The Never-Ending Story: On the Need for Continuous Privacy Impact Assessment

Laurens Sion
imec-DistriNet, KU Leuven
Heverlee, Belgium
laurens.sion@cs.kuleuven.be

Dimitri Van Landuyt
imec-DistriNet, KU Leuven
Heverlee, Belgium
dimitri.vanlanduyt@cs.kuleuven.be

Wouter Joosen
imec-DistriNet, KU Leuven
Heverlee, Belgium
wouter.joosen@cs.kuleuven.be

***Abstract*—The importance of privacy by design has increased with initiatives such as the General Data Protection Regulation (GDPR). While static, design-level assessment of privacy aspects provides considerable benefits in the creation of privacy-preserving software-intensive systems, operational aspects that are difficult to predict at design-time also play a key role. This is particularly true in the instance of privacy impact or privacy risk: while existing approaches succeed fairly well in assessing the overall risk from a static design context, they are not well suited to capture risk elements that are dynamic and often impossible to foresee.**

In this position paper, we highlight this problem at the basis of a number of realistic motivational scenarios and outline our vision towards continuous privacy impact assessment and risk management.

***Index Terms*—privacy by design, GDPR, DevOps, continuous privacy assessment, privacy risk, DevPrivOps**

1. Introduction

The rising number of personal data breaches has led to an increased attention towards data protection and privacy. The significance of privacy is further confirmed with the introduction of legislation and guidelines such as the EU's General Data Protection Regulation (GDPR) [1], the OECD Privacy Guidelines [2], the Global Privacy Standard [3], and the Generally Accepted Privacy Principles (GAPP) [4].

The effective realization of the principles of Privacy by Design (PbD) and Data Protection by Design (DPbD) is strongly rooted in the notion of a *privacy impact assessment*, which in turn is concretized into *privacy risk*, and thus requires effective privacy risk assessment and management approaches. Indeed, countermeasures should be informed and motivated by an extensive characterization of the involved risk. Additionally, legal reasoning should take into account and correctly weigh these risks (i.e. a *risk-based approach* [1] is required).

In its essence, this is the risk of harming the data subjects' fundamental rights, as delineated in the regulatory frameworks discussed above [1]. In practice however, this overall risk is approximated as a combination of different risk factors such as the perceived likelihood and impact of data breaches, or of other factors that influence the risk such as the value and nature of the involved data sets, the nature of the involved data subjects (e.g., their age, their societal status in case of celebrities or politicians), etc.

Many approaches have been proposed to estimate the overall privacy risk through assessing individual risk factors. These approaches are however mainly focused on the static realization of this principle at development and design time, i.e. when determining the means of data collection and processing activities, and thus many of the relevant factors have to be estimated beforehand.

These approaches however ignore a fundamental operational or dynamic aspect of the privacy risk equation: many of the privacy risk factors commonly considered are in essence based on operational metrics which cannot accurately be determined from a system's design blueprint (for example, the actual volume of data collected, the actual user base of the system once it is brought to market, etc.).

In this position paper, we highlight and motivate this problem. First, we present an overview of the risk factors used in existing approaches (Section 2). Then, we illustrate in real-world examples how a number of these risk factors are difficult or even impossible to predict at design or development time (Section 3). Finally, we discuss the potential role of a number of privacy-enhancing infrastructure components that will enable continuous monitoring of these privacy risk factors (Section 4).

2. State of the art on privacy risk assessment

Adopting a risk-based approach in its essence implies that any data collection and processing activity should be accompanied with a characterization of the associated *privacy risk*, i.e. the risk of harm to the fundamental rights of data subjects, and explicit approaches to properly manage such risks.

In practice, this overall notion of risk is difficult to quantify, and many practical approaches resort to approximating the overall risk as a combination of more tangible risk factors such as, for example, the projected impact and estimated likelihood [15] of specific data leaks. In this section, we discuss the state of the art in privacy risk assessment and prioritization with specific emphasis on highlighting concrete privacy risk factors that are used to quantify the overall privacy risk. These are summarized in Table 1, which shortly describes each of the identified risk factors and summarizes how they are estimated in practice.

In PRIAM [5], the risk factors are derived from an information collection phase in which the system is described together with the risk sources, privacy weaknesses, feared events, and privacy harms.

This research is partially funded by the Research Fund KU Leuven, the PRiSE KU Leuven-C2 research project, and the Flemish Research Programme Cybersecurity.

TABLE 1. MAPPING SUMMARY OF THE DIFFERENT RISK FACTORS USED IN STATE-OF-THE-ART PRIVACY RISK MODELS TO QUANTIFY THE OVERALL PRIVACY RISK ASSOCIATED TO A DATA COLLECTION OR PROCESSING ACTIVITY.

Privacy risk factor	Description	How quantified	Used in
Data type sensitivity	The sensitivity level of the collected data impacts the potential harm to data subjects	Decided when determining the means of data collection, at design time.	[5]–[7]
Data volume	The volume of data collected and processed is an indication of (i) the inherent value of a data set, (ii) the impact on data subjects in case of breaches.	Estimated at design time.	[5]–[7]
Data subject type	The nature of the data subjects affected (e.g. minors, data subjects with a public profile) has an impact on the degree of harm in case of data breach or misuse.	Decided at design-time, considering the user base of a service or product.	[6], [7]
Data subject scale	The amount of involved data subjects amplifies the overall harm in case of data breaches or misuse.	Decided at design-time, considering characteristics of the user base of a service or product.	[6]
Data retention	The time window in which personal data is kept before active deletion indicates the time window of exposure to harm.	Decided at design time, enforced through policies at runtime.	[6]–[9]
Threat probability	Expresses how easy is it to pose a certain threat to a system. These may be security threats (cf. STRIDE [10] or privacy threats, LINDDUN [11] threats).	Design-time assessment (at the basis of system design models).	[6], [7], [12], [13]
Threat Vulnerability	Combination of the strengths of countermeasures in place and the probability of insider access.	Analysis of security architecture or security posture/assessment of existing countermeasures.	[6], [7], [12], [13]
Data locality	Selecting a specified storage location impacts the potential privacy harm (e.g., private data center vs. public cloud).	At design or deployment time, based on data placement decisions.	[8], [9]
Collection intent	The nature of the data collection activity and the means of obtaining data (voluntarily provided by data subjects, observing and/or recording activities, direct questioning or probing for information) impacts the potential privacy harm.	Design-time, when determining the means of data collection.	[9], [12]–[14]
Processing means	The way in which information is stored and used after initial collection (aggregation, identification, insecurity, secondary use, exclusion [14]) and the level of invasiveness (intrusion or decisional interference).	Design-time, when determining the means of data processing.	[7], [9], [12]–[14]
Intended disclosures	Breaching confidentiality, disclosing information to third parties, amplifying the accessibility, threatening to disclose, or distorting information.	Based on assessment of the main functionality of the system under design.	[7]–[9], [12]–[14]
Data subject awareness and control	Insufficiently informing data subjects about collection and processing is considered harmful as data subjects will have no means to exert their fundamental rights (to be forgotten, etc.)	Awareness is realized using data subject dashboards, privacy policies and consent forms established at design time.	[7], [9], [13]

The risk model proposed by Sion et al. [6] extends FAIR [15] for privacy threat modeling, thus exclusively in the context of requirements and architecture.

The RFC6973 [9] provides privacy guidelines for protocols and explicitly lists the privacy harm sources in function of the nature of the collection and processing activities, data locality, intended disclosure and unawareness (exclusion).

Solove's taxonomy of privacy [14] delves into the legal notion of privacy risk and identifies sources of privacy harm that stem from the means of collection, the nature of the processing and its impact on the data subject, and the disclosures of the obtained information to third parties.

Cronk [12] also relies on the risk components from FAIR [15] but uses the taxonomy of privacy from Solove [14] to assess the potential privacy harm.

Hong et al. [7] present a privacy risk model for ubiquitous systems using a set of questions to elicit privacy risks which are subsequently prioritized.

The NIST Privacy Risk Assessment Methodology [13] applies the NISTIR 8062 risk model [16] to identify and prioritize privacy risks.

Complementary risk modeling and management methods and frameworks [17]–[20] put emphasis on the appropriate treatment of privacy risk from an organizational risk management perspective and as such have been excluded from this summary, as they do not prescribe concrete risk factors to perform the risk assessment.

The summary table shows that many of these approaches are intended for the early stages of the development of a system (conception, requirements, architecture), and thus many of the risk factors in practice have to estimated, based upon future projection, for example of the intended user base of the system. In following section, we discuss a number of real-world scenarios in which such static risk assessment falls short as later evolutions that are difficult to anticipate drastically change the risk profile of the described data collection/processing operation.

3. Motivational scenarios

In this section, we discuss some concrete, real-world scenarios that illustrate the shortcomings of one-shot or infrequent risk assessments. Throughout the section, we refer to the affected *privacy risk factors* from Table 1.

Changing User Bases. A social network is a textbook example of a software service that strongly benefits from the network effect (formulated in Metcalfe's law [21] as the value of a network increasing with the square of its users). The user bases of such Internet services have the ability to grow suddenly and drastically, and this in turn impacts the '*data subject scale*' risk factor. Furthermore, an increase in popularity of a social network amongst, for example, teenagers (i.e. minors) can shift the dominant *data subject types* in the user base and may again be a necessary trigger for re-assessment of the overall risk. Finally, specific items posted in a social network may unexpectedly go

viral (i.e. popularity increases substantially in a short time frame), and this in turn may amplify the factor of '*intended disclosures*' or may pose risk in terms of '*data subject awareness and control*' as this level of exposure might not have been originally intended or foreseen by the user.

Changing Application Usage. Another typical source of unanticipated change is related to effective usage of a service. A messaging application may, for example, initially be intended and used for personal communication between citizens (*data type sensitivity, processing*). However, with evolutions in the *data subject scale* as outlined above, the service may be increasingly used by, for example, government employees [22] or physicians, which means that nature of the information being processed and transferred (*processing*) by the service may significantly change (*data type sensitivity*) and thus a re-assessment of the overall privacy impact is in order.

Changing Legal Context. The laws and regulations that apply to data processing operations are not static or constant. For example, determining the exact legal context that applies in a specific case depends upon: (i) the location of the company processing the data; (ii) the location of the processing operations themselves (*data locality*); (iii) the nationality of the data subject (*data subject type*); and (iv) changes in the interpretation of the laws due to court opinions. These factors influence the risk to the data subjects (*data subject type, data type sensitivity*) as they can modify which risk components (*threat probability, threat vulnerability*) will play a dominant role in the risk assessment. To illustrate, one of the most visible effects of the GDPR entering into force is that controllers and processors have to systematically consider the privacy impact posed by their processing operations.

Changing Threat Context. When the user base of a system changes (*data subject scale, data subject type*), as mentioned above, new types of users such as members of political campaigns or activists, may lead to the system becoming a more desirable and high-impact target for specific adversaries (*threat probability, threat vulnerability*).

Not only the adversaries, but also improvements in adversarial techniques (e.g., data analysis techniques) can later increase the impact of previously-collected or -published (pseudonymized) data sets (*threat vulnerability*). An illustration of this issue is the case of the successful de-anonymization of a data set released by Netflix [23]. A different example is the development of techniques to infer personal information from demand-response systems [24] such as appliance load monitoring [25] to reveal information about the types of devices installed in homes with smart energy meters to record energy consumption.

Changes in Countermeasure Effectiveness. As illustrated with the large numbers of issues in vulnerability databases [26], [27], the effectiveness of security countermeasures in existing software systems frequently changes over time (*threat vulnerability*). Hence, any up-front assessment that does not take into account changes in the effectiveness because of newly-discovered vulnerabilities, ignores these evolutions and severely underestimates the risk. An illustration of this is the Equifax breach which

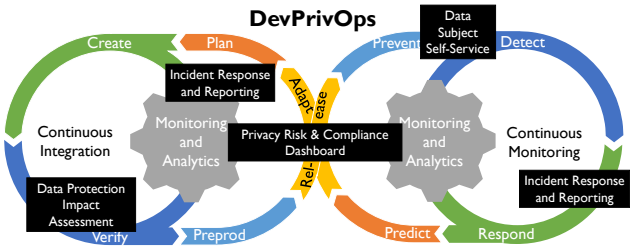

Figure 1. Overview of the DevOps life cycle overlaid with a number of enablers for continuous risk assessment (*DevPrivOps*).

was caused by a six-month-old unpatched vulnerability in the Apache Struts framework [28], [29].

4. Enablers for Continuous Risk Assessment

Contemporary systems have the ability to scale to proportions that can have a considerable impact on the privacy of their users. For example, Facebook [30] holds privacy-related code to a higher standard than code just related to functionality—illustrating how an operational metric, i.e. the amount of users (*data subject scale*) of the service, influences the development practices.

Tight integration of the notion of data protection risk into the practices of continuous development and integration (DevOps) leads to the vision of *DevPrivOps*. Figure 1 overlays the DevOps life cycle with a number of enabling systems that in our view will contribute to more reactive and continuous privacy risk assessment and management. In the remainder of this section, we outline these systems and their role in this context.

Privacy Risk & Compliance Dashboard. This is a dashboard offered to developers and operators that provides a continuous, run-time inspection view on the system, users, processed information, and overall compliance status. Such a dashboard system will provide developers and operators with accurate operational metrics about the user base, which in turn will enable continuous risk estimation. As such, these metrics will support developers with the prioritization of privacy efforts during development.

Not only does such a dashboard provide consistency and accuracy of data, it becomes almost trivial to monitor evolutions of the different risk factors over time and implement appropriate heuristics to conclude that operational context has shifted significantly to warrant an in-depth re-assessment of the overall risk [31].

Self-Service for Data Subject Rights. A second enabling component can be found in data subject self-service dashboards that allow data subjects to manage their risk and to exercise their data subject rights. These systems can generate dynamic and accurate descriptions of which data is processed and for which reasons, thereby ensuring compliance with transparency and user notification requirements (*data subject awareness and control*). Furthermore, they provide an effective interface to exercise a data subject's rights (i.e. access, rectification, etc.), as all the necessary infrastructure to keep track of all this information is already in place. For example, keeping track of data subject consent, which can be revoked any time, is essential to ensure the legality of the processing operations.

Dashboards tailored to data subjects enable further integration with existing risk assessment and decision support systems [8], [32], [33] aimed at data subjects.

Incident Response and Reporting. This support system is tightly integrated with incident and intrusion detection systems (IDSs) and focuses on the aspects of responding and reporting. The response to a security incident often involves, besides breach notification, a detailed assessment to determine the actual privacy impact of such a breach. An incident report can be generated, based on an assessment of the impacted risk factors to automatically determine the scope and impact of a data breach.

Furthermore, it can assist in the notification process to the affected data subjects (or authorities), by coordinating the notifications and ensuring that the appropriate countermeasures, such as invalidating the passwords of the impacted users, are executed properly.

5. Conclusion

In this position paper, we start from the observation that the privacy risk and privacy impact assessment approaches in the current state of the art are mainly intended to be used in a static requirement, design, or implementation context. We argue that this is inherently problematic, as a number of the risk factors essentially require prediction or estimation whereas we have shown many of these to be subject to unanticipated change. As such, we highlight the need for (and lack of) methods and tools to continuously revise the privacy risk in the context of an operational system. We argue that such methods are required to more reactively tune the actions of privacy hardening a system to its operational reality.

This vision is strongly aligned to the principles of continuous development and privacy-oriented DevOps (DevPrivOps) as it relies extensively on metrics and information obtained from the operational, run-time context of the system and uses these to steer further development.

References

[1] European Union, "Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016," *Official Journal of the European Union*, vol. 59, no. L 119, pp. 1–88, may 2016.

[2] OECD, "The OECD Privacy Framework," *Organisation for Economic Co-Operation and Development*, 2013.

[3] A. Cavoukian, "Creation of a Global Privacy Standard," 2006.

[4] American Institute of Certified Public Accountants Inc. and Canadian Institute of Chartered Accountants, "Generally Accepted Privacy Principles," no. August, pp. 1–84, 2009.

[5] S. J. De and D. Le Métayer, "PRIAM: a privacy risk analysis methodology," in *Data Privacy Management and Security Assurance*. Springer, 2016, pp. 221–229.

[6] L. Sion, D. Van Landuyt, K. Wuyts, and W. Joosen, "Privacy risk assessment for data subject-aware threat modeling," in *2019 IEEE Security and Privacy Workshops (SPW)*. IEEE, 2019.

[7] J. I. Hong, J. D. Ng, S. Lederer, and J. A. Landay, "Privacy risk models for designing privacy-sensitive ubiquitous computing systems," *Proceedings of the 2004 conference on Designing interactive systems processes, practices, methods, and techniques*, 2004.

[8] S. J. De and D. L. Metayer, "Privacy risk analysis to enable informed privacy settings," in *2018 IEEE European Symposium on Security and Privacy Workshops (EuroS PW)*, April 2018, pp. 95–102.

[9] A. Cooper, H. Tschofenig, B. Aboba, J. Peterson, J. Morris, M. Hansen, and R. Smith, "RFC 6973: Privacy considerations for Internet protocols," *IETF*, 2013.

[10] S. Hernan, S. Lambert, T. Ostwald, and A. Shostack, "Threat Modeling: Uncover Security Design Flaws Using The STRIDE Approach," *MSDN Magazine*, vol. 6, nov 2006.

[11] M. Deng, K. Wuyts, R. Scandariato, B. Preneel, and W. Joosen, "A privacy threat analysis framework: supporting the elicitation and fulfillment of privacy requirements," *Requirements Engineering*, 2011.

[12] R. J. Cronk, *Strategic Privacy By Design*, 2018.

[13] NIST, "NIST Privacy Risk Assessment Methodology (PRAM)," Feb. 2019.

[14] D. J. Solove, "A Taxonomy of Privacy," *University of Pennsylvania Law Review*, vol. 154, no. 3, pp. 477–564, 2005.

[15] J. Freund and J. Jones, *Measuring and managing information risk: a FAIR approach*. Butterworth-Heinemann, 2014.

[16] S. Brooks, M. Garcia, N. Lefkovitz, S. Lightman, and E. Nadeau, "An introduction to privacy engineering and risk management in federal systems," NIST, Tech. Rep. NIST IR 8062, Jan. 2017.

[17] M. S. Lund, B. Solhaug, and K. Stølen, *Model-driven risk analysis: the CORAS approach*. Springer Science & Business Media, 2010.

[18] S. Abu-Nimeh and N. R. Mead, "Privacy risk assessment in privacy requirements engineering," *2009 2nd International Workshop on Requirements Engineering and Law, RELAW 2009*, pp. 17–18, 2009.

[19] S. J. De and D. Le Métayer, "A Refinement Approach for the Reuse of Privacy Risk Analysis Results," in *Privacy Technologies and Policy*, E. Schweighofer, H. Leitold, A. Mitrakas, and K. Rannenberg, Eds. Cham: Springer International Publishing, 2017, pp. 52–83.

[20] NIST, "NIST Privacy Framework: A Tool for Improving Privacy through Enterprise Risk Management," NIST, Tech. Rep., Jan. 2020. [Online]. Available: https://www.nist.gov/system/files/documents/2020/01/16/NIST%20Privacy%20Framework_V1.0.pdf

[21] A. Endres and H. D. Rombach, *A handbook of software and systems engineering: Empirical observations, laws, and theories*, 2003.

[22] "Eu commission to staff: Switch to signal messaging app," https://www.politico.eu/article/eu-commission-to-staff-switch-to-signal-messaging-app/, Feb 2020.

[23] A. Narayanan and V. Shmatikov, "Robust De-anonymization of Large Sparse Datasets," in *2008 IEEE Symposium on Security and Privacy (sp 2008)*, May 2008, pp. 111–125, iSSN: 2375-1207.

[24] M. A. Lisovich, D. K. Mulligan, and S. B. Wicker, "Inferring personal information from demand-response systems," *IEEE Security Privacy*, vol. 8, no. 1, pp. 11–20, Jan. 2010.

[25] G. W. Hart, "Residential energy monitoring and computerized surveillance via utility power flows," *IEEE Technology and Society Magazine*, vol. 8, no. 2, pp. 12–16, Jun. 1989.

[26] "Common Vulnerabilities and Exposures (CVE)," Available from MITRE, 2019. [Online]. Available: https://cve.mitre.org/index.html

[27] NIST, "National Vulnerability Database (NVD)," 2019. [Online]. Available: https://nvd.nist.gov/

[28] "The apache software foundation confirms equifax data breach due to failure to install patches provided for apache® struts™ exploit," https://blogs.apache.org/foundation/entry/media-alert-the-apache-software, Sep 2017.

[29] "CVE-2017-5638," MITRE, 2017. [Online]. Available: https://cve.mitre.org/cgi-bin/cvename.cgi?name=CVE-2017-5638

[30] D. G. Feitelson, E. Frachtenberg, and K. L. Beck, "Development and Deployment at Facebook," *IEEE Internet Computing*, 2013.

[31] A. Barth, A. Datta, J. C. Mitchell, and H. Nissenbaum, "Privacy and contextual integrity: framework and applications," in *2006 IEEE symposium on security and privacy (IEEE S&P'06)*, May 2006.

[32] S. J. De and A. Imine, "To reveal or not to reveal: balancing user-centric social benefit and privacy in online social networks," in *Proceedings of the 33rd Annual ACM Symposium on Applied Computing*. ACM, 2018.

[33] F. Karegar, N. Gerber, M. Volkamer, and S. Fischer-Hübner, "Helping john to make informed decisions on using social login," in *Proceedings of the 33rd Annual ACM Symposium on Applied Computing*, 2018.