www.cyber.st.dhs.gov/docs/IRC_Hard_Problem_List.pdf
15. ISO 14721:2003. Space data and information transfer systems - Open Archival Information System Reference model (OAIS), <http://www.iso.org>
16. Miles, S., Groth, P.T., Munroe, S., Jiang, S., Assandri, T., Moreau, L.: Extracting causal graphs from an open provenance data model. *Concurrency and Computation: Practice and Experience* 20(5), 577–586 (2008)
17. MoReq2 specifications. Model Requirements for the management of electronic records Update and Extension (2008), <http://www.moreq2.eu>
18. Muniswamy-Reddy, K.K., Holland, D.A., Braun, U., Seltzer, M.: Provenance-aware storage systems. In: USENIX Annual Technical Conference, General Track, pp. 43–56 (2006)
19. NF Z42-013. Electronic archival storage-Specifications relative to the design and operation of information processing systems in view of ensuring the storage and integrity of the recording stored in these systems, <http://www.boutique.afnor.org>
20. Sar, C., Cao, P.: Lineage file system. Technical Report (January 2005), <http://crypto.stanford.edu/~cao/lineage>
21. Simmhan, Y., Plale, B., Gannon, D.: A survey of data provenance in e-science. *SIGMOD Record (Special section on scientific workflows)* 34(3), 31–36 (2005)
22. ISO Standard for using PDF format for the long-term archiving of electronic documents ISO-19005-1 - Document management - Electronic document file format for long-term preservation - Part 1: Use of PDF 1.4 (PDF/A-1), <http://www.pdfa.org>
23. Moreau, L., Plale, B., Miles, S., Goble, C., Missier, P., Barga, R., Simmhan, Y., Futrelle, J., McGrath, R.E., Myers, J., Paulson, P., Bowers, S., Ludaescher, B., Kwasnikowska, N., Van den Bussche, J., Ellkvist, T., Freire, J., Groth, P.: The Open Provenance Model (v1.01) specifications. *Future Generation Computer Systems* (2009)