

Towards Privacy Compliance: A Design Science Study in a Small Organization

Ze Shi Li^a, Colin Werner^a, Neil Ernst^a and Daniela Damian^a

^aUniversity of Victoria

ARTICLE INFO

Keywords:
requirements engineering
continuous software engineering
privacy
GDPR
design science

ABSTRACT

Context: Complying with privacy regulations has taken on new importance with the introduction of the EU's General Data Protection Regulation (GDPR) and other privacy regulations. Privacy measures are becoming a paramount requirement demanding software organizations' attention as recent privacy breaches such as the Capital One data breach affected millions of customers. Software organizations, however, struggle with achieving privacy compliance. In particular, there is a lack of research into the organizational practices and challenges involved in compliance, particularly for small and medium enterprises (SMEs), which represent a sizeable portion of organizations. Many SMEs use a continuous software engineering (CSE) approach, which introduces additional adoption and application challenges. For example, the fast pace of CSE makes it harder for SMEs that are already more resource constrained to prioritize non-functional requirements such as privacy.

Objective: This paper aims to fill a gap in the under-researched area of continuous compliance with privacy requirements in practice, by investigating how a continuous practicing SME dealt with GDPR compliance.

Method: Using design science, we conducted an in-depth ethnographically informed study over the span of 16 months and iteratively developed two artifacts to help address the organization's challenges in addressing GDPR compliance.

Results: We identified 3 main challenges that our collaborating organization experienced when trying to comply with the GDPR. To help mitigate the challenges, we developed two design science artifacts, which include a list of privacy requirements that operationalized the GDPR principles for automated verification, and an automated testing tool that helps to verify these privacy requirements. We validated these artifacts through close collaboration with our partner organization and applying our artifacts to the partner organization's system.

Conclusions: We conclude with a discussion of opportunities and obstacles in leveraging CSE to achieve continuous compliance with the GDPR. We also highlight the importance of building a shared understanding of privacy non-functional requirements and how risk management plays an important role in an organization's GDPR compliance.

1. Introduction

In 2016, the European Union (EU) passed a privacy law known as the General Data Protection Regulation (GDPR) that came into full force in the EU in 2018 [23]. The enactment of the regulation created an urgent need for software organization to comply with the requirements of the regulation or risk penalties. Unlike any prior privacy regulations, the GDPR was a trailblazing regulation that directed stringent requirements upon organizations regarding how they process and/or collect personal data. Organizations outside the EU, including in countries such as Canada must also comply as long as they deal with EU citizen data.

As a *privacy* regulation, the GDPR represents one type of non-functional requirements (NFRs). NFRs, such as privacy and security, focus on the quality of the software, as opposed to specific functions [30]. NFRs are profoundly important for organizations, particularly as NFRs affect a software's architecture, yet, NFRs are also difficult to document and validate [2].

At the same time, many modern software organizations are increasingly adopting continuous software engineering practices (CSE) [25, 27] as a form of software development practice in search of promised benefits such as rapid feed-

back and fewer integration problems [27]. As part of CSE, these organizations adhere to the principles of fast and frequent builds, and automated and frequent deployments [27]. The fast moving and competitive landscape in software development makes it ideal for organizations to adopt CSE to align software with customer feedback. However, as organizations leverage rapid feedback to improve its product and find a suitable market, non-functional requirements (NFRs) are often neglected [31]. Managing NFRs in a CSE environment is not without challenges [67], in particular with respect to verifying an NFR through automated means.

In a bid to reduce the risk of penalty and obey the GDPR, organizations embarked on achieving compliance, which can require adjusting their software and information processing practices to satisfy with the new requirements imposed by the GDPR. This compliance is costly, but large enterprises can likely absorb this, or already face similar compliance challenges with other regulations. For instance, a report in 2018 by Veritas claimed that the average organization expected to spend 1.3 million Euros on GDPR compliance initiatives ¹. Such amount is much less palatable for small- and medium-sized enterprises (SMEs) that are more resource constrained. Moreover, a mistake

¹<https://www.veritas.com/content/dam/Veritas/docs/infographics/gdpr-infographic-en.pdf>

with respect to requirements may be particularly costly for a small, agile organization [3]. SMEs are likely to have fewer resources than a large organization to direct towards compliance and development, and may experience more difficulty with the GDPR [57], and hiring a dedicated team of lawyers and privacy experts may not be feasible.

To understand the reasons behind non-compliance, studies about the GDPR have often relied on surveys and interviews to study general compliance challenges [57, 51]. While some research has investigated tool assisted GDPR compliance [49, 46, 34], to the best of our knowledge there are no empirical insights on how organizations respond to GDPR compliance challenges. In particular, no study investigated how SME organizations adopt GDPR in their system and process, or challenges they face in their CSE practices.

In this paper, we report on an in-depth 16-month long investigation, using design science (based on Hevner et al. [36]) about an organization's journey towards GDPR compliance. Our collaborating organization, Gamma², is a SME (initially a startup) that has a large number of EU-based users. In addition, our collaborator makes extensive use of CSE practices such as daily builds, automated builds, and automated deployments, enabling us to also investigate GDPR compliance challenges specific to CSE environments. Prior to this work, we have conducted a couple of other studies with this collaborator that pertains to their continuous practices and how they deal with NFRs [67, 68]. While those previous studies share similarities with this study with respect to CSE and NFRs, the focal point of this study pertains to our collaborator's GDPR compliance. Through this study we bring the following contributions:

- we describe three GDPR compliance challenges based on a detailed exploration of a SME's GDPR adoption process; item we describe the affect of CSE on a SME's GDPR compliance;
- we develop a list of privacy requirements that are derived from GDPR principles important to our collaborating organization. The requirements are automatically verified to help identify GDPR violations; an automated testing tool to help test the list of privacy requirements is also presented as it served as the means to realize verification of the privacy requirements;
- we present an empirically-grounded discussion on the opportunities and obstacles of leveraging CSE to achieve continuous compliance with the GDPR;
- we discuss how our collaborating organization tries to mitigate GDPR risk; we detail three strategies our collaborator uses to offload GDPR risk to other parties to manage risk.

²Real name and some identifying details have been changed for confidentiality.

We first introduce the background and related work in Section 2. In Section 3 we describe our design science methodology, and in particular, our process for acquiring an in-depth understanding about our collaborating organization in order to develop and evaluate our design science artifacts (described in Section 5). In Section 4 we describe our collaborator's challenges with compliance, before we discuss the relationship between continuous compliance and continuous software engineering, building continuous compliance, establishing shared understanding of the GDPR, and managing GDPR risk in a small organization. Subsequently, we describe the limitations of our study in Section 7. Finally, we summarize our research study and conclude this paper in Section 8.

2. Background and Related Work

2.1. General Data Protection Regulation (GDPR)

The GDPR is a comprehensive set of privacy regulations that grants privacy rights and protections to individuals and correspondingly prescribe restrictions on how organizations must treat personal data [23]. When the EU enacted the GDPR, an organization that operates in the EU or collects and/or processes data from EU customers must be GDPR compliant. As a result, any organization that deals with customer data from the EU must abide by the GDPR or potentially face severe consequences. The GDPR prescribes a maximum fine of the higher of either 20 million Euros or 4% of annual revenue based on the severity of non-compliance. Due to the scope of the GDPR requiring *any* organization collecting personal data to adhere to regulations, even organization not in the EU must also comply. The GDPR lists six main data protection principles including: lawfulness, fairness, and transparency, purpose limitation, data minimization, accuracy, storage limitation, and integrity and confidentiality. Overall, the GDPR is a complex law that consists of 99 articles and 173 recitals [23], written in legal speak.

Organizations were initially given a two year grace period to prepare for the GDPR. As the replacement to the 1995 EU Data Protection Directive, the GDPR united privacy regulations in the EU under one umbrella regulation [22]. A unified regulation means that an organization can streamline its treatment of privacy and just focus on one central privacy regulation in the EU.

Other governments have passed laws that mimic and further ratchet the GDPR's stringent regulations. For example, the CCPA³ and SHIELD Act⁴ are laws from two different United States states. Therefore, studying GDPR compliance in software organizations and staying on top of privacy compliance should help prepare organizations for future privacy regulations.

2.2. Non-Functional Requirements (NFRs) in CSE

NFRs, also known as quality attributes or architecturally significant requirements, can profoundly affect a system's ar-

³https://leginfo.ca.gov/faces/billTextClient.xhtml?bill_id=201720180AB375

⁴<https://www.nysenate.gov/legislation/bills/2019/s5575>

chitecture [6]. The GDPR, in essence, is a privacy NFR that severely impacts software organizations that collect and/or process data from EU users. Apart from privacy, other common NFRs include performance, scalability, and availability. While many studies have tried to characterize NFRs [6] [29] [30], NFRs are often still difficult to enforce and validate in practice [7]. As manual methods are the default strategy often employed to test NFRs [10], organizations often suffer from a lack of consistency in testing and significant demands for time [62].

Even privacy by design (PbD), a concept proposed to increase privacy in software [11], is not without criticism for “vagueness” [35] and difficulty of implementation [59]. Moreover, in agile startup organizations, it has been found that NFRs are frequently neglected [31]. However, late testing of NFRs is not ideal and identified as less successful than early testing [50]. In addition, delayed treatment and testing of NFRs may have serious consequences as rapid releases without consideration for NFRs may allow “resource and performance creep” [55].

While testing in continuous integration (CI), where CI is a CSE practice, is rather understudied [69]; although Yu et al.’s [69] recent systematic literature review reveals that CI may assist in testing of NFRs. The authors focused on nine different NFRs and found that CI is underutilized for testing NFRs, despite CI’s potential benefits. Yet, there are several challenges that may limit CI’s usefulness for testing NFRs, including the lack of supporting CI tools that are trivial to integrate and natural difficulty of testing and automating some NFRs as NFRs are often cross-cutting and cannot be broken into small modules [69].

We previously studied the management of NFRs in CSE [67, 68], specifically, we partnered with our collaborator in these studies. While those studies related to CSE and managing NFRs, they are separate from this study nor do they share data. We conducted separate data collection for those studies from this study. The goal of this study is to investigate the GDPR compliance of our collaborator through an in-depth study of the company.

David Farley [16], a pioneer of the continuous software engineering movement, also advocates for a particular CSE practice known as continuous compliance [26]. Farley argues that CSE is paramount to an organization’s regulatory compliance because CSE offers the ability to keep an orderly audit trail of code changes and managing environments and testing with automation [16]. While research on continuous compliance is still in its infancy, there have been efforts to model continuous compliance [61]. Web-based continuous compliance testing tools were suggested as a strategy to manage risks especially in software infrastructure [61]. The five aspects of the continuous compliance model proposed by Steffens et al. [61] include compliance rules, compliance tests, software components, software system, and continuous compliance testing.

2.3. Current GDPR Research

While numerous surveys have been used to gauge the level of GDPR compliance in organizations [48, 13], research into how small CSE practicing organizations comply with the GDPR is limited. We discuss related work regarding research about the GDPR including state of compliance in organizations, privacy approaches developed before and after the GDPR, and takeaways for organizations complying with the GDPR.

2.3.1. State of compliance in organizations:

Numerous organizations are not GDPR compliant post GDPR deadline [48]. In fact, some organizations may never be fully compliant [13]. In particular, smaller organizations that did not previously take appropriate security and privacy measures, may feel burdened by GDPR compliance [57]. To the best of our knowledge, the GDPR makes it clear that organization must be “fully” compliant for as long as they are either collecting or processing data. The GDPR does not specify a detailed breakdown on the “levels” of compliance. However, it is reasonable to assume that the more comprehensive and complete an organization’s software compliance to the GDPR, the less likelihood they will experience penalties if at all.

While numerous privacy frameworks and approaches for the GDPR have been developed [4, 5, 9, 17, 40, 43, 49, 53, 63, 65, 66], there is little empirical evidence indicating that there exists a robust, comprehensive approach that is universally adopted by software organizations. In particular, there is little evidence that shows continuous practicing SMEs are adopting these frameworks. The adoption of privacy frameworks and tools are often limited by factors such as relevance to the GDPR, existing as a prototype, or lack of practical implementation details for an organization. Although some privacy tools may capture significant elements of the GDPR, the magnitude of resources required to adopt these tools may not be possible for a small, resource constrained organization. Below, we describe some of the proposed frameworks and tools.

2.3.2. Privacy approaches developed prior to the GDPR:

Huth et al. [40] compared 8 privacy engineering approaches developed prior to the existence of the GDPR to gauge their support for the GDPR. Privacy approaches reviewed include the influential and highly cited works of Deng et al. [18] and Spiekermann et al. [60]. Deng et al.’s LINDDUN methodology aims to identify privacy threats in a system through analysis of the system’s data flow diagram in consideration of seven privacy properties [18], but the methodology was developed prior to the GDPR [18] and omits several GDPR privacy properties [40].

2.3.3. Model driven compliance frameworks and tools:

There are also theoretical frameworks and models for tool assisted legal compliance [49, 46, 34]. Palmirani et al. [49] proposed a framework to check compliance by modeling GDPR text into legal concepts and analyzing an or-

ganization's business process modeled by business process model and notation (BPMN). However, BPMN may not be feasible for resource constrained SMEs, as organizations must model their business process using this notation. Another framework monitors the logs of a system to determine GDPR compliance [46]. While monitoring logs may help detect violations of privacy, the framework is limited by its coverage as the framework only checks for actions that occurred in the system, but neglects other potential vulnerabilities in the system [46]. A related framework prescribes a methodology to specify, enforce, and check privacy policies for data intensive applications [34], albeit it is not specific to the GDPR. Likewise, a model based compliance approach models both the GDPR and an organization's legal and technical documents, such as privacy policies and requirements specifications [65]. However, all four theoretical frameworks [49, 46, 34, 65] are still considered proof of concepts and not deployed in a real world setting.

2.3.4. Takeaways and challenges for organizations complying with the GDPR:

In addition to frameworks and tools, several studies analyzed the GDPR's implications and challenges [1, 64, 33]. Holistically analyzing the GDPR, Tikkinen-Piri et al. [64] found twelve ramifications that an organization must be cognizant of, such as the need for an organization to designate a data protection officer if the organization conducts systematic monitoring of users or uses special categories of data. More importantly, Tikkinen-Piri calls for more empirical studies "conducted among personal data intensive companies" and especially in contexts such as a SME [64].

The GDPR prescribes multiple rights for users, but supporting and complying with these rights may be complicated and challenging. For instance, in an interview study focused on GDPR data subject rights, researchers formulated suggestions for organizations to adequately adhere to GDPR data subject rights [1]. From the study, Altorbq et al. [1] identified twelve challenges (e.g. service availability and storage location, flexibility and standard agreements, roles, responsibilities, and expectations, and understanding and engagement) and developed fourteen recommendations to help mitigate these challenges, demonstrating that just adhering to data subject rights granted by the GDPR is non-trivial. In our research, we also noted the challenge of 'awareness and knowledge', which is similar to understanding and engagement. However, our research differs from this study as our focal point is not pinpointed on data subject rights compliance, but rather an organization's complete compliance with emphasis on GDPR data processing principles. Issues with awareness and knowledge does seem like a common problem in regulatory compliance, in a systematic literature review on security and privacy in electronic health records, it was found that a only a few articles indicated the necessity of employee training in security and privacy [24]. Furthermore, similar to our collaborating organization, a sizable contingent of the articles in the study mention the use of pseudo-anonymization [24].

2.3.5. Developing requirements for organizations:

Regulatory language is known to be ambiguous [8], and several works have investigated methods for better defining requirements in health care contexts [8, 45]. Moreover, there are suggestions for leveraging automated test suites to ensure that health care systems adhere to regulatory frameworks [47]. There are some proposed GDPR-specific privacy frameworks to guide compliance [9] and help elicit requirements from an organization [5]. The appropriateness of the requirements were validated with privacy experts, but the requirements lacked clear cut measurables for validation and framework steps. Moreover, the framework steps were high level with little implementation details. Ringmann et al. [53] defined technical requirements that served to help make a software GDPR compliant, albeit with limitations. The requirements are relatively generic as the authors wanted the requirements to apply to as many organizations as possible [53]. Similarly, Hjerpe et al. performed a single case study with a service oriented SME and identified GDPR requirements based on constraints that applied to the SME [38]. Subsequently, Hjerpe et al. implemented these requirements by modifying their collaborating organization's architecture. However, this work provides little additional empirical insight on the organization's GDPR compliance adoption process, such as encountered practices and challenges. To the best of our knowledge, there are no other empirical studies that investigated the GDPR adoption journey in CSE practicing SMEs.

2.3.6. Research gap addressed by our work:

Hence, our research fulfills several research gaps in the current literature. We answer Tikkinen-Piri et al.'s [64] call for more empirical research into GDPR compliance and challenges encountered by SMEs, which is an understudied context. We conducted an in-depth empirical study with a CSE practicing SME to identify GDPR compliance practices and challenges. In addition, literature on the GDPR has thus far neglected CSE practicing contexts. This research gap is concerning given the widespread usage of CSE in software organizations and CSE's propensity to impact all levels of an organization including planning, development, testing, and deployment. Finally, we add to the scarce empirical evidence on the actual use of privacy-related artifacts and supporting tools to achieve compliance in software organizations.

3. Research Methodology

To conduct our research we used an ethnographic influenced design science methodology where we worked in close collaboration with Gamma, our collaborating SME that employs CSE practices. Design science [56, 36] was relevant and most appropriate for our investigation since our research approach was influenced by a problem solving mentality. Our design science research was done in two methodological stages: an initial problem characterization stage and the development and evaluation of design science artifacts stage.

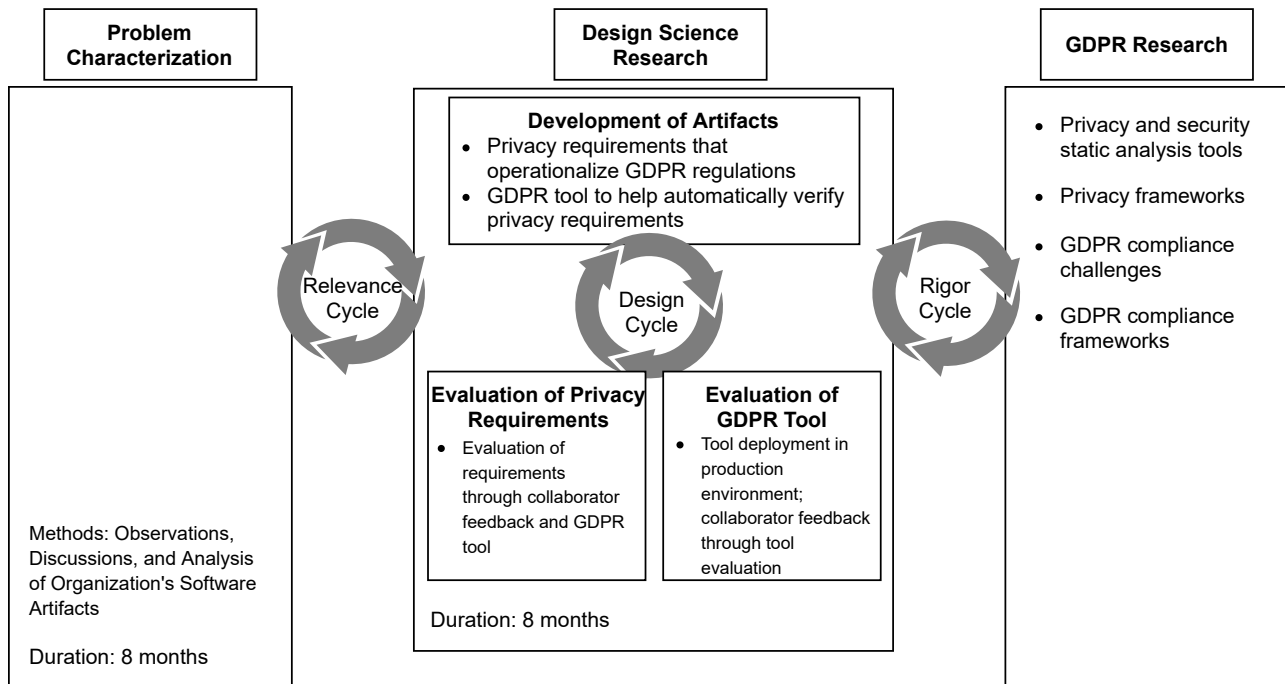


Figure 1: Design Science Methodology, based on Hevner et al. [36]. Design science is a research process emphasizing characterizing relevant problems (left side), iteratively developing and evaluating artifacts that potentially solve those problems (middle), and ensuring the findings and artifacts align with relevant literature (right side).

With operations in the EU, our collaborator had the urgency and desire to achieve GDPR compliance despite being a Canadian company. To facilitate our gathering of in-depth knowledge and understanding about our collaborator, the organization agreed to a long (16 months) research involvement for our study. To achieve a comprehensive understanding about our collaborator, co-authors of this paper, especially the first co-author physically spent extensive amounts of time at Gamma over an 18 month period to learn, observe, study, and advise our collaborator. For example, we physically observed the organization's sprint planning, demo, and retrospective meetings to understand how tasks are allocated and distributed.

Our methodological steps, including data collection and analysis methods are illustrated in Fig. 1. The design science artifacts produced in our research must be relevant to our collaborator and rigorously evaluated [36]. The design science artifacts produced from our research were validated in two ways. First, internal validation with the partner organization, and second, external validation through an analysis of existing literature regarding the GDPR, and methodologies and tools designed to handle privacy. In particular, we ensured external validation by reviewing existing literature on GDPR implications for organizations, and privacy frameworks and methodologies designed to help achieve compliance. As shown by Fig. 1, we consulted relevant GDPR research in our bid to develop the two novel artifacts. We describe our methodology in detail, after we introduce details about our partner organization.

3.1. Research Setting

Gamma⁵ is a data gathering and analysis startup founded within the last decade. Gamma's primary business is to develop in-house software leveraging cloud-based platforms; these platforms contribute to their primary source of revenue. Gamma implements CSE to a higher degree, including automated builds and testing (including feedback) and varying levels of automated deployment based on the component. Gamma has grown from a small startup into a mature, established leader, and is implementing the recommended CSE best-practices to a high degree.

The emphasis on CSE is important as Gamma analyzes millions of data points each day, with a significant percentage coming from EU users, and must quickly react to changes in the data. Prior to the enactment of the GDPR, the organization dedicated some effort towards privacy NFRs, but the overall treatment of privacy was not yet adequate per GDPR regulations. Since the GDPR mandates compliance from any organization that collects personally identifiable data from EU citizen for commercial purposes, our collaborator was obliged to attain GDPR compliance. However, the employees in the organization did not have any GDPR specific training prior to the movement towards GDPR compliance. The data collected by our collaborator are immediately pseudonymized upon collection as an initial precautionary measure to protect privacy. To assist with the organization's compliance, our collaborator also hired

⁵Real name and some identifying details have been changed for confidentiality.

several different outside lawyers and consultants. While these external consultants sometimes provided conflicting answers to the same question, they provided a second opinion that the organization can use as a reference to compare. During our 16-month long study, our collaborator experienced immense growth. In particular, the number of employees in the organization grew from a dozen to several dozen.

Our collaborator's business approach is agile and fast moving with feature iteration essential to long-term survival. Our collaborator uses CSE practices, as advocated by Fowler [28], including automating the build and deployment, and keeping the build fast. Our collaborator also leverages tools, such as Jenkins, to automate software build and deploy software to production. Once code is committed and pushed to source control, our collaborator's deployment pipeline builds the code and runs automated tests against the code, if pertinent tests exist. Our collaborator heavily relies on cloud-hosted solutions, not only for data storage and analysis, but also for storing and managing its burgeoning infrastructure. Our collaborator has multiple categories of partners: 1) partners who receive data from our collaborator; 2) third-party services who provide infrastructure to collect, store, and process data; 3) partners who facilitate data collection.

Over the course of our research, the first author became deeply acquainted with our collaborator's processes, tools, and practices, by becoming part of the team and its activities. During the 16 month study, he spent at least 1 to 2 days per week immersed in our collaborator's offices observing and conversing with our collaborator employees and learning about the organization's business, software, and processes. To acquire a reasonable perspective of our collaborator's work, he participated in meetings, such as planning and retrospective meetings, and performed tasks, such as creating data flow documentation. We also received access to our collaborator's source control repositories, project management tools, and infrastructure hosted in the cloud.

When our research began, our collaborator was much smaller in size (i.e., a dozen of employees before tripling to a few dozen) and Gamma's employees often had a multitude of responsibilities and roles. The balancing of multiple responsibilities and roles frequently led to time pressures for Gamma employees.

3.2. Design Science: Problem Characterization

The first major step of design science is learning about the problem affecting the partner organization. Hence, we began by learning about how our collaborator prepared for the GDPR and the problems the organization experienced during the preparation, including the organizational and business contexts that made compliance challenging.

For this particular methodological phase, over a period of eight months, the first author spent 1 to 2 days per week at the organization.

We participated in weekly and monthly meetings,

observed numerous GDPR compliance-related discussions, and conversed with nearly every Gamma employee (i.e. more than 90%). The only reason we were unable to speak to every single employee is that a small portion of the workforce was not co-located. The first author kept a journal log of the events observed and insights learned during the time spent in the company. Since the first author was given free rein to directly ask any question to the employees, the first author could store clarification answers directly in the daily logs as well. In these daily logs, the first author recorded details including which employees met which employees, what tasks were assigned to employees, who was in meetings, and which roles a employees played. Moreover, the logs included answers we asked employees about the nature of their work and their role in compliance. The first author grouped these answers by main themes and referred to the contextual observations for more details. The second author also having in-depth knowledge about Gamma's organization and processes, provided a sanity check of the grouping of themes.

We learned about our collaborator's continuous practices – we analyzed source code and tools, to understand how our collaborator planned work, developed code, tested software, types of tools used to support our collaborator's work and, most importantly, the amount of preparation conducted for GDPR compliance.

Our initial interactions with the employees included prodding the employees to learn about the nature of their work and roles. We typically started out by asking each employee about their role, their day-to-day tasks, and the types of tasks there were undergoing in preparation for the GDPR compliance. Our observations of the employees occurred as they carried out their work, which included meetings conducted with other employees. For example, we observed the employees discussions surrounding the allocation of tasks during the sprint planning. In addition, we observed team meetings and project meetings between various members of the development members. Our observation of our collaborator was ongoing in nature and did not suddenly abruptly end when the problem characterization stage ended. We carried on the practice of keeping a journal log every time we were at the company.

In particular, we closely studied nine employees at our collaborator. We choose these employees as their work included development in some level of capacity as well as having direct access to these employees to ask for clarification if we ever had any questions. We had extensive access to these employees and could observe and converse with these employees in their work. Table 1 lists more details about these study participants. Among the details include their role, experience in our collaborator, and overall industry experience. For ethical and privacy considerations, each study participant is anonymized. In general, while a manager represents someone whose primary focus is managing developers or other employees, managers may still perform development and testing tasks. It is typical for an employee to balance multiple responsibilities in a startup organization like our

collaborator. In contrast to a manager, a developer represents someone who mostly works in development, testing, or operations. Through extensive observation of these nine employees and our collaborator as a whole, we identified three main challenges to our collaborator's GDPR compliance. To help increase rigor of the observation data gathered by the first co-author who became a close member of our collaborator during the study, the second co-author also spent substantial time collecting observational and discussion data from our collaborator. In addition, the second co-author validated the data collected by the first co-author. We elaborate on these findings in Section 4.

3.3. Design Science: Development and Evaluation of Artifacts

A key component of design science is to help alleviate the identified challenges once the (properly contextualized) industry problem is characterized. To this end we developed two design science artifacts that were influenced by our understanding of challenges at our collaborator: a set of privacy requirements that allowed the identification of GDPR exposures, and a tool to assist with automatically testing these privacy requirements (referred to as privacy NFRs henceforth).

These artifacts were developed and evaluated in the design cycle iterations in full collaboration with our collaborator's employees over a period of eight months. To assist with mitigating compliance challenges in our collaborator, it was paramount that our collaborator provided guidance and feedback in the evaluations of our artifacts. In total, we conducted four design cycle iterations with our collaborator each lasting 6-8 weeks. For each design science iteration, we spent several weeks developing the privacy requirements (artifact 1) and its corresponding automated tests in our tool (artifact 2) and then collecting feedback for 4-6 weeks. Specifically, after each design science iteration, we presented our artifacts to at least one of collaborator's main employees tasked with GDPR compliance and elicited feedback from them. Based on direct feedback from the meeting with our collaborator's GDPR compliance main correspondences, and insights gathered from applying our artifacts to Gamma's production software system, we incorporated Gamma's suggestions to the next iteration of our artifacts. During the development and evaluation of artifacts step of our research, we visited our collaborator 1 or 2 times per week similar to the problem characterization stage. Our continuous interactions with our collaborator enabled us the opportunity to converse and meet with our collaborator's employees to gather feedback as needed in addition to our regular meetings with our collaborator's main GDPR compliance correspondents. Our design science artifacts are described in more details in Section 5.

4. Problem Characterization: Understanding GDPR Compliance Challenges at Gamma

Below we describe the main findings that we observed in the problem characterization stage of our research as shown

Table 1

Participant's Role and Experience. Note: All participants have limited or non-existent legal experience.

Id	Role	Role Exp.	Exp. in Org.	Industry Exp.
P1	Developer	Medium	< 5 years	5-10 years
P2	Developer	Senior	> 5 years	5-10 years
P3	Manager	Senior	< 5 years	10-15 years
P4	Manager	Senior	> 5 years	5-10 years
P5	Developer	Medium	> 5 years	5-10 years
P6	Developer	Senior	< 5 years	5-10 years
P7	Developer	Medium	< 5 years	5-10 years
P8	Developer	Senior	< 5 years	> 15 years
P9	Manager	Senior	> 5 years	5-10 years

by Figure 1. Specifically, we describe in detail the three main compliance challenges experienced by Gamma. After describing the challenges of complying to the GDPR, we detail the affects of CSE on Gamma's GDPR compliance, which to Gamma's benefit were mostly positive. Finally, we explain how Gamma uses various strategies to try to mitigate GDPR risk.

The identified challenges with GDPR compliance are:

1. reliance on manual GDPR tests,
2. limited awareness and knowledge of privacy requirements, and
3. balancing GDPR compliance in a competitive data business.

To ground these challenges in the specific organizational context at Gamma, we outline in Table 2 a number of contextual factors that contributed to one or more of these challenges. The contextual factors are explained in more details in the following subsections. Each factor represents an aspect that contributes to the primary challenges experienced by our collaborator while trying to comply with the GDPR. Similarly, in Table 3, we list the study participants who directly reported or experienced these challenges. The quotes listed in subsequent sections are responses to questions we directly asked the study participants listed in Table 1. The questions in combination with our observations at the organization were recorded by us in daily logs. As a result, the results in Table 2 and 3 were analyzed based on daily logs that we recorded in our collaborating organization from observing employees and directly asking them questions. Below we describe the challenges in detail. In Section 4.4, we explain our findings regarding the benefits of CSE with respect to GDPR compliance at Gamma and in Section 4.5 we describe several factors that help our collaborator mitigate GDPR risk.

4.1. Challenge 1: Reliance on Manual GDPR Tests

A significant challenge that repeatedly hindered our collaborating organization's GDPR efforts was testing that privacy regulations prescribed by the GDPR were met. The difficulty of testing for GDPR compliance was exacerbated by the complexity of the GDPR that consequently resulted

Table 2

Relationship between observed GDPR challenges and organizational context of our collaborator. Contextual factors (rows) *contributes to* one or more specific GDPR challenges in the organization (columns). The challenges are grouped into 3 major categories: testing, awareness and knowledge, and business and workflow (i.e. subsections in §4). Note: these contextual factors were identified from our close observations and questions and answers with our collaborators. This data was recorded in daily logs.

	Challenges				
	Reliance on Manual GDPR Tests	Limited Awareness and Knowledge of Privacy Obligation	Awareness of Knowledge of Privacy Obligation	Balancing Compliance in a Competitive Business	GDPR in a Data Business
Number of GDPR regulations		X			
Ambiguity of GDPR		X			
Lack of legal training		X			
Lack of privacy experience		X			
Conflicting advice from experts		X		X	
Nature of business		X		X	
Size of organization	X	X		X	
Lack of time	X			X	
Context Increased growth of infrastructure and data	X			X	
Lack of shared understanding				X	

Table 3

Primary challenges from Table 2 observed by us or reported by each participant. Note: observation and response data was collected in Gamma and recorded in daily logs.

P#	Reliance on Manual GDPR Tests	Limited Awareness and Knowledge of Privacy Obligation	Balancing GDPR Compliance in a Competitive Data Business
P1	X	X	
P2	X	X	
P3		X	X
P4		X	X
P5		X	X
P6	X	X	
P7		X	
P8			X
P9	X	X	X

in our collaborator frequently relying on manual tests to fill the void. Manually testing for GDPR compliance is not a trivial task especially considering that our collaborator was a rapidly growing company whose developers were often hard pressed for time. With respect to GDPR compliance, employees often reflected that they have limited time to complete the tasks “*I would ... but I have no time*” (P2) and “*I wish I had more time*” (P6).

Manually testing for GDPR compliance was a rather

difficult experience for those employees tasked with the job as they had to conduct the compliance work on top of their typical responsibilities. Checking the compliance of Gamma’s system, especially its vast infrastructure, was primarily delegated to a few specific developers. During the compliance effort, these developers incurred the vast majority of work related to checking for GDPR compliance in the system. These developers had to perform this GDPR compliance work at the same time as conducting their other responsibilities and tasks. Therefore, time constraints and manual GDPR checks became a more prominent challenge for these employees, particularly for P1 and P2. We can see from Table 3’s “Testing” column, that this situation was reflected in our observations and self reported challenges from these employees. In contrast, the other developers and managers who had a much fewer number of tasks related to manual inspection of Gamma’s system for GDPR compliance, experienced less difficulty with respect to testing for GDPR compliance.

A critical factor for the difficulty of automating GDPR tests is that it is hard to convert GDPR regulations into automated tests. In particular, GDPR regulations sometimes affect multiple facets of a software system that make developing automated tests onerous. For instance, if our collaborator decided to cease collecting a data parameter or a data parameter is no longer deemed acceptable, a developer would need to manually check Gamma’s databases to verify the data parameter was no longer collected by our collaborator. Complying with the GDPR’s data processing principles relied primarily on manual enforcement and verification. In addition to GDPR data processing regulations, Gamma’s man-

ual process for handling compliance tests also extended to supporting data subject rights granted by the GDPR: “When a user sent a request to opt out to us, the emails come to me and I have to tell them how to opt out.” (P9) P9 would manually contact the user and provide instructions on how to remove themselves from Gamma’s data collection and processing. This task may not be GDPR specific, but it is still important to our collaborator’s GDPR compliance because failure to comply with the users’ requests may expose the organization to GDPR fines. Ultimately, such user request can quickly add up over time if the number significantly increases.

In contrast to automated tests, manual tests are laborious, error prone, and time consuming [19]. It is very easy for a developer to check the wrong database or run the wrong query. Erroneously checking for compliance provides no benefit to an organization as decisions should not be based on inaccurate data. Hence, Gamma’s manual GDPR can hinder the organization’s compliance and result in rework or retesting of the software. In theory, to check that the software is compliant with privacy requirements, after every change to the database, a developer would have to conduct the same type of manual test after every code change. Due to the ad-hoc nature of the manual GDPR compliance tests that our collaborator conducted, it is difficult to gauge the total number of tests. While it is difficult to precisely quantify the exact number of manual tests that our collaborating organization used for GDPR compliance, we can reliably estimate based on observations and discussions with our collaborators that manual tests represent the vast majority of GDPR tests. We refer to any test for verifying GDPR compliance as a *GDPR test*. In particular, checking that third party libraries and frameworks adhered to GDPR principles was all manually conducted as well as checking that the organization’s infrastructure was secure and private. To our knowledge, very few automated existed and/or was used during the compliance process.

Our collaborator is a rapidly growing company who adheres to CSE principles and emphasizes fast-paced development to get its product to its customers, but we observed that the organization’s fast moving nature has negative effects on the completeness of its testing. While fast paced development provides the advantages of quicker time to market and faster feedback, we observed that it can be disadvantageous in our collaborator’s case because at the time of our study they did not have a complete test suite for its software. In particular, we saw that our collaborator had few tests related to verifying GDPR compliance. We refer to any test for verifying GDPR compliance as a *GDPR test*. When our collaborator began to strive towards GDPR compliance, our collaborator was challenged as the organization had few GDPR relevant tests in its disposal that could help check for compliance. Notwithstanding whether or not relevant GDPR tests exist prior to the compliance effort, our collaborator tasked several employees (mostly developers) to delve into Gamma’s software and infrastructure and check that the system is GDPR compliant. However, as we described in the

first paragraph of this section, our collaborator’s employees suffered from time constraints and limitations caused by predominantly manual means of verification with respect to checking GDPR compliance.

The growth of the number of our collaborator’s employees reflects the growth of our collaborator’s business, infrastructure and data. Yet, this fast pace growth of infrastructure and data further increased the difficulty of verifying compliance for our collaborator. For example, a rapid growth of 10 new databases in one week means that the organization has 10 extra databases that it must consider with respect to privacy safeguards. Among the many cloud services that are integral to Gamma’s business and software, our collaborator relies on a multitude of third-party services like Amazon Web Services (AWS), Google Cloud Platform (GCP), and Azure. our collaborator hosts more than fifty databases and over one hundred servers on a single third-party cloud service. However, without many automated GDPR tests, it was arduous for a developer to manually review all these databases.

4.2. Challenge 2: Limited Awareness and Knowledge of Privacy Obligations

Our collaborator struggles to properly identify privacy problems, largely due to the complexity and magnitude of the GDPR and inexperience dealing with privacy regulations. To make matters worse for long term privacy compliance, our collaborator must stay on top of upcoming privacy regulations and any updates to prior privacy regulations.

As illustrated by Table 3, almost every single closely studied participant reported or was observed to experience challenges with awareness and knowledge about the GDPR. The only exception was P8, who appeared to have zero to almost no work assigned towards GDPR compliance nor significantly impacted by the GDPR in his role. In contrast, every other participant 1-9 ostensibly experienced significant difficulty becoming familiar with GDPR regulations.

It is difficult to quantify exactly how the limited awareness and knowledge of privacy obligations related to the GDPR manifested within the organization. However, we observed several factors that could have played a role. First, we noted that the organization did not provide GDPR specific training for all the employees in the organization. Second, we observed there is an insufficient collective shared understanding on the expectations of privacy.

The employees lacked experience with privacy and did not have the processes in place for sharing privacy knowledge when they learned something new. As there was no systematic approach to disseminating knowledge, new insights acquired by developers regarding privacy would be lost. Without adequate communication and documentation, lack of shared understanding could manifest in undesirable rework of tasks [68]. We observed that when lack of shared understanding is present, developers working in cross-functional teams may develop features that are incompatible with each other.

In theory, to sufficiently address the GDPR, our collaborator

erator's employees should become reasonably knowledgeable about the GDPR, but attaining an adequate understanding was difficult in practice. Neither Gamma's developers nor managers are well-versed in legal language and do not have privacy specific training. In consequence, our collaborator's employees do not have the experience and training to draw clear-cut conclusions based on demands of the GDPR, which is not unique to this setting [15, 8]. Several participants expressed the difficulty of completing GDPR compliance tasks as “[Evaluating GDPR compliance of third party libraries and tools] is difficult because I am not an expert in the GDPR” (P1) and “Interpreting the rules and regulations [was challenging]. The rules weren't clear on what can be collected and what is considered private” (P9).

Much of the GDPR regulation is ambiguous and like all regulations, depends on subsequent court cases to establish precedent for proper interpretation. At least three of our participants specifically agreed that compliance is hard in the face of conflicting answers provided by external consultants, hired by our collaborator, to the same question.

This difficulty is compounded even further with disagreements between different domain experts such as developers and lawyers. We observed conflicting arguments made regarding whether our collaborator's system had adequate privacy measures or not. One group of external consultants would feel strongly that the organization had robust privacy considerations, whereas another group of consultants would list off a series of recommendations that the organization must adopt. In particular, one point that garnered frequent debate was what types of data are allowed to be collected and/or processed under the GDPR. Without concrete guidance from the GDPR on what types of data is collectable in certain circumstances, deciding whether to proceed with collecting a type of data becomes onerous.

However, Gamma's GDPR compliance is not only about the present, but also making amends and preparations for the future. In particular, our collaborator must recognize upcoming regulations and anticipate amendments or new precedence setting interpretations of current privacy laws. Taking a proactive approach to long term privacy compliance, our collaborator should “stay up to date with the regulations. Put efforts in research and implement the changes” (P7). Our study participants recognized that the GDPR was only the beginning of a growing trend of privacy regulations widely advocated by regulators and the general public: “No [not aware of any new regulations], but US will probably adopt something similar to the GDPR”, yet, none of our participants could definitively describe nor list an upcoming privacy regulation. Solely staying up to date with the GDPR is already quite difficult, “[A large challenge is knowing] changes to the GDPR. Especially minor changes [and amendments] can be difficult for companies to find out” (P4). As there were no explicit requirements that came with the GDPR, asking developers to interpret ambiguous laws may lead to missing, vague, or even wrong requirements.

Justifying data collecting practices, and educating users

on our collaborator's data collection purposes, is also strenuous. As explained by P9 “a user needs to be educated on why we are collecting data”. Without sufficiently explaining the purpose of Gamma's data collection, a user may refuse to grant permission for Gamma's data collection. In the worst case, a user may even report our collaborator to a data protection agency. User concern is important to our collaborator's business. Effectively educating and communicating our collaborator's collection purpose to users is critical for our collaborator.

4.3. Challenge 3: Balancing GDPR Compliance in a Competitive Data Business

Properly valuing GDPR compliance is a challenge. GDPR compliance is only one part of Gamma's business concerns. In particular, as a SME with many competitors, our collaborator must constantly contemplate its business needs, such as developing new features and releasing updates based on customer feedback. To ensure survival and increase business, our collaborator must “stay competitive in terms of [volume of data] collected and presented, while respecting privacy concerns of anonymization” (P5). To stay competitive against other companies, our collaborator needs to continue increasing the amount of data collection, while also considering demands of the GDPR.

The feeling of underscoring the importance of balancing the business needs and GDPR compliance is high in priority for managers, but this feeling is less so experienced in developers. In reference to Table 3, all of the participants who have a primary role as manager reported “business and workflow” as a GDPR compliance challenge. In contrast, most developers have less direct impact and relationship with the business. The difference in roles is reflected by Table 3 as two developers reported “business and workflow” as a challenge. Simply put, developers do not need to have a holistic view of the business and be mindful of the impact of GDPR on the business on a continual basis.

However, P9, a manager, explains that earning the trust of users and receiving consent is paramount to the success of the organization, but even if a user consents, “there may be a regulator who says we can't collect this data”.

After all, at the time that we conducted our study, the GDPR was still in its infancy, and there was little precedence regarding the boundaries of privacy enforcement from the authorities. A study participant characterized the waiting as “we are waiting for if there will be some sort of litigation in the industry”. Our collaborator may present legitimate grounds for data collection and take adequate steps to safeguard its data and systems. Nonetheless, a privacy regulator can ultimately rule that our collaborator must rescind the data collected and modify the types of data collected in the future. Such ruling can have devastating effects on the competitiveness of our collaborator's business. Additionally, more onerous regulations may be enacted that our collaborator must comply.

Since our collaborator's system already exists, becoming GDPR compliant meant re-designing large aspects of

the system to comply with the GDPR. Some aspects of our collaborator's system already existed for years. At this stage, significant modifications to the system architecture is not trivial. P4 specifically acknowledged that revamping Gamma's legacy systems for the GDPR is quite difficult. Our collaborator feels that managing privacy concerns is easier in new projects beginning with an elevated priority for privacy, i.e., "build privacy in" [12]. In addition to managing Gamma's own systems for compliance, our collaborator must also vet its partners to satisfy the GDPR's emphasis on shared responsibility between controllers and processors. P3, a manager, expounded "[it's challenging] making sure that partners who receive data are compliant". Checking that partners are compliant is even more a hassle as the number of partners increase.

4.4. Windfall: CSE affect on Gamma's Compliance

Gamma's use of CSE practices resulted in advantages for our collaborator that are typical for an organization adopting CSE, namely quick release [37] and feedback [54]. While these advantages are desirable for organizations, we found that the advantages also extend to assisting with GDPR compliance. In particular, CSE's rapid feedback facilitates acquisition of usage information regarding the GDPR and quick release supports our collaborator making necessary changes if any GDPR related issues require immediate updates. Moreover, as alluded to by the first challenge identified in our problem characterization: reliance on manual GDPR tests, our collaborator does not have a large breadth of automated tests. While the organization has some automated tests, the tests are not yet comprehensive nor abundant.

Study participants support our findings: "Through CI, [our software] can be generated, modified, and fixed within a couple of hours" (P5) and "[allows involvement] with external stakeholders" (P9). However, the compliance benefits are contingent on employees possessing a reasonable amount of GDPR knowledge. For a developer to quickly implement changes and make sense of the feedback, the developer must first understand the expectations of the GDPR.

Without discounting the aforementioned advantages afforded by CSE, we also noticed a challenge exerted by CSE. In theory, CSE should break down silos between developers [39], but we noticed the existence of a lack of shared understanding in our collaborator with respect to privacy. This challenge relates to the second challenge identified in our problem characterization: limited awareness and knowledge of privacy obligations. The potential dearth of shared understanding of NFRs in CSE was previously reported by Werner et al. [68]. At Gamma, we observed the lack of shared understanding often manifesting in unsafe assumptions. For example, we saw instances of developers *assuming* processes and system elements not directly tied to them are secure and compliant. A developer explained that GDPR compliance was not a significant concern because their work dealt with data that was already pre-processed. The devel-

oper assumed that prior processes contained safeguards and checks that would ensure the data is GDPR compliant. However, the developer's assumption implies that the organization has mechanisms in place to ensure this assumption is accurate and traceable, which our collaborator does only partially. While one strategy to at least partially alleviate this privacy risk is ensuring that pre-processed data is GDPR compliant, a developer managing large-scale data processing should still have a reasonable understanding of the potential privacy implications.

4.5. Mitigating GDPR Risk

However, we observed several factors that our collaborator believes helps the organization mitigate the risks from the GDPR. Our collaborating organization is a small software organization based in Canada. While the organization has some customers based in the EU, most of the organization's customers are outside of the EU. Consequently, our collaborator believes that it has less exposure than a EU based organization. Moreover, our collaborator removed parts of its EU operations after the GDPR to better comply with regulations.

Another aspect that Gamma uses to lower GDPR risk is through offloading risk on other parties as part of managing GDPR risk. As described in earlier sections, Gamma relies heavily on third-party cloud services. Without third-party cloud services, Gamma would not experience such significant growth to its business. Gamma felt that these services provide an added benefit of GDPR compliant privacy safeguards. Likewise, some of Gamma's employees felt a sense of security due to extensive partnerships with large clients who should be adequately compliant. Gamma believes that these clients would have the resources to become compliant and also extend support to their smaller partners like Gamma and help with their compliance. During our study, Gamma also hired external consultants to conduct an external GDPR review. The external consultants were hired to provide a sanity check for Gamma that their GDPR compliance was adequate. To demonstrate that the organization passed the review, the consultants presented Gamma with a compliance certificate. For our collaborators, the certificate marked a stamp of approval from legal experts.

5. Artifacts of the Design Science Research

Recall that our design science approach (cf. Fig. 1) began with understanding the problem context at Gamma (Section 4). We identified three challenges with GDPR compliance: limited awareness and knowledge of privacy obligations, reliance on manual GDPR tests, and difficulty balancing GDPR compliance with the business and other work. In assisting the organization in addressing these challenges on its journey for GDPR compliance, we developed two artifacts: privacy requirements and an automated tool that tests for these privacy requirements. Altering Gamma's business as part of the third challenge is not within the scope of our research, and we therefore focused on mitigating the first two challenges with our two artifacts. Our first artifact (privacy

Table 4
Privacy Requirements Derived from GDPR Principles

RQ Number	Privacy Requirement	GDPR Principle
REQ1	A database must be encrypted for integrity	Integrity and Confidentiality
REQ2	Each server must exist with a purpose	
REQ3	Each server without purpose must be removed	
REQ4	Each server must have a corresponding cloud firewall	
REQ5	Each server storage must be encrypted	
REQ6	Each server storage must exist for a purpose	
REQ7	Each cloud firewall must use secure protocols inbound and outbound	
REQ8	Each cloud firewall must limit access to reliable sources	
REQ9	Each cloud firewall must limit outbound communication to reliable sources	
REQ10	Each load balancer must use end to end encryption	
REQ11	Each load balancer must use secure protocols	
REQ12	Each cloud storage resource must be encrypted	
REQ13	Each cloud storage resource must limit access from unapproved sources	
REQ14	Each cloud storage resource must limit modification and deletion from unapproved sources	
REQ15	Each access management resource must not grant all permissions	Data Minimization
REQ16	Each access management resource must not grant permissions to infrastructure resources	
REQ17	Each router must limit outbound communication to unapproved sources	
REQ18	Each database must not collect personal data types outside an organization's data collection purpose	Storage Limitation
REQ19	Each database tuple must not live indefinitely	

requirements derived from GDPR principles) mitigates the challenge of limited awareness and knowledge of privacy obligations. The artifact does this by developing a set of testable and measurable privacy requirements that are pertinent and important to our collaborator. Our second artifact (automated testing tool of privacy requirements) mitigates the challenge of reliance on manual GDPR tests by providing a tool to automatically execute tests so our collaborator does not need to dedicate additional overhead for testing compliance.

As per our methodology, Section 3.3, we completed four iterative development and evaluations stages for our artifacts over eight months. The evaluation stage for each iteration lasted roughly 4-6 weeks after several weeks of development. For each artifact, we first describe the artifact, followed by the process of its iterative development and evaluation as per our design science methodology.

5.1. Privacy Requirements Derived from GDPR Principles

Our first artifact of the design science research is a list of privacy requirements derived from GDPR principles. This artifact is intended to deal with challenges of GDPR awareness (by making GDPR regulations explicit), and acts as the crucial first step to automating the testing of compliance.

5.1.1. Creating the Artifact (Privacy Requirements)

Drawing from our problem characterization, we understood that our collaborator struggled to make sense of the

GDPR regulations, which impeded Gamma's compliance. In an ideal context, our collaborator is knowledgeable of *all* intricacies of the GDPR, but this was extraordinarily idealistic and impractical. To satisfy Gamma's compliance, we needed to focus on aspects of the GDPR that matter to our collaborator. Therefore, we developed a list of privacy requirements drawn from the GDPR and pertinent to our collaborator software and infrastructure. In addition, recognizing that our collaborator would not have extra resource capacity to manually verify these privacy requirements, we needed to consider the plausibility of automating testing of the requirements.

We used three properties as criteria to develop the privacy requirements: 1) a requirement is derived from a GDPR principle, 2) a requirement is important and relevant to our collaborator, and 3) a requirement can be verified using an automated testing tool of privacy requirements and through testing, it identifies potential GDPR non-compliance. Table 4 lists the GDPR principles we considered and the privacy requirements we developed. In describing these requirements, we describe the GDPR principles and give examples of the associated privacy requirements that operationalized these principles. While there are numerous possible GDPR requirements, not all requirements are relevant for our collaborator. Our collaborator acknowledges that it does not need to consider privacy principles that are out of scope for the company. Hence, the requirements listed in table 4 are all relevant for our collaborator.

The GDPR has six main data processing principles 1)

lawfulness, fairness and transparency; 2) purpose limitation; 3) data minimization; 4) accuracy; 5) storage limitation; and 6) integrity and confidentiality [23]. Based on input from our collaborator and our own observations, we scoped our initial effort to three GDPR principles: integrity and confidentiality, storage limitation, and data minimization.

The purpose of *integrity and confidentiality* is ensuring an organization adequately handles personal data, and safeguard that data from malicious attacks or accidental misappropriation. One example of a privacy requirement developed based on this GDPR principle is that databases must be encrypted for integrity. This was explained to us by two different participants: “*I added more encryption to the databases*” (P6) and “*[I worked on] disk and storage encryption*” (P2). REQs 1-17 are attributed to this GDPR principle as this principle encompasses almost every element of Gamma’s system. Prior to our study, Gamma’s compliance approach was rather ad-hoc and employees did not have a structured understanding of GDPR expectations. With our REQs 1-17 to clarify some “expectations” from the integrity and confidentiality with respect to our collaborator.

Our second operationalized principle, *storage limitation* represents the idea of keeping data no longer than necessary. An organization must ensure that it has a process to remove a datum after a period of time. For example, a datum is automatically removed after a year. Our collaborator collects a plethora data and the data should automatically be removed after a specified time frame.

Data minimization is our third operationalized principle. It prescribes that personal data should only be collected if necessary and relevant to an organization’s data collection purpose. Our collaborator collects a large assortment of data and data types. It is onerous for a developer to manually verify whether the organization is collecting more personal data than allowed.

Drawing on feedback from our collaborator, we chose not to operationalize three principles (i.e. lawfulness, fairness, and transparency, accuracy, and purpose limitation). Notwithstanding less relevance for our collaborator, these principles are also more subjective in nature. For instance, the accuracy principle prescribes that personal data must be kept up to date and inaccurate personal data is fixed or erased [23]. However, data collected by our collaborator is pseudonymized. Moreover, our collaborator has no desire and minimal ability to identify any data subject. If data were inaccurate for any particular reason, data subjects should experience minimal impact as there is little possibility for our collaborator to even identify a data subject, and there are no implications for vital interests nor monetary exchange between data subjects and our collaborator. Finally, we are not dismissing the importance of these three principles, we had to operate within the scope of our collaborator and what they deemed important. We acknowledge that a similar study in a different organization may prioritize a different set of GDPR principles.

5.1.2. Iterative Development and Evaluation of Privacy Requirements Derived from GDPR Principles

The privacy requirements in Table 4 were initially created from input from our collaborator and the three principles in focus. We then iterated the definitions of the privacy requirements with ongoing input from our collaborator. We used our second artifact, the automated testing tool of privacy requirements, as a research instrument to iteratively evaluate and refine the list of privacy requirements. Specifically, we checked whether automated testing the requirements could identify GDPR-compliance violations. We had continuous access to our collaborator’s employees throughout our research, which provided us with ample opportunities to discuss and acquire suggestions for refinement.

The first property of our criteria used to develop the privacy requirements was crucial because the goal is improving awareness and knowledge of GDPR privacy obligations in Gamma. If the developed requirements were not derived from the GDPR, then the requirements failed to meet the goal of making GDPR regulations explicit and raising our collaborator’s GDPR awareness. For instance, the GDPR suggests that encryption is a suitable organizational measure to protect privacy. The privacy requirement, “A database must be encrypted for integrity” was pertinent to the GDPR’s suggestions for suitable security measures. The example requirement also fulfilled the criteria’s second property as the requirement was about a crucial infrastructure element for our collaborator. Our collaborator relies on a substantial number of databases for its business. Therefore, our collaborator fully agreed with and supported this example requirement.

However, a privacy requirement derived from a GDPR principle did not guarantee that the requirement was relevant to our collaborator and important within the organization’s context. An representative example would entail any requirement related to something that our collaborator does not possess or use. Similarly, a requirement that lacked enough “importance” for our collaborator or occurred infrequently would fail to satisfy our criteria even if it was derived from a GDPR principle and relevant to our collaborator. For example, the requirement “Each load balancer must *only* use secure protocols” was derived from GDPR. The requirement should have been highly relevant to our collaborator because the organization uses load balancers and the automated testing tool flagged some load balancers initially verifying this requirement.

Yet, the organization did not agree with the requirement. The organization agreed with the requirement in theory as relying on secure protocols seem prudent and sensible. However, from a practical perspective, the organization could not support the requirement as it would have required changes that were not possible at the time. Our collaborator agreed with the requirement in principle, but viewed the requirement, in its original form, as something to maybe revisit at a later time. To satisfy the requirement’s relevancy and im-

portance to our collaborator, we revised the requirement by dropping the word ‘only’: “Each load balancer must use secure protocols”. This accounted for cases where a load balancer listened to both HTTP and HTTPS traffic. Essentially, our collaborator accepted the risks of load balancers that it perceived as otherwise secure. We discuss this acceptance of risk and highlight the organization’s risk management in more details later in the Discussion Section of the paper.

To satisfy the third property of our privacy requirements, each requirement had to be testable using our automated testing tool of privacy requirements. The tool provided a useful vehicle to test and evaluate our privacy requirements as the tool performed automatic verification of the requirements. When the tool is executed, it helps to affirm whether or not a requirement can be tested automatically. Additionally, if a requirement failed when the tool executed, it indicates an area of non-compliance within our collaborator. For instance, an employee once exclaimed, “*It is peculiar that [redacted]...that should have all been fixed a while ago!*” Our understanding of implications of the GDPR evolved throughout the research. Apart from our collaborator, we also refined our privacy requirements based on lessons learned from external events. For example, when the Capital One breach (a massive data breach affecting millions of Canadian and US banking customers whose social security/insurance and bank account numbers were compromised) occurred [14], that largely originated from misconfigurations of cloud infrastructure, we developed requirements that applied to access and modification rights (i.e. REQs 13-17).

5.2. Automated Testing Tool of Privacy Requirements

Our second artifact of the design science research, an automated testing tool that automates verification of privacy requirements, primarily fulfills our goal of helping reduce our collaborator’s manual testing of GDPR compliance. As the tool executes on its own without manual intervention, our collaborator is not further burdened with manual testing. We also developed a tool that could help Gamma build awareness of GDPR requirements as developers see the test log in Jenkins, a CI tool. In particular, the results of the tool help our collaborator directly identify where GDPR violations may exist in the system and isolate the problem. For our collaborator, the tool assists with raising awareness of defects within the organization’s system that otherwise may be overlooked. Finally, since our implementation produces a list of issues, we believe this will help with risk management.

5.2.1. Creating the Artifact

The tool is a series of Python scripts tailored for Amazon Web Services (AWS) and customized to verify the privacy requirements listed in Table 4. Therefore, the tool was designed to perform automated testing of these privacy requirements. When we modified or added a new privacy requirement in Table 4, correspondingly, we made relevant changes to the tool to test for the new privacy requirement. When

the tool is executed, the scripts use AWS’ API to acquire a detailed breakdown of Gamma’s entire infrastructure hosted on AWS. Every infrastructure that our collaborator stores on AWS is captured, including everything from databases, servers, and cloud storage. Upon acquiring the list of infrastructure resources, the tool iterates through the details of each resource and compares with the conditions prescribed by each relevant privacy requirement. Therefore, the tool compares each database to the privacy requirements that pertain to databases. If an infrastructure resource fails a condition, the tool stores the condition, the corresponding privacy requirement, and the resource in question. The failed condition is deemed an area of non-compliance that warrants further investigation from our collaborator. The tool does not compare an infrastructure resource to a non-related privacy requirement as it would not be useful. For example, checking that a database does not violate the conditions for REQ10 (i.e. load balancer) as the requirement does not apply to databases. After iterating through all the infrastructure resources and comparing with the conditions of their related privacy requirements, the tool generates a detailed report of all potential non-compliance areas. The report contains a breakdown of the non-compliance areas by various metrics such as location and infrastructure type. The report is stored by Jenkins, a CI tool, and is publicly available to any Gamma employee.

Ultimately, the tool provided a vehicle for us to execute automated tests for GDPR compliance in practice and gather empirical evidence of continuously applying a tool to verify GDPR compliance. The tool ran without requiring a human to trigger an execution, as Jenkins triggers the tool to run as often as Gamma desired, which our collaborator originally set at once per week.

Moreover, the tool helped validate whether Table 4’s privacy requirements are automatically verified by flagging non-compliance issues as reflected by these privacy requirements. Whenever potential GDPR issues in our collaborator’s infrastructure and code are found, our automated tool produces a list with detailed information about each issue, such as location, name, ID, type of resource, and pertinent GDPR principle, which allows a developer to investigate the problem in more detail.

5.2.2. Iterative Development and Evaluation of Automated Testing Tool of Privacy Requirements

To iteratively develop and evaluate the tool, we collected and acted upon feedback from our collaborator while analyzing the results produced by the tool. We iterated and analyzed four versions of the tool. The feedback helped evaluate and improve the accuracy and efficiency of the tool. For instance, if the tool found 11 load balancers of a specific type, we manually verified that there were indeed 11 load balancers. Since the tool served as a means to perform automated testing of the privacy requirements from our first artifact of the design science research, the tool is a mirror of those requirements and reflects any changes to those

requirements. In essence, the tool checked for potential non-compliance and provided meaningful details that can help an employee investigate the non-compliance problem.

Table 5 displays the average number of defects identified by our tool for each iteration 1-4. They are grouped by various resource types, which are linked to the resource types described by our privacy requirements in Table 4. After each iterative development of the tool and successive execution of our tool on Gamma's production system, we presented and discussed the tool execution results with at least one experienced Gamma employee tasked with realizing GDPR compliance. In particular, we verified each defect identified by our GDPR tool to identify whether our tool and in essence our requirements helped to capture legitimate issues with our collaborator's system. Our collaborator previously conducted ad-hoc manual testing of their GDPR compliance, but had no systematic approach. Together with our collaborator we went through each defect to determine whether it was an issue or not. Table 5 shows all the confirmed issues our tool was able to identify in each iteration. The issues identified by our tool frequently surprised our collaborators as they had no idea that they had such issues in their systems.

Since our tool executed on our collaborator's production system, the tool's findings represent real issues found in the organization's system. We collated and discussed its results with our collaborator.

Moreover, our collaborators also provided suggestions for further enhancement of the tool during this discussion. Although the tool brought more visibility of GDPR compliance issues by identifying and relaying problematic areas of the system, our collaborator did not consistently create tasks after each sprint to address identified problems found by our tool (as of this writing). It may be that employees were currently limited by time as the organization was undergoing a rapid transformation period and felt the potential GDPR exposures were not "severe" enough to cause a drastic penalty if temporarily not investigated and resolved. However, when we inquired, our collaborator agreed that the elements identified as potentially not GDPR compliant should have been added to the organization's backlog.

The bottom-up approach to resolving potential GDPR compliance problems that we settled on involved 3 steps. First our tool checks and detects potential violations. For example, one of our primary contacts at Gamma once said this about the findings of our tool, "*These are peculiar findings that should be addressed in subsequent sprints*". Once the tool identifies violations, the problem is added to the organization's backlog and prioritized. Subsequently, the organization works to fix the problem and removes the ticket from its backlog. In theory, executions of the tool thereafter will not find the same problem.

6. Discussion and Implications

This research study surfaced several interesting issues with respect to the use of CSE and privacy NFRs. Our design science research methodology allowed us to obtain in-

Table 5

The table lists the average number of defects identified by our GDPR tool at our collaborating organization across the four iterations. The defects are grouped by resource type.

Resource Type	It. 1	It. 2	It. 3	It. 4
Database	120	208	206	207
Server	32	30	27	34
Route Gate				
Load Balancer Type 1	47	0	0	0
Load Balancer Type 2	0	2	2	2
Router	39	0	0	0
Cloud Storage	0	0	68	62
Cloud Firewall	369	370	368	373
Server Storage	138	138	131	112
Cloud Network	0	0	0	0
Access Management	0	0	0	128
Total Average				

depth knowledge of Gamma's challenges and practices when trying to become GDPR compliant. In reflecting on our empirical insights, we first discuss the role of **continuous compliance** on GDPR in our collaborator and how it may help similar organizations become continuously compliant. Not only do we touch on the importance of just enough compliance engineering, but also the importance of continuous practices in building a **shared understanding** of the privacy non-functional requirement (NFR). We conclude the discussion with an analysis of how **risk management** played an important role in GDPR compliance at Gamma.

6.1. Continuous Privacy Compliance and Continuous Software Engineering

CSE, a core engineering practice at Gamma, emphasized a quick feedback loop through faster iteration cycles [28, 26]. In particular, small, startup organizations embrace CSE as the ability to get their product quickly to market, as well as quickly adapting to change, which are crucial factors to business.

Despite the difficulty of defining and testing NFRs [52, 42], it has been shown that CSE may help verify NFRs [69]. CSE emphasizes the importance of quantifying, acquiring useful feedback, and monitoring of an NFR. These are critical to satisfying organizational and regulatory privacy requirements.

As demonstrated by the first artifact in our design science methodology, privacy regulations may be operationalized into privacy NFRs. However, many GDPR compliance initiatives are the antithesis of continuous: they emphasize up-front modeling [49], complex legal analysis [65], or periodic (and expensive) outside consultants [58]. What we observed at Gamma instead was a bottom-up privacy approach that focused on practical and achievable outcomes, incremental improvements to privacy compliance. We discuss in the following section how this *risk management* approach works; it is important to realize that since the GDPR is relatively untested, such approaches may end up being badly

misguided.

Consistent testing and monitoring of software for NFRs pertinent to regulatory compliance is not new [21]. Continuous compliance is described as automatically checking regulatory compliance after each sprint as opposed to conducting compliance checks after extended intervals [26]. The regression testing aspect of continuous compliance aims to prevent any non-compliance issue from recurring. As privacy regulations, such as the GDPR, prescribe that an organization must comply at all time, “continuous” verification of an organization’s compliance can minimize the duration of non-compliance issues. The concept of “at all times” is ambiguous, but one interpretation is that it means issues must either not exist, or be fixed as soon as the organization is aware [41]. As described above, with the help of our artifacts, our collaborator in an ideal scenario would: Execute our tool on Gamma’s production system. Use our tool’s results to create issues in the organization’s backlog. The collaborator would then prioritize these issues manually. Finally, the privacy tasks are assigned a high priority and resolved in subsequent sprints. These steps would help our collaborator to identify and resolve GDPR compliance problems. Unfortunately, as described in Section 5.2.2, at the time that our study finished, our collaborator did not consistently add identified issues to the organization’s backlog. The process to add issues to the organization’s backlog is not automated. To add to the backlog, a developer would need to manually create a backlog issue or the tool must do so automatically. This manual addition to the backlog is a limitation of our tool and ideally should be automated. Our organization was fast moving in nature and growing rapidly at the time so adding backlog tasks from the tool results became lower priority than their other duties.

Right before our study completed, we were nevertheless encouraged to see our collaborator begin dedicating more resources to interpret the results by assigning someone to go through the weekly results. Adhering to CSE practices, Gamma had a continuous pipeline and automated tests, but the total number of automated tests was limited. In particular, with respect to GDPR tests, there was a slight disconnect between our automated GDPR tests and the organization’s backlog. Following the process of detecting issues, and then creating a backlog task was still a work in progress for our collaborator. The organization acknowledged the importance of consistent testing and monitoring, but has more to improve. The organization needs to link up the feedback loop of our automated tests together with future work tasks so that issues identified in the system can be resolved in the near short term.

As previously discussed in Section 4.5, our collaborator underwent an external GDPR review with consultants during our study. While the external study provides valuable information about the state of the organization’s compliance, its short comings are also apparent. External reviews are extremely costly for a small organization to bear and the annual nature of the reviews exposes an organization to compliance problems in the interval between reviews.

6.2. Becoming Continuously Compliant

For an organization to move to a continuous compliance model, our study at Gamma sheds light on a potential few steps. First, operationalization of GDPR principles into privacy NFRs turns abstract compliance goals into concrete and testable requirements. The privacy NFRs then become the criteria to enforce each compliance check. Next, the development of the automated testing tool helps automate the verification of these privacy NFRs. As the tool is automatically executed on a weekly basis, the tool produced an actionable list of GDPR exposures after each sprint.

However, we observed several limiting factors to the effectiveness of achieving privacy compliance in CSE that other organizations should be aware of when following these guidelines. One concern is whether complete operationalization of the GDPR is possible. We operationalized several GDPR data processing principles in our study, but there are three other GDPR data processing principles, along with data subject rights that we did not operationalize. Other works have focused on automating other individual aspects of the GDPR such as checking and enforcing privacy policies [65], and detecting GDPR violations based on actions on a system [46].

The second challenge is about building a culture that prioritizes compliance NFRs. In smaller organizations in particular, NFRs are neglected until they become a problem [31]. The lowered priority of NFRs is exacerbated by CSE’s emphasis on quick release of features. In consequence, more emphasis is placed on functional requirements in CSE [68].

As described in previous sections, we observed the focus of features as opposed to NFRs as our collaborator did not regularly add tasks to their backlog based on the findings of the tool. The frequency of updating the list of privacy requirements and likewise the corresponding automated tool is at the organization’s discretion as they can chose to make updates on a weekly or monthly basis. Based on experience at Gamma, in the end, an organization must decide for itself the interval and time allocation for compliance work that is acceptable for the organization. In the case of our collaborator, our study occurred during a tumultuous time when the organization underwent massive growth and lots of upheaval occurred in terms of work assignments and employee onboarding. Employees are constantly busy and finding time to translate GDPR tool results into tasks and subsequently working on such tasks is overly time consuming. Our collaborator agreed that the tool’s results should have been added to the organization’s backlog, but employees have been busy with other work. We comment on reasons for this in the following section.

6.3. Building Shared Understanding of the GDPR With CSE

We attribute part of this problem to a lack of shared understanding of the importance of GDPR compliance. In our study, we observed and described in our problem characterization that our collaborator did not conduct systematic employee training on the GDPR, nor enact explicit GDPR

policies. Employees had to conduct individual research on the GDPR and knowledge was disseminated on an ad-hoc basis. Since there was no specific requirement for how much knowledge about the GDPR an employee must master, the level of familiarity with the GDPR varied greatly between employees. For instance, we observed an employee in our collaborator who was not aware of the existence of the GDPR. An employee who is not adequately versed in the GDPR can accidentally create privacy risks, particularly in the fast-paced CSE environment at Gamma. One minor configuration mishap in a system's infrastructure can result in the exploitation of millions of users' sensitive data [14].

As described in our section about CSE's affect on Gamma's compliance, we did observe that one benefit of CSE is that such issues are at least surfaced and discussed. Thus, CSE may act as a driver to improving shared understanding, particularly with its focus on quick turnaround and shared, global repository awareness of such issues. We are planning to investigate the extent to which CSE itself may improve awareness of these NFRs, once they are verified through automated tests.

6.4. GDPR Risk in a Small Company

Earlier we mentioned that technical compliance checking of the GDPR at Gamma was managed bottom-up, emphasizing automation and testing. This was one prong of the organization's overall risk management approach to the GDPR. The GDPR has rightly been characterized as a landmark in privacy regulation of technology, with big implications for IT management and engineering. As a small company, based in Canada, Gamma's management thus takes a prudent, fiduciary approach to the GDPR: balancing non-compliance risk (i.e. lawsuits and fines) against the cost (time, resources, etc.) of compliance.

We comment on two aspects of this risk management: 1) how our collaborator focused on balancing time and compliance, 2) how our collaborator offloads part of its risk using its corporate partners, third-party providers, and outside consultants.

6.4.1. Balancing compliance and non-compliance

Our study corroborated the many prior studies indicating difficulty in achieving GDPR compliance. Many of our closely studied participants described the challenges that they experienced while working towards compliance. For instance, we described in our problem characterization that an employee was tasked with verifying that each tool and service that the organization used complies with the GDPR. This process was particularly challenging as there was a plethora of items that the employee had to verify. In addition, the employee did not have any significant GDPR expertise, thus creating even more difficulty in the verification. Employees were aware of the costs of GDPR non-compliance, at least in theory, and the engineering challenges involved (presumably one reason they partnered with academic researchers!).

Despite the potential for burdensome compliance costs and heavy fines, we observed Gamma taking a balanced ap-

proach. The organization's overall response was delayed to witness the initial batch of GDPR litigation. A couple of heavy fines were handed out in the first year of the GDPR, albeit there were not many fines overall⁶.

Our study suggests it is important to be mindful of the context. Gamma is a smaller company in a highly competitive environment, where time to market is critical. Focusing on the wrong requirements could easily cause the entire business to failure. These concerns are immediate and ever-present, unlike the more abstract problems of the GDPR. As a non-EU company, there may be some perception that EU regulations are less relevant. However, Gamma already removed a portion of its business in the EU, and strove to remove personally identifying information from the initial data collection wherever possible. Thus the GDPR compliance effort was only one aspect of the company's overall risk management and business strategy.

6.4.2. Offloading GDPR Risk to Other Parties

Another risk management strategy we observed at Gamma was to *offload* GDPR risk on other parties wherever possible in three ways. Essentially, offloading risk to third-parties is a practice to shift some responsibilities and liabilities to one or more third-parties [67] through strategies such as partnerships or service agreements.

The three ways we observed in this study include: 1) using presumed-compliant technical partners, 2) relying on deeper-pocketed clients, 3) using external privacy consultants.

Use of third-party cloud services, such as AWS, Azure, and Google Cloud Platform, was an important aspect of the organization's CSE infrastructure. Services like AWS provide state of the art cloud infrastructure, which naturally should have excellent GDPR compliance safeguards that could help protect an organization's data.

Of course, this strategy is not without problems, as even industry leading third-party services may have vulnerabilities. Third-party services may not be fully GDPR compliant and organizations typically still need to self manage the configuration of its cloud infrastructure. Instead of having full control over its GDPR compliance, an organization may be unintentionally exposing itself to the compliance of its third-party providers. Purely migrating and deploying one's infrastructure on a third-party service is insufficient, an organization must take active steps to manage the infrastructure to correct vulnerabilities. As a result, we note that an organizations practicing CSE should increasingly ensure that adequate attention is devoted to configuration management, which is viewed as facilitating achievement of CSE's benefits [39]. In addition, an organization may experience the negative impacts of vendor lock-in [32], a phenomenon that has existed for several decades. Switching to a new service may require significant work to ensure privacy is maintained.

We also observed some implicit offloading risk on other parties, which includes Gamma's clients, much bigger corporations. As users of the services Gamma provides, these

⁶<https://www.enforcementtracker.com/>

clients have liability to GDPR violations as well. We noticed that some at Gamma relied, at least informally, on those clients ensuring the compliance was sufficient.

We speculate that the larger clients' compliance is linked to our collaborator and they have exposure to any compliance issues affecting our collaborator. Furthermore, if any compliance problems surface, the belief is that these clients have large budgets at their disposal that could put together expert teams to help mitigate those problems. Contributing to this belief is the lack of significant fines, for the most part, in the first year of the GDPR. Small companies were mostly spared from penalty, where as large corporations such as British Airways instead received the brunt of the penalties⁷

The final aspect of offloading GDPR risk on other parties we observed was to use external privacy consultants. The main role was to act as a sanity check and potential legal cover (due diligence) in any lawsuit. For instance our collaborator hired privacy consultants to review compliance and suggest improvements. As described by P4, “[*Reputable consultant*] reviewed our compliance and he was impressed” (and hence, P4 saw no need to do anything else). However, solely relying on the positive review of a consultant may provide a false sense of assurance and decrease motivation to further improve the organization's compliance. After all, even privacy consultants often have contrasting views on the same issue or miss aspects in a review [20]. Research has also shown that building a shared understanding between consultants and domain experts is auspicious for privacy compliance [44]. Finally, the compliance certificate that privacy consultants provide is perceived as reputable, but it may be misleading. Yet, the concept of performing a intermittent audit seems like the antithesis of the CSE principles. CSE stresses quick feedback loops and iterations to reduce problems, but we observed that external compliance audits still occur to satisfy the need to provide expert guidance on compliance. However, as of this writing, there is not yet a universal standard method or framework to conduct a GDPR compliance audit. Ultimately, compliance may not be determined until faced with regulatory or legal action.

7. Limitations

For internal validity, we note that the nature of the design science cycle — from problem to solution and back to validation in practice — ensures that the solution has relevance to (at least) our partner. However, we also ensured credibility and analyzability of the data by:

- the primary researcher embedded with our collaborator maintained a researcher diary of observations, with entries of each observation day;
- expanded the relevant challenges with our observational data and diary notes, and then validated these challenges for relevance with our collaborators;

- we conducted iterative, ongoing member checking with three members of the company, as we developed our analysis;
- we member checked the final conclusions, described in this paper, with the primary contact at Gamma.

To ensure credibility of the report, we use thick descriptions of our research approaches. However, our confidentiality agreement with our partner limits our ability to be completely transparent. At Gamma we studied participants in a variety of roles to ensure we had a valid sample. The observation and discussion data we collected from our participants may be limited due to participants performing or talking a certain way because they know they are being watched (the observer effect). We removed any themes that did have corroborating support from multiple participants. As part of our iterative, design science approach, we validated each challenge and operationalized GDPR requirement with employees in our collaborator organization. We triangulated our observational data and limited analysis of code and issue tracker artifacts. We acknowledge that our list of privacy requirements does not encapsulate our collaborator's entire software system neither do the requirements represent every GDPR principle. However, these requirements were developed in close collaboration with our collaborator within our research time frame and draws from our collaborator's advice on which software components were critical.

Another limitation is that our results may be biased because we gifted our collaborator with our artifacts and our collaborator may be biased as a result. However, the fact that the artifacts are developed for the organization's use should also mean that the organization is serious about the development and evaluation of the artifacts. This is especially true since our collaborator permitted the artifacts to be implemented in our collaborator's production software, presumably forcing them to ensure that the artifacts are suitable and adequate.

Our collaborator follows standard CSE practices and employs a continuous integration pipeline for activities such as building and testing software. Our design science artifacts did not get to fully reap the benefits of the organization's automated testing framework. The results of our GDPR tool were not consistently fed back into the backlog. Overall, the organization was not devoid of automated testing, but the organization's automated tests were not comprehensive. The fact that the organization did not have a robust testing framework in place was a limitation of our work. Our study could have incurred different results if the organization already had a fully fledged automated testing framework that it heavily relied upon. In contrast, one possible cause that our collaborator did not habitually convert our tool results into tasks was that they were not accustomed to automated testing.

The interpretation of research results may be subject to researcher bias as one co-author has extensive knowledge about our collaborator, as part of our in-depth design science approach. In our view, the extensive knowledge merely served to provide context about our collaborator, not to bias

⁷<https://www.enforcementtracker.com/>

any inferences or conclusions of results. We secured institutional ethical review approval prior to our study. We also reminded participants that we were not at the company to judge or find blame, that they would be anonymous, and our research goal was focused on the company and GDPR, not individuals.

As with any single company case study research, generalizability of insights to other organizations is not the focus of the research. However, the characteristics of Gamma are important: a small organization (several dozen employees), Gamma operates in a data gathering business, with a reliance on cloud infrastructure. We expect that organizations of similar size and context (e.g. GDPR-applicable, cloud-based, CI-practicing) to encounter similar challenges as Gamma, and find the details about Gamma's journey to address GDPR compliance useful in guiding their own efforts, or making changes to avoid some of Gamma's challenges.

Finally, we acknowledge that the privacy requirements we developed were operationalized from only a portion of the GDPR, and do not claim to achieved a full operationalization of the GDPR. We also recognize that our privacy requirements do not represent all the possible GDPR related requirements pertinent to Gamma's system.

8. Conclusion

Empirical research studying the practices and challenges of GDPR compliance in SMEs using CSE is still a relatively lacking area of research. In our design science research, we spent significant time studying and understanding the organization, and identified challenges that the organization faced in its efforts for GDPR compliance: significant reliance on manual GDPR tests, limited awareness and knowledge of privacy requirements, and balancing GDPR compliance in a competitive data business. For example, developers tasked with compliance work often have limited knowledge and experience working with privacy requirements, especially regulatory frameworks written in legal language. Efforts to comply with the GDPR are further exacerbated by the complexity of the GDPR and the difficult nature of automating testing for GDPR compliance. Relying on manual GDPR tests is time consuming and laborious. For a small organization like our collaborator, relying on manual GDPR tests was burdensome.

The silver lining we found was that our collaborator's use of CSE provided positive benefits with respect to GDPR compliance. CSE emphasizes quick release and feedback, two desirable qualities for our collaborator who wish to swiftly correct areas of non-compliance. Managing GDPR compliance is complex and our collaborator attempts to mitigate and reduce GDPR risk. In particular, our collaborator offloads GDPR risk on other parties, which shifts some liability and responsibility to those third-parties away from our collaborator.

While the GDPR ranked as a highly important regulatory for our collaborator, we also found that GDPR compliance

is just one challenge that a small organization must be cognizant about. Given the resource constraints of a small organization, the organization must balance the trade-offs with continuing and growing the business with GDPR requirements.

To help address these compliance challenges, we iteratively developed a list of privacy requirements. The privacy requirements are testable and measurable, and provide concrete obligations for our collaborator regarding the GDPR. Moreover, we developed a GDPR tool for our collaborator to conduct automated testing of these privacy requirements on the organization's system. The GDPR tool helps alleviate the challenge of relying on manual testing and provides a vehicle to test our privacy requirements.

Finally, we discussed the efficacy of utilizing CSE to achieving GDPR compliance in our collaborating organization, specifically through continuous compliance.

From these insights, we also offer several recommendations to assist organizations with GDPR compliance. First, upon working towards GDPR compliance, an organization is well served to develop procedures for communicating and documenting knowledge regarding the GDPR. Although systematic privacy training for all employees is ideal, our study suggests that such training is not a reasonable expenditure for small or resource constrained organization. Instead, an organization could strive to disseminate privacy knowledge as much as possible to limit the challenges understanding the GDPR. Second, we recommend that organizations develop privacy requirements and corresponding automated tests for verifying GDPR compliance. Automated tests allow an organization to perform repeatable and frequent verification of their system adhering to those requirements. We acknowledge that developing requirements and automated tests is not trivial so our recommendation for organizations is that they start small with a few tests. Even a few automated tests is better than zero tests. Third, we suggest that organizations evaluate their GDPR risk, particularly if they are a small organization. In our study our collaborator was a small Canadian software organization that tried taking a balanced approach of managing and offloading GDPR risk, such as employing external consultants to do their due diligence and increase reputability of their compliance. An organization can assess their own situation and context to determine their level of GDPR risk and whether some offloading of GDPR risk is appropriate.

To complement our research, future studies should gather more empirical evidence on utilizing continuous privacy compliance in both small and large organizations, and research research deriving and implementing additional privacy requirements from the GDPR. Future empirical studies should also be mindful of the circumstances of each examined organization. As our study clearly suggests, organizational context highly influences an organization's compliance response from deriving pertinent requirements to privacy risk management.

References

- [1] Altorbaq, A., Blix, F., Sörman, S., 2017. Data subject rights in the cloud: A grounded study on data protection assurance in the light of GDPR, in: 2017 12th International Conference for Internet Technology and Secured Transactions (ICITST), pp. 305–310.
- [2] Ameller, D., Ayala, C., Cabot, J., Franch, X., 2012. How do software architects consider non-functional requirements: An exploratory study, in: 2012 20th IEEE International Requirements Engineering Conference (RE), pp. 41–50.
- [3] Aranda, J., Easterbrook, S., Wilson, G., 2007. Requirements in the wild: How small companies do it, in: 15th IEEE International Requirements Engineering Conference (RE 2007), IEEE, Delhi, India, pp. 39–48.
- [4] Ataei, M., Degbelo, A., Kray, C., Santos, V., 2018. Complying with Privacy Legislation: From Legal Text to Implementation of Privacy-Aware Location-Based Services. *ISPRS International Journal of Geo-Information* 7, 442.
- [5] Ayala-Rivera, V., Pasquale, L., 2018. The Grace Period Has Ended: An Approach to Operationalize GDPR Requirements, in: 2018 IEEE 26th International Requirements Engineering Conference (RE), pp. 136–146. ISSN: 2332-6441, 1090-705X.
- [6] Bellomo, S., Ernst, N., Nord, R., Kazman, R., 2014. Toward Design Decisions to Enable Deployability: Empirical Study of Three Projects Reaching for the Continuous Delivery Holy Grail, in: 2014 44th Annual IEEE/IFIP International Conference on Dependable Systems and Networks, pp. 702–707.
- [7] Borg, A., Yong, A., Carlshamre, P., Sandahl, K., 2003. The bad conscience of requirements engineering: An investigation in real-world treatment of non-functional requirements, in: Third Conference on Software Engineering Research and Practice in Sweden (SERPS'03).
- [8] Breaux, T., Anton, A., 2008. Analyzing Regulatory Rules for Privacy and Security Requirements. *IEEE Transactions on Software Engineering* 34, 5–20.
- [9] Brodin, M., 2019. A Framework for GDPR Compliance for Small and Medium-Sized Enterprises. *European Journal for Security Research* 4, 243–264.
- [10] Caracciolo, A., Lungu, M.F., Nierstrasz, O., 2014. How do software architects specify and validate quality requirements?, in: European Conference on Software Architecture, Springer, pp. 374–389.
- [11] Cavoukian, A., 2010. Privacy by design: the definitive workshop. *Identity in the Information Society* 3, 247–251.
- [12] Cavoukian, A., Fisher, A., Killen, S., Hoffman, D.A., 2010. Remote home health care technologies: how to ensure privacy? Build it in: Privacy by Design. *Identity in the Information Society* 3, 363–378.
- [13] Chantzou, I., 2019. GDPR Turns 1: Many Companies Still Not Ready. URL: <https://www.symantec.com/blogs/expert-perspectives/gdpr-turns-1-many-companies-still-not-ready>.
- [14] CloudSploit, . A Technical Analysis of the Capital One Hack - CloudSploit. URL: <https://blog.cloudsploit.com/a-technical-analysis-of-the-capital-one-hack-a9b43d7c8aea>.
- [15] Cool, A., 2019. Impossible, unknowable, accountable: Dramas and dilemmas of data law. *Social Studies of Science* 49, 503–530.
- [16] Dave Farley, 2019. Continuous Compliance. URL: <http://www.davefarley.net/?p=285>.
- [17] De Hert, P., Papakonstantinou, V., Malgieri, G., Beslay, L., Sanchez, I., 2018. The right to data portability in the GDPR: Towards user-centric interoperability of digital services. *Computer Law & Security Review* 34, 193–203. URL: <https://linkinghub.elsevier.com/retrieve/pii/S0267364917303333>, doi:10.1016/j.clsr.2017.10.003.
- [18] Deng, M., Wuyts, K., Scandariato, R., Preneel, B., Joosen, W., 2011. A privacy threat analysis framework: supporting the elicitation and fulfillment of privacy requirements. *Requirements Engineering* 16, 3–32.
- [19] Dustin, E., Rashka, J., Paul, J., 1999. Automated Software Testing: Introduction, Management, and Performance. Addison-Wesley Longman Publishing Co., Inc., Boston, MA, USA.
- [20] Elga, A., 2007. Reflection and disagreement. *Noûs* 41, 478–502.
- [21] Ernst, N., Bellomo, S., Nord, R.L., Ozkaya, I., . Enabling Incremental Iterative Development at Scale: Quality Attribute Refinement and Allocation in Practice , 35.
- [22] European Commission, 2019. Data protection in the EU. URL: https://ec.europa.eu/info/law/law-topic/data-protection/data-protection-eu_en. [Online; accessed 2019-04-06].
- [23] European Parliament, 2016. Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/46/EC (General Data Protection Regulation) (Text with EEA relevance). URL: <http://data.europa.eu/eli/reg/2016/679/oj/eng>.
- [24] Fernández-Alemán, J.L., Senor, I.C., Lozoya, P.A.O., Toval, A., 2013. Security and privacy in electronic health records: A systematic literature review. *Journal of Biomedical Informatics* 46, 541–562. URL: <https://www.sciencedirect.com/science/article/pii/S1532046412001864>, doi:<https://doi.org/10.1016/j.jbi.2012.12.003>.
- [25] Fitzgerald, B., Stol, K.J., 2014. Continuous Software Engineering and Beyond: Trends and Challenges, in: Proceedings of the 1st International Workshop on Rapid Continuous Software Engineering, ACM, New York, NY, USA, pp. 1–9. Event-place: Hyderabad, India.
- [26] Fitzgerald, B., Stol, K.J., 2017. Continuous software engineering: A roadmap and agenda. *Journal of Systems and Software* 123, 176–189.
- [27] Fowler, M., 2006. Continuous integration. URL: <https://martinfowler.com/articles/continuousIntegration.html>.
- [28] Fowler, M., Foemmel, M., 2006. Continuous integration. URL: <https://web.archive.org/web/20200522100132/https://martinfowler.com/articles/continuousIntegration.html>. [Online; accessed 2020-05-21].
- [29] Glinz, M., 2005. Rethinking the notion of non-functional requirements. *Third World Congress for Software Quality* .
- [30] Glinz, M., 2007. On non-functional requirements, in: International Conference on Requirements Engineering, pp. 21–26.
- [31] Gralha, C., Damian, D., Wasserman, A.I.T., Goulão, M., Araújo, J., 2018. The Evolution of Requirements Practices in Software Startups, in: Proceedings of the 40th International Conference on Software Engineering, pp. 823–833.
- [32] Greenstein, S.M., 1997. Lock-in and the Costs of Switching Mainframe Computer Vendors: What Do Buyers See? *Industrial and Corporate Change* 6, 247–273.
- [33] Grundstrom, C., Väyrynen, K., Iivari, N., Isomursu, M., . Making Sense of the General Data Protection Regulation—Four Categories of Personal Data Access Challenges , 10.
- [34] Guerriero, M., Tamburri, D.A., Di Nitto, E., 2018. Defining, enforcing and checking privacy policies in data-intensive applications, in: Proceedings of the 13th International Conference on Software Engineering for Adaptive and Self-Managing Systems, Association for Computing Machinery, Gothenburg, Sweden, pp. 172–182.
- [35] Gürses, S., Troncoso, C., Diaz, C., 2011. Engineering privacy by design. *Computers, Privacy & Data Protection* 14, 25.
- [36] Hevner, A.R., March, S.T., Park, J., Ram, S., 2004. Design Science in Information Systems Research. *Management Information Systems Quarterly* , 32.
- [37] Hilton, M., Tunnell, T., Huang, K., Marinov, D., Dig, D., 2016. Usage, costs, and benefits of continuous integration in open-source projects, in: 2016 31st IEEE/ACM International Conference on Automated Software Engineering (ASE), IEEE, pp. 426–437.
- [38] Hjerpe, K., Ruohonen, J., Leppänen, V., 2019. The General Data Protection Regulation: Requirements, Architectures, and Constraints. arXiv:1907.07498 [cs] ArXiv: 1907.07498.
- [39] Humble, J., Farley, D., 2010. Continuous Delivery: Reliable Software Releases through Build, Test, and Deployment Automation (Adobe Reader). Pearson Education.
- [40] Huth, D., Matthes, F., 2019. “Appropriate Technical and Organizational Measures”: Identifying Privacy Engineering Approaches to Meet GDPR Requirements. *AMCIS 2019 Proceedings* .
- [41] ICO, 2020. Accountability and governance. URL: <https://ico.org.uk/for-organisations/guide-to-data-protection/guide-to-the-general-data-protection-regulation-gdpr/accountability->

- and-governance/.
- [42] Leppänen, M., Mäkinen, S., Pagels, M., Eloranta, V.P., Itkonen, J., Mäntylä, M.V., Männistö, T., 2015. The highways and country roads to continuous deployment. *IEEE Software* 32, 64–72.
- [43] Lindqvist, J., 2017. New challenges to personal data processing agreements: is the GDPR fit to deal with contract, accountability and liability in a world of the Internet of Things? *International Journal of Law and Information Technology* 26, 45–63. URL: <https://doi.org/10.1093/ijlit/eax024>, doi:10.1093/ijlit/eax024, arXiv:<https://academic.oup.com/ijlit/article-pdf/26/1/45/23695982/eax024.pdf>
- [44] Massey, A.K., Antón, A.I., 2008. A Requirements-based Comparison of Privacy Taxonomies, in: *2008 Requirements Engineering and Law*, pp. 1–5.
- [45] Massey, A.K., Otto, P.N., Hayward, L.J., Antón, A.I., 2010. Evaluating existing security and privacy requirements for legal compliance. *Requirements Engineering* 15, 119–137.
- [46] de Montety, C., Antignac, T., Slim, C., 2019. GDPR Modelling for Log-Based Compliance Checking, in: Meng, W., Cofta, P., Jensen, C.D., Grandison, T. (Eds.), *Trust Management XIII*, Springer International Publishing, Cham. pp. 1–18.
- [47] Morrison, P., Holmgreen, C., Massey, A., Williams, L., 2013. Proposing regulatory-driven automated test suites for electronic health record systems, in: *2013 5th International Workshop on Software Engineering in Health Care (SEHC)*, pp. 46–49.
- [48] Narendra, M., 2019. Almost a third of EU firms still not GDPR compliant. URL: <https://gdpr.report/news/2019/07/22/almost-a-third-of-eu-firms-still-not-gdpr-compliant/>.
- [49] Palmirani, M., Governatori, G., 2018. Modelling Legal Knowledge for GDPR Compliance Checking, in: *JURIX*.
- [50] Poort, E.R., Martens, N., van de Weerd, I., van Vliet, H., 2012. How architects see non-functional requirements: Beware of modifiability, in: Regnell, B., Damian, D. (Eds.), *Requirements Engineering: Foundation for Software Quality*, Springer Berlin Heidelberg, Berlin, Heidelberg. pp. 37–51.
- [51] Poritskiy, N., Oliveira, F., Almeida, F., 2019. The benefits and challenges of general data protection regulation for the information technology sector. *Digital Policy, Regulation and Governance ahead-of-print*.
- [52] Ramesh, B., Cao, L., Baskerville, R., 2010. Agile requirements engineering practices and challenges: an empirical study. *Information Systems Journal* 20, 449–480.
- [53] Ringmann, S.D., Langweg, H., Waldvogel, M., 2018. Requirements for Legally Compliant Software Based on the GDPR, in: Panetto, H., Debruyne, C., Proper, H.A., Ardagna, C.A., Roman, D., Meersman, R. (Eds.), *On the Move to Meaningful Internet Systems. OTM 2018 Conferences*, Springer International Publishing. pp. 258–276.
- [54] Rossi, C., Shibley, E., Su, S., Beck, K., Savor, T., Stumm, M., 2016. Continuous deployment of mobile software at facebook (showcase), in: *Proceedings of the 2016 24th ACM SIGSOFT International Symposium on Foundations of Software Engineering*, pp. 12–23.
- [55] Savor, T., Douglas, M., Gentili, M., Williams, L., Beck, K., Stumm, M., 2016. Continuous deployment at Facebook and OANDA, in: *Proceedings of the 38th International Conference on Software Engineering Companion - ICSE '16*, ACM Press, Austin, Texas. pp. 21–30.
- [56] Sedlmair, M., Meyer, M., Munzner, T., 2012. Design Study Methodology: Reflections from the Trenches and the Stacks. *IEEE Transactions on Visualization and Computer Graphics* 18, 2431–2440.
- [57] Sirur, S., Nurse, J.R.C., Webb, H., 2018. Are we there yet? Understanding the challenges faced in complying with the General Data Protection Regulation (GDPR). arXiv:1808.07338 [cs] ArXiv: 1808.07338.
- [58] Smith, O., 2018. The GDPR Racket: Who's Making Money From This \$9bn Business Shakedown. URL: <https://www.forbes.com/sites/oliversmith/2018/05/02/the-gdpr-racket-whos-making-money-from-this-9bn-business-shakedown/>.
- [59] Spiekermann, S., 2012. The challenges of privacy by design.
- [60] Spiekermann, S., Cranor, L.F., 2009. Engineering privacy. *IEEE Transactions on Software Engineering* 35, 67–82.
- [61] Steffens, A., Lichter, H., Moscher, M., . Towards Data-driven Continuous Compliance Testing , 7.
- [62] Taipale, O., Kasurinen, J., Karhu, K., Smolander, K., 2011. Trade-off between automated and manual software testing. *International Journal of System Assurance Engineering and Management* 2, 114–125.
- [63] Tamburri, D.A., 2020. Design principles for the General Data Protection Regulation (GDPR): A formal concept analysis and its evaluation. *Information Systems* 91, 101469.
- [64] Tikkinen-Piri, C., Rohunen, A., Markkula, J., 2018. EU General Data Protection Regulation: Changes and implications for personal data collecting companies. *Computer Law & Security Review* 34, 134–153.
- [65] Torre, D., Soltana, G., Sabetzadeh, M., Briand, L.C., Auffinger, Y., Goes, P., 2019. Using Models to Enable Compliance Checking Against the GDPR: An Experience Report, in: *2019 ACM/IEEE 22nd International Conference on Model Driven Engineering Languages and Systems (MODELS)*, IEEE, Munich, Germany. pp. 1–11.
- [66] Urquhart, L., Sailaja, N., McAuley, D., 2018. Realising the right to data portability for the domestic internet of things. *Personal and Ubiquitous Computing* 22. doi:10.1007/s00779-017-1069-2.
- [67] Werner, C., Li, Z., Lowind, D., Elazhary, O., Ernst, N., Damian, D., 2021. Continuously managing nfrs: Opportunities and challenges in practice. *IEEE Transactions on Software Engineering* PP, 1–1. doi:10.1109/TSE.2021.3066330.
- [68] Werner, C., Li, Z.S., Ernst, N., Damian, D., 2020. The lack of shared understanding of non-functional requirements in continuous software engineering: Accidental or essential?, in: *2020 28th IEEE International Requirements Engineering Conference (RE)*, IEEE.
- [69] Yu, L., Alégroth, E., Chatzipetrou, P., Gorschek, T., 2020. Utilising ci environment for efficient and effective testing of nfrs. *Information and Software Technology* 117, 106199.