Security and Privacy Aspects of Semantic Data

Sabrina Kirrane

Synonyms

Security and Privacy for the Resource Description Framework.

Definitions

Access Control is a mechanism used to restrict access to data or systems, based on rules that grant *subjects* (e.g. individuals, groups, roles) *access rights* to *resources* (e.g. data or systems) (Sandhu and Samarati 1994). Enforcement is usually broken into two stages *authentication* and *authorisation*. Authentication involves the verification of the subjects identity or attributes. Whereas, authorisation is a mechanism used to determine if the requester (i.e. the subject) has the access rights necessary to carry out the request.

Encryption is an effective means of ensuring the confidentiality and integrity of information stored locally or transferred over a network (Menezes et al 1996). Encryption involves the translation of data into an unintelligible form through the use of a secret key. Decryption is the process of restoring data to its original form through the use of a key (which may or may not be the same as the key used to encrypt the data). Encrypted data is referred to as cipher or cipher text, whereas unencrypted data is commonly known as plain text.

Trust mechanisms are used to verify the validity of a claim (e.g. the identity/attributes of an individual, or the correctness of data). The most widely used trust mechanisms include *policies* and *reputation* (Artz and Gil 2007). Policies are used to govern the exchange of credentials that are often certified by trusted third parties. While, reputation mechanisms take the form of provenance information and metrics that are calculated from previous actions and behaviors. Where no such data is available trust may be established via referral from other trusted parties.

Anonymisation involves the removal of personally identifiable information from datasets. One of the most well know anonymisation techniques k-anonymity, involves the use of suppression (i.e. masking sensitive data) and generalisation (i.e. choosing broader classification terms for sensitive data), in order to group individuals into equivalence classes, whereby each individual in a class is indistinguishable from *k-1* other individuals (Samarati and Sweeney 1998).

Sabrina Kirrane
Vienna University of Economics and Business, Austria e-mail: sabrina.kirrane@wu.ac.at

Overview

The Resource Description Framework (Manola and Miller 2004) is designed to facilitate data integration and reuse by representing distributed data in a machine readable format. RDF *vocabularies* (otherwise known as *ontologies*) are collections of RDF triples that can be used to describe both schema and instance data. Each triple, which is composed of a *subject-predicate-object* expression, asserts a binary relationship between two pieces of information. *Internationalised Resource Identifiers (IRIs)* and *literals* are used to represent information, which can be either physical or abstract in nature. The *RDF-Schema* (RDFS) ontology (Brickley and Guha 2014) is composed of a set of classes and properties commonly used to describe RDF data. RDFS does not describe the structure of an RDF graph, but rather provides a framework that can be used to denote classes, properties and relations. Vocabularies are often placed in a common *namespaces*, referenced via *prefixes*. In the examples that follow the default prefix `:` is used to denote an example enterprise ontology `<http://example.org/ex/>`. In addition, well known `rdf` and `foaf` prefixes are used for the RDF built-in vocabulary and FOAF social network ontology respectively. *Example* 1 demonstrates how RDF can be used to represent information pertaining to `Joe Bloggs`.

Example 1 (RDF triples). The following triples states that the entity `:JBloggs` is a person whose first name is `Joe`, last name is `Bloggs` and salary is `60000`:

```
:JBloggs rdf:type foaf:Person.
:JBloggs foaf:givenName "Joe".
:JBloggs foaf:lastName "Bloggs".
:JBloggs :salary 60000.
```

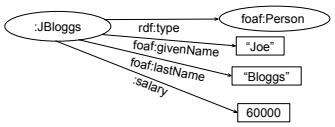

Fig. 1 Triples represented as an RDF graph

An *RDF graph* is a finite set of RDF triples, with subjects and objects represented as nodes and predicates represented as edges. *Figure* 1 demonstrates how the triples in *Example* 1 converge to form a graph, with IRIs represented as ovals and literals represented as rectangles. A collection of RDF graphs, which can include a default graph and one or more named graphs is known as an RDF dataset.

In a recent survey by Fernandez Garcia et al (2016) the authors analysed topics appearing in papers that were published in Semantic Web conference proceedings and journals from 2006 to 2015 inclusive. The results of the conducted text analysis confirmed that traditional Semantic Web topics, such as knowledge representation, data management, system engineering, searching, browsing and exploration, and data integration, dominated the field up to 2015.

According to Fernandez Garcia et al (2016), although topics relating to security and privacy have shown a minor increase over the years, the topics remain under represented in comparison to traditional topics.

Much of the early research on security and privacy in the context of the Semantic Web focused on using RDF to represent existing access control models and standards, and demonstrating how

the technology could be used to develop general policy languages. Later the focus moved to the development of access control strategies for RDF and the Semantic Web. Other popular topics over the years include demonstrating how existing encryption mechanisms can be used to protect RDF data and establishing trust mechanisms for the Semantic Web. More recently, the landscape has broadened to include the encryption and anonymisation of RDF data.

Key Research Findings

The goal of this section is to introduce the reader to key research findings in relation to Semantic Web security and privacy, and as such it focuses on the predominant topics within the community, namely, access control, encryption, trust and anonymisation.

Access Control

Access Control (AC) for the RDF data model has predominately focused on using patterns to specify authorisations, enabling inference based on semantic relations between policy entities and demonstrating how RDF can be used to form general policy languages.

Reddivari et al (2005) demonstrate how access control rules can be used to manage access to an RDF store. Two predicates `permit` and `prohibit` are used to grant and deny access rights based on common database actions (e.g. `INSERT, DELETE, SELECT`) to one or more triples using triple patterns (cf. *Example* 2). A triple pattern is composed of an RDF triple with optionally a variable (denoted by a ?) in the subject, predicate and/or object position.

Example 2 (AC Rules with triple patterns). The following rule states that a subject `Alice` can create instances of any class (denotes as `?y`) if there is an assertion that subject `Alice` created that class.

```
permit(INSERT(Alice,
(?x,rdf:type,?y)))
:- createdNode(Alice,?y)
```

Jain and Farkas (2006) build on the approach proposed by Reddivari et al (2005), by demonstrating how RDFS entailment rules can be used to derive authorisations for inferred triples. While, Kirrane et al (2013) demonstrate how authorisations based on quad patterns (where the fourth element denotes the named graph) can be used to enforce Discretionary Access Control (DAC), whereby users can pass their access rights on to other users. Like Jain and Farkas (2006) the authors derive access rights for derived data using RDFS entailment rules.

When it comes to access control enforcement, typical enforcement strategies involve filtering unauthorised data based on access control policies and executing queries against the filtered dataset, or using query rewriting techniques to inject access control filters into queries.

Dietzold and Auer (2006) and Gabillon and Letouzey (2010) demonstrate how graph patterns (i.e. sets of triple patterns) constrained by a `WHERE` clause can be used to created a new dataset that only contains data the subject is permitted to access. The authorised dataset

is created using SPARQL the standard query language for RDF (Seaborne and Prud'hommeaux 2008). Essentially, authorisations contain filters that refer to sparql `CONSTRUCT` queries that are used to generate the authorised dataset. In Gabillon and Letouzey (2010) a rule such as `Permit(Alice, SELECT, foafview.txt)` can be used to permit subject `Alice`, access right `SELECT` on resource `foafview.txt`. The resource `foafview.txt` simply contains a `CONSTRUCT` query such as that presented in *Example* 3). When a requester submits a query, a new dataset is created based on the matched authorisations. The query is executed against the new dataset, which only contains data that the requester is permitted to access.

Example 3 (Construct view). The following query creates a dataset that contains all data relating to people with `Bloggs` as a lastname.

```
CONSTRUCT {?x ?p ?y}
WHERE {
?x ?p ?y .
?x foaf:lastName "Bloggs"}
```

An alternative enforcement strategy proposed by Abel et al (2007) uses query rewriting to create bindings for variables that are subsequently added to the query `WHERE` clause. In the case of negative authorisations the bindings are added to a MINUS clause, which is appended to the query. A simple SPARQL `SELECT` query is presented in *Example* 4 and sample rewritten queries containing positive and negative filters are presented in *Example* 5 and *Example* 6, respectively. When a requester submits a query, the enforcement framework rewrites the query according to the matching authorisations, and the rewritten query is subsequently executed against the new dataset, ensuring that the requester is only returned data that they are permitted to access.

Example 4 (SELECT query). The following query returns all data.

```
SELECT *
WHERE { ?s ?p ?o }
```

Example 5 (Positive filter). The following query, which contains a positive filter, only returns the information for `:JBloggs`.

```
SELECT *
WHERE { ?s ?p ?o .
FILTER ( ?s = :JBloggs ) }
```

Example 6 (Negative filter). The following query, which contains a negative filter, returns everything except the `:salary` information.

```
SELECT *
WHERE { ?s ?p ?o .
MINUS { ?s ?p ?o .
FILTER ( ?p = :salary ) }}
```

In addition to the access control mechanisms described above there have been a number of standardisation initiatives that could be used to limit access to RDF data. *Web Identity and Discovery (WebID)* (Sporny et al 2011) is a mechanism that can be used to uniquely identify and authenticate a person, company, organisation or other entity, by means of a Uniform Resource Identifier (URI). While, *Web Access Control (WAC)* W3C (n.d.) is an RDF vocabulary and access control framework, that can be used for policy specification and enforcement. Both Villata et al (2011) and Sacco and Passant (2011) extend WAC to cater

for access control over the RDF data model. Using the extended vocabularies, it is possible to associate access control with individual RDF resources (subjects, predicates and objects) and also collections of RDF resources (named graphs). In addition, the authors extend the vocabulary to cater for a broader set of access privileges.

An alternative policy language, called the Open Digital Rights Language (ODRL) (Iannella and Villata 2018), is a general rights language that can be used to define rights for limiting access to digital resources. When it comes to ODRL and RDF, primary research efforts to date focus on demonstrating how ODRL can be used to express a variety of access policies (Steyskal and Polleres 2014; Steyskal and Kirrane 2015) and using ODRL vocabularies to specify RDF licenses (Cabrio et al 2014).

A comprehensive survey of existing access control strategies for RDF is presented in Kirrane et al (2017).

Encryption

Encryption techniques for RDF have received very little attention to date, with work primarily focusing on the partial encryption of RDF data, the querying of encrypted data and the signing of RDF graphs.

Giereth (2005) demonstrate how public-key encryption can be used to partially encryption RDF fragments (i.e. subjects, objects, or predicates). The ciphertext and the corresponding metadata (algorithm, key, hash etc...) are represented using a literal that they refer to as an encryption container.

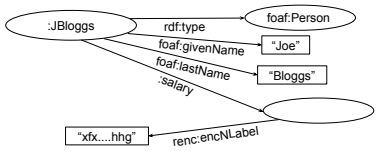

Fig. 2 Partially Encrypted RDF graph

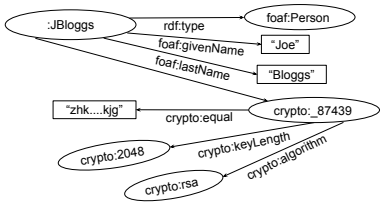

Fig. 3 Partially Encrypted RDF graph and Metadata

When only the object is encrypted, the object part of the triple is replaced with a blanknode (i.e. an anonymous resource) and a new statement is created with the blanknode as the subject, the encryption container as the object and a new `renc:encNLabel` predicate (cf. *Figure* 2). The treatment of encrypted subjects is analogous. The encrpytion of predicates is a little more difficult, as reification (a technique used to make statements about resources) is needed to associate the new blanknode with the relevant subject, object and encryption container.

Rather than simply storing the encrypted data and metadata in a literal, Gerbracht (2008) discuss now the metadata can be represented using multiple triples using their crypto ontology. The encrypted element or subgraph is replaced with a new unique identifier and new statements are added for the encrypted data and the corresponding metadata (cf. *Figure* 3).

Kasten et al (2013) in turn focus on querying encrypted data. In the

proposed framework each triple is encrypted eight times according to the eight different triple pattern binding possibilities. The proposed approach allows for graph pattern queries to be executed over the ciphertext, at the cost of storing multiple ciphers for each statement. An alternative approach by Fernández et al (2017) demonstrates how functional encryption can be used to generate query keys based on triple patterns, whereby each key can decrypt all triples that match the corresponding triple pattern. While, other work by Kasten et al (2014) investigates enabling the signing of graph data at different levels of granularity.

Trust

In 2007, Artz and Gil (2007) conducted a survey of existing trust mechanisms in computer science in general and the Semantic Web in particular. The authors highlight that traditional approaches focused primarily on authentication via assertions by third parties, however in later years the topic evolved to include historical interaction data, the transfer of trust from trusted entities, and decentralised trust mechanisms (e.g. voting mechanisms or other consensus decision making mechanisms).

Existing work on trust and semantic data focuses primarily on demonstrating how existing trust metrics can be applied to Semantic Web Data, the development of policy languages to support trust and negotiation and the identification of trust architectures and frameworks.

Ding et al (2003, 2005) discuss how various trust mechanisms can be combined in order to determine the reliability of information published on the Semantic Web. The proposed trust mechanism combines historical data, information obtained directly from trusted semantic agents and information based on referrals from trusted agents.

The PeerTrust policy language and framework (Gavriloaie et al 2004) demonstrates how together semantic annotations and access control rules can be used to support automated trust negotiation and access control. An alternative policy language called Protune is described in Bonatti and Olmedilla (2005, 2007)). Although Protune is in fact a general policy language, the authors focus primarily on trust negotiation and policy explanations.

Bizer and Oldakowski (2004) propose a trust architecture that combines reputation, content and context based trust mechanisms. The *Information Integration Layer* aggregates data from several sources and adds the relevant provenance metadata. The *Repository Layer* is used to store the information and associated metadata in named graphs. The *Query and Trust Evaluation Layer* uses trust policies to make trust decisions. Here the authors rely on a query language TriQL.P that is used to return the query results together with a justification tree that can be used to understand how the query results fulfil the trust requirement. Finally the *Application and Explanation Layer* receives requests and provides the trust decision together with the relevant explanations.

More recently, Laufer and Schwabe (2017) propose a framework that can be used to describe the trust process. Inputs to be considered include the data and associated metadata, contextual information relating to the action that needs to be taken, together with trust policies speci-

fied by the agent. Like (Ding et al 2003) the trust process relies on historical data, along with direct and indirect sources of information.

Anonymisation

The anonymisation of RDF data has recently emerged as a popular research topic, with work to date focusing on the application of k-anonymity (Samarati and Sweeney 1998) or differential privacy (Dwork 2006) to RDF data.

Radulovic et al (2015) propose a framework called k-RDFanonymity, which includes an anonymisation model, generalisation and suppression operations and distortion metrics, that are specifically tailored for the anonymisation of RDF data. The authors highlight the fact that RDF differs from tabular data as identifiers, quasi identifiers, and sensitive attributes can appear in the subject, predicate and object positions position. Additionally the anonymisation needs to be able to handle data represented as literals and IRIs. In the proposed model generalisation involves the replacement of resources (i.e. literals or IRIs) with more general resources based on domain hierarchies. While, suppression involves either the removal or replacement of resources.

Heitmann et al (2017) build on this work to ensure protection against neighbourhood attacks. The proposed approach, which is known as k-RDF-Neighborhood anonymity ensures that one-hop neighbours of an anonymised resource are indistinguishable from k-1 one-hop neighbours of other resources.

Other work in relation to RDF anonymisation include adopting graph or statistical database approaches. Lin (2016) take inspiration from existing graph isomorphism-based anonymisation techniques, discussing their suitability for RDF data from both a security and a computational complexity perspective. While, Silva et al (2017) explore the application of existing differential privacy mechanism to RDF data and propose a framework that can be used to compute differential privacy parameters.

Examples of Application

Semantic Web technologies have a solid foundation in open standards as evidenced by the various World Wide Web Consortium (W3C) recommendations, however the layers of the Semantic Web technology stack (Berners-Lee 2000) that relate to security and privacy (i.e. unifying logic, proof, trust and cryptography) are still very immature. Although the application of the key research findings described in the previous section are still very exploratory, several of the articles are guided by real world uses cases and practical applications.

For instance, the Protune policy language (Bonatti and Olmedilla 2005, 2007)), which was developed by the Research Network of Excellence on Reasoning on the Web, known as REWERSE, was tasked with building the foundations for the advanced of Web systems and applications by developing inter-operable reasoning languages.

Fernández et al (2017) are motivated by a real word use case that involves the combination of open and closed data in a data market scenario. In order to demonstrate the suitability of the

proposed encryption mechanism the authors conduct a performance evaluation over two real world datasets: Jamendo a large dataset containing licensed music and the AEMET metereological dataset.

Although the initial evaluation of the trust framework proposed by Ding et al (2003) was conducted using simulated data, the authors later discussed how trust mechanisms could be used in the context of homeland security, in order to identify suspicious individuals, relationships, activities or events (Ding et al 2005). Similarly, Laufer and Schwabe (2017) describe how the proposed trust framework can be used to evaluate the trustworthiness of claims in relation to political agents in Brazil coming from a variety of public sources (e.g. news stories, tweets, social media postings).

Existing work on anonymisation appears to be less applied that the other topics with authors simply motivating their work by referring to privacy concerns in domains, such as healthcare and energy (cf. (Radulovic et al 2015)).

Future Directions for Research

From a community perspective, it is well know that privacy is a multidisciplinary research area, which brings with it the need for closer collaboration between computer scientists, humanities and social scientists and legal scholars. Although initiatives such as the *Society, Privacy and the Semantic Web - Policy and Technology* (PrivOn) workshop, which was collocated with the International Semantic Web Conference (ISWC) from 2013 to 2017, provides a forum for multidisciplinary research, stronger collaboration between different research communities is still needed.

From a technical perspective, there is a need for more applied work and a focus on attacker models across all privacy and security topics. Additionally there are many open research questions concerning the topics presented in this paper, several of which are outlined below.

When it comes to information security there is still no standard access control strategy for the Semantic Web. Considering the array of access control specification and enforcement mechanisms proposed to date, a necessary first step is to develop a framework that can be used to evaluate existing offerings in terms of correctness, completeness and robustness.

As for encrypted RDF, it is still not possible to execute complex queries and computations over encryption RDF data. One interesting avenue for future work is the application of Homomorphic encryption to RDF, however it brings with it performance and scalability issues that still need to be tackled. Another open research topic is the simplification of key management for multiple datasets, federated querying over encrypted data and providing support for the revocation of existing keys.

In terms of trust, a recent article by Beek et al (2016) highlights several issues with respect to the quality of existing data and datasources, claiming that the Semantic Web is neither traversable nor machine-processable, and consequently arguing that the Semantic Web needs centralisation. A counter argument, that is more in keeping with the goals of the Semantic Web, would be to argue for the application of trust mechanisms into the fabric of the Semantic Web, which could be

brought about by the realisation of the upper layers and vertical layers of the Semantic Web technology stack.

Anonymisation is a relativey new topic within the Semantic Web community with works primarily focusing on k-anonymity. However, it is well known that k-anonymity is prone to homogeneity and background knowledge attacks. Common extensions mechanisms include l-diversity (Li et al 2007), which ensures sensitive attributes within an equivalence class are suitably different, and t-closeness, which ensures that the distribution of each equivalence class is representative of the distribution of the entire dataset (Machanavajjhala et al 2006).

Other promising privacy and security research directions that remain underdeveloped and as such have not been presented in this article include usage control, which is defined as an extension of access control that enables data publishers to dictate not only who can access their data but also what they are permitted to do with this data (Bonatti et al 2017). Related topics include transparency, which involves being open with respect to data processing and sharing, and accountability, which involves making data consumers responsible for their actions. Here, interesting avenues for future work include the adoption and extension of non-repudiation and fair exchange protocols.

Cross-References

Big Data for Cyber Security, Privacy aware identity management, Privacy-Preserving Data Analytics, Privacy-preserving Record Linkage.

References

Abel F, De Coi J, Henze N, Koesling A, Krause D, Olmedilla D (2007) Enabling advanced and context-dependent access control in rdf stores. In: The Semantic Web, Springer Berlin Heidelberg, Lecture Notes in Computer Science, vol 4825, pp 1–14, URL http://dx.doi.org/10.1007/978-3-540-76298-0_1

Artz D, Gil Y (2007) A survey of trust in computer science and the semantic web. Web Semantics: Science, Services and Agents on the World Wide Web 5(2):58–71

Beek W, Rietveld L, Schlobach S, van Harmelen F (2016) Lod laundromat: Why the semantic web needs centralization (even if we don't like it). IEEE Internet Computing 20(2)

Berners-Lee T (2000) Semantic web - xml2000. Accessed 13 January 2018, available at https://www.w3.org/2000/Talks/1206-xml2k-tbl/slide10-0.html

Bizer C, Oldakowski R (2004) Using context- and content-based trust policies on the semantic web. In: Proceedings of the 13th international World Wide Web conference on Alternate track papers & posters, ACM, pp 228–229

Bonatti P, Olmedilla D (2005) Driving and monitoring provisional trust negotiation with metapolicies. In: Policies for Distributed Systems and Networks, 2005. Sixth IEEE International Workshop on, pp 14–23

Bonatti P, Kirrane S, Polleres A, Wenning R (2017) Transparent personal data processing: The road ahead. In: International Conference on Computer Safety, Reliability, and Security, Springer, pp 337–349

Bonatti PA, Olmedilla D (2007) Rule-based policy representation and reasoning for the semantic web. In: Proceedings of the Third International Summer School Conference on Reasoning Web, Springer-Verlag, RW'07, pp 240–268, URL http://dl.acm.org/citation.cfm?id=2391482.2391488

Brickley D, Guha R (2014) RDF Schema 1.1. W3C Recommendation, available at http://www.w3.org/TR/2014/REC-rdf-schema-20140225/Overview.html, W3C

Cabrio E, Aprosio AP, Villata S (2014) These are your rights a natural language processing approach to automated rdf licenses generation. In: The Semantic Web: Trends and Challenges, Springer, pp 255–269

Dietzold S, Auer S (2006) Access control on rdf triple stores from a semantic wiki perspective. In: Proceedings of the ESWC'06 Workshop on Scripting for the Semantic Web

Ding L, Zhou L, Finin TW (2003) Trust based knowledge outsourcing for semantic web agents. In: Web Intelligence, pp 379–387

Ding L, Kolari P, Finin T, Joshi A, Peng Y, Yesha Y (2005) On homeland security and the semantic web: A provenance and trust aware inference framework. In: AAAI Spring Symposium: AI Technologies for Homeland Security, pp 157–160

Dwork C (2006) Differential privacy. In: Proceedings of the 33rd International Conference on Automata, Languages and Programming - Volume Part II, Springer-Verlag, Berlin, Heidelberg, ICALP'06, pp 1–12, DOI 10.1007/11787006_1, URL http://dx.doi.org/10.1007/11787006_1

Fernández JD, Kirrane S, Polleres A, Steyskal S (2017) Self-enforcing access control for encrypted rdf. In: European Semantic Web Conference, Springer, pp 607–622

Fernandez Garcia JD, Kiesling E, Kirrane S, Neuschmid J, Mizerski N, Polleres A, Sabou M, Thurner T, Wetz P (2016) Propelling the potential of enterprise linked data in austria. roadmap and report. URL https://www.linked-data.at/wp-content/uploads/2016/12/propel_book_web.pdf

Gabillon A, Letouzey L (2010) A view based access control model for sparql. In: Network and System Security (NSS), 2010 4th International Conference on, pp 105–112

Gavriloaie R, Nejdl W, Olmedilla D, Seamons KE, Winslett M (2004) No registration needed: How to use declarative policies and negotiation to access sensitive resources on the semantic web. In: ESWS, Springer, pp 342–356

Gerbracht S (2008) Possibilities to Encrypt an RDF-Graph. In: Proc. of Information and Communication Technologies: From Theory to Applications, pp 1–6

Giereth M (2005) On Partial Encryption of RDF-Graphs. In: Proc. of International Semantic Web Conference, vol 3729, pp 308–322

Heitmann B, Hermsen F, Decker S (2017) krdf-neighbourhood anonymity: Combining structural and attribute-based anonymisation for linked data. In: Proceedings of the 5th Workshop on Society, Privacy and the Semantic Web - Policy and Technology (PrivOn2017) (PrivOn), URL http://ceur-ws.org/Vol-1951/#paper-03

Iannella R, Villata S (2018) ODRL Information Model 2.2. W3C proposed recommendation, W3C, available at https://www.w3.org/TR/odrl-model/

Jain A, Farkas C (2006) Secure resource description framework: An access control model. In: Proceedings of the Eleventh ACM Symposium on Access Control Models and Technologies, ACM, SACMAT '06, pp 121–129, URL http://doi.acm.org/10.1145/1133058.1133076

Kasten A, Scherp A, Armknecht F, Krause M (2013) Towards search on encrypted graph data. In: Proc. of the International Conference on Society, Privacy and the Semantic Web-Policy and Technology, pp 46–57

Kasten A, Scherp A, Schauß P (2014) A Framework for Iterative Signing of Graph Data on the Web, Springer International Publishing, Cham, pp 146–160. DOI 10.1007/978-3-319-07443-6_11, URL https://doi.org/10.1007/978-3-319-07443-6_11

Kirrane S (2015) Linked data with access control. PhD thesis, INSIGHT Centre for Data Analytics, National University of Ireland, Galway, URL https://aran.library.nuigalway.ie/handle/10379/4903

Kirrane S, Abdelrahman A, Mileo A, Decker S (2013) Secure manipulation of linked data. In: The Semantic Web - ISWC 2013, Springer Berlin Heidelberg, Lecture Notes in Computer Science, vol 8218, pp 248–263, URL http://dx.doi.org/10.1007/978-3-642-41335-3_16

Kirrane S, Mileo A, Decker S (2017) Access control and the resource description framework: A survey. Semantic Web 8(2):311–352, URL http://www.semantic-web-journal.net/system/files/swj1280.pdf

Laufer C, Schwabe D (2017) On modeling political systems to support the trust process. In: Proceedings of the 5th Workshop on Society, Privacy and the Semantic Web - Policy and Technology (PrivOn2017) (PrivOn), URL http://ceur-ws.org/Vol-1951/#paper-07

Li N, Li T, Venkatasubramanian S (2007) t-closeness: Privacy beyond k-anonymity and l-diversity. In: Data Engineering, 2007. ICDE 2007. IEEE 23rd International Conference on, IEEE, pp 106–115

Lin Z (2016) From isomorphism-based security for graphs to semantics-preserving security for the resource description framework (rdf). Master's thesis, University of Waterloo

Machanavajjhala A, Gehrke J, Kifer D, Venkitasubramaniam M (2006) l-diversity: Privacy beyond k-anonymity. In: Data Engineering, 2006. ICDE'06. Proceedings of the 22nd International Conference on, IEEE, pp 24–24

Manola F, Miller E (2004) RDF Primer. W3C Recommendation, available at http://www.w3.org/TR/rdf-primer/, W3C

Menezes AJ, Van Oorschot PC, Vanstone SA (1996) Handbook of applied cryptography. CRC press

Radulovic F, García Castro R, Gómez-Pérez A (2015) Towards the anonymisation of rdf data. DOI 10.18293/SEKE2015-167, URL https://doi.org/10.18293/SEKE2015-167

Reddivari P, Finin T, Joshi A (2005) Policy-based access control for an rdf store. In: Proceedings of the Policy Management for the Web workshop, pp 78–83

Sacco O, Passant A (2011) A privacy preference ontology (ppo) for linked data. In: Linked Data on the Web, CEUR-WS, URL http://ceur-ws.org/Vol-813/ldow2011-paper01.pdf

Samarati P, Sweeney L (1998) Protecting privacy when disclosing information: k-anonymity and its enforcement through generalization and suppression. Tech. rep., Technical report, SRI International

Sandhu RS, Samarati P (1994) Access control: principle and practice. IEEE communications magazine 32(9):40–48

Seaborne A, Prud'hommeaux E (2008) SPARQL Query Language for RDF. W3C recommendation, available at http://www.w3.org/TR/rdf-sparql-query/, W3C

Silva RRC, Leal BC, Brito FT, Vidal VMP, Machado JC (2017) A differentially private approach for querying rdf data of social networks. In: Proceedings of the 21st International Database Engineering & Applications Symposium, ACM, New York, NY, USA, IDEAS 2017, pp 74–81, DOI 10.1145/3105831.3105838, URL http://doi.acm.org/10.1145/3105831.3105838

Sporny M, Inkster T, Story H, Harbulot B, Bachmann-Gmr R (2011) WebID 1.0 - Web Identification and Discovery. W3C working draft, W3C, available at http://www.w3.org/2005/Incubator/webid/spec/

Steyskal S, Kirrane S (2015) If you can't enforce it, contract it: Enforceability in policy-driven (linked) data markets. In: SEMANTiCS (Posters & Demos), pp 63–66

Steyskal S, Polleres A (2014) Defining expressive access policies for linked data using the odrl ontology 2.0. In: Proceedings of the 10th International Conference on Semantic Systems, ACM, pp 20–23

Villata S, Delaforge N, Gandon F, Gyrard A (2011) An access control model for linked data. In: On the Move to Meaningful Internet Systems: OTM 2011 Workshops, pp 454–463

W3C (n.d.) Webaccesscontrol. Accessed 13 January 2018, available at https://www.w3.org/wiki/WebAccessControl