Provenance as a Security Control

Andrew Martin John Lyle Cornelius Namilkuo

Department of Computer Science
University of Oxford

(firstname.lastname@cs.ox.ac.uk)

Abstract

Much has been written about security and provenance. Although both have their own large areas of concern, there is a very significant intersection. One is often brought to bear upon the other, in the study of the security of provenance. We discuss through a series of examples how provenance might be regarded as a security control in its own right. We argue that a risk-based approach to provenance is appropriate, and is already being used informally. A case study illustrates the applicability of this line of reasoning.

1 Introduction

Much has been written about security and provenance. Although both have their own large areas of concern, there is a very significant intersection.

Classic conceptions of information security refer to it comprising *confidentiality*, *integrity*, and *availability*. Provenance is very strongly bound up with the *integrity* part, so much so that many of the high-level concerns of security are often almost indistinguishable from provenance concerns. For example, the question of whether (or how) the content of a particular electronic document has changed since its creation is very much an integrity question in the security world; it is equally well a proper question for the analysis of the document's provenance.

In developing tools and reasoning methods for approaching these topics, the fields diverge somewhat. Here, however one is often brought to bear upon the other, in the study of the security *of* provenance. This is an important theme in its own right [1, 4], and also gives rise to studies in means of achieving security of provenance in a variety of other contexts, including databases [13], cloud computing [6, 8], and SOA [12]. Clearly, the value of provenance data varies as the value of the primary data varies, and increasingly valuable digital assets will require increasingly strong provenance security in this way.

Our purpose here, however, is to discuss how provenance might be regarded as a security control in its own right. Formal models for this have been described by Cheney [2], and provenance has been described as a security measure in some particular application domains [10]. Our perspective is that such controls are often important *even when* the provenance data is itself poorly-controlled, and lacks the integrity assurances or formal structures which might be desirable. A strong theme in both the study of security and of provenance would be that these disciplines, when skilfully applied, relate to the "art of the possible" rather than any sense of perfection or an ideal solution.

Section 2 describes the nature of security controls, and how provenance may find a home among them; Section 3 offers a small case study in this area, and Section 4 presents some tentative conclusions.

2 Appropriate Security Controls

It is well-known that 'perfect' security in modern software systems is unachievable. As a result, means of determining what is 'good enough' is absolutely essential, if an unbounded amount of effort is not to be expended in the protection of security. When viewed as part of the security analysis, decisions about the collection of provenance (both its extent, and the controls upon its quality and accuracy) must be subject to the same judgement.

Pragmatic security analysis — particularly in the business world — is frequently founded upon the concepts of ISO 27002 [5] wherein the notion of a *control* is described as a *means of managing risk, including policies, procedures, guidelines, practices or organizational structures*. Such a notion of controls encompasses tools as diverse as operating systems access permissions and the vetting of critical staff prior to employment. The selection of such controls is necessarily based on an analysis of risk, in a form of likelihood-weighted cost-benefit analysis: risk combines an estimate of the likelihood of a (unwanted) event with the cost (or impact) of its consequence.

Such an approach sits in some contrast to the world of cryptography, where the notion of *work factors*, calculations based on information theory and entropy, and reduction to known hard problems, guides the development of cryptosytems.

Provenance clearly has potential as a security control — most chiefly in the protection of data integrity, by providing mechanisms through which unwanted or unexpected changes to data (or software, or other items) can be detected (and hence discouraged). The selection of tools and techniques for provenance, the amount and extent of the provenance, and so on — is naturally the subject of risk analysis. The cost of collecting, storing, and processing the provenance will be weighed against the benefit gained from it (the extent to which 'bad' or missing provenance is avoided, and the extent to which a bad decision-making process is avoided). The security controls which need to be applied *to* that provenace similarly must be determined by the threats (and related risks) which arise in its context of use.

As with many security controls, we must be aware that instrumentation of a system in this way is usually subject to feedback: the socio-technical system responds to a change in its dynamics. Increased reliance upon provenance data will inevitably motivate fresh attacks. Therefore, although we will generally assume that for security some provenance is better than none, we must be aware of the limitations of this view, and the danger of a false sense of security.

Although our main focus is on provenance for the protection of integrity, in areas such as *post hoc* access control, provenance helps to protect confidentiality also, as we discuss below. Moreover, loosely-coupled provenance graphs can assist in the design of systems for fault tolerance — and so promote the availability security property also. The latter also serves to illustrate that we have in mind the full richness of *provenance*, rather than 'mere' logging solutions.

2.1 Status Quo

To the immense chagrin of those who look for massive work factor advantage over their adversaries, 'real world' security often adopts a distinctly pragmatic stance, with few strong guarantees. A commonplace example is reliance upon plain email as a form of evidence, even sometimes in courts of law. Email headers may carry elements of provenance — identities of senders, servers, and intermediate delivery points — all of these are readily forged (as is the message content). Perhaps the forger must take care that the headers have not (or, even, the entire message has not) been logged by one of the intermediate servers — and today's regulatory compliance might make this an unwelcome bet. The diverse and loosely-linked forms of provenance serve to strengthen the over all security.

This use of email only throws into relief a range of other poorly-evaluated sources of provenance in everyday life — letters on 'official letterhead', a modern arms-race with degree certificates and examination transcripts, and more. The security of such provenance is generally poor, but it is nevertheless *good enough* (wisely or otherwise) for a variety of security-critical applications. The risk-based approach asks not how idealized provenance can be collected and processed, but how achievable changes in design and practice can lead to reduced risk.

2.2 Kinds of access controls

Provenance is most clearly related to the *integrity* property, but in access control — which may cover *read* as well as *write* and other methods — it plays a part in the protection of confidentiality.

Many security decisions are necessarily made 'live' — that is, synchronously — on the basis of evidence at hand at the time of the decision measured against a carefully-planned static policy. This is appropriate in cases where a failure to control the distribution of data (or to protect its integrity) will lead quickly to a catastrophic outcome. In the situations described later, social pressure or the threat of retribution helps to prevent the catastrophe, but for some scenarios (such as the protection of certain national secrets) that threat may be insufficient, or inefficient.

It is well-known that in other cases, a highly permissive policy makes more sense. In retail banking, processing of identity documents (passport office etc.), and in emergency medicine, it is most practical to give all authenticated users access to all records. In this instance, it is necessary to keep audit logs of all such accesses, and use anomaly detection either (a) to aim to detect every unusual behaviour, or (b) to do enough checks that users have an incentive (strong likelihood of discovery, and then punishment) to act in accordance with policy.

A middle path lies with 'break the glass' security policies [9] — where a static policy controls everyday behaviour, but trusted individuals are permitted to invoke an emergency override. The override comes with an obligation/overhead of enhanced logging and audit - so that the actions taken within that mode can be studied carefully, and individuals given an incentive to act appropriately.

In the latter two cases, logging and audit is an essential security control. This is problematic if every fine-grained access is logged: the volume of data becomes large, as does the privacy overhead. We see here the merit of semantically rich provenance information, over and above simplistic logs. Knowing the *context* of an access, and the use to which the data is put (whether in a query, consolidated report, individual record, physically printed or emailed, etc.) serves to eliminate false positives.

2.3 Scientific and Workflow integrity

Much of the community's interest in provenance arises of course in the collection and processing of scientific data — particularly in workflows and in database processing. Provenance plays two obvious (and related) security roles here: one in helping to prevent (or discover) scientific fraud; the other in enabling records to satisfy with integrity the requirements of regulatory authorities.

Although a design objective is often to make the collection as painless as possible for the user through instrumenting systems for transparent collection, few of these means of collecting provenance are tamper-proof. We may see here a clear application (perhaps unwittingly) of a risk-based approach: the anti-circumvention measures are just as strong as they need to be. Other work has described how those measures might be considerably strengthened through the use of technologies such as trusted computing [7].

Where provenance information is attached to billing for utilities — phone, electricity, etc. — we see the beginnings of an approach to *non-repudiation*, a secondary form of integrity and one of the more difficult general problems in information security. Where one party wishes to claim that a particular event did (or did not) take place, additional provenance information about the context (the configuration of the relevant systems at the time, the preceding and following events, and more) may be invaluable — even if falling short of 'proof'.

Many modern security controls themselves depend today on collection and processing of log data. It may be a stretch to describe this as provenance-based security, but approaches such as IF-MAP [11] are beginning to illustrate the value of semantically-rich multi-faceted data in making decisions around intrusion detection and reaction, for example.

2.4 Controlling for Privacy

Ensuring privacy is a common application of security controls. One of the key requirements in some concepts of privacy — including that embodied in the European Data Protection Directive [3] — is the control of *disclosures* of personal information. Using the techniques of provenance allows appropriate tracking of the situations in which personal information has been disclosed, and so provenance as a control enables compliance with the law in this setting.

Plainly, privacy gives rise to a significant countervailing concern: it is widely accepted that the best means to protect privacy is to collect and retain *as little data as possible*. No matter how strong the privacy controls, errors and flaws will lead to unwanted disclosures. Moreover, rich provenance information will frequently enable the actors to be identified, even if their explicit identities are hidden or discarded. Privacy is a common application of security technologies, and this tension with the collection of provenance seems unavoidable: it is clearly a necessary topic for further study. In our present context, we may observe that where the collection of additional provenance information entails more processing of personal data, the *increased* use of provenance may lead to an *increased* information security risk, in contrast to the general case.

3 Case Study: Forming and Maintaining a Relationship

Consider the common situation where collaboration occurs across organisational boundaries — perhaps for the creation of an *ad hoc* project, or during the due diligence process prior to an acquisition or merger, or in an emergency. The common security mechanisms which might support such ventures are, by their nature, heavyweight and hard to deploy quickly and for a limited period. Harmonising policies to the extent necessary to create mutually-trusted authentication, say, is observed to be enormously time-consuming. Few security mechanisms work well without a single master control point; a 'root of trust'. Access control based upon attributes possessed by the actors becomes untenable if the assertion of those attributes is not authoritative.

Such collaborations often work in practice due to relatively straightforward agreements — legal documents — and goodwill on the part of the participants, backed by simplistic record keeping. Large-scale activities present a challenge because the scale is not well-supported by such simple controls.

The techniques of provenance — perhaps incomplete, varying in quality and semantics, but with well-defined procedures behind them — have much to offer in this situation. As we noted in Section 2.2 access control can be effected in many cases (especially when parties are intending to collaborate) through record-keeping, recording context, and keeping account not merely of access to data, but of what happens later (disclosures and further processing). If parties to a collaboration agreement trust one another's provenance recording and reporting (in whole or in part) then they can work together despite the absence of an over-arching domain controller or elaborate enterprise rights management system.

As with much else of what we have described, such an approach must be subject to risk assessment. If participants in the collaboration have too great an incentive to fail to play by the rules, the relatively weak collection of provenance may not save the situation. On the other hand, if the risk profile is such that the majority of actors will be conscientious, but simply need to give evidence of compliance to another party, well-placed provenance collection may be completely suitable. This is noteworthy because the integration of strong security systems (and indeed, IT systems in general) is known to give rise rather often to collaboration failures.

Related comments apply in a variety of other federated IT systems — whether or not multiple actors are involved. One may desire a unified security solution to span a variety of cloud and web service endpoints, but reality is heterogeneous. Semantically rich provenance has potential to unify the security controls in a suitably loose way — a control suitable to the realistic risk in many contexts.

4 Conclusion

We have argued that in most contexts, when security matters, some provenance is better than none. A risk-based approach to provenance appears appropriate — stronger provenance controls (those which are harder to circumvent, harder to spoof) are generally to be preferred over weaker ones, but even the weakest ones appear to have considerable value. The security *of* such controls may be strengthened by a variety of technical means, which are explored in the literature. Even in the absence of such strengthening, provenance has considerable value as a security measure, playing a role in protecting *both* integrity and confidentiality.

We have observed that many security controls are based upon 'live' (or synchronous) decision-making, based on instantaneous assessments of a system's state. That state information plays a part in protecting the integrity of data; its own integrity is often harder to assure, and in turn this depends on the integrity of the systems and software which created them. Such second-order and third-order concerns are (ironically) difficult to handle in a security context and are frequently simply *assumed*, or assured by alternative means such as separate architectural or design information. A risk-based approach to provenance offers a much more integrated or holistic conceptual view of such issues — and lends itself to use as a security control for dynamic heterogeneous systems.

5 Acknowledgments

Parts of this work are supported by the *Trust Domains* project funded by UK TSB and EPSRC, reference TS/I002634/1. We are very grateful for the comments of the anonymous referees, whose advice we have imperfectly attempted to follow.

References

[1] Braun, U., Shinnar, A., and Seltzer, M. Securing provenance. In *HOTSEC'08: Proceedings of the 3rd conference on Hot topics in security* (Berkeley, CA, USA, 2008), USENIX Association, pp. 1–5.

[2] Cheney, J. A formal framework for provenance security. In *Computer Security Foundations Symposium (CSF), 2011 IEEE 24th* (June 2011), pp. 281 –293.

[3] European Parliament and Council. Directive 95/46/EC on the protection of individuals with regard to the processing of personal data and on the free movement of such data, Oct. 1995.

[4] Hasan, R., Sion, R., and Winslett, M. Introducing secure provenance: problems and challenges. In *StorageSS '07: Proceedings of the 2007 ACM workshop on Storage security and survivability* (New York, NY, USA, 2007), ACM, pp. 13–18.

[5] ISO. Information technology — security techniques — code of practice for information security management. International Standard ISO27002, ISO, 2005.

[6] Lu, R., Lin, X., Liang, X., and Shen, X. S. 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* (New York, NY, USA, 2010), ASIACCS '10, ACM, pp. 282–292.

[7] Lyle, J., and Martin, A. Trusted computing and provenance: better together. In *Proceedings of the 2nd conference on Theory and practice of provenance* (Berkeley, CA, USA, 2010), TAPP'10, USENIX Association, pp. 1–1.

[8] Muniswamy-Reddy, K.-K., Macko, P., and Seltzer, M. Provenance for the cloud. In *FAST '10: Proceedings of the 8th USENIX conference on File and storage technologies* (2010), USENIX, pp. 15–14.

[9] Rissanen, E., Firozabadi, B., and Sergot, M. Towards a mechanism for discretionary overriding of access control. In *Security Protocols*, B. Christianson, B. Crispo, J. Malcolm, and M. Roe, Eds., vol. 3957 of *Lecture Notes in Computer Science*. Springer Berlin / Heidelberg, 2006, pp. 312–319.

[10] Sultana, S., Bertino, E., and Shehab, M. A provenance based mechanism to identify malicious packet dropping adversaries in sensor networks. In *Distributed Computing Systems Workshops (ICDCSW), 2011 31st International Conference on* (june 2011), pp. 332 –338.

[11] Trusted Computing Group. TNC IF-MAP Binding for SOAP. Specification, TCG, July 2010.

[12] Tsai, W., Wei, X., Chen, Y., Paul, R., Chung, J.-Y., and Zhang, D. Data provenance in soa: security, reliability, and integrity. *Service Oriented Computing and Applications 1* (2007), 223–247. 10.1007/s11761-007-0018-8.

[13] Zhang, J., Chapman, A., and LeFevre, K. Do you know where your data's been? - tamper-evident database provenance. In *Secure Data Management* (2009), W. Jonker and M. Petkovic, Eds., vol. 5776 of *Lecture Notes in Computer Science*, Springer, pp. 17–32.