Easy Access to Provenance: an Essential Step Towards Trust on the Web

Tom De Nies, Sam Coppens, Ruben Verborgh, Miel Vander Sande, Erik Mannens, Rik Van de Walle, Danius Michaelides, and Luc Moreau

In: METHOD 2013: The 2nd IEEE International Workshop on Methods for Establishing Trust with Open Data, 2013.

Easy Access to Provenance: an Essential Step Towards Trust on the Web

Tom De Nies*, Sam Coppens*, Ruben Verborgh*, Miel Vander Sande*,
Erik Mannens*, Rik Van de Walle*, Danius Michaelides†, Luc Moreau†

*Ghent University - iMinds
Electronics and Information Systems, MMLab
Gaston Crommenlaan 8 bus 201
B-9050 Ledeberg-Ghent, Belgium
{tom.denies,sam.coppens,ruben.verborgh,
miel.vandersande,erik.mannens,rik.vandewalle}@ugent.be

†University of Southampton
Electronics and Computer Science
Southampton SO17 1BJ
United Kingdom
{dtm,l.moreau}@ecs.soton.ac.uk

Abstract

This paper describes the mechanisms involved in accessing provenance on the Web, according to the new W3C PROV specifications, and how end-users can process this information to make basic trust assessments. Additionally, we illustrate this principle by implementing a practical use case, namely Tim Berners-lee's vision of the "Oh, yeah?" button, enabling users to make trust assessments about documents on the web. This implementation leverages the W3C PROV specification to provide user-friendly access to the provenance of Web pages. While the extension described in this paper is specific to one browser, the majority of its components are browser-agnostic.

1. Introduction

In the research community, provenance has always been viewed as an essential component of establishing trust for information resources. It includes all information about entities, activities, and people involved in producing a piece of data or thing [1]. Recently, PROV [1], the standard for provenance by the W3C Provenance Working Group, has been published as a Proposed Recommendation. This means that now, users and applications can produce and consume provenance in a standard model and an array of standardized serializations. In this paper, we present our ideas to enable making basic trust assessments of information on the Web, based on the availability of its provenance. Our main goal is to present the provenance of information on the Web in such a way that a non-expert user can easily understand it, and make a decision whether or not to trust the information.

The remainder of the paper is structured as follows. In Section 2, we provide a brief overview of related work on using provenance for trust assessments. Next, in Section 3, we describe how provenance can be accessed and retrieved in a standard way, followed by a description of the use of a validation service in Section 4, ensuring that the retrieved provenance complies with the standard. In Section 5, we explain how this information can be used to generate human-understandable trust assessments. In Section 6, we describe the implementation of a use case, incorporating these ideas into a browser extension. Finally, in Section 7, we discuss the remaining challenges and future work.

2. Related Work

In literature, a significant amount of work is available on provenance as well as trust assessment on the Web. Here, we will limit ourselves to those works describing a combination of both, as those are most relevant to our goals. For a detailed survey on trust in computer science and on the Web, we refer to [2].

Most approaches in literature agree that *reputation* is an essential component for making trust assessments. In [3], a system is proposed that generates recommendations for content, based on the trust a user has in the agents that produced and/or published the content. The added value of this system, is that it not only makes use of the general reputation of a person or organization, but also of the trust relationships in the

user's FOAF[1] profile, resulting in a very personalized trust assessment. The FOAF profile lists the user's contacts, and therefore provides some context for the trust assessment, as the user is likely to trust someone he or she considers a friend. This is an important consideration to keep in mind for future implementations of our proposed approach.

However, as stated in [4], reputation estimation alone is not enough. The authors of [4] present an approach to compute the trustworthiness of user-annotated tags in a video corpus, based on a combination of reputation and provenance specified in W3C PROV. Their main goal is to provide reasoning and information retrieval software with automatically generated information regarding the trustworthiness of data. Another combination of reputation and provenance is described in [5], where the trustworthiness of sensor network data is assessed based on the reputation of network nodes in the provenance trace of this data.

In [6], events are identified that invoke distrust for a user, and it is described how these events relate to the provenance of the distrusted information. In Section 5, we apply a similar logic by checking for indicators that might generate these distrust events in the PROV associated with a Web resource.

We observe that most these approaches focus on reputation of the source, in two cases in combination with provenance. To the best of our knowledge, there is currently no system that directly maps the level of availability of provenance associated with information on the Web to a trustworthiness level. This is why we propose our approach, in which we access the provenance of a Web page, and make trust assessments based on this information.

3. Accessing Provenance

The Provenance Working Group has published a note, stating the recommended methods to associate provenance to a document. For the full specification, we refer to PROV-AQ [7]. Here, we will summarize the essentials.

In the specification, three mechanisms are proposed for a provenance provider to supply information that may assist a provenance consumer to locate the provenance descriptions related to a document: the HTTP *Link* header, the HTML *<link>* tag, and *RDF(a)*. Provenance descriptions for a resource can be provided in two ways: either by using a *provenance resource* that contains a set of provenance descriptions about the resource, or by using a *provenance query service*, where provenance for the resource can be retrieved.

1. Friend Of A Friend: http://www.foaf-project.org/

For a resource accessible using HTTP, the provenance descriptions may be linked from the Link header included in the HTTP response to a GET or HEAD request, as specified in [8]. The `has_provenance` and/or `has_query_service` link relation types may be used, as follows:

```
Link: <provenance-URI>; rel=
 "http://www.w3.org/ns/prov#has_provenance";
 anchor="target-URI",
 <provenance-service-URI>; rel=
 "http://www.w3.org/ns/prov#has_query_service";
 anchor="target-URI"
```

Here, the `provenance-URI` is used to indicate the provenance descriptions associated with the document, in which the document itself is referred to as `target-URI`. If no anchor parameter is provided, the `target-URI` is assumed to be the URI of the requested resource in the HTTP request. The `provenance-service-URI` refers to a service description that provides the consumer with the necessary information to submit a query to retrieve the provenance descriptions for the `target-URI`. Multiple `has_provenance` Link header fields are permitted.

For resources represented as HTML, a provenance resource may be linked to by appending a `<link>` element to the HTML `<head>` element of the document. Three link relation types are defined: `has_provenance`, `has_anchor` and `has_query_service`. The `provenance-URI`, `target-URI` and `provenance-service-URI` have the same meaning as specified above.

```
<html>
<head>
 <link href="provenance-URI" rel=
 "http://www.w3.org/ns/prov#has_provenance">
 <link href="target-URI" rel=
 "http://www.w3.org/ns/prov#has_anchor">
 <link href="provenance-service-URI" rel=
 "http://www.w3.org/ns/prov#has_query_service">
</head>
<body> ... </body>
</html>
```

Finally, a resource identified by a `resource-URI` and represented as RDF (in any syntax, including RDFa) may contain triples that relate the resource to its own provenance. Therefore, the link relations `has_provenance`, `has_anchor` and

has_query_service may also be used as RDF properties to create these triples.

```
@prefix prov: <http://www.w3.org/ns/prov#>.
<resource-URI>
 prov:has_provenance <provenance-URI>;
 prov:has_anchor <target-URI>;
 prov:has_query_service <provenance-service-URI>;
```

In our use case (described in Section 6), a provenance extractor searches the document the user is viewing for provenance URIs specified using these three methods, and aggregates the results into a list of provenance resources. Note that currently, no provenance query services are available, which could be used to retrieve provenance from an external source. However, once they do become available, supporting query services is a matter of searching for these query service URIs in addition to provenance URIs. Since these query services will build up a reputation over time, additional trust assessments regarding the provenance retrieved through them will become possible.

4. Validation

In order to make easy, quick assessments regarding the trustworthiness of a document on the Web, a user needs more information than just the location of its provenance resource(s). The user needs to know whether the specified provenance resources actually exist, who created them, and whether they can be considered valid. While trust cannot directly be derived from it, validation of provenance does provide the user with an indication that the asserter of the provenance put effort into remaining compliant with the standard, and that the provenance is at least more likely to be trustworthy than invalid provenance.

4.1. URI Existence and Source

The first step in the validation process is to check whether the linked provenance URIs actually exist and if the provenance resources can be retrieved. This is done by sending an HTTP request for each provenance URI. If the headers of the specified provenance URI can be retrieved, the URI exists and is passed to the validator; otherwise, the URI is flagged as non-existent. Additionally, the location of the provenance is compared to the location of the original document, and it is stored regardless if they are the same or not. This happens because a user might prefer to trust a document whose provenance is stored in a trusted repository, instead of at the same location as the document itself.

4.2. PROV Validator

The concept of valid provenance is defined in PROV Constraints [9]. This document states that valid PROV descriptions satisfy certain definitions, inferences, and constraints to provide a measure of consistency checking and reasoning over provenance. While validation is no guarantee for trustworthiness, it does guarantee that the supplied provenance can be consumed by all applications compliant to the standard, and therefore, it is a valuable property.

PROV-Constraints defines 56 distinct definitions, inferences, and constraints. This, in addition to the various PROV serializations, makes implementing a validator for PROV a non-trivial task. A comprehensive, publicly available validation service[2] was developed. We will not discuss the details and inner workings of this validator here, and describe the use of its public API instead.

The API for the validator is used by sending an HTTP POST request to http://provenance.ecs.soton.ac.uk/validator/validation/submit with the following parameters:

```
validate : 'Validate'
url : <provenance-URI>
file : <file upload>
statements : <the provenance statements>
```

Note that for our use case, only the fields validate and url will be set. No content type is specified. Content negotiation allows the following PROV representation types to be validated: text/turle, text/prov-notation, rdf/xml, application/provenance+xml and text/json. When the process is complete, the validator refers to an XML document, which contains the validation result. This XML document contains a child element for each validation error, and two additional elements for the provenance of the validation result itself. This means that if no error elements are present, the provenance was valid, and it can be labeled as such. To help the user understand this validation result, we also link to the (temporarily available) detailed validation report at the validation website.

2. http://provenance.ecs.soton.ac.uk/validator/

5. Basic Trust Assessment

Providing users with access to the provenance of Web resources is an important step to allow them to make trust assessments, but this might be difficult for users who are not experts in the field of provenance or computer science in general. Therefore, an interpretation of the available information is presented to the user, in a way that he/she understands. Based on the trends in literature (as seen in Section 2), we define several criteria to make trust assessments based on provenance associated with a Web resource:

1) **availability**: whether there is provenance available for the resource;
2) **validity**: whether the provenance is well-structured and valid;
3) **provenance of provenance**: the source of the provenance, who asserted it, etc.;
4) **consistency**: whether the provenance is consistent with alternative sources;
5) **correctness**: whether the provenance corresponds with what's actually in the content;
6) **reputation**: the reputation of the agents and sources mentioned in the provenance.

Instead of returning a single trust score to the end-user, we choose to provide conclusions regarding each of these criteria to the end user. While a trust-score is valuable information for a machine agent making decisions on filtering or retrieving content, a human user might not understand the meaning of this score, and will possibly misinterpret it. Providing information to the user about each of the above criteria is aligned with the vision detecting distrust events, as described in [6], and will be more usable for non-expert users.

In our use case, described in Section 6, we provide information on criteria 1, 2, 3, and 6 to the user, because they are directly computable from the provenance associated with a Web resource using existing technology. Criteria 4 and 5 require more advanced processing of the content and provenance statements. Incorporating these two criteria is part of our future work, as mentioned in Section 7.

6. Use Case: The "Oh, Yeah?"-button

In 1997, Tim Berners-Lee proposed that each browser should have a button marked "Oh, yeah?" [10], that a user can press when he/she loses the feeling of trust when viewing a document. Upon pressing the button, information is shown about why the user should (dis)trust the document. In this section, we describe a browser extension that constitutes a first implementation of this "Oh, yeah?" button. When the button is clicked, the browser acquires the provenance resources linked to by the document the user was looking at and displays the URIs, whether the provenance is valid and whether the URI actually exists, in addition to a number of automatically derived statements regarding the trustworthiness of the document. The extension is available for download at the Google Chrome Web Store[3].

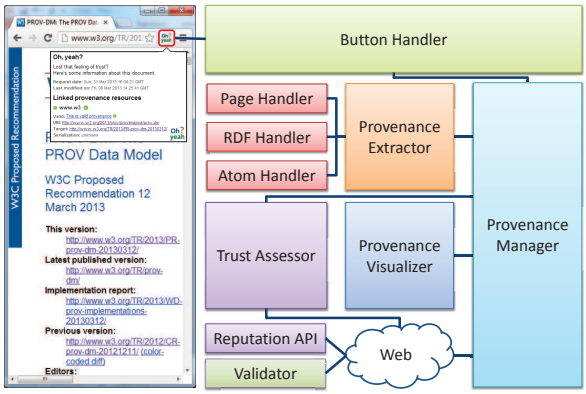

Figure 1. Overview of the browser extension

6.1. Overview of the Extension

In Figure 1, an overview of our application is shown. The "Oh, yeah?" button is located at the upper right corner of the browser window, and is only enabled when provenance is associated with the page being viewed. Upon pressing the button, the content and headers of the document are passed to the Provenance Manager, which then processed the information in four steps.

1) The Provenance Extractor extracts the provenance from the document, using a suitable method for each supported document type.
2) The linked provenance resources are fetched (if they exist), and validated using an web-based validation service.
3) All information is interpreted by the Trust Assessor, using a web-based reputation API.

3. http://chrome.google.com/webstore/detail/oh-yeah/nniblclgdefhcllookkjmfaoboggkcef

4) The results are summarized and visualized in a pop-up by the Provenance Visualizer.

In the next sections, we will explain each of these steps in detail.

6.2. Implementation Choices

Our application aims to bring provenance to a broad audience – not only to experts in the field of provenance – through a lightweight and understandable visualization. The application is written in Javascript, and therefore should be usable in most browsers. However, as explained in Section 3, provenance may be specified in the headers of HTTP requests, and our application must be able to intercept these requests. Therefore, we opted to build an extension specifically for one browser (Google Chrome), due to its easy access to Web requests (specifically, through the `chrome.webRequest` module). However, all other components of the application are browser-agnostic, and we plan to create extensions for Mozilla and Safari in the near future.

6.3. Criteria for Trust and Distrust

As mentioned in Section 5, our implementation of the "Oh, yeah?"-button considers four criteria from which a distrust event can be derived. The following rules are applied:

- **Provenance availability:** If provenance is linked to, the application checks the existence of the provenance URI(s) and relays this information to the user. If no provenance is linked to, the button is not activated.
- **Provenance validity:** If the provenance is successfully validated, an icon indicating this is valid provenance is displayed, if not, a warning is shown.
- **Provenance of provenance:** The location of the provenance linked to the resource is displayed to the user, indicating whether it is hosted at the same location as the resource, or at an external source. It can be argued whether either one is more cause for trust than the other. Ideal is when both a local and an external provenance record are present, as this provides a reference for consistency checking.
- **Reputation:** The domain names in all URLs referring to agents and derivation sources are extracted from the provenance statements, and their reputation is assessed by an external API, specifically the Web of Trust[4] (WOT) API. Web of Trust returns a numerical reputation score, which translates to a human readable rating, ranging from "very poor" to "excellent". These are the ratings shown to the user, as well as a clarification of the confidence, the estimated reliability of the reputation value. This way, generating unnecessary distrust events is avoided.

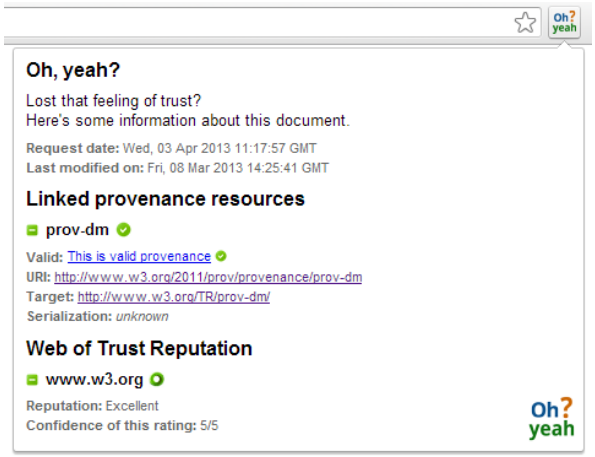

Figure 2. Visualization of trust assessments

6.4. Visualization

The information acquired is relayed back to the user, by showing an unintrusive pop-up window above the document, right under the "Oh, yeah?"-button. This pop-up displays the timing details of the document (when it was requested and when it was last modified), and a list of provenance resources. Each of the items in this list contains details of the selected provenance resource. These details include: the provenance URI, whether it could be retrieved, its source, the validation result and the serialization used. Additionally, the reputation of the domains mentioned in the provenance is relayed to the user. In Figure 2, this visualization is illustrated for the provenance associated with the PROV-DM Proposed Recommendation document at W3C[5].

4. http://www.mywot.com/wiki/API
5. http://www.w3.org/TR/2013/PR-prov-dm-20130312/

7. Conclusions & Future Work

The implementation of the "Oh, yeah?"-button illustrates that thanks to the finalization of the PROV standard, we are a few steps closer to bringing trust assessments to the Web. Enabling easy access to the provenance of Web resources adds value to both consumers and providers of these resources. Consumers gain access to additional input to make an informed decision when deciding to trust the information on the Web, and providers gain an incentive to assert and publish the provenance of their resources.

In future work, we aim to research a finer-grained analysis of the provenance linked to Web resources. More specifically, this would allow us to generate statements regarding the consistency and correctness of the provenance information, based on cross-checking of the information in multiple provenance records and the content. Furthermore, the disadvantage of the centralized approach for reputation assessment is that the central service decides which sources are trustworthy, usually based on crowdsourcing. As explained in [3], a more personalized approach is beneficial, where the preferences and relations of the user are taken into consideration when calculating the reputation assessment.

While the current version of our browser extension supports most features of PROV-AQ, query services remain to be implemented. As the usage of PROV becomes more wide-spread, various provenance query services will appear, and we will adapt our application to use these services to acquire the provenance of Web resources. Note that in this scenario, specifying links to provenance might be unnecessary, and our application can discover it independently through a provenance repository that can be freely queried. Finally, provenance might be missing or non-existent, especially for documents predating the PROV standard. Therefore, it is important that methods are implemented for reconstruction of provenance based on the content [11], [12].

Acknowledgment

The research activities in this paper were funded by Ghent University, iMinds (a research institute founded by the Flemish Government), the Institute for Promotion of Innovation by Science and Technology in Flanders (IWT), the FWO-Flanders, and the European Union. We would also like to thank the members of the W3C Provenance WG for their hard work on PROV.

References

[1] L. Moreau, P. Missier (Eds.), and W3C Provenance Working Group, "PROV-DM: The PROV Data Model. W3C," 2012.

[2] D. Artz and Y. Gil, "A survey of trust in computer science and the semantic web," *Journal of Web Semantics: Science, Services and Agents on the World Wide Web*, vol. 5, no. 2, pp. 58–71, 2007.

[3] J. Golbeck and A. Mannes, "Using trust and provenance for content filtering on the semantic web," in *Proceedings of the Models of Trust for the Web Workshop*, 2006.

[4] D. Ceolin, P. Groth, W. R. van Hage, A. Nottamkandath, and W. Fokkink, "Trust evaluation through user reputation and provenance analysis," in *8th International Workshop on Uncertainty Reasoning for the Semantic Web*, 2012, p. 15.

[5] H.-S. Lim, Y.-S. Moon, and E. Bertino, "Provenance-based trustworthiness assessment in sensor networks," in *Proceedings of the Seventh International Workshop on Data Management for Sensor Networks*. ACM, 2010, pp. 2–7.

[6] X. Li, T. Lebo, and D. L. McGuinness, "Provenance-based strategies to develop trust in semantic web applications," in *Provenance and Annotation of Data and Processes*. Springer, 2010, pp. 182–197.

[7] G. Klyne, P. Groth (Eds.), L. Moreau, O. Hartig, Y. Simmhan, J. Myers, T. Lebo, K. Belhajjame, and S. Miles, "PROV-AQ: Provenance Access and Query. W3C," 2012.

[8] M. Nottingham, "Web linking (RFC 5988)," 2010.

[9] J. Cheney, P. Missier, L. Moreau (Eds.), and T. De Nies, "Constraints of the PROV Data Model. W3C," 2012.

[10] T. Berners-Lee, "Cleaning up the User Interface, Section-The "Oh, yeah?"-Button," 1997.

[11] S. Magliacane, "Reconstructing provenance," in *The Semantic Web–ISWC 2012*. Springer, 2012, pp. 399–406.

[12] T. De Nies, S. Coppens, D. Van Deursen, E. Mannens, and R. Van de Walle, "Automatic discovery of high-level provenance using semantic similarity," in *Provenance and Annotation of Data and Processes*. Springer, 2012, pp. 97–110.