Privacy Patterns for Internet of Things

A Handbook

PUBLISHED BY LEANPUB

Brand names, logos and trademarks used herein remain the property of their respective owners. This listing of any firm or their logos is not intended to imply any endorsement or direct affiliation with the author.

How to cite this book
Charith Perera, Lamya Alkhariji, Atheer Jeraisy, Omer Rana Privacy Patterns for Internet of Things: A Handbook , Leanpub Publishers, 2020

Lead Editor: Charith Perera
Editors: Omer Rana
Authors: Charith Perera, Nada Alhirabi, Lamya Alkhariji, Atheer Jeraisy

Acknowledgements

Version 1.0 (Expected), September 2020

Contents

Preface

This book is not an introduction to Internet of Things (IoT), software engineering, or privacy. There are many books that provides sufficient content on those topics. This book assumes that you are reasonably proficient in software design and development. We also expect you to have some understanding on IoT. However, expertise in IoT or privacy is not expected and will provide series of use cases throughout the book to support the discussions.

On the other hand this books is not an advanced technical computer science text book. It is a book of privacy patterns that can be useful in designing privacy aware IoT application. Therefore, specific knowledge on privacy preserving techniques or algorithms are not expected. This book share some similarities with the popular object oriented design patterns book *"Design patterns : elements of reusable object-oriented software"* by Erich Gamma, John Vlissides, Ralph Johnson, and Richard Helm. This book inspired by their effort and the value it brought to the software engineering community.

Designing and developing IoT applications is much more complicated than designing and developing desktop, mobile, or web applications. First, IoT applications require both software and hardware (e.g., sensors and actuators) to work together on multiple different type of nodes (e.g., micro-controllers, system-on-chips, mobile phones, miniaturized single-board computers, cloud platforms) with different capabilities under different conditions. Secondly, IoT applications development requires different types of software engineers to work together (e.g., embedded, mobile, web, desktop). This complexity of different software engineering specialists collaborating to combine different types of hardware and software is compounded by the lack of integrated development stacks that support the engineering of end to end IoT applications.

> **— Who Should Read This Book.** This book is primarily aimed at following audiences:
>
> - Are you a software engineering who is building Internet of Things solutions ?
> - Are you a undergraduate student, masters student, PhD student, or a researcher interested in Internet of Things

> and privacy implications ?
>
> -
>
> This book is for you.

1. Introduction

This chapter introduce you to the Internet of things (IoT), its history, why IoT has become a buzz word, current IoT marketplace and the major weaknesses in IoT. If you are well aware of IoT, you may directly move to Section **??**. This chapter aims to create a foundation for upcoming chapters.

1.1 History

Before we investigate the IoT in depth, it is important to look at its evolution. In the late 1960s, communication between two computers was made possible through a computer network [28]. More specifically, in 1969, The first message is sent over the ARPANET, the predecessor of the Internet. The first patent for a passive, read-write RFID tag was granted in 1973. A year later, in 1974, a Universal Product Code (UPC) label is used to ring up purchases at a supermarket for the first time.

In the early 1980s the TCP/IP stack was introduced. Then, commercial use of the Internet started in the late 1980s. Later, the World Wide Web (WWW) became available in 1991 which made

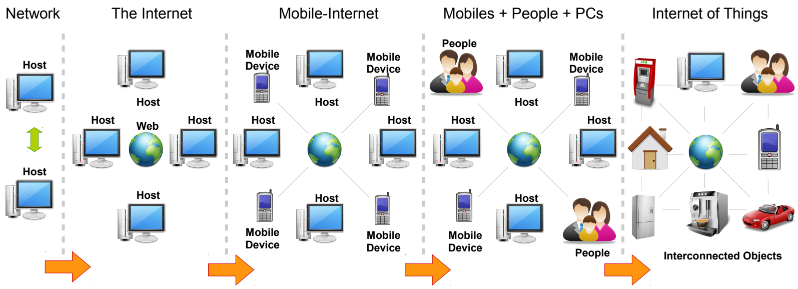

Figure 1.1: Evolution of the Internet of Things (IoT)

the Internet more popular and stimulate the rapid growth. Web of Things (WoT) [17], which based on WWW, is a part of IoT. Later, mobile devices connected to the Internet and formed the mobile-Internet [8]. With the emergence of social networking, users started to become connected together over the Internet. The next step in the IoT is where objects around us will be able to connect to each other and communicate via the Internet [12]. Figure 1.1 illustrates five major phases in the evolution of the Internet of Things.

The term *'Internet of Things'* was coined by Kevin Ashton executive director of the Auto-ID Center in 1999 [5]. Therefore, the term itself is over a decade and half old. However, the ideas of connected devices are more than three decades older [26]. Pervasive computing and ubiquitous computing are the term commonly used at the time.

> **— History of the Internet of Things.** In-depth historical reviews are presented here: A look back at the history of the Internet of Things [6], History of the Internet of Things [26], Why it is called Internet of Things [29], A Very Short History of The Internet of Things [30].

1.2 Internet of Things

The Internet of Things (IoT) does not have a well accepted definition. Instead, IoT has been described and defined by many different parties from many different perspectives. In this section, we will introduce you to a wide variety of definitions.

Definition — (1). Things have identities and virtual personalities operating in smart spaces using intelligent interfaces to connect and communicate within social, environment, and user contexts [25].

Definition — (2). The Internet of Things allows people and '*things*' to be connected Anytime, Anyplace, with Anything and Anyone, ideally using Any path/network and Any service [16].

Definition — (3). Internet of Things is the network of physical devices, vehicles, buildings and other items embedded with electronics, software, sensors, actuators, and network connectivity that enable these objects to collect and exchange data [21].

Definition — (4). Sensors and actuators embedded in physical objects are linked through wired and wireless networks, often using the same Internet Protocol that connects the Internet. [26].

Definition — (5). The Internet of Things is a network of networks where, typically, a massive number of objects / things / sensors / devices are connected through communications and information infrastructure to provide value-added services [19].

In parallel to the term *Internet of Things (IoT)*, Cisco has been driving the term *Internet of Everything (IoE)*. Intel initially called it the *Embedded Internet*. Some other terms used are M2M (Machine to machine) communication Web of Things, Industry 4.0, Industrial internet (of Things), Smart systems, Pervasive computing, Intelligent systems [26]. These terms are interrelated to each other as summarized in Figure 1.2.

— Machine-to-Machine (M2M). The term Machine to Machine (M2M) has been in use for more than a decade, and is well-known in the Telecoms sector. M2M communication had initially been a one-to-one connection, linking one machine to another. But today's explosion of mobile connectivity means that data can now be more easily transmitted, via a system of IP networks, to a much wider range of devices [22].

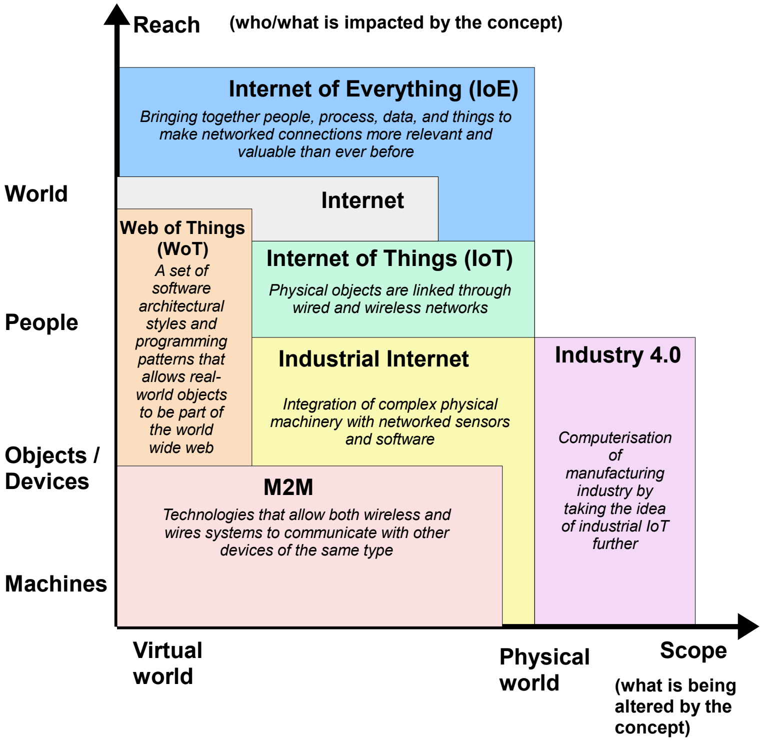
Reach
(who/what is impacted by the concept)
World
People
Objects / Devices
Machines
Internet of Everything (IoE)
Bringing together people, process, data, and things to make networked connections more relevant and valuable than ever before
Internet
Web of Things (WoT)
A set of software architectural styles and programming patterns that allows real-world objects to be part of the world wide web
Internet of Things (IoT)
Physical objects are linked through wired and wireless networks
Industrial Internet
Integration of complex physical machinery with networked sensors and software
Industry 4.0
Computerisation of manufacturing industry by taking the idea of industrial IoT further
M2M
Technologies that allow both wireless and wires systems to communicate with other devices of the same type
Virtual world
Physical world
Scope
(what is being altered by the concept)

Figure 1.2: Concepts Related to IoT. Reproduced from [26].

> **— Sensor Networks.** Wireless sensor networks (WSN), sometimes called wireless sensor and actuator networks (WSAN), are spatially distributed autonomous sensors to monitor physical or environmental conditions, such as temperature, sound, pressure, etc. and to cooperatively pass their data through the network to a main location [2].

> **— Industrial Internet of Things.** The term industrial internet is strongly pushed by General Electrics. It goes beyond M2M since it not only focuses on connections between machines but also includes human interfaces.

> **— Internet.** In the above graph, the internet is a fairly small box. In its core it connects only people.

— Web of Things (WoT). The Web of Things is much narrower in scope as the other concepts as it solely focuses on software architecture.

— Internet of Everything (IoE). Still a rather vague concept, IoE aims to include all sorts of connections that one can envision. The concept has thus the highest reach.

— Industry 4.0. The term Industry 4.0 that is strongly pushed by the German government is as limited as the industrial internet in reach as it only focuses on manufacturing environments. However, it has the largest scope of all the concepts. Industry 4.0 describes a set of concepts to drive the next industrial revolution. That includes all kinds of connectivity concepts in the industrial context. However, it goes further and includes real changes to the physical world around us such as 3D-printing technologies or the introduction of new augmented reality hardware.

— More IoT Definitions and Descriptions. Definitions collected and synthesized by the IEEE Internet of Things community are documented here [27].

IoT Devices (*'Things'*) on the Internet of Things

As you may have understood by now, *'Things'* play a significant role in Internet of Things paradigm. There isn't any formal definition to describe a *'Thing'* in IoT paradigm. We have illustrated variety of different *'Things'* that can be part of IoT paradigm in Figure 1.3.

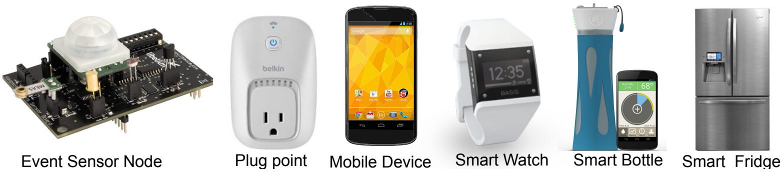

Figure 1.3: A *'Thing'* can be any object around us from refrigerators to bottles to watches to mobile phones to electrical plugs to sensors.

It is important to note that terms such as objects, smart objects, internet connected objects (ICOs), nodes, devices, IoT devices, smart devices are also used interchangeably in IoT related documentations in order to refer to *'Things'*.

Let us now explore the major characteristics of a *'Thing'*. First, it is important to understand that, any object can become part of the IoT. One major characteristics is computational capability. Each *'Thing'* show have some kind of computational capabilities. Next, each *'Things'* should be be able to communicate with the Internet. This does not mean that each object should have a direct or permanent connection to the Internet. For example, a *'Thing'* may communicate with a near-by mobile phone using Bluetooth and the phone may forward the data to the Internet using its WiFi capabilities. In another example, a *'Thing'* may connect to the Internet using its GPRS communication capability once a week. In Figure 1.4, we illustrate how an everyday object may be converted into an IoT device in IoT. Typically, IoT devices have both sensing and actuation capabilities as well.

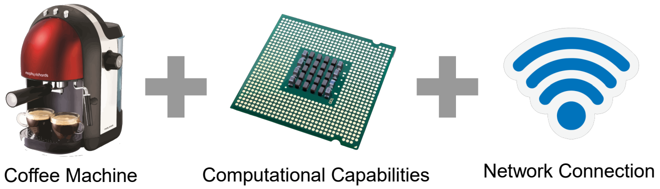

Figure 1.4: An everyday object embedded with some amount of computational and network communication capabilities can be identified as an IoT Device

Common Internet of Things Solutions Architecture

There is no consensus on constitutes a suitable architecture for an IoT solution. Systems have varying requirements that affect the choice of architecture. For example, in a centralised architecture, sensors lie on the periphery and are only concerned with data acquisition. These peripheral devices feed data to a centralised entity, which processes, analyses, stores and disseminates the data. This architectural pattern has many well documented benefits including reliability, scalability and interoperability [31]. This is in contrast to a distributed IoT architecture where processing occurs on the periphery at the device level, and data may or may not then be sent to a centralised server or other peripheral devices. The distributed approach still has many issues that needs to be addressed but provides more fine grain control over the data produced. We can categorise different types of IoT solutions architecture into four segments [31]: 1) centralised, 2) collaborative, 3) connected intra-net of Things, and 4) distributed IoT.

Out of these architectures, centralised architecture is the most widely used in IoT solutions. The centralised architecture, as shown in Figure 1.5, consists of three components: 1) IoT devices, 2) Gateway devices, and 3) IoT cloud platforms. Today, there are many different vendor who provides both hardware and software components in order to support rapid IoT solutions development. We can see these components in the IoT solutions marketplace as well though though they may not be clearly visible to the end-users.

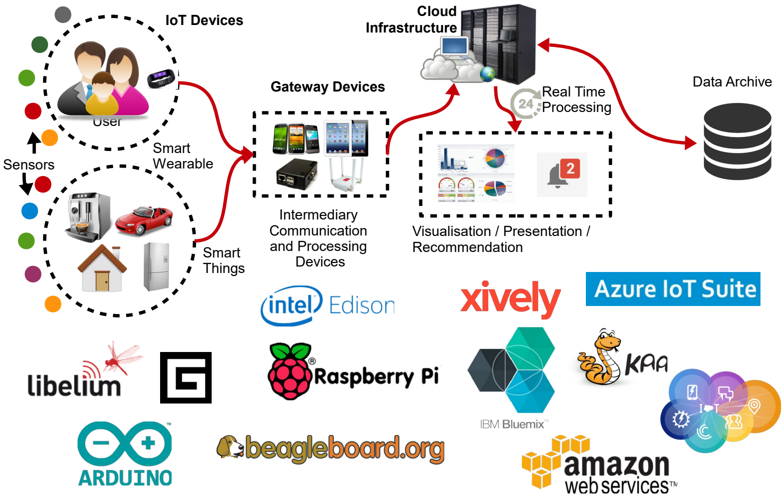

Figure 1.5: Common Internet of Things solutions architecture comprises with three components: 1) IoT devices, 2) Gateway devices, and 3) IoT cloud platforms.

1.3 Privacy

The Internet of Things: making the most of the Second Digital Revolution [34]

The Guide to Data Protection [20]

Securing the Internet of Things [18]

Protecting Data and Privacy in the Cloud [33]

Privacy-Enhancing Technologies [36]

Privacy by Design in the Age of Big Data [9]

Privacy by design in big data [1]

Privacy Bridges [35]

Privacy and Data Protection by Design from policy to engineering [10]

PRIME white paper [24]

Opinion 8/2014 on the on Recent Developments on the Internet of Things [4]

Internet of Things Privacy and Security in a Connected World [15]

Internet of Things IoT Governance, Privacy and Security Issues [13]

Handbook on European data protection law [14]

Getting smarter about smart cities: Improving data privacy and data security [23]

Direct Marketing Association Data Guide [11]

Consumer Perceptions of Privacy in the Internet of Things What Brands Can Learn from a Concerned Citizenry [3]

A Practical Guide to the Data Protection Act [37]

1.4 Privacy by Design

1.5 Patterns

> **Definition — Software Design Pattern.** In software engineering, a software design pattern is a general reusable solution to a commonly occurring problem within a given context in software design. It is not a finished design that can be transformed directly into source or machine code. It is a description or template for how to solve a problem that can be used in many different situations. (wikipedia.org)

1.6 Pattern Template

The template we used to organise and present privacy patterns in this book was proposed by Romanosky et al. [32]. It is also a simplified version of the Pattern-Oriented Software Architecture (POSA) outline and developed by Bushman et al. [7].

Privacy Pattern Name

The name provides a short descriptive title or active phrase that generally illustrates the solution.

Context

The context describes the general situations and assumptions under which the problem occurs. It describes the scope, market, user or other conditions that, if changed, would alter the problem or solution.

Problem

Describes the problem that repeatedly occurs and the forces that are in conflict for the given context. The forces can arise from tensions or conflicts from users, computing systems, corporations, the natural environment, legal regulations, etc.

Solution

This is the fundamental solution that best resolves and balances the forces. The better the forces are balanced, the better the solution. The discussion provides a guideline or strategy for implementing the solution and should allow the reader the freedom to craft the solution in the most appropriate way.

Constraints and Consequences

C Consequences describe both the benefits and liabilities of the pattern because solutions are not always able to resolve

[0] The POSA format was originally developed for software engineering patterns and so also describes other sections such as Dynamics, Implementation and Variations that we will not cover in this paper.

each of the forces. Therefore, any conflicts not resolved or limitations of the solution should be listed.

Motivating Scenario

Scenario This is an example scenario to describe ■

Classification

Know Uses and Related Work

> — **Name of a Usage.** A true pattern will have many real-world implementations. Without these, the pattern is only a potentially great idea. The better a pattern can demonstrate actual uses, the better it is and the more useful it will be to others.

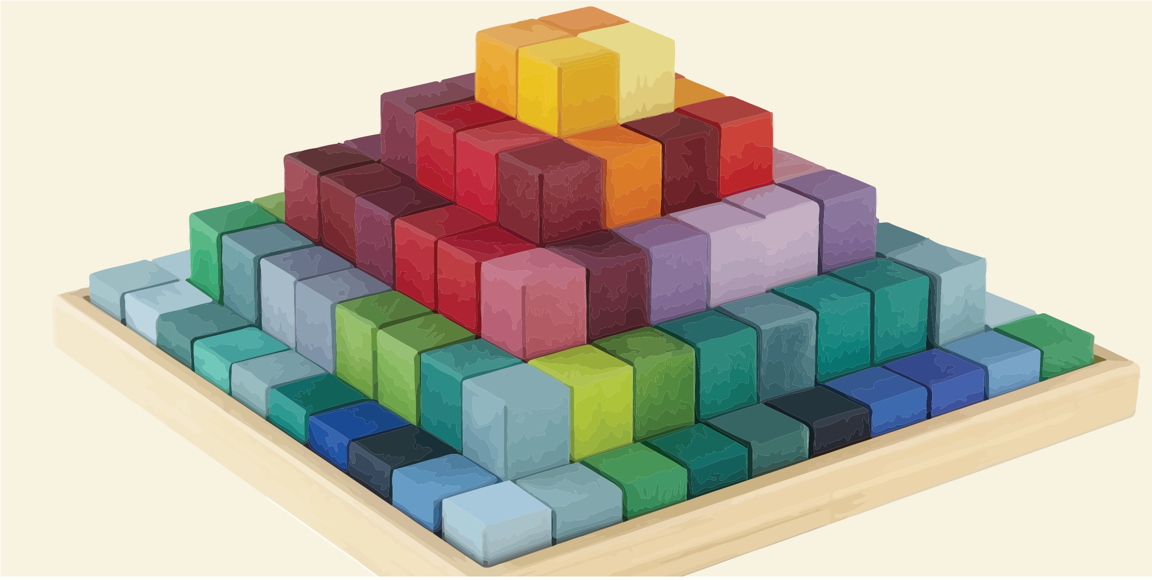

2. Privacy Patterns Catalogue

In this chapter, we describe each privacy pattern using the template introduced in the Section 1.6. It is important to note that the order we present the following patterns does not the their relative importance at any means. We discuss how these patterns can be used to design privacy-aware IoT applications in Chapter 3.

2.1 Protection Against Tracking

Summary

This pattern avoids the tracking of visitors of websites via cookies. It does this by deleting them at regular intervals or by disabling cookies completely.

Context

This pattern is applicable when personal identifiable information is tracked through software tools, protocols or mechanisms such as cookies and the like.

Problem

With every single interaction in the web you leave footmarks and clues about yourself. Cookies for example enable webservers to gather information about web users which therefore affects their privacy and anonymity. Web service providers trace user behavior, which can lead to user profiling. Also, providers can sell the gathered data about users visiting their pages to other companies.

Goal

G Restricting a website to not be able to track any of the user's personal identifiable information.

Solution

Restricting usage of cookies on the client side by deleting cookies on a regular basis e.g. at every start-up of the operating system or enabling them case-by-case by deciding if the visited website is trustworthy or not and by accepting a cookie only for the current session. At the highest level of privacy protection cookies are disabled, but as a consequence web services are restricted. Another solution could be that cookies are exchanged between clients, so that sophisticated user profiles emerge.

Constraints and Consequences

C With cookies disabled there is no access to sites that require enabled cookies for logging in.

C Other tracking mechanisms for user fingerprinting may still work even when cookies are disabled.

Motivating Scenario

> **Scenario** Alice wants to buy shoes and she wants to shop online. She heads to an online shop and searches for shoes but can't decide which ones she wants, so she buys neither of them. The next day she finds a couple of emails in her inbox, giving her suggestions for other shoes and alerting her that the viewed shoes are now on sale. ■

Know Uses and Related Work

> **— Onion Routing.** Junkbuster is an old proxy filtering between web server and browser to block ads and cookies, but it is no longer maintained. A program named CookieCooker (`http://www.cookiecooker.de/`) provides protection for monitored user behaviour and interests by exchanging cookies with other users or using a random selection of identities. Unfortunately, this project also seems to be not maintained anymore. There is also the Firefox Add-on Self-Destructing Cookies which deletes cookies of tabs as soon as they are closed.

Categories

- Anonymity
- Unlikability
- Tracking
- Cookies
- Minimize
- Exclude
- Hide
- Aggregate

Related Patterns

- P Masked Online: Need to explain

- P Strip Metadata: Need to explain

Supporting Patterns

- P Onion Routing: Need to explain

Sources

- `https://privacypatterns.org/patterns/Protection-against-tracki`
- `https://privacypatterns.eu/#/patterns/protection-against-track`
  `0-0-0-0-0-0-0-0-0-0-0-0-0-0-0-0`

2.2 Location Granularity

Summary

Support minimization of data collection and distribution. Important when a service is collecting location data from or about a user or transmitting location data about a user to a third-party.

Context

When a service is collecting location data from or about a user or transmitting location data about a user to a third-party.

Problem

Many location-based services collect current or ongoing location information from a user in order to provide some contextual service (nearest coffee shop; local weather; etc.). Collecting more information than is necessary can harm the user's privacy and increase the risk for the service (in the case of a security breach, for example), but location data may still need to be collected to provide the service. Similarly, users may want the advantages of sharing their location from your service to friends or to some other service, but sharing very precise information provides a much greater risk to users (of re-identification, stalking, physical intrusion, etc.).

Accepting or transmitting location data at different levels of granularity generally requires a location hierarchy or geographic ontology agreed upon by both services and a more complex data storage model than simple digital coordinates.

Truncating latitude and longitude coordinates to a certain number of decimal places may decrease precision, but is generally not considered a good fuzzing algorithm. (For example, if a user is moving in a straight line and regularly updating their location, truncated location information will occasionally reveal precise location when the user crosses a lat/lon boundary.) Similarly, using "town" rather than lat/lon may occasionally reveal more precise data than expected when the user crosses a border between two towns.

Goal

G

Solution

Since much geographic data inherently has different levels of precision (see geographic ontologies, for example) – like street, city, county, state, country – there may be natural divisions in the precision of location data. By collecting or distributing only the necessary level of granularity, a service may be able to maintain the same functionality without requesting or distributing potentially sensitive data. A local weather site can access only the user's zip code to provide relevant weather without ever accessing precise (and therefore sensitive) location information.

A similar pattern is location fuzzing which uses an algorithm to decrease the accuracy of location data without changing its lat/lon precision. This may be useful if the application only functions on latitude/longitude data but can be vulnerable to attack.

In some cases, less granular data may also better capture the intent of a user (that tweet was about Sproul Plaza in general, not that particular corner) or be more meaningful to a recipient ("Nick is in Berkeley, CA" means more to my DC relatives than the particular intersection). For more along these lines, see, for example, the Meaningful Location Project.

Constraints and Consequences

C

Motivating Scenario

Scenario 1. Fire Eagle location hierarchy

Yahoo! Fire Eagle allows user to provide location infor-

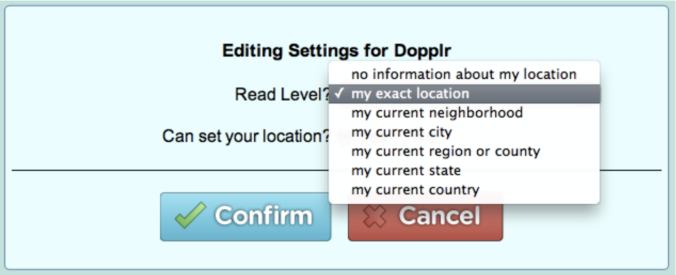

Editing Settings for Dopplr

Read Level? ✓ my exact location

Can set your location?

Figure 2.1: Fire Eagle location hierarchy

mation to applications using eight different "levels" of granularity in their hierarchy:

- No information
- As precise as possible
- Postal code
- Neighborhood
- Town
- Region
- State
- Country

Fire Eagle specifically requires that recipient applications be written to handle data at any of the levels and allows updating the user's location at any level of granularity.

2. Twitter "place" vs. "exact location" Twitter allows users to tag a tweet with either exact coordinates, a Twitter "place" (a town, neighborhood or venue) or both.

3. Geode
One of the fore-runners to the W3C Geolocation API, Firefox's experimental Geode feature allowed JavaScript access to the current location at four different levels of granularity.

Know Uses and Related Work

> Many online organizations provide signals to their customers. Often, they are publicly and freely available, but can also be purchased by third parties. The online auction site, eBay, for example, uses a reputation system to assist other buyers in feeling more comfortable purchasing from an unknown seller. Many other ecommerce sites (such as Amazon) rely heavily on the reputation and referral systems in order to help customers make a more informed decision.
>
> Websites are more commonly publishing their privacy policies in order to assuage the privacy concerns of their users [ECC2005]. Users are also stating that they would be more comfortable interacting online if the site had displayed the TRUSTe or BBBOnline symbols, or had a privacy policy [CRA1999].

Categories

- Location
- Minimization
- Abstract
- Group

Related Patterns

P

Supporting Patterns

P

Sources

- https://privacypatterns.org/patterns/Location-granularity

2.3 Minimal Information Asymmetry

Summary

Context

Users frequently interact with controllers whose services (or products) they have not used before. At this point the knowledge the user has about the controller and its practices, especially regarding privacy, is typically nonexistent. The controller as a whole has a much clearer understanding of its policies. It also begins to know a lot about the user in a short time period, if not already well informed. The user needs to put in sufficient effort to investigate the controller to know about its practices to provide valid consent. The controller needs this valid consent to lawfully process the user's information.

Problem

Controllers have far more information than the users who utilize their services, which makes the users vulnerable to exploitation.

Information asymmetry is generally described as one party having more or better information about a transaction than the other. In order for a healthy consumer relationship to ensue, users should know close to as much about the controller's practices as it would be expected to itself.

Forces and Concerns

- Users sometimes want to use services of an unknown party, and are cautious about what it might do with their data
- Users may not want to provide any more information than necessary, but want the services to function properly
- Controllers want users to understand the intentions behind the data they collect, and be content with how they use it
- Controllers need to ensure that users understand purposes and means for processing before their consent will be valid

Goal

G

Solution

Require minimal information from the user, so that only as much personal data as is required, explained, and consented to, is processed. Further reduce the imbalance of policy knowledge by writing clear and concise policies rather than, or in addition to, complex and verbose ones.

Implementation

Limit the amount of data needed to provide the services necessary to the users, and where appropriate, prefer less sensitive data to do so. Give users the option to opt in to features which require more data, but keep it minimal by default. If the amount of data needed is minimized, then users have less they need to understand, and less to disagree with. This also allows for more simple policies.

Making policies more clear and concise is also crucial, as users will not want to sift through long-winded texts to understand what would happen with their data. Highlight important aspects for users themselves, rather than allowing them to become cluttered with legal jargon, detail, and complexity. While certain elements cannot be explained adequately without doing so at length, not all aspects are relevant at once. Some elements may be summarized without the detail, so that users may better understand the current focus. The full detail should still exist however, and be easily located.

Constraints and Consequences

C

Motivating Scenario

Scenario ■

Know Uses and Related Work

> Many online organizations provide signals to their customers. Often, they are publicly and freely available, but can also be purchased by third parties. The online auction site, eBay, for example, uses a reputation system to assist other buyers in feeling more comfortable purchasing from an unknown seller. Many other ecommerce sites (such as Amazon) rely heavily on the reputation and referral systems in order to help customers make a more informed decision.
>
> Websites are more commonly publishing their privacy policies in order to assuage the privacy concerns of their users [ECC2005]. Users are also stating that they would be more comfortable interacting online if the site had displayed the TRUSTe or BBBOnline symbols, or had a privacy policy [CRA1999].

Categories

- Inform
- Provide

Related Patterns

Privacy Mirrors

Personal Data Table

Supporting Patterns

Sources

- https://privacypatterns.org/patterns/Minimal-Information-Asymmetry

2.4 Informed Secure Passwords

Summary

Context

Credentials are required by numerous services (and products) in order to ensure that only authenticated and authorized users have access to certain features. Controllers typically provide authentication mechanisms in the form of usernames and passwords. Although these provide a weak form of security when used incorrectly, they are more convenient for users than many less popular and more secure alternatives. Controllers often try to circumvent the shortcomings of passwords by encouraging users to change them frequently, use stronger variations, check them, and prevent disclosure and reuse. However, users make use of many services, and use many passwords, thus discouraging proper application. This misapplication can result in personal data being accessed by unauthorized persons.

Problem

Users must regularly maintain many strong passwords, remember them, and protect them, but are not well equipped to do so. So instead many choose weak ones and reuse them.

Forces and Concerns

- Users do not want to remember many long or complex passwords
- Users do not want others to access their secured services
- Controllers want to protect accounts from unauthorized access
- Controllers do not want to apply too much pressure to their users to maintain and protect strong and unique passwords

Goal

G

Solution

Provide users with assistance in understanding and maintaining strong passwords which are easier to remember.

Structure

Assistance is typically provided in the following ways:

- Passive mechanisms (e.g. help button)
- Static mechanisms (e.g. pop-ups)
- Dynamic mechanisms (e.g. dynamically adjusting message) [the] method that is most noticed by the users and therefore also most helpful

Implementation

Short passwords, those at character lengths which are feasible for brute forcing, are not secure. The difficulty to brute force is affected by the known complexity, such as using a variety of character types. However, complexity affects password memorability more than strength. It is more important that passwords are long enough, and varied enough. This does not mean they should be difficult to remember. A couple unrelated words strung together can be a very secure yet memorable password.

These aspects can be weighted together to provide the user with a strength meter, as well as the explanations behind it. Examples of secure passwords should also be provided, but not accepted as the actual password. Do not enforce the use of special characters and numbers. Length, along with sufficient variation, should be the deciding factor in password strength.

Given enough resources and time however, state of the art character lengths can be overcome. It is as such useful to change them more regularly than the time it would take to brute force them. Otherwise, the longer that a password remains unchanged, the more likely it is that the password has been compromised.

Therefore a mechanism should also be provided to remind a user when it is time to start thinking of a new password. Based on how strong the original was this may be more or less often. Unusual Activities may also justify a more frequent update.

> When verifying whether a user used the same password in a second field, to prevent mistypes, simply indicate whether the fields match with a recognizable affirmation. Typically this uses a green theme, and may use a check mark.

Constraints and Consequences

C Secure passwords are very important in [an interconnected world]. Users generally tend to use familiar words such as names of pets and family members and no special [characters] when creating a password. These passwords can hence be easier hacked using social engineering than longer [and more complex passwords]. Secure passwords are a necessary step towards personal security. Using the above approach, the user obtains more feedback on the safety of the entered password and is therefore able to create safe passwords that can be remembered.

Motivating Scenario

Scenario ■

Know Uses and Related Work

> Strongpasswordgenerator.com both provides explanation on state of the art approaches to secure passwords in a layperson friendly manner and helps generate them.

Categories

- Inform
- Explain

Related Patterns

P Unusual Activities

P Informed Credential Selection

Ⓟ Appropriate-Privacy-Icons

Ⓟ Icons-for-Privacy-Policies

Ⓟ Privacy-color-coding

Supporting Patterns

Ⓟ

Sources

- https://privacypatterns.org/patterns/Informed-Secure-Passwords

2.5 Awareness Feed

Summary

An Awareness Feed warns the user about the potential consequences of sharing their personal data. It does so before that data is collected or used, and continues to do so whenever a change in context is detected. This change may include newly provided information by the user, and changes in the environment in which the controller (i.e. provider) operates or processes personal data.

This pattern allows users to make informed decisions regarding if, when, and how they share their personal data. As more information is collected, the user may become more identifiable, and the data relating to them may become more invasive. Awareness Feed keeps users aware of both the short-term and long-term repercussions in their data sharing choices.

Context

In a situation where user data is collected or otherwise processed, particularly personal data, many users are concerned about the potential repercussions of their actions. Controllers (e.g. organizations), which have dynamic and evolving services (or products) which users interact with, may share this concern. This may be for legal, ethical, or public appearance reasons.

These controllers also care about the monetary implications of a solution, often including the opportunity cost of informed users against the risks and profits of over-sharing. For-profit organizations regularly want to bolster their market share by overcoming competition with state of the art technologies. These changes may have important consequences, unintentional or otherwise, for users of the system. Controllers want to limit the exposure of these risks to their userbase, even if from a third party, as they are responsible for their data.

Such controllers may already have in place a Privacy Dashboard, seeking to complement it, or wish to maintain awareness through various other services. They likely consider Lawful Consent and thus seek to ensure that their users are properly and priory informed before making regrettable decisions. They would nonetheless need to prevent notification fatigue if they were Preventing Mistakes or Reducing Their Impact like in this pattern.

Problem

Users are often unaware of the privacy risks in their data sharing activities, especially risks which are indirect or long-term. How can we best ensure that users become aware of these risks?

This problem is agitated by the organizational aim to provide novel and competitive services while keeping users informed. The difficulty of this is frequently underestimated. The pitfalls controllers face as a consequence manifest both in taking shortcuts and in unexpected long-term effects.

Forces and Concerns

- Users do not necessarily realize the effects of their information sharing, but often want to use new or interesting features
- Some users are discouraged from sharing as they do realize that they are not informed about risks to their privacy, but cannot reasonably change that themselves
- Controllers aim to provide or utilize novel and or competitive services, but explaining potential risks to privacy in those services is often non-trivial and generates a fear of upsetting the userbase and endangering trust
- ome controllers wish to empower users by informing them, but do not want to jeopardize their business model, or ability to process in a timely fashion

Shortcuts

The appeal of convenience features may sway controllers into flawed implementations which undermine user privacy. Automated decisions, influenced by past actions or by other potentially inaccurate metrics, may result in sharing decisions which users do not approve of. The same holds for features which are not adequately assessed. While a controller might intend all the necessary tools for informed decisions to be present, short-sighted process flows may violate user trust all the same.

Long-term Effects

Over time, supposedly harmless data may amass into more revealing information, especially when paired with the right metadata. Being able to link user activity to other sources of information may also result in far more exposing situations than expected.

Not only are users often unaware of the potential consequences of their actions, even controllers themselves regularly fail to anticipate how revealing their services can be. While some users approach this uncertainty with caution, others will risk their privacy in hopes of using the services. Though the uncertainty might not prevent their participation, it may still jeopardize their trust in the system.

Goal

G

Solution

Warn users about potential consequences before collecting or otherwise processing personal data, early enough to be appreciated and late enough to be relevant.

This information should be provided before the point where privacy risks could materialize. If there is some delay before further processing after collection, the user has some time to review the risks. Until the user accepts them however, that further processing should not take place.

This pattern is a compound pattern, one in which multiple patterns work together to address a broader problem. It combines the following patterns:

- Impactful Information and Feedback;
- Increasing Awareness of Information Aggregation;
- Privacy Awareness Panel;
- Appropriate Privacy Feedback; and
- Who's Listening.

Rationale

It is not likely enough that users are informed prior to being provided a service, nor is it reasonable to expect that consent acquired in bulk is properly informed. Consent is not necessarily freely given, either, if the lack of consent presents a wall to a service that the user wants or believes to need.

A concerted effort needs to be made to present the user with unintimidating information relevant to their privacy risks for a service. Providing too much information lessens the chances that the user will read it, while too little information may not properly inform the user. Informing the user too late also puts the user at unnecessary risk.

By making this effort, the controller avoids accusations of negligence in informing their users.

Implementation

Every service which makes use of personal information should be investigated by its creators during its creation, or retrospectively if already available. The controller in question is responsible for this. Not only will this affect the user's understanding once presented to them in layperson terms, but it will also allow the controller to realize the privacy impact of their services. This may encourage them to improve the services to be more respecting of privacy. A good solution composes of accessibility, as well as transparency and openness.

Accessibility

There needs to be a balance between the user effort required both to use a service and maintain their privacy. Information about the risks should not be deceptive, or difficult for laypersons to comprehend. Meeting this balance may also be challenging, as fully comprehending the risks involved might require a certain understanding of the system itself.

In order to reduce the quantity of the presented information, only the contextually significant information need be presented. Furthermore, the information should be available in the level of detail sought by the user: in both concise and detailed variants. A short description may be used in Preventing Mistakes or Reducing Their Impact. A more in depth variation may give them confidence that even if they cannot comprehend it, someone would speak out if something were amiss. In a similar vein, detailed descriptions should be comprehensible enough to avoid accusations of being deliberately complex or misleading.

One way in which to explain the risks involved in a process is through example. This is particularly useful in the case of information aggregation. Visualizing the publicity of data is also useful, users can see how visible information would be, or is, to the outside world. Similar decisions by those who choose to set examples may also help in influencing informed sharing behaviour.

Transparency and Openness

Users need to be able to trust that a system does not pose unnecessary risks. Fostering a familiarity with openness and transparency about the processes involved may garner this trust. It allows those who invest time an opportunity to be certain, and those who trust in public perception to be at ease.

Constraints and Consequences

C The solution of this pattern will cause users to have a better understanding of the potential consequences of information they share. It may empower some to share knowing they may do so safely, though it will cause less activity overall, as many users will be more careful about what they put online. This is not necessarily a negative consequence, though, since regretted decisions merely garner mistrust and prevent future activity. The controller will be able to introduce new services or update old ones with confidence that users are given the opportunity to consider their decisions in full light of the service's potential consequences.

C In addition to lower adoption of risky services, due to public consequences, there will be more cost involved in reworking them. Unattended and system-wide process changes, which negatively affect consequences, will be more difficult to perform. They will not be possible without first disabling the affected services. This is similar to the way some controllers (e.g. Google) handle changes in privacy policies.

C Due to more privacy-minded implementations, the system will not anticipate users as easily, though for many this will be a worthwhile tradeoff. While there is cost in creating good solutions, the long-term cost of bad ones (especially in good faith) can often be higher.

Constraints

By informing users prior to the activation of any services which use personal data, many aspects of a system are less fluid and thus require additional forethought. Instead of quick integration into the system, which may have come with many privacy oversights, users will be exposed to consequences that they might not have otherwise realized. Care will need to be taken to ensure that these users do not become overwhelmed. As a consequence of better informed users, however, questionable services are more open to scrutiny and thus many shortcuts will no longer be viable.

Motivating Scenario

> **Scenario** Full adoption of this pattern is not yet commonplace, yet there exist examples of feedback loops to users about activities corresponding to them. This includes notifications such as 'user X wants to access Y', or retrospectively, 'user X accessed Y'. There also exist services which require opt-in, accompanied

> by explanations of their effects. Conversely telemetry is often opt-out, but occasionally explains what information is at stake. ■

Know Uses and Related Work

Categories

- Inform
- Provide

Related Patterns

Building Trust and Credibility

Preventing mistake or Reducing Their impact

Lawful Consent

Selective Disclosure

Privacy Aware Wording

Layered Policy Design

Privacy Aware Network Client

Icon for Privacy Policies

Appropriate Privacy Icons

- Privacy Labels
- Privacy Color Coding
- Task-based Processing
- Trust Evaluation of Services Sides
- Privacy Dashboard
- Impactful Information and Feedback
- Increasing Awareness of Information Aggregation
- Privacy Awareness Panel
- Appropriate Privacy Feedback
- Who’s Listening

Supporting Patterns

- Layered Policy Design

Sources

- https://privacypatterns.org/patterns/Awareness-Feed

2.6 Encryption with user-managed keys

Summary

Use encryption in such a way that the service provider cannot decrypt the user's information because the user manages the keys.

Enable encryption, with user-managed encryption keys, to protect the confidentiality of personal information that may be transferred or stored by an untrusted 3rd party.

Supports user control, cloud computing and mobile.

Context

User wants to store or transfer their personal data through an online service and they want to protect their privacy, and specifically the confidentiality of their personal information. Risks of unauthorized access may include the online service provider itself, or third parties such as its partners for example for backup, or government surveillance depending on the geographies the data is stored in or transferred through.

Problem

How can a user store or transfer their personal information through an online service while ensuring their privacy and specifically preventing unauthorized access to their personal information?

Requiring the user to do encryption key management may annoy or confuse them and they may revert to either no encryption, or encryption with the online service provider managing the encryption key (affording no protection from the specific online service provider managing the key), picking an encryption key that is weak, reused, written down and so forth.

Some metadata may need to remain unencrypted to support the online service provider or 3rd party functions, for example file names for cloud storage, or routing information for transfer applications, exposing the metadata to risks of unauthorized access, server side indexing for searching, or de-duplication.

If the service provider has written the client side software that does the client side encryption with a user-managed encryption key, there can be additional concerns regarding whether the client software is secure or tampered with in ways that can compromise privacy.

Goal

G

Solution

Encryption of the personal information of the user prior to storing it with, or transferring it through an online service. In this solution the user shall generate a strong encryption key and manage it themselves, specifically keeping it private and unknown to the untrusted online service or 3rd parties.

Constraints and Consequences

C

C

Motivating Scenario

Scenario

- Spider Oak: online backup, sync, sharing enabling user managed personal information in zero knowledge privacy environment
- Least Authority: secure off-site backup system with client side encryption
- LastPass: encrypted credentials and personal information

database with user managed encryption keys

Some have used the term "zero-knowledge" to describe this pattern; however, "zero-knowledge proof" is a cryptographic term with a distinct meaning.

Know Uses and Related Work

Categories

- Encryption
- Control
- Mobile
- Cloud
- Hide
- Restrict

Related Patterns

P

Supporting Patterns

P

Sources

- https://privacypatterns.org/patterns/Encryption-user-managed-keys

2.7 Federated Privacy Impact Assessment

Summary

The impact of personal information in a federation is more than the impact in the federated.

Context

Identity Management scenarios (that is, when the roles of the Identity Provider and the Service Provider are separated).

Problem

Identity Management solutions were introduced to decouple the functions related to authentication, authorization, and management of user attributes, on the one hand, and service provision on the other hand. Federated Identity Management allows storing a data subject's identity across different systems. All together, these form a Federation that involves complex data flows.

Federated Management solutions can be used to improve privacy (e.g. by allowing service providers to offer their services without knowing the identity of their users). However, the complexity of data flows and the possibility of collusion between different parties entail new risks and threats regarding personal data.

Goal

G Deal with privacy risks associated from the federation of different parties in an Identity Management solution.

Solution

A Privacy Impact Assessment is conducted by all the members of the federation, both individually and in conjunction, so as to define shared privacy policies, prove they are met, and demonstrate the suitability of the architecture, in the benefit of all the members.

Constraints and Consequences

C The consequences depend on the results of the privacy-impact analysis.

Motivating Scenario

Scenario An Identity Provider issues pseudonyms to authenticate users at third-party Service Providers, which can in turn check the authenticity of these pseudonyms at the Identity Provider, without getting to know the real user identity. However, the Identity Provider knows all the services requested by the users, which discloses personal information to the Identity Provider and allows it to profile the users. ■

Know Uses and Related Work

The New Federated Privacy Impact Assessment (F-PIA). Building Privacy and Trust-enabled Federation. Information and Privacy Commissioner of Ontario & Liberty Alliance Project, January 2009.

Categories

- Risk management
- Procedure
- Enforce
- Create

Related Patterns

P Obligation-management

Supporting Patterns

P

Sources

- https://privacypatterns.org/patterns/Federated-privacy-impact-
- https://privacypatterns.eu/#/patterns/federated-privacy-impact
  0-0-0-1-0-0-1-0-0-0-1-0-1-0-0-0

2.8 Use of dummies

Summary

This pattern hides the actions taken by a user by adding fake actions that are indistinguishable from real.

Context

This pattern is applicable when it is not possible to avoid executing, delaying or obfuscating the content of an action.

Problem

When users interact with ICT systems their actions reveal a lot of information about themselves. An option would be for users to not perform such actions to protect their privacy. However, this is not possible since users cannot completely avoid executing these actions because they need to perform them to achieve a goal (e.g., search for a word on the Internet, send an email, search for a location).

Goal

G To hinder the adversary's ability to infer the user behavior, as well as her preferences.

Solution

Since the action must be accurately performed, an option to provide privacy is to simultaneously perform other actions in such a way that the adversary cannot distinguish real and fake (often called dummy) actions.

Constraints and Consequences

C This pattern entails the need for extra resources to perform the dummy actions, both at the side of the user that must repeat the action, and at the server side that must process several actions. Sometimes it may degrade the quality of service since the service provider cannot personalize services. It has been demonstrated that generating dummies that are perfectly indistinguishable from real actions (in terms of content, timing, size, etc...) is very difficult

Motivating Scenario

Scenario Alice wants to search for an abortion clinic on Google, but she does not want to reveal her intentions of abort to an adversary that may be eavesdropping this search (e.g., ISP provider, system administrator of her workplace, etc). ■

Know Uses and Related Work

The use of this pattern has been proposed to protect privacy in location based services (the user reveals several locations to the service provider so that her real location is hidden), anonymous communications (the user sends fake messages to fake recipients to hide her profile), web searches (the user searches for fake terms to hide her real preferences).

Categories

- Hide
- Obfuscate

Related Patterns

P

Supporting Patterns

P

Sources

- https://privacypatterns.org/patterns/Use-of-dummies
- https://privacypatterns.eu/#/patterns/use-of-dummies/0-0-0-0-0-0-0-0-0-0-0-0-0-0-0-0

2.9 Who's Listening

Summary

Context

Users of a service regularly share its usage with other users. Sometimes these are users they know personally, an sometimes these are anonymous, unauthenticated persons. This occurs particularly in shared or collaborative environments where content is generated. Knowledge of the contributions of other users contributes to additional or refined content in general. Controllers facilitating this interaction therefore encourage the users to form groups or allow public access. Though when the amount of disclosure is high it is difficult to keep track of attribution and modification.

Problem

Users do not know if the content they are accessing or have disclosed has been accessed or modified by others, nor if it is someone they know.

Forces and Concerns

- Users want to know who can access their disclosures and those of others
- Users want to know that specific other users have accessed or modified content
- Controllers do not want users to be unaware of who can see their disclosures
- Controllers want to log access to prevent abuse

Goal

G

Solution

> Provided that users know their access is not private, inform them of other users, even unauthenticated, which are also access-

ing the content in question.

Implementation

Ensure that it is made clear to the user that the content they are about to view is accessed in a shared and public manner. Their access will be visible to others, and may be recorded by the system (if applicable) for historic views, or for preventing abuse.

The implementation of the system prior to this will likely only require the addition of UI elements to indicate the access state as the system already perceives it. Each user may be shown using some identifier easily recognizable by other users, such as a randomly selected avatar (e.g. Gravitar), initial(s), username, or profile picture. The same may identify unauthenticated users as 'anonymous'.

Where historic views are provided, the same consistent identifier can be shown next to differential changes along with timestamps. The ability to edit, remove, or anonymize a contribution may also be available if desired. Details of these extra features, or justification for the lack of user ability to perform these actions, should be provided prior to usage.

Constraints and Consequences

C This pattern will only work, if the users trust the system that provides the information and log in personally. In web based systems that don't require personal login, it is not possible to [reliably] detect, who is visiting the site (even cookies [and browser fingerprints] do not [necessarily] reveal information about the users' identities). This is problematic for [attribution], but it ensures that the users can control their privacy.

Motivating Scenario

Scenario The [MIME] protocol provides an option so that receivers of the [message] are asked to confirm the message. It is defined in [RFC 8098 (Hansen & Melnikov 2017)]

The BSCWshared workspace system (Bentley, Horstmann, and Trevor 1997) [logs] accesses to the shared [content]. The event log can be queried [by users] and for each document stored in the shared workspace, the users can define notification patterns. By these means, it is possible for an author of a document to find out who read the document (and when).

Various collaborative environments, like Google's Docs, or chat rooms, instant messaging, and other immediate content sharing mediums frequently provide lists of currently online users.

These can also indicate a number of anonymous users who have not authenticated, but have reduced privileges.

Know Uses and Related Work

Categories

- Inform
- Provide

Related Patterns

P Awareness Feed

P Privacy Awareness Panel

P Appropriate Privacy Feedback

P Buddy List

P Reciprocity

Supporting Patterns

P

Sources

- https://privacypatterns.org/patterns/Whos-Listening

2.10 Privacy Policy Display

Summary

Context

Privacy policies are an important element in the processing activities of a controller. They not only relay to data subjects, the users, crucial aspects about the processing in question, but also adhere to the laws which mandate those policies. Balancing the accessibility of these policies however with the legal comprehensiveness needed is nontrivial. As such users do not naturally familiarize themselves with privacy policies as they need to be verbose, and often complex, to comply with the law. It is therefore necessary that controllers ensure that users are indeed informed before soliciting their consent.

Problem

Whenever the user's information is requested, it must be clear to them exactly what information is needed, who requests it, and what will be done with it.

Forces and Concerns

- Users do not want to read extensive policies, but they do want to understand any relevant risks
- Controllers need users to understand specific policy elements in order to legally process their data
- Users would rather be provided with relevant and ideally concise information than all of it at once
- Controllers want users to trust that they are not trying to hide the risks of using the system

Goal

G

Solution

> As requests for personal data are made, state clearly what information is needed by whom, for which purposes, and by what

means, prior to soliciting consent.

Implementation

The Article 29 Working Party of the Data Protection Directive of the European Union have set out recommendations regarding the distribution of policy into a layered format. They suggest three tiers, each providing additional detail. Users should have clearly visible access to successive detail upon the controller's request of the related personal data.

The first tier, 'short notice', shall offer core information necessary for users to understand the purposes and means of processing. It should provide a clear mechanism to obtain further detail. This tier is aimed towards maximum user understanding.

The second tier, 'condensed notice', includes a summary of pertinent information as required by Article 13 of the General Data Protection Regulation (GDPR), the successor to the Directive. This non-exhaustively includes additional information regarding contact details of applicable entities, legal basis or obligation, legitimate interests, recipients, retention, data subject rights, and whether automated decision making is in use.

The third tier, 'full notice', presents all remaining information required by the GDPR in addition to the previous information. This is the variation which expresses the full detail of the policy which best holds up the legislative requirements.

Constraints and Consequences

C By constantly reminding users what it really means to share their information, they will better contemplate the personal data they choose to provide. However, users may also become fatigued or otherwise desensitized by frequent reminders and begin to overlook privacy policies. As such it is important to balance the levels of visibility and implicit severities of the information conveyed.

Motivating Scenario

Scenario Alice wants to search for an abortion clinic on Google, but she does not want to reveal her intentions of abort to an adversary that may be eavesdropping this search (e.g., ISP provider, system administrator of her workplace, etc).

Know Uses and Related Work

Categories

- Inform
- Provide

Related Patterns

- Dynamic Privacy Policy Display
- Policy Matching Display
- Privacy Aware Wording
- Layered Policy Design
- Platform for Privacy Preference

Supporting Patterns

Sources

- https://privacypatterns.org/patterns/Privacy-Policy-Display
- http://privacypatterns.wu.ac.at:8080/catalog/

2.11 Layered Policy Design

Summary

Split privacy policies into nested, successively refined versions. Leave the legalese to the lawyers.

Context

As the law in various parts of the world requires a number of considerations, policies tend to be long, complex documents which are difficult to understand. The same holds true for privacy, which supplies its own legislative concerns, particularly regarding data protection. The [data] controller in these instances, provides users (data subjects) with services (or products) to which privacy policies apply. These suffer the same detail rich and superfluous content pitfalls as other policies, though are legally required to be available to users in a manner which is both understandable and complete.

A data controller offers detailed, legal explanations of their privacy and data protection policies.

Problem

The controller needs to balance comprehension and comprehensiveness in their privacy policies in order to ensure that users choose to inform themselves. If they do not, then processing their information is unlawful.

Forces and Concerns

- Users do not want to read complex and long policies, and most will simply not read them unless they are very concise
- Users still want to understand any important distinctions which might cause them risks they would rather not take
- Controllers want to comply with legal requirements to avoid punitive measures as well as bad publicity
- Controllers also want users to know what they are signing up for when using a service, without being unpleasantly surprised

Privacy policies may be difficult to understand and hard to read. What was initially conceived as an instrument to inform users is now almost useless, as they have become riddled with legalese and all sort of extraneous details. As a consequence, users do not read the privacy policies, for being long and cumbersome.

However, privacy policies are legally binding documents, which makes it difficult to get just rid of these legal aspects.

Goal

G Make users really understand what they can expect about their personal data from a data controller (in terms of which data is managed, for which purposes, etc.)

Solution

Extract the most crucial aspects of the privacy policy, which users are most likely to read, to the foreground. Nest successive detail levels within these components so that users can quickly find information that is relevant to them.

Implementation

A short notice may provide a summary of the practices that deal with personal data, highlighting those which may not be evident to the data subject. Then, a longer policy may provide specific information, split into sections, detailing any uses of personal data. And finally, the whole legal text of the privacy policy can be specified.

Constraints and Consequences

C [Helps users] understand what they can expect about their personal data from a data controller (in terms of which data is managed, for which purposes, etc.) Also fosters simplicity, transparency and choice.

However, [multiple] versions of the privacy policies [need to] coexist, which may introduce potential contradictions; in particular, the data controller must ensure that updates are performed in parallel and coherently.

The use of this pattern fosters simplicity, transparency and choice. However, two versions of the privacy policies coexist, which may introduce potential contradictions; in particular, the data controller must ensure that updates are performed in parallel and coherently.

Motivating Scenario

Scenario See examples at Terms of Service Didn't Read. The average user would take 76 work days to read the privacy policies they encounter each year.

Know Uses and Related Work

- An early example of layered privacy policy by TRUSTe and its mobile version, which are discussed in Pinnick, T. Layered Policy Design. TRUSTe Blog, 2011
- There are several sites that use this pattern nowadays, albeit not always with that name. One example is Banksia Villages, which provides a Simplified Privacy Policy as well as an Extended one
- It is recommended by British Information's Commissioner Office in its Privacy Notices Code of Practice (p.55)
- This concept is quite similar to the Creative Commons license layers in the field of copyright management.

Categories

- Inform
- Explain

Related Patterns

P Awareness Feed

P Appropriate Privacy Icons

P Icons for Privacy Policies

P Privacy Labels

P Privacy Color Coding

P Abridged Terms and Conditions

P Privacy Aware Wording

P Privacy Policy Display

P Impactful Information and Feedback

P Dynamic Privacy Policy

P Policy Matching Display

Supporting Patterns

P

Sources

- https://privacypatterns.org/patterns/Layered-policy-desig
- https://privacypatterns.eu/#/patterns/layered-policy-desi
  0-0-0-0-0-0-0-0-0-2-0-0-1-0-0-0

2.12 Discouraging Blanket Strategies

Summary

Context

Socially oriented services on the Internet allow their often diverse userbase to share content. These masses of users and shared content are also varied enough to discourage individual attention. Controllers prefer to protect themselves from additional complexity and investment into features which provide them with less data. Users, however, feel in need of privacy settings to distinguish their personal risk appetite from that of the norm. They each have their own ideas about the sensitivities of their information, which makes sufficient controls difficult to implement.

Problem

Overly simplified privacy settings following all or nothing strategies could result in over-exposure, self-censoring, and unsatisfied users.

These all or nothing strategies could refer to privacy settings which holistically apply to all content, or to binary (or otherwise deficient) choices for public visibility.

Forces and Concerns

- The level of invasiveness depends upon the context of the content being shared
- Users have differing levels of sensitivity attributed to contexts, and thus different levels to provide their content
- They trust the controller to different extents
- Controllers do not want to violate user privacy
- They want to protect themselves from blame if a user's privacy is violated
- They want users to produce content which is at least somewhat valuable. (For instance, when nobody can see the content there won't be an impact in collaborative environments)

Goal

G

Solution

Provide users with the possibility to define a privacy level for content being shared with the controller, or with other users. Give them a range of visibilities, so that they can decide the access-level of the content being shared according to different users, or service-defined groups.

Implementation

Provide users easily-recognizable visual elements to define the privacy level for each content submission. Use controls, such as (drop-down) lists, combo boxes, etc. to provide a range of possible privacy levels.

The privacy levels could be defined in terms of social group of the users in question to the user who is sharing content. For instance, family, closest friends, colleagues, acquaintances, everybody. This is in line with Reasonable Level of Control and Selective Access Control, where a user might be given the opportunity to define their own groups or set individual privacy levels.

The privacy controls themselves also need to be designed in such a way that it is very clear to users what each setting does and what it means for their privacy.

Constraints and Consequences

C

Benefits

Grants users complete control over the privacy of the content being shared, which may lower the bar for them to share certain data they otherwise would not.

Liabilities

Users could find having to set privacy settings every time they share content to be tedious. It would be necessary to define reasonable defaults for privacy settings (least effort for minimal sharing).

Motivating Scenario

Scenario

- Facebook
- Google Plus

Know Uses and Related Work

Categories

- Control
- Choose

Related Patterns

P

Supporting Patterns

P

Sources

- https://privacypatterns.org/patterns/Discouraging-blanket-strategies

2.13 Reciprocity

Summary

Context

In services where users may either socially or collaboratively contribute, participation may be a foundation for the service's business model. In these situations the quality and frequency of content affects the success of the service, and thus users have a large impact on its survival. Whether any single user contributes, or not, plays a role in profitability, which puts the controller in a position to encourage or enforce equal participation. Users may respond to such ideas negatively, however, especially if they do not see potential gains worthy of their effort and personal risks to privacy.

Problem

Equal participation does not always result in equal rewards. In some cases, participants do not need to contribute at all to benefit from the content generated by the group. Any who feel slighted are then likely to contribute less, eventually jeopardizing results for the group.

Forces and Concerns

- Users may feel uncomfortable due to unfairly spread workload, and it could generate problems with the overall tasks' fulfillment. Some might decide to leave the group altogether
- Inequality may additionally generate a tense work or social environment
- Controllers want group dynamics to work so that content generation continues

Goal

G

Solution

Limit the benefits gained from the group effort to the amount of effort contributed. All contribution should be afforded proportionate gains.

Structure

Ensure that all group members' activities result in an improved group result that is beneficial for all group members again. Prohibit people to benefit from group results if they are not willing to help the group in return.

Implementation

Prior to completing designs on functionality, determine the benefits as opposed to efforts or costs on all possible user activities. Weigh these, with input from any necessary stakeholders. Any feature which does not affect more than one user does not need to be assessed.

For user groups that are able to affect one another within a feature or functionality, consider them each a case for a collaboration mode. If a user within this group performs an activity, they are expected to reciprocate on any benefits (or gain from costs). This is in proportion to the weighted effort of the feature determined earlier.

The way in which users reciprocate is up to specific implementation. It may include required effort (satisfied by certain activities) before their activity's resulting benefit is realized.

Alternatively, it may prevent additional beneficial activities until they contribute. It may also make their discrepancy public, allowing the users to determine tolerable thresholds. In all these cases it is useful to keep track of each user's ratio within each collaboration mode they feature.

It is important that any use of user data is done so under the explicit and properly obtained permissions required. Deriving value from participation rewards users for providing personal information, and thus they must be informed about how their data may be used.

Constraints and Consequences

Benefits

Finding the inequalities in the design phase involving all stakeholders can reduce the objections for participating to the system since the benefit is made explicit to the end-user. Using this pattern minimizes reasons for groupware applications failure.

Liabilities

The pattern is only needed in situations, where the critical mass of participation can only be reached with most users participating. If the community is very large (e.g. a news group), it can succeed with a small number of active participants and a larger number of inactive participants (free riders, lurkers). Consent given by users needs to be freely-given, which is a requirement easily overlooked as controllers are tempted to coerce participation. As with sunk cost, emotional investment can pressure users into choices they do not truly consent to. Therefore, Lawful Consent should be used in this pattern.

Motivating Scenario

Scenario

- Reddit
  - Gold: those who never ’gild’ are set apart from those that do, as this fact is made clear within a profile.
  - Upvotes: while individual votes are not publicized, the votes received as part of contributing are.
- Facebook Friends; LinkedIn Contacts; etc.: these relationships are one-to-one, to have a contact is to be a contact.
- TUKAN;
- Buddy Lists in Instant Messaging systems;
- Bulletin Board Systems.

Know Uses and Related Work

TUKAN: The collaborative programming environment TUKAN introduced the concept of modes of collaboration (MoC) to ensure reciprocity. A MoC is a lightweight mode, which defines

> possible collaborative activities. It combines a specific level of privacy (cf. Masquerade) with the right of receiving information about other users. It thus provides a set of predefined Attention Screens. This combination ensures that a user can only utilize information from other users at a privacy level on which he is also willing to reveal personal information.

Categories

- Control
- Update

Related Patterns

P Incentivized Participation

P Pay Back

P Masquerade

P Buddy List

P Lawful Consent

Supporting Patterns

P

Sources

- https://privacypatterns.org/patterns/Reciprocity

2.14 Asynchronous Notice

Summary

Context

Many sensor related or other recurring forms of data collection are important for improving service (or product) quality, but occur in a manner which is not apparent to the user. Even where the user is informed of such processing, the nature of that processing may cause it to occur within contexts the user would not consent to. Users are also subject to forgetfulness. The controller processing this information therefore seeks to ensure that consent is retained. Some interfaces necessitate more restrictive use of screen real estate, however, and as such can not accommodate extensive information or persistent elements.

Problem

Users being tracked and monitored may not consent to processing they had previously consented to, as the context surrounding that processing is subject to change.

Also, initial consent may have been forged by an attacker or have been provided by another user of a shared device – if synchronous notice is only provided at the time of consent, a user may inadvertently distribute personal information over a long period of time after having lost control of their device only momentarily.

Forces and Concerns

- Users may change their minds or forget about consent they have given
- Users may not realize the processing they consented to is currently in effect, potentially allowing collection of information they would not want collected
- Controllers do not want to collect data for which consent is uncertain, where users may feel violated or otherwise let down
- Controllers cannot remind users of their consent all the

time

- Providing an asynchronous notice requires a reliable mechanism to contact the user (a verified email address or telephone number, for example). Care should be taken to ensure that the mechanism can actually reach the person using the device being tracked. (For example, notifying the owner of the billing credit card may not help the spouse whose location is being surreptitiously tracked.)
- In contrast to the common privacy practice of providing consistent and reliable systems, you may wish to provide random asynchronous notice. If there is a concern that a malicious user may have opted-in the user without their knowledge, a notice that is sent once a week at the same time each week may allow the attacker to borrow the device at the appointed time and clear the notice
- Many repeated notices may annoy users and eventually inure them to the practice altogether. Take measures to avoid unnecessary notices and some level of configuration for frequency of notices. This must be balanced against the concerns of an attacker's opting the user in without their knowledge

Goal

G

Solution

Whenever there is a context switch, sufficient duration, or random spot check, provide users with a simple reminder that they have consented to specific processing. The triggers and means for contacting the user may be chosen by the user themselves, who should be able to review and if necessary retract their consent.

Implementation

Implementation depends on the medium chosen for conveying the notification, and also on the service facilitating collection.

For mediums with less space, shorter messages should be provided, but even in more traditionally long-winded options such as email, brevity should be favored. The user should be able to obtain more information by a linking mechanism, also dependent on the medium. The most important information to provide is the fact that they have consented to specific data for specified purposes, and that a context change, spot check, or specified duration has triggered the reminder. Context changes are most notable, as these are most likely to affect the consent. Note that changes to purposes and means instead require new consent, not merely notification.

Asynchronous notices may also include a summary of the data recently collected (since the last notice, say) in order to provide clarity (and reminders) to the user about the extent of collection. By ensuring that users aren't surprised, asynchronous notice may increase trust in the service and comfort with continued disclosure of information.

Constraints and Consequences

C

Motivating Scenario

Scenario 1. Google Latitude reminder email

Google Latitude users can configure a reminder email (see below) when their location is being shared with any application, including internal applications like the Location History service.

This is a reminder that you are sharing your Latitude location with the following application(s):

Google Location History You may disable these applications at any time by going to `https://www.google.com/latitude/apps?hl=en]`

Do more with Latitude Go to `https://www.google.com/`

latitude/apps on your computer and try the following:

Google Location History lets you store your history and see a dashboard of interesting information such as frequently visited places and recent trips. Google Talk Location Status lets you post your location in your chat status. Google Public Location Badge lets you publish your location on your blog or site.

You are receiving this reminder once a week. To change your reminder settings, go to: `https://www.google.com/latitude/apps?hl=en&tab=privacyreminders`

2. Fire Eagle My Alerts

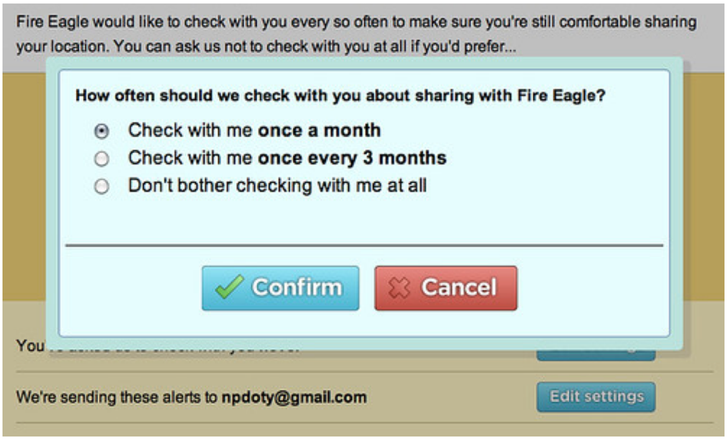

Fire Eagle would like to check with you every so often to make sure you're still comfortable sharing your location. You can ask us not to check with you at all if you'd prefer...

How often should we check with you about sharing with Fire Eagle?

- ● Check with me **once a month**
- ○ Check with me **once every 3 months**
- ○ Don't bother checking with me at all

You

We're sending these alerts to **npdoty@gmail.com**

Figure 2.2: Fire Eagle My Alerts

Know Uses and Related Work

Categories

- Location
- Notice
- Mobile
- Inform

- Notify

Related Patterns

P Preventing Mistakes or Reducing Their Impact

P Impactful Information and Feedback

P Ambient Notice

P Buddy List

P Single Point of Contact

Supporting Patterns

P

Sources

- https://privacypatterns.org/patterns/Asynchronous-notice
- http://privacypatterns.wu.ac.at:8080/catalog/

2.15 Abridged Terms and Conditions

Summary

Context

Controllers which provide services (or products) to users have various policies, including those which affect user privacy, which need to relay to the user. If users do not have knowledge of the risks, rights, and responsibilities relevant to them, this is the fault of the controller. Keeping users (the data subjects) informed, especially prior to acquiring consent, is a legal requirement. As such, controllers need to ensure that this is the case. It is however difficult to keep user attention on such matters, as they are often more willing to spend time on other things, including actually using a system. Efforts to hold attention by force also face active resistance.

Problem

Users often overlook Terms and Conditions when presented with them in their entirety before the use of a service.

Forces and Concerns

- Controllers need to write Terms and Conditions in a manner which will hold up to scrutiny from the law, but this is not accessible to users
- Controllers want to ensure that users are fully aware of the risks of using the system before using it, for both legal and image purposes
- Users want to get to using the service without being blockaded by walls of text, but the also do not want to be blindsided about policy
- Users want to understand the risks in as little time as necessary, at a granularity most suitable to their value of it

Goal

G

Solution

Summarize the legally sufficient Terms and Conditions into concise and relevant variations which suit the user's level of interest and attention. At first use of a service, users should be able to investigate further, but not have to read much to understand the risks involved.

Implementation

Prepare the concise Terms and Conditions according to a user perspective, focusing on matters which are most important to them. Aspects which do not affect them should not be included in summarized variations. Where areas of potential interest are easily bundled, group them under a general summary with option to expand further. Using titles in this regard is less helpful if they do not summarize the policies involved, as expanding should not be necessary unless the user notes an area worth their concern. Aim towards a page or less of information, as the inclusion of a scrollbar may dissuade the user.

The full, legally sound version should also be available, and should not contradict the summarized information. This applies at first use as well, as a user should be allowed to review detail prior to being subjected to it.

Constraints and Consequences

C The appropriate and concise summarization of the Terms and Conditions will allow users to get a sufficient idea of the rights, risks, and responsibilities relevant to them. As it should be brief, only the most carefree users will overlook them. It will not therefore be guaranteed that users are fully informed, and this should be taken into account.

C Due to the fact that the [Terms and Conditions] of an application are condensed to a size that is easily comprehensible, a user's trust in the application can be increased. [Additionally, this] ensures a greater transparency to the user since possible implications for the user, which may result through the usage of the application, can be recognized more easily beforehand.

Motivating Scenario

Scenario ■

Know Uses and Related Work

> - Support-U: An example of an abridged TAC is given in fig. 3. The figure shows the results of the abridged TAC pattern used for the Support-U application
> - Connect-U: The user has to sign a license agreement of the size of one page in A4 format. On this page the agreement about the data usage is described in clear detail
> - Meet-U: The key points of TAC that affect the user's privacy the most, are displayed on one screen. Hence, the gathering and processing of data are addressed and summarized briefly. The long version of the TAC is linked. The user has to agree on that before continuing with the application

Categories

- Inform
- Explain

Related Patterns

P Privacy Aware Wording

P Layered Policy Design

P Privacy-Aware Network Client

Supporting Patterns

P

Sources

- https://privacypatterns.org/patterns/Abridged-Terms-and-Conditions

2.16 Policy Matching Display

Summary

Give one careful thought to your privacy needs, then be always able to swiftly apply what you decided.

Context

Controllers have policies written in a manner appropriate for legal evaluation, as it is the legal compliance which warrants them in the first place. Users tend to not be able to comprehend such language, and do not typically care to spend the time and effort required to parse it. However, much of the content in these policies is consistent throughout the services they use.

Users value using a service (or product) without having to go through repetitive and verbose policy detail. However, these users must still understand the policies which apply to them in order to not be blindsided. Controllers need to avoid this as keeping users happy is integral to a sustainable business model.

A user wants to start using a new service, which lets the user configure several privacy-related parameters. The user often does the same with new, different service providers.

Problem

Users may get overwhelmed by the complexity of policies impacting privacy when using a service, compromising the validity of their informed consent.

Forces and Concerns

- Users do not want to have to read privacy policies, but do want to know about relevant and important distinctions from their personal preferences
- Controllers need to have policies which are tailored to legal compliance, but also need users to understand risks and responsibilities.
- Users may not like the default values chosen by controllers

for application settings, even if those defaults are privacy friendly
- Controllers would like users to use a service immediately, with as little in the way and as little potentially discouraging as possible.

Goal

G Allow users to provide a consistent privacy-related behavior, while reducing their cognitive workload every time they enroll in a new service.

Solution

Retrieve user policy preferences and use these to highlight contradictions with the privacy policy. Where possible, configure application settings to the values which best adhere to these preferences.

Implementation

User policy preferences may be collected and managed by a controller, exposed by their user agent, or at a well-known URI. They may be highlighted through an overlay of elements or handled in-line where context plays an important role. In either case these notifications should not encourage users to apply settings which do not match their preferences in order to remove them.

On the other hand, if the notification is not noticeable, the user may overlook an important policy distinction. Notifications which are persistent or ubiquitous may quickly desensitize users, and should also be used with care.

Before contracting a service, the service provider retrieves the user preferences (exposed by their user agent, or at a well-known URI), and presents the user a comparison between their preferences and the privacy policies applied by default by the service operator, which in turn automatically adapts any configurable values to the user’s declared preferences.

Constraints and Consequences

C Allows users to provide a consistent privacy threshold while reducing cognitive workload as they use services.

Constraints

Expressing and comparing the policies requires a consistent machine-readable format. There however numerous approaches to this. The Platform for Privacy Preferences pattern addresses this through eXtensible Markup Language.

This pattern requires sharing a machine-readable format to express and exchange definitions of privacy policies between the user agent and the service providers. Several such formats exist, yet they are not always supported by either user agents or by service providers. Besides, not all the privacy policy nuances can be expressed in existing privacy policy languages.

Motivating Scenario

Scenario ■

Know Uses and Related Work

- For an academic discussion, see Graf, C., Wolkerstorfer, P., Geven, A., &Tscheligi, M. (2010, November). A pattern collection for privacy enhancing technology. In PATTERNS 2010, The Second International Conferences on Pervasive Patterns and Applications (pp. 72-77)
- For a discussion of privacy languages see Kumaraguru, P., Cranor, L., Lobo, J., &Calo, S. (2007, July). A survey of privacy policy languages In SOUPS'07: Proceedings of the 3rd Symposium on Usable Privacy and Security. and Becker, M. Y., Malkis, A., & Bussard, L. (2010)
- A related, classic initiative was W3C's The Platform for Privacy Preferences 1.1 (P3P1.1) Specification, however, the matching was performed at the client's side
- A more recent example is available at S4P: A generic language for specifying privacy preferences and policies. Microsoft Research

- ... and Sacco, O., & Passant, A. (2011, March). A Privacy Preference Ontology (PPO) for Linked Data. In LDO

Categories

- Inform
- Provide
- control

Related Patterns

P Privacy Policy Display

P Platform for Privacy Preferences

P Trust Evaluation of Services Slides

P Dynamic Privacy Policy Display

P Icons for Privacy Policies

P Appropriate Privacy Icons

P Privacy Color Coding

P Privacy Aware Network Client

Supporting Patterns

P

Sources

- https://privacypatterns.org/patterns/Policy-matching-disp
- https://privacypatterns.eu/#/patterns/policy-matching-dis 0-0-1-1-0-1-0-0-0-0-0-1-1-0-0-0
- http://privacypatterns.wu.ac.at:8080/catalog/

2.17 Incentivized Participation

Summary

Users of a system have varying privacy concerns, and different sensitivities associated with their personal information. These users need ways to contribute without leaking sensitive details, or to perceive a worthwhile tradeoff for those details. This can be achieved through social encouragement (i.e. participation and shared trust), direct value exchanges (discounts and giveaways), or some other derived value (e.g. positive reinforcement).

Context

A data controller derives various values from the participation of its users (i.e. data subjects). The more that these users participate, explicitly providing context and implicitly providing metadata (e.g. statistics and telemetry), the better the controller fares in a number of respects. Despite this key relation, over-sharing can greatly infringe upon a user's right to privacy. Many controllers therefore aim to respect this right when benefiting from user interactions.

As the controller should recognise the necessity of specific, informed, and freely given Lawful Consent, users are made aware of the pitfalls of such a system. As they are informed, perhaps through a combination of a Privacy Dashboard or Awareness Feed, users may balance the privacy related tradeoffs.

This minimises the privacy risks taken according to the user's personal informed choices, and protects the controller from inadvertently undermining the user's privacy. The controller still desires participation, however, and may therefore make additional concessions or provisions to help make the tradeoff worthwhile or non-existent. The controller may complement its strategy with more granular choice in order to achieve this, such as with Selective Disclosure and Selective Disruption.

Problem

Controllers which gain from user activity want to push for participation, but this can negatively affect users.

Users have varying degrees of concern about their privacy, and do not respond to different forms of encouragement the same way. By penalising under-sharing and inactivity, or being misleading, users become alienated and distrusting of the system. As such this problem has multiple elements. These include asymmetric returns on investment, and the standard incentive deficiency, where users lack the encouragement to participate.

Asymmetric Returns on Investment In many situations, some benefit more than others. In extreme cases, users may benefit through minimal participation and thus contribute very little to the system's derived value. Those who do not perceive an acceptable value despite considerable contribution may then withdraw.

An example of this behaviour might be seen in dating sites where users with only a flattering picture may succeed more than those with detailed profiles. Similar cases can be made for other social media, as well as with asymmetric bandwidth on peer to peer sharing. With torrent technologies, this is often referred to as 'leeching'.

Standard Incentive Deficiency

Users which provide limited or vague information due to privacy concerns may have less opportunity for participation. Another way this occurs is when they are not driven by positive social reinforcement. The lack of friends, followers, potential matches, etc. leads to user inactivity.

Forces and Concerns

- Controllers want to encourage users to grow their networks and further participate, though this may increase their exposure to privacy concerns and bad user experience
- User concerns over privacy may cause for adverse reactions to unwelcome changes or discoveries in policy, especially

attempts to goad or guilt-trip users into activity (e.g. "This user does not participate")

- Where controllers sanction users for inactivity, or undesirable activity, they affect user experience, for better or for worse
- Controllers may wish to lock secretive users out of certain services (or products), but this is likely to alienate them
- Users may want to use a limited set of functions which do not undermine their privacy, whereas controllers derive less value from these users

Goal

G

Solution

Privacy concerns need to be met with valid reassurances about issues which matter to the user. Firstly, users should know that the system holds their preferences in high regard. Secondly, they should perceive real value in their participation. Finally, if desired, users should be assisted in a smooth transition into the ecosystem.

Rationale

Implementation

The three elements of the solution are elaborated on in the following sections.

Adherence to Preferences

Users need to know they are able to participate without the system undermining their personal preferences. This should apply from the very first usage of a system. Everything the system does globally must adhere to privacy friendly defaults. Any service which cannot uphold these expectations should be deactivated for new users, and only be enabled once these users consent to the additional processing. Attempts to solicit this should not be invasive.

Value Perception

With privacy concerns at ease, encouraging equal participation entails reciprocity. These can be in both social and financial forms.

All participation should result in value derivation (social or otherwise) for all participants, and not just individuals. As a consequence of this mindset, the derived benefits of users who do not participate are limited. This secondary effect may also be the primary mechanism for reciprocity, though positive sum approaches will be met with more support.

As an additional measure, or where equality is not feasible, provide an alternative incentive. If not social, then financial incentives (discounts, waived fees) can be provided to active users. This can be limited to those who the system can identify as receiving poor value returns on their contributions. User retention examples tend to be less frequent, with the odd website sending ’come back’ emails which promise fee reductions.

Note that attempting to pinpoint users based on activity levels may reveal more sensitive characteristics, and as such should in any case require their prior informed consent.

Another approach is the explicit provision of virtual currency necessary to benefit from the system, those who contribute will then have more currency at their disposal. They may opt to be applauded for their efforts publicly, but again should not be forced to.

Examples of social value perception are the Facebook like/reaction, Google’s +1, Reddit Gold, and Twitter’s reposting. These approaches are enacted by the users instead of the system, and are therefore less intrusive.

The system should permit participation itself without a risk to privacy. Those with low identifiability should not be barred from participation.

Transition Assistance

In terms of computer aided pairing, greater participation may be achieved if users consent to intelligent nudging. Shümmer suggests that users are more likely to participate when their interests are shared with others, and thus, a system would help users identify those with similar interests. This is another solution which relies on prior consent, however. By encouraging interaction between these users, a system would derive more activity and therefore further value.

On the other hand, Shümmer points out that mixing dissimilar users may also result in unexpected activity. It may allow the system to discover notions about its users which were not previously apparent. This is yet another way to increase value, though will likely be far more intrusive than the former. Users should be properly informed of the possible consequences of 'mixing it up'.

In order to ensure that recommendations made by the system do not have increasingly negative side effects, the system should learn from ineffective suggestions. This is limited to where it has permission to do so. Where user activity drops, a system should aim not continue in the same fashion as before. When it climbs, however, the system should improve whatever characteristics likely resulted in that climb.

This can be made more explicit by soliciting feedback from the user themselves, as also suggested by Shümmer. Even this, though, is subject to negative reactions. However, acclimating users to an environment of openness and transparency will also build trust - potentially resulting in the use of services users would not have used otherwise.

Constraints and Consequences

C Applying the concepts represented in this pattern may have certain trade-offs associated.

Some users respond negatively to being nudged into participation, it is also intrusive Learning about interests may be considered invasive, and should require an opt-in, along with assurances for privacy concerns. Furthermore, any assumptions a system makes about a user due to incomplete or misleading data may lead to further reduced activity. Suggesting unappealing interactions might cause the user to seek alternative social media or withdraw from sharing altogether.

Where the system is quite large, it is more resilient against inactive users, and thus can sustain a considerable amount of inequality. However, if an adequate balance is not maintained it may result in an unpredictable ecosystem. Social media giants such as Twitter have struggled to turn a profit, while Facebook's stock price has climbed continually.

Constraints

Isolating users, and learning from their actions, based on feedback loops requires prior informed and explicit consent, as potentially invasive conclusions may be derived.

Motivating Scenario

Scenario ■

Know Uses and Related Work

Categories

- Control
- Choose
- Consent

Related Patterns

P Pay Back

P Reciprocity

Supporting Patterns

P

Sources

- https://privacypatterns.org/patterns/Incentivized-Participation

2.18 Outsourcing (with consent)

Summary

Context

Controllers often do not have the means to feasibly or sufficiently process the data they oversee to the extent they desire. In these cases, they seek an external processor or third party to handle the process. This typically conflicts with their already obtained consent from their users (their data subjects), as further processing by a third party is not necessarily compatible with the agreed upon purposes. In these situations, the controller does not have legally obtained consent for this processing and will be liable if they carry it out.

Problem

Third party processors do not inherent user consent granted to a controller, but need each user's consent before they may process their information. The processor cannot contact the necessary users as they have no lawful access to any means to identify them.

Forces and Concerns

- Controllers wish to outsource processing when it is not feasible or viable to do so themselves
- Third party processors want to process information efficiently without needing to address other considerations
- The controller does not want to be liable, or damage their reputation
- Outsourcing has a strong impact on the security and privacy requirements of organizations. A contract will bind both parties
- The outsourced third party will be obliged by all data protection principles to which the controller is, in addition to stricter measures imposed on processors

Goal

G

Solution

Obtain additional (Lawful Consent) [Lawful-Consent] for the specific purposes needed from each user before allowing the third party to process their data. Do not process the data of users who do not consent.

The consent can be seen as a contract establishing what and how data will be processed by the [third party]. The [controller] must also ensure, preferably by a written agreement, that the [third party] strictly follows all conditions relating to data processing that were imposed on [them].

Implementation

Before outsourcing data processing, it is necessary to obtain consent from the user and create an agreement between the controller and the third party. The consent itself needs to be freely given, informed, specific, and explicit. It should indicate purposes and means (physical or informational) regarding the controller and the third party. There is also an execution dependency between the controller and the user.

Figure 2(b) shows an SI* model explaining the solution of Compagna et al. (2007)

Constraints and Consequences

C

Benefits

The pattern solves the problem of granting [the consent] necessary to perform out-sourced data processing by assuring [users that their information is] processed according to the contract.

Liabilities

The [controller] may want assurance that the [third party] does not repudiate the data processing agreement and the [user] does not repudiate the consent. As such the controller may decide to use the Non-repudiation pattern.

Motivating Scenario

> **Scenario** The scenario described by Compagna et al. (2007) features a Health Care Centre (data controller) and a user (data subject), Bob, who needs constant supervision. The subcontractor, a Sensor Network Provider (third party supplier), installs and maintains the network responsible for automated monitoring of Bob's health. This subcontractor needs additional specific, informed, explicit, and freely given consent from Bob. ■

Know Uses and Related Work

Categories

- Control
- Consent

Related Patterns

P Lawful Consent

P Non-repudiation

Supporting Patterns

P

Sources

- https://privacypatterns.org/patterns/Outsourcing-[with-consent]

2.19 Ambient Notice

Summary

Context

Modern sensor systems process massive quantities of data, performing complex and silent operations which users typically overlook. The tracking of user information is used to improve the quality of these services (or products), and typically users wish to benefit from this. This is particularly evident in consumption, location, and physical activity tracking. While these users do not want to be exposed to extensive and otherwise intimidating details, this processing must be done under the user's informed and explicit consent. Once the consent for it has been obtained, processing may occur regularly or in real-time. While users are to be informed of this, in-progress readings still happen in a manner which is streamlined and not inherently noticeable. Users are also capable of forgetting sometime after consent was given.

Problem

Users are frequently unaware of the sensors currently tracking them. It is important that they understand that their personal data is being further collected in order for their informed consent to remain valid. This should be unobtrusive, however, so as to avoid notification fatigue or desensitization.

A user may not realize that an application given permission to access [sensor data] is doing so continuously or repeatedly, or may not remember that explicit permissions given in the past allow a service to access data again later. In some cases, past explicit permission may not have been provided by the current user of the device (but instead by a spouse, parent or even an ex-spouse or stalker who temporarily had control of the device or the account). If notice is provided only at the time of consent, a user may inadvertently distribute personal information over a long period of time after having lost control of their device only momentarily.

Forces and Concerns

- Users are capable of forgetting or reconsidering their consent, affecting the legitimacy of any processing under it
- Users may overlook processing which is not made apparent to them, allowing sensors to record data they would not otherwise
- Controllers aim to ensure that consent is retained, they want to avoid collecting data against the user's wishes
- Controllers want to prevent users from inadvertently sharing personal data which they regret being processed
- A tray full of ambient notices may annoy or confuse users and inure them to ongoing practices. Take measures to avoid unnecessary notice. This must be balanced against the concerns of an attacker's opting the user in without their knowledge

Goal

G

Solution

Provide an unobtrusive but clearly visible notification while sensors are in use, without interrupting the flow of user activity. This notification should be interactive in order to prevent, delay, or further explain the data collection.

Implementation

The best place to provide transparency is the place where data is collected. The sensor is the first component noticed of a complex system, because the user is directly confronted with the sensor during collection. Hence, provided information in this place is easier to access by the user. Because of that, sensors should be equipped with a user interface for instant check of the collected data. Such an interface can consist of a simple display or message box. In more complex environments, optical codes and links can refer the user to more elaborated information sources.

Constraints and Consequences

C

Motivating Scenario

Scenario
- Location services icons: Mac OS X, Google Chrome, et al

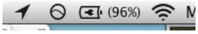

Figure 2.3: Fire Eagle location hierarchy

Mac OS X Lion adds an ambient location services icon (a compass arrow) which appears in the task bar momentarily when an application is accessing the device's location.

Figure 2.4: Fire Eagle location hierarchy

Chrome adds a cross-hair icon to the location bar when a web site accesses the device location via the W3C Geolocation API. Clicking on the icon provides potential actions: clearing the saved consent for this site and accessing settings.

Similar examples exist in at least Android, iOS and Windows.

Know Uses and Related Work

> QR-code based information access, smart meter display.

Categories

- Notice
- Mobile
- User-Interface
- Inform

Related Patterns

P Asynchronous Notice

P Informed Implicit Consent

P Single Point of Contact(SPoC)

P Preventing Mistakes or Reducing Their Impact

P Impactful Information and Feedback

Supporting Patterns

P

Sources

- https://privacypatterns.org/patterns/Ambient-notice
- http://privacypatterns.wu.ac.at:8080/catalog/

2.20 Dynamic Privacy Policy Display

Summary

Context

Controllers are mandated by various laws and regulations to ensure that users, their data subjects, are adequately informed before requesting consent. Failing this, the consent loses legitimacy and the controller may face repercussions. However, the main mechanism for supplying this information resides with privacy policies, which must also conform to legislative norms. The language this necessitates is long and complex, which jeopardizes the chances of users understanding it. This information can be summarized, and otherwise reworded to make the content more accessible to users, though typically the length of such summaries are still quite long.

Problem

Not all contexts are suitable for extensive privacy policy information, yet users often still need to be able to obtain additional data without breaking those contexts.

Forces and Concerns

- Controllers need informed consent for collection, sometimes with limited screen space available, yet users do not typically read privacy policies on their own
- Users do not want to spend time and effort reading through privacy policies
- Controllers want users to actually use their service (or product), but users will not do so if the cost of doing so entails disproportionate effort
- Users want to be able to get to using the service quickly, without needing to visit multiple policy pages

Goal

G

Solution

Provide the user with additional relevant policy information on hover or tap, by way of ’tooltips’, to best inform them given contextual limitations. In a mobile setting these tooltips may unobtrusively become available to tap when the relevant control is most in focus (i.e. selected, centered, or occupies most of the screen).

This information may highlight the nature and potential consequences of the disclosure, and should be displayed consistently.

Implementation

The information should be provided to the user where it is needed. Therefore the tooltip should appear on demand (i.e., need of information). This could be for example in a login dialog as soon as the user navigates the mouse into the concerning part of the interface. The tooltip should then be made visible to the user and contain all necessary information for making an informed decision.

Constraints and Consequences

C By displaying the relevant [information pertaining to] privacy policies whenever they apply to what the user is currently doing or about to do, the user’s awareness of what will happen with the information they’re about to share is increased.

However, users may also happen to not trigger the tooltip, especially when using a mobile device. As such it is important that they are aware of its existence, and its importance, given the current context.

Motivating Scenario

Scenario When a user needs to login and is given numerous options, with limited space provided, each option can have an assigned tooltip. These can appear on hover, tap, or scroll, where necessary appearing with less opacity until the user taps the tooltip itself. It can also lead to further detail through ’(see more)’ in a recognizable blue underlined hyperlink format. To encourage use of this a variant may either scroll through detail or show a visible scroll bar. Not needing the user to leave the application or webpage will require less effort on their part. ■

Know Uses and Related Work

The PrimeLife HCI Pattern Collection (V2) features a prototype using tooltips to convey policy information on hover.

Categories

- Inform
- Notify

Related Patterns

P Policy Matching Display

P Platform for Privacy Preferences

P Privacy Policy Display

Supporting Patterns

P

Sources

- https://privacypatterns.org/patterns/Dynamic-Privacy-Policy-Display
- http://privacypatterns.wu.ac.at:8080/catalog/

2.21 Privacy Labels

Summary

Context

Users use a variety of services (or products) for which there are different effects on their privacy. The providers of these services have varying policies around that usage, and thus affect privacy differently. Typically, the differences appear in a privacy policy document, or set of documents. Services encourage users to read this information, which can be quite extensive and involved. Users do not typically have the time or patience to investigate this information on their own.

Problem

Due to the effort required, users often do not investigate the various privacy policies of the services they use, leaving them uninformed about the potential consequences of their consent and choices. Services tend to have overly complex policies, and present them inconsistently, which agitates this issue.

Forces and Concerns

- Users want to know how much personal data they must share to use a service, without unnecessary or disproportionate effort
- Users want to quickly determine which services provide the functionality they seek with the privacy trade offs they can best accept
- Controllers want users to realize what data they use, and how they use it, so that they do not process it without informed consent
- Controllers also want users to understand the options they have in privacy preferences, and the advantages of opting into further sharing

Goal

G

Solution

Present the user with an standardized privacy ’nutritional’ label to quickly summarize policy information.

Structure

Putting a box around the label identifies the boundaries of the information, and, importantly, defines the areas that are “regulated” or should be trusted. This is a common issue when the label is placed in close proximity to other information, but may not be as significant an issue online.
Using bold rules to separate sets of information gives the reader an easy roadmap through the label and clearly designates sections that can be grouped by similarity.

Providing a clear and boldfaced title, e.g., Privacy Facts, communicates the content and purpose of the label specifically and assists in recognition.

Finally, we have defined a maximum width of 760px for this label and all following designs in this paper. One important consideration was that the privacy label design be printable to a single page and viewable in the standard width of today’s internet browsers.

Implementation

The tabular format can be filled in automatically if a site uses [Platform for Privacy Preferences].

Privacy Labels use four colored squares to help convey information quickly:

- Dark Red Square: we will collect and use your information in this way
- ’opt out’ Red Square: by default, we will collect and use your information in this way unless you tell us not to by opting out
- Light Blue Square: we will not collect and use your information in this way
- ’opt in’ Blue Square: by default, we will not collect and use your information in this way unless you allow us to by

opting in

In the short table variation, the label omits any rows (information types) which are entirely light blue (no collection or use). Instead this information gets summarized in text below the label using short natural-language format. Similar rows are merged into combined statements for brevity.

Constraints and Consequences

C The Privacy Label authors conducted a study where they assessed respondents' (n=764) attention to presented policies. They were able to determine how long respondents looked at each policy and where that affected their opt-out and further investigation decisions in the study. These were randomly divided between Privacy Labels (n=188), short table version (n=167), short text version (n=169), the full original policy document (n=162), and Layered Policy Design (n=78). Privacy Labels tested best among the respondents, followed by short table and text variations. Layered Policy Design was not found to perform any better than the full text when not additionally rephrasing policies.

Motivating Scenario

Scenario ■

Know Uses and Related Work

Privacy Labels are currently implemented using Privacy Bird and Privacy Finder Their source code is also available.

Categories

- Visualize
- User-Interface
- Inform
- Explain

Related Patterns

P Impactful Information and Feedback

P Layered Policy Design

P Privacy Aware Wording

P Privacy Color Coding

P Privacy-Aware Network Client

P Impactful Information and Feedback

P Informed Secure Passwords

P P3P (Platform for Privacy Preferences)

P Awareness Feed

P Trust Evaluation of Services Sides

Supporting Patterns

P

Sources

- https://privacypatterns.org/patterns/Privacy-Labels

2.22 Data Breach Notification Pattern

Summary

This pattern assures that a certain minimum data breach notification delay is not exceeded.

Context

Controllers of services (or products) provided to users collect mass amounts of data, a lot of it personal, to improve the quality and user experience of that service. This is all to be done under the informed consent of the user, who should properly understand the risks involved for their data. One such risk is that of unauthorized access, modification, removal, or sharing of data. If such a data breach occurs, notification is required. Any controller within (or providing services or products within) the EU must notify the supervisory authority of their main establishment or representative. This must occur within 72 hours unless justified. Notifying users is dependent on whether they are sufficiently affected.

This pattern is applicable in any environment where PII is stored and that allows monitoring of specific events.

Problem

When data breaches occur, numerous risks become apparent for multiple parties, these parties need to be notified and the risks need to be mitigated. Subsequent instances should be prevented through lessons learned.

Forces and Concerns

- Users want to know if anything has happened to compromise their data, their security, or their privacy
- Users want the controller to mitigate the risks before and after a breach to the best of their ability
- Controllers want to prevent risks from materializing and place measures against breaches happening in future
- Controllers also want to prevent users from suffering consequences from the breach, or from ignorance of the breach

In case a data breach has occurred, i.e. Personally Identifiable Information (PII) has leaked, the data owner must be notified. The notification process in turn may not work correctly, so it has to be monitored.

Goal

G The pattern goal is to constantly ensure a minimum delay of notification, should a data breach have occurred, and in case a notification exceeds the allowed delay, to indicate this by appropriate means.

Solution

Detect and react to data breaches quickly, notifying the supervisory authority of details, particularly risk mitigation, in order to establish whether users must also be informed. Properly handling these events will strengthen user trust rather than weaken it.

Implementation

A monitoring system logs access to [personal data] along with a time-stamp. A notification process continuously verifies that only authorized access is listed in this log file, and in case of unauthorized access notifies the data owner and logs the notification action in the log file, again accompanied by a time-stamp. A notification monitoring process finally continuously checks that $t_n - t_l <= max_n p$ (t_n denoting the time of notification, t_l the time of data leakage, $max_n p$ the maximally allowed period of notification). In case $t_n - t_l > max_n p$ it alerts the [associated] Incident Manager. A formal model can be found here.

In the event of a breach, the controller should first notify the supervisory authority within 72 hours of it's discovery, and no later without sufficient justification. The processor of personal data, where not also the controller, should notify the controller immediately.

Notification to the authority should include the nature and extent of the personal data affected, the contact for follow up, likely consequences, and the measures proposed or taken to mitigate the breach's effects. If absolutely necessary these details can be provided as they become available. Any breaches should also be documented for future review.

Where users are affected in a manner which risks their personal rights and freedoms, they shall also be informed of at least the contact, consequences, and measures to be taken, without undue delay. This is not the case if disproportionate effort would be needed, the data remains protected, or the risk is already sufficiently mitigated. The supervisory authority shall assist in determining whether informing users is necessary.

Note that associations or other representative bodies may prepare codes of conduct for data breach notifications. These notifications may also be affected by binding corporate rules, or guidelines, recommendations, and best practices from the board, to promote consistency.

A monitoring system logs access to clients' PII along with a timestamp. A notification process continuously verifies that only authorized access is listed in this log file, and in case of unauthorized access notifies the data owner and logs the notification action in the log file, again accompanied by a timestamp. A notification monitoring process finally continously checks that $t_n - t_l <= max_n p$ (t_n denoting the time of notification, t_l the time of data leakage, $max_n p$ the maximally allowed period of notification). In case $t_n - t_l > max_n p$ it alerts the PII Incident Manager.

Constraints and Consequences

In order to [detect the breach], the [controller] must have in place an access control mechanism and a monitoring mechanism [for personal data]. The pattern cannot ensure that [the relevant] Incident Manager will take adequate actions, hence this process has to be established and controlled by other means.

Motivating Scenario

Scenario Assume a [company] stores all employees' data [through a controller's service]. There is a contractual agreement between [them] that each data leakage is reported within one hour. Now Bob, an employee of [the controller] and not authorized to read [the company's] data, succeeds in circumventing [the] access control mechanisms and reads [personal] data. This represents a data breach of which [the company] has to be notified within an hour. ■

Know Uses and Related Work

This pattern is based on the privacy principle "Accountability" specified in ISO/IEC 29100 that is also used in Annex A of ISO/IEC 27018. More specifically, it addresses A.9.1 Notification of a data breach involving Personally Identifiable Information (PII). Uses of the pattern as a concrete instantiation of A.9.1 are not known.

Categories

- Inform
- Notify
- Control

Related Patterns

P Unusual Activities

Supporting Patterns

P

Sources

- `https://privacypatterns.org/patterns/Data-breach-notification-`
- `https://privacypatterns.eu/#/patterns/data-breach-notification`
  `0-0-0-0-0-0-0-5-1-0-0-1-0-0-0-0`

2.23 Pseudonymous Messaging

Summary

A messaging service is enhanced by using a trusted third party to exchange the identifiers of the communication partners by pseudonyms.

Context

This pattern can be used for online communications by email, through message boards, and newsgroups.

Problem

Messaging includes all forms of communication through emails, articles, message boards, newsgroups etc. This information could be stored and used to build sophisticated user profiles. Sometimes it can also be used to prosecute people.

Goal

G The goal of this pattern is to prevent unforeseen ramifications of the use of online messaging services.

Solution

A message is send by a user to the server, which exchanges the sender's address with a pseudonym. Replied messages are sent back to the pseudonymous address, which will then be swapped back to the original.

Constraints and Consequences

C

Motivating Scenario

Scenario Alice is a political activist and tries to organize a political demonstration. Since her government does not like free speech, her communication channels are intensely monitored

> and one day, she simply disappears into a labor camp and is never seen again. ■

Know Uses and Related Work

> Users can communicate more freely. Pseudonym servers can be misused to send offensive messages, for spam mails or by criminals for illegal activities. Under those circumstances it could be necessary to revoke the pseudonymity of the corresponding parties.

Categories

- Hide
- Categories
- Messaging
- Email
- Pseudonymity
- Dissociate

Related Patterns

pseudonymous-identity

masked-online-traffic

Supporting Patterns

strip-metadata

protection-against-tracking

Sources

- `https://privacypatterns.org/patterns/Pseudonymous-messaging`
- `https://privacypatterns.eu/#/patterns/pseudonymous-messaging/0-0-0-0-0-0-0-0-0-0-0-0-0-0-0-0`

2.24 Onion Routing

Summary

This pattern provides unlinkability between senders and receivers by encapsulating the data in different layers of encryption, limiting the knowledge of each node along the delivery path.

Context

A system in which data is routed between different nodes.

Problem

When delivering data, the receiver has to be known. If the system provides the functionality that the receiver of data should be able to answer, than the receiver should also know the address of the sender. When forwarding information over multiple stations then, in a naive implementation, each station on the delivery path knows the sender and the final destination.

Goal

G The goal of this pattern is to achieve unlinkability between senders and receivers.

Solution

The solution is to encrypt the data in layers such that every station on the way can remove one layer of encryption and thus get to know the immediate next station. This way, every party on the path from the sender to the receiver only gets to know the immediate successor and predecessor on the delivery path.

Constraints and Consequences

C If there are too few hops, the anonymity set is not big enough and the unlinkability between sender and receiver is at risk. The same problem occurs when there is too few communication going on in the network. The multiple layers of encryption will bloat up the data and consume bandwidth. If all nodes on the delivery path collaborate in deducing the sender and the receiver, the pattern becomes useless.

Motivating Scenario

Scenario Alice is a whistleblower and tries to forward data to Bob who works at the press. She sends the corresponding documents as an e-mail-attachment. Eve monitors the traffic and can see who sent this mail to whom. The next day, police raids Alices apartment and sends her to jail. Bobs mail account gets seized. ■

Know Uses and Related Work

The TOR-browser, a web-browser specifically designed to ensure anonymity makes heavy use of onion routing.

Categories

- Routing
- Anonymous-Communication
- Hide
- Mix

Related Patterns

P anonymous-reputation-based-blacklisting

Supporting Patterns

P strip-metadata

P protection-against-tracking

Sources

- https://privacypatterns.org/patterns/Onion-routing
- https://privacypatterns.eu/#/patterns/onion-routing/0-0-0-0-0-0-0-0-0-0-0-1-0-0-0
- http://privacypatterns.wu.ac.at:8080/catalog/

2.25 Strip Invisible Metadata

Summary

Strip potentially sensitive metadata that isn't directly visible to the end user.

Context

When a service requires a user to import data from external sources (eg. pictures, tweets, documents) different types of metadata may be transmitted. Users may not be aware of the metadata as it can be automatically generated or not directly visible. Services might be inadvertently responsible for exposing private metadata, or going against users' expectations.

Problem

Users are not always fully aware of the various kinds of metadata attached to files and web resources they share with online services. Much of this data is automatically generated, or not directly visible to users during their interactions. This can create situations where, even though users share information explicitly with services, they may be surprised to find this data being revealed. In certain cases where the data is legally protected, the service could be held responsible for any leakage of sensitive information.

How should services that need users to share data and upload files treat additional metadata attached with files? In case of uploading documents and images, which parts of the metadata can be treated as explicitly shared information.

Goal

G

Solution

Stripping all metadata that is not directly visible during upload time, or during the use of the service can help protect services from leaks and liabilities. Even in cases where the information is not legally protected, the service can protect themselves from surprising their users and thus alienating them.

Additionally when users share data with services, they can be presented with a preview of the data obtained by the service, including any metadata [[Preview Shared Data]]. This allows users to be more aware of information that they are sharing with the services, and in many cases with other entities on the Internet.

To summarize: user metadata that can not be made visible to users clearly should be stripped to avoid overstepping the users' expectations.

Constraints and Consequences

C

Motivating Scenario

Scenario

- Uploading images to twitter.com
  Twitter.com removes EXIF data from images uploaded to their image sharing service. Previously there have been many breaches of personal location by using EXIF data shared by image sharing services.

- Hiding EXIF data on Flickr.com
  In certain cases services might build features based on metadata, or the metadata sharing could be an important part of the community of users. Flickr.com allows users to hide their EXIF data from public display, and also provides an interface for users to easily see whether they are sharing location

Know Uses and Related Work

Categories

- Metadata
- Minimization
- ExiF
- Location
- Media
- Minimize
- Strip

Related Patterns

P

Supporting Patterns

P

Sources

- https://privacypatterns.org/patterns/Strip-invisible-metadata
- http://privacypatterns.wu.ac.at:8080/catalog/

2.26 Pseudonymous Identity

Summary

Hide the identity by using a pseudonym and ensure a pseudonymous identity that can not be linked with a real identity during online interactions.

Context

This pattern can be used for systems in which users are identified by public identities.

Problem

Many kinds of sensitive information are released through web interactions, email, data sharing or location-based systems, which can contain the name of a user or header information in packets. Another problem could be to interact anonymously in a forum. However too much interaction in a forum with an anonymous identity can be dangerous in the sense that the relation between original identity and a pseudonymous identity can be exposed.

Goal

Hide the identity of the participants.

Solution

> Initiate a random pseudonym, that can not be related to the original, so that the identity is hidden. Furthermore a pseudonym depends on concealment, so the pseudonym allocation needs protection.

Constraints and Consequences

The real identity of a user is hidden. In certain scenarios there is a need for additional space to store the pseudonym-identity mapping. Extensive Usage of the same pseudonym can weaken it.

Motivating Scenario

Scenario Assuming some students are writing an exam and they have to fill out a form about their identity, where there is an optional field for a chosen pseudonym. This way the result can be released under the chosen pseudonyms and the identity of each student is hidden. But by being observant, some students might be able to figure out which identity belongs to which pseudonym and so the confidentiality of the identity is compromised. ■

Know Uses and Related Work

Anonymizer are well-known tools for anonymous web interactions. They work for example by using a proxy between a request sender and a recipient to strip header information like $HTTP_{U}SER_{A}GENT$ in packet headers because they contain metadata about packet senders. The Mixmaster is an anonymous remailer that hides the sender and recipient identity by stripping its name and assigning a pseudonym. Some data sharing systems with a privacy-preserving focus make use of pseudonyms so that identifying information such as names and social security numbers are hidden. For example various electronic healthcare systems are using pseudonyms for the storage of e-health records.

Categories

- Anonymity
- Pseudonymity
- Hide
- Dissociate

Related Patterns

P pseudonymous-messaging

Supporting Patterns

P strip-metadata

P protection-against-tracking

Sources

- https://privacypatterns.org/patterns/Pseudonymous-identity
- https://privacypatterns.eu/#/patterns/anonymity-set/0-0-0-0-0-0-0-0-0-0-0-0-1-0-0-0

2.27 Personal Data Store

Summary

Subjects keep control on their personal data that are stored on a personal device.

Context

The pattern is applicable to any data produced by the data subject (or originally under his control) as opposed to data about him produced by third parties.

Problem

Data subjects actually lose control over their data when they are stored on a server operated by a third party.

Goal

G Enhance the control of the subjects on their personal data.

Solution

A solution consists in combining a central server and secure personal tokens. Personal tokens, which can take the form of USB keys, embed a database system, a local web server and a certificate for their authentication by the central server. Data subjects can decide on the status of their data and, depending on their level of sensitivity, choose to record them exclusively on their personal token or to have them replicated on the central server. Replication on the central server is useful to enhance sustainability and to allow designated third parties (e.g. health professionals) to get access to the data.

Constraints and Consequences

C Data subjects need to be equipped with a personal data store.

Motivating Scenario

> **Scenario** Patients want to keep control over their health data but also to grant specific access to some health professionals. ■

Know Uses and Related Work

> It has even been deployed for certain types of services, in particular, in the health sector.

Categories

- Control
- Separate
- isolate

Related Patterns

P user-data-confinement-pattern

P sticky-policies

Supporting Patterns

P

Sources

- `https://privacypatterns.org/patterns/Personal-data-store`
- `https://privacypatterns.eu/#/patterns/personal-data-store/0-0-0-0-0-0-0-0-0-0-0-0-0-0-0-0`

2.28 Trust Evaluation of Services Sides

Summary

Context

When using a service (or product) offered by a controller, the level of trust held by users is crucial. Without sufficient trust, the users would seek alternatives or generate bad publicity. They will use a system more cautiously, regardless of whether it is necessary. In many systems this lessens the quality of service offered, not only to the user in question, but holistically.

Problem

Users want to have reason to trust that a service does not undermine their personal privacy requirements. They do not want to have to take controllers, and third parties, at their word alone.

Forces and Concerns

- Controllers, as well as third parties, want to show that they are provably trustworthy and reliable
- Less confident entities will not make this effort alone
- Users want to verify claims which controllers and third parties make without having to do so themselves
- Users benefit from a standardised way of indicating trust, as it is easier and quicker to look into if done consistently and often

Goal

G

Solution

Supply a function which informs users of the trustworthiness and reliability of services, and that of the third parties connected to those services. These qualities may be determined, and assured, through independent evaluation of given criteria.

Structure

Information regarding a service's trustworthiness and reliability needs to be clearly indicated to the user prior to or during collection. It may therefore be brought up along with obtaining informed consent. This ensures that the user does not make misinformed or uninformed decisions, especially as this can seriously jeopardise trust.

A visual highlight which succinctly asserts this quality may also be displayed in persistent manner, or where otherwise contextually relevant.

Implementation

A trust evaluation function should be based on suitable parameters for measuring the trustworthiness of communication partners and for establishing reliable trust.

[Trust] in a service provider can be established by monitoring and enforcing institutions, such as data protection commissioners, consumer organisations and certification bodies. Privacy seals certified by data protection commissioners or independent certifiers (e.g., the EuroPrise seal, the TRUSTe seal or the ULD Gütesiegel) therefore provide especially suitable information for establishing user trust. Such static seals can be complemented by dynamic seals conveying assurance information about the current security state of the system and its implemented privacy and security (PrimeLife) functions. Further information sources by independent trustworthy monitoring organisations that can measure the trustworthiness of services sides can be blacklists maintained by consumer organisations or privacy alert lists provided by data protection commissioners.

[Also,] reputation metrics based on other users' [ratings] can influence user trust. Reputation systems, [for instance] in eBay, can however often be manipulated by reputation forging or poisoning. Besides, the calculated reputation values are often based on subjective ratings by non-experts, [through which privacy-friendliness may be difficult to judge].

A trust evaluation function should in particular follow the following design principles:

- Use a multi-layered structure for displaying evaluation results
- Make clear who is evaluated
- Inform the user without unnecessary warnings
- Use a selection of meaningful overall evaluation results

Constraints and Consequences

C Users will be able to better justify the trust they place in controllers who measure high levels of trustworthiness and reliability, and will know of greater risks in lower trust. A familiarity with the approach will also cause a healthy skepticism of controllers who do not participate, or have low confidence evaluations.

Motivating Scenario

Scenario Determine an appropriate metric for evaluating trustworthiness of partners of the service who will receive personal data as third parties. This can be simple, such as meeting expectations, failing them, or exceeding them. PrimeLife suggests 'poor', 'fair', and 'good', with fair evaluations having neither negative nor positive influences. Blacklists or alert lists make for a poor evaluation regardless of positive aspects.

These evaluations are shown to users prior to their related parties having consent for access. The notification is not shown too frequently, as extensive warnings may be misleading to users. While they should be aware of neutral or unevaluated parties, it may not be desired to worry them without cause. There should also be just enough information to raise awareness, allowing the user to investigate further if desired. A notification for a fair evaluation may be 'we have not found any issues with this partner' for example, with a neutral colour which matches the rest of the interface. Poor evaluations could be yellow or red (alarming colours), with good evaluations green or blue (positive colours).

■

Know Uses and Related Work

Categories

- Transparency
- Access
- Inform
- Provide

Related Patterns

P Policy Matching Display

P Awareness Feed

P Icons for Privacy Policies

P Privacy Color Coding

P Appropriate Privacy Icons

Supporting Patterns

P

Sources

- https://privacypatterns.org/patterns/Trust-Evaluation-of-Services-Sides

2.29 Aggregation Gateway

Summary

Encrypt, aggregate and decrypt at different places.

Context

A service provider gets continuous measurements of a service attribute linked to a set of individual service users..

Problem

> The provision of a service may require detailed measurements of a service attribute linked to a data subject to adapt the service operation at each moment according to the demand load. However, these measurements may reveal further information (e.g. personal habits, etc.) when repeated over time.

Goal

G Let the service provider have reliable access to the aggregated load at every moment, so as to fulfil its operating requirements, without letting it access the individual load required from each specific service user.

Solution

> A homomorphic encryption (e.g. Paillier) is applied at the metering system, using a secret shared with the service provider (generated by applying e.g. Shamir's Secret Sharing Scheme) The encrypted measurements from a group of users are transmitted to an independent yet trusted third party. This third-party cannot know about the content of each measurement (as it is encrypted), but it can still operate on that data in an encrypted form (as the encryption system is homomorphic). There are different trusted third parties for each group of users. In order to improve the privacy resilience, each user may belong to several groups at the same time.

The trusted third-party aggregates the measurements from all the users in the same group, without accessing the data in the clear at any time.

The service provider receives the encrypted, aggregated measurement and decrypts it with the shared secret.

A feeder metering system can be added as a measuring rod which introduces a comparison for each group of meters.

Constraints and Consequences

C There is a need to deploy trusted third parties that compute the aggregations over each group of users. Note that they need to be honest (i.e., they cannot collude with the other parties involved), but they need not respect confidentiality (as they only have access to encrypted contents). Smart meters are needed that have computation resources to apply secret generation and homomorphic encryption procedures (note that this is trivial when dealing with the use of computational resources, but it does not have to be always available in the case of e.g. smart grid systems). The potential range of measured values must be large enough to avoid brute force attacks. Robust homomorphic encryption schemes introduce a large computational load.

Motivating Scenario

Scenario An electric utility operates a smart grid network with smart meters that provide measurements of the instantaneous power consumption of each user. The utility employs that information to adapt the power distribution in a dynamic fashion, according to the user demand at each moment.

Know Uses and Related Work

- Lu, R., Liang, X., Li, X., Lin, X., & Shen, X. (2012). Eppa: An efficient and privacy-preserving aggregation scheme for secure smart grid communications.Parallel and Distributed Systems, IEEE Transactions on, 23(9), 1621-1631
- Rottondi, C., Verticale, G., & Capone, A. (2013). Privacy-

> preserving smart metering with multiple data consumers. Computer Networks, 57(7), 1699-1713
> - Kursawe, K., Danezis, G., & Kohlweiss, M. (2011, January). Privacy-friendly aggregation for the smart-grid. In Privacy Enhancing Technologies (pp. 175-191. Springer Berlin Heidelberg

Categories

- Aggregate
- Hide
- Restrict
- Separate

Related Patterns

- trustworthy-privacy-plug-in
- added-noise-measurement-obfuscation

Supporting Patterns

- user-data-confinement-pattern
- anonymity-set

Sources

- `https://privacypatterns.org/patterns/Aggregation-gateway`
- `https://privacypatterns.eu/#/patterns/aggregation-gateway`
  `0-0-0-1-0-1-0-1-1-0-0-1-1-1-0-0`

2.30 Privacy Icons

Summary

A privacy policy which is hard to understand by general audience is summarized and translated into commonly agreed visual icons. A privacy icon is worth a thousand-word policy.

Context

This pattern can be applied to any system which collects end user data. It can be presented in an interactive web page but also as part of a physical product which can collect data (e.g. fitness tracker).

Problem

Many organizations provide privacy policies which are too lengthy and hard to understand by the general audience. These policies are oriented as legal disclaimers for legal issues, rather than to inform end users so they can consent to the organization practices after being clearly informed of the collected data, its purpose, and the processing and potential sharing with third parties.

Goal

G Truly inform customers of the privacy policy of a system/organization.

Solution

Include within the service/device a very accessible and visual explanation of the privacy policy. Icons are a great complement to written text, as they may convey much information at a glance through a different modality (images). Standardized icon sets may thus be added to the privacy policy.

Constraints and Consequences

C Users may understand, at first glance, what are the potential risks of consenting of a privacy policy. In order to be useful, the icons must be well known and understood by the majority of the potential users before being used. A common meaning of the icon needs to be shared by the community. Educational material can be built upon the implications of each of these icons.

Motivating Scenario

Scenario Alice buys a fitness tracker and she is aware that the device collects her location, and sends it to a central web service in order to provide her with her fitness statistics (her fitness routes, the time spent...). The device provider aggregates this data and provides a business analytics service to third parties.

Alice is totally unaware of this secondary use of her data and may not agree to it. But accessing this policy involves accessing a website and going through a lengthy and legally oriented document.

Know Uses and Related Work

- The current version of the forthcoming EU Data Protection Regulation includes a set of privacy icons that should be used within European services and organizations
- https://disconnect.me/icons
- https://wiki.mozilla.org/Privacy_Icons
- http://yale.edu/self/psindex.html
- http://www.privacybird.org/
- https://netzpolitik.org/2007/iconset-fuer-datenschutzer
- http://knowprivacy.org/policies_methodology.html
- http://www.privicons.org/
- TheEU-fundedPrimeLifeprojectalsoproposedasetofprivacyic
  Holtz,L.E.,Zwingelberg,H.,&Hansen,M.(2011).Privacypolic
  //link.springer.com/chapter/10.1007%2F978-3-642-20317-6
  15)InPrivacyandIdentityManagementforLife(pp.279-285)
  .SpringerBerlinHeidelbergandHoltz,L.E.,Nocun,
  K.,&Hansen,M.(2011).Towardsdisplayingprivacyinformation
  InPrivacyandIdentityManagementforLife(pp.338-348)
  .SpringerBerlinHeidelberg.

> - • The Use of Privacy Icons and Standard Contract Terms for Generating Consumer Trust and Confidence in Digital Services CREATe Working Paper 2014/15 (October 2014)
>
>   Currently, most of these are only applied by client-side solutions. See also the Privacy Icons entry at Ideas for a Better Internet (kind of a pattern repository by the Berkman Center for Internet and Society in Harvard).

Categories

- Privacy-Policy
- Inform
- Explain

Related Patterns

P privacy-aware-network-client

P privacy-color-coding

P layered-policy-design

Supporting Patterns

P

Sources

- `https://privacypatterns.org/patterns/Privacy-icons`
- `https://privacypatterns.eu/#/patterns/privacy-icons/0-1-1-1-0-0-0-1-1-4-0-0-1-0-0-1`
- `http://privacypatterns.wu.ac.at:8080/catalog/`

2.31 Privacy-Aware Network Client

Summary

A privacy policy which is hard to understand is in an automated way converted into a more easy to read format.

Context

This pattern is limited to the web browsing domain.

Problem

Privacy policies are typically written to satisfy legal requirements ahead of conveying concise and easily understandable information to users. This makes users less informed overall.

Forces and Concerns

- Users do not want to read through long policy documents on each site they visit while browsing the Internet
- Users want to understand any notable risks or tradeoffs of using a site, preferably before being subject to them
- Controllers want to adhere to the many requirements of law, in a manner that is balanced and best reduces risks for them
- Controllers also want users to have a pleasant user experience without enduring shocking revelations or underhanded agendas

  Many websites have privacy policies which are hard to understand for the general audience. Many people enter websites with different intentions like shopping, research, etc. At the same time those responsible collect, use and release information about a user by explaining it through statements called privacy policies. These privacy policies are not easy to read and to understand.

Goal

G Expand the awareness of the user towards privacy policies of a website.

Solution

Provide a privacy preserving proxy which securely parses and interprets the privacy policies of controllers, supplying users with standardized and easily understood summaries of those policies.

Structure

Figure 1 [of the paper] shows a class diagram for the relationships between the user, the server, and the proxy. Each server can publish many policies and each user can be made aware of many policies at a time through the proxy. In Figure 2 [of the paper], a user wishes to access some information or interact with files on the server, which publishes its privacy Policy. The access occurs in the following sequence:

- The User interacts with the Server through a network Client
- The Client consults the Proxy for privacy policies
- The Proxy discovers the correct Policy (or Policies) made available by the Server, for the information or files in question
- The Proxy displays a user-friendly screen to the User requesting approval of the Policy, prior to allowing access to the information or permitting the interaction
- The User makes a decision after reviewing the Policy

Implementation

Design and implement a proxy able to parse and interpret privacy policies written in some standard language. This proxy could be built as a specialized version of the Proxy pattern [by the Gang of Four]. The proxy could be able to interpret several privacy languages or just one of them. Successful use of the pattern requires that the proxy can understand the server's privacy language.

Design and implement a secure communication channel between network clients and their proxies. This is necessary to avoid interception of the user choices by malicious [actors].

The problem can be solved by designing and implementing a privacy proxy that can parse and interpret policies. Afterwards it translates the policies into a human-readable format to present them in a user-friendly way.

Constraints and Consequences

C The user's awareness of the privacy policy rises so that more informed decisions can be made. The proxy is able to automatically detect changes of the privacy policy. A separate secure connection is needed for the proxy for every access to an area which is secured by a privacy policy. Policy constraints need to allocate local storage in the client. An attack on this could lead the user to decisions which he would otherwise not do. If there are any breaches of privacy it can be blamed on the client if he did use a privacy-aware client for a particular access.

Motivating Scenario

Scenario Alice uses several web services but is not aware of the their privacy policies. Even when she reads the policies, she is still not aware of the actual implications of the legal description. In the absence of other solutions, she does not read the policies and does not understand the ramifications. ■

Know Uses and Related Work

JRC P3P Proxy Version 2.0 is a P3P user agent acting like an intermediary. Depending on the specified privacy preferences of a user, it controls the access to web servers. Another known P3P user agent is AT&T Privacy Bird. Privacy Bird is a tool warning users if privacy policies of visited websites are not matching with their invidual privacy preferences.

Categories

- Notice
- Mobile
- User-Interface
- Inform
- Notify
- Privacy-Policy
- Proxy
- P3p
- Explain

Related Patterns

- policy-matching-display
- Appropriate Privacy Icons
- Icons for Privacy Policies
- Privacy Labels
- Privacy Color Coding
- Abridged Terms and Conditions
- Privacy Aware Wording
- Awareness Feed
- Impactful Information and Feedback
- Platform for Privacy Preferences
- Privacy-Aware Network Client

Supporting Patterns

- privacy-color-coding

Sources

- https://privacypatterns.org/patterns/Privacy-aware-networ
- https://privacypatterns.eu/#/patterns/privacy-aware-netwo
  0-0-0-0-0-0-0-2-0-0-0-0-1-0-0-0

2.32 Sign an Agreement to Solve Lack of Trust on the Use of Private Data Context

Summary

Context

Users do not inherently trust controllers who provide services (or products), as they do not have assurances as to what the controller's or their processor's true intentions are. Controllers and processors typically aim to make profit, but this might be at the expense of users if those users do not consider their privacy needs. The controller might have reasonable defaults or levels of control, but users also need to feel reassured that their choices are being honored. This is especially true of what they do or do not provide Lawful Consent for.

Problem

The controller does not necessarily have the trust of its users, and needs this trust for its services to process their data.

Forces and Concerns

- The controller wishes to provide services to the user, but needs their trust and consent to do so
- Processors want to manipulate data without having to worry about whether the data contains consented information or not
- Users want to use services, but not at the risk of their own personal privacy requirements being undermined
- Users want to know what they can safely provide to the controller and what information might be revealed about them if they use the service
- Users need to feel that the controller will honor any decision taken about their personal data

Goal

G

Solution

The service should provide the user with a contractual agreement (featuring privacy policy) which binds the controller to their word, provided that the user consents to the processing of data needed for specific purposes. The agreement should also bind any representative of the controller. It should be straightforward and clear enough for the user to comprehend.

Implementation

The service should feature a mechanism (e.g. landing page or unavoidable introduction) prior to collection, which stipulates the need for user consent. There should be a reasonable effort to prevent users from bypassing this mechanism.

The specific purposes for which their data will be processed should be made clear. The service should, at the same time, outline the contractual obligations it will be held against should the user consent. The user should be able to seek further detail about these obligations without first needing to consent.

If users decide to consent, they can make this clear by interacting with a mechanism (e.g. button) which clearly represents their agreement to the contract.

A further implementation could additionally allow the user access to a subset of the service which does not require any data, in order to help justify their consent. This would also alleviate the user's potential apprehension about the time taken to review and inform themselves about their decision.

Constraints and Consequences

C

Benefits

The controller, any of their representatives, and their users are tied to the terms of the contract and the legal implications it holds. Any disputes will involve both contract law and privacy law.

Liabilities

Users may be discouraged to use a service if they are made aware of the risks to their privacy, or introduced to the ways in which their data can be used to reveal information. They may also be tempted to consent without reading about the contract or how their data may be used. Therefore it is useful to not force an immediate decision, as this can invalidate the consent as not freely given or uninformed.

Motivating Scenario

Scenario ■

Know Uses and Related Work

Categories

- Control
- Consent

Related Patterns

P Sign an Agreement to Solve Lack of Trust on the Use of Private Data Context

P Lawful Consent

P Obtaining Explicit Consent

Supporting Patterns

P

Sources

- https://privacypatterns.org/patterns/Sign-an-Agreement-to-Solve-Lack-of-Trust-on-the-Use-of-Private-Data-Context

2.33 Single Point of Contact

Summary

Context

Many controllers make use of a storage platform (i.e. ’cloud’ facilities), such as e-Health services that keep their sensitive patient data in a distributed online storage. The sensitivity of this information raises concern and garners a need for special care. The storage medium in this case rules out typical security approaches.

Problem

Effective distributed storage services require specialized privacy management. The deficiencies of traditional means may be expressed through the following:

- traditional security mechanisms are platform dependent;
- typically they are difficult to federate or distribute;
- compliance with protocol can be cumbersome; and
- as such they are often inflexible.

Forces and Concerns

- Controllers wish to protect the sensitive or otherwise personal data they are charged with
- They want to acquire genuine Lawful Consent in a streamlined fashion
- They need this process to be facilitated, supervised, and provably sound

Goal

G

Solution

Single Point of Contact adopts a claim-based approach for both authentication and authorization similar to a super-peer design, also acting as a (Resource) Security Token Service, an Identity and Attribute Provider, and a Relying Party. It features a tried and proven expressive e-consent language, and can communicate with other SPoCs in a Circle of Trust

Rationale

Overcoming the inflexibility of traditional security mechanisms is partly solved by claim-based identity, which provides a platform-independent way of presenting identity information.

Structure

A SPoC is essentially a security authority, which protects patients' privacy in e-Health applications by providing a claim-based authentication and authorisation functionality (Baier et al. 2010), and facilitating secure communication between an e-Health service and its clients.

SPoC shares characteristics with a Central Medical Registry (CMReg), which performs authentication and manages identifying access to anonymised medical documents in a central repository. SPoC additionally facilitates secure e-Health service development and integration. It is able to share Electronic Health Records (EHRs) through a peer-to-peer network as an overarching, claim-based, super-peer-like representative of the e-Health community. Multiple SPoCs may also communicate, constituting a Circle of Trust.

See Fan et al. (2012) Figure 1 for a visual depiction."
The SPoC features a Domain Ontology for providing vocabulary towards claims and policies, a Policy Engine for consent syntax using natural language and pseudonym storage, and an Interface Service. The interfaces provided include Authentication, Authorisation, and Pseudonym Resolution.

Implementation

A SPoC is able to issue security tokens as a Security Token Service (STS), authenticate local domain users as an Identity Provider, certify attributes as an Attribute Provider, and accept external claims as a Relying Party. When in a Circle of Trust, the SPoC can also translate the claims of other SPoCs as a Resource

STS.

SPoCs' implementation of e-consent features the following levels, based on Coiera et al. (2004):

- general consent [with or without specific exclusions];
- general denial [with or without specific consents];
- service authorisation;
- service subscription; and
- investigation.

As with Pruski's (2010) e-CRL, SPoCs' e-consent also considers specific grantees, operations, purposes and period of validity.

For more information see Fan et al. (2012).

Constraints and Consequences

C The SPoC ensures that the privacy of sensitive medical data is protected, and that it is distributed securely and only to the people who are allowed to access the data. However, it requires a reliable credential-based authentication system to be able to validate requests.

Motivating Scenario

Scenario ■

Know Uses and Related Work

Categories

- Control
- Choose

Related Patterns

P Lawful Consent

Supporting Patterns

P

Sources

- https://privacypatterns.org/patterns/Single-Point-of-Contact

2.34 Informed Implicit Consent

Summary

Context

Processing of user (data subject) information, particularly that which potentially identifies a user or group, requires their explicit informed consent. Inaction is not considered valid consent. However, not all instances make this feasible. As such there are circumstances in which legitimate interests of the controller may justify collection without first obtaining a clear statement of permission to do so. Security footage around a controller's premises, or fraud detection, for example, cannot reasonably be made optional to users of the service (or product). What constitutes legitimate interests in these contexts depends on the relationship and reasonable expectations between the controller and user. As such, sensitive data, or special categories of data, are more difficult to justify.

Problem

A controller needs to collect and otherwise process reasonable information to fulfill their legitimate interests regarding a user, but cannot feasibly acquire each user's explicit consent.

Forces and Concerns

- Users should not have to frequently and explicitly consent for regular, everyday, ubiquitous services which are expected and acceptable for legitimate interests
- Users do not want to have certain data processed, and need a way to avoid implicitly consenting to it
- Controllers do not want to have to obtain explicit consent in real-time bulk for expected and acceptable legitimate interests
- Controllers want to ensure that legitimate consent exists before processing

Goal

G

Solution

Provide clear and concise notice that by using the service, the user implicitly consents to the processing necessary to fulfill legitimate interests. Ensure that this notice is perceived prior to the application of the effects it describes.

Implementation

Ensure that users are informed sufficiently prior to any processing with clear and concise notice, the complete detail of which should also be accessible. In digital mediums, this is straightforward, working similarly to Cookie Walls on websites. Users should be given the opportunity to choose not to use the service and therefore not be subject to the processing it requires.

In physical instances it is more difficult to be sure that users take note of this. On devices, lights have often been used to convey a recording state. This, while clear once already subject to processing, is not sufficient however. Instead, large signs are commonplace to indicate the use of data collection. The most familiar example would be "Smile, you're on camera". Of course, this is less clear than "Our premises is recorded for security purposes, by entering you consent to this processing. See more info at [address]". These signs should be posted, visible prior to recording, at all entrances or otherwise where applicable.

Constraints and Consequences

C Users will be informed before implicitly providing consent to reasonable processing for legitimate interests of the controller.

Motivating Scenario

Scenario Given a Sensor Network, Provider, and Controller, collected data is delegated by the Controller through the Provider to the Sensor Network. The Sensor Network collects some data with explicit consent, but this data may also be personal for a user who has not given such consent. This data may be potentially identifying, and thus the user should be informed prior to its processing. The Controller must ensure that the Provider of the Sensor Network provides any potential users with unambiguous warning of the collection, and that individual consent is infeasible. This may make use of a clear and legible warning sign. The Sensor Network itself should also be visible and obvious, clearly indicating when collection is ongoing. Societal norms may dictate this, such as security cameras in some contexts (commercial areas where valuables may be stolen) needing little warning. ■

Know Uses and Related Work

Categories

- Inform
- Notify

Related Patterns

P Ambient Notice

P Asynchronous Notice

P Preventing Mistakes or Reducing Their Impact

Supporting Patterns

P

Sources

- https://privacypatterns.org/patterns/Informed-Implicit-Consent

2.35 Enable/ Disable Functions

Summary

Context

Users frequently have data collected about them, often in situations where it needn't be. Many of these cases are due to good intentioned, expansive, functionality. Not all users seek to take advantage of all functions, however. Some controllers aim to consider this in their designs.

Problem

Not all users desire or benefit from all functionality.

Consider users living in an Ambient Assisted Living environment: these users are surrounded by various sensors such as video cameras, motion sensors or electrical current sensors that are used to monitor the actual situation of a person. Another example are the acceleration sensors included in smartphones. A [service (or product)] can recommend places of interest to the user by considering the gathered [data]. With regard to these examples it becomes obvious that [services] often unobtrusively collect highly critical and personal context data of users.

Forces and Concerns

- Informational self-determination: The pattern considers a user's basic right of informational self-determination. This is due to the fact that a user is able to explicitly agree or disagree to a certain function depending on the context data needed by the function. Therefore, the user has direct control of the context data collection process. This satisfies the principles of necessity, transparency, giving consent and responsibility. They are part of the user's right of informational self-determination and are described in detail in [Kuner] and [Hornung & Schnabel].
- Trust: The pattern increases a user's trust in the [service] by offering the possibility to prevent the collection and inference of certain personal context data. Hence, [users]

can be sure that personal data that is critical to [them] is not gathered, stored or further processed by third parties.

- Transparency: The pattern provides transparency to the user by giving an overview, which function needs which personal context data of a user to work properly. For this reason a user is aware of the context data that is gathered

Goal

G

Solution

Enable users to choose which functions they do not consent to using, nor wish to provide the required data for.

Implementation

A solution is given if the user can explicitly agree or disagree to certain functions. For this purpose, the [service] has to display every function and its required context data. A possible way of displaying these functions and the used context data may be the use of the privacy consent form, which is included in every application.

Constraints and Consequences

C By enabling the user to explicitly agree or disagree to certain functions, a context aware application like Support-U might not be able to provide all of its possible functionalities to the user anymore. However, the usage of this pattern in the development process of context-aware applications might additionally strengthen the user's confidence in the usage of UC systems.

Motivating Scenario

Scenario Support-U

In the shown privacy consent form each function, which utilises personal context information, is listed. Furthermore, the user is able to activate or to deactivate the functions, e.g., to enable a live stream or to enable predicting her next context.

Meet-U

Meet-U provides several functions that make use of localization mechanisms and the personal data the user supplies. That includes the user's interests, buddy list and [their] preferred means of transportation. For indoor navigation a RFID sensor attached to the user is exploited. The user can now switch off the navigation function so that neither the indoor nor the outdoor localization continue to operate. The user's preferences concerning transportation will be no longer available. Further functions can be disabled correspondingly. Turning off, for example, the advanced search engine would stop using the user's interests.

■

Know Uses and Related Work

- Support-U: An example of an abridged TAC is given in fig. 3. The figure shows the results of the abridged TAC pattern used for the Support-U application
- Connect-U: The user has to sign a license agreement of the size of one page in A4 format. On this page the agreement about the data usage is described in clear detail
- Meet-U: The key points of TAC that affect the user's privacy the most, are displayed on one screen. Hence, the gathering and processing of data are addressed and summarized briefly. The long version of the TAC is linked. The user has to agree on that before continuing with the application

Categories

- Control
- Update

Related Patterns

P Negotiation of Privacy Policy

P Lawful Consent

Supporting Patterns

P

Sources

- https://privacypatterns.org/patterns/Enable-Disable-Functions

2.36 Privacy Colour Coding

Summary

In a social networking site a user gets direct visual cues which privacy settings apply on which shared elements.

Context

The numerous policies and settings around privacy for each service (or product) used by a user would be quite complex and time consuming if such a user endeavored to investigate them. Policies are written for legal compliance and settings are often configured for best experience rather than privacy. Even in the instances where privacy friendly defaults are used, they may cripple the usability of the system, or otherwise disable desirable features. Some settings can also be difficult to consider due to overly brief and vague descriptions.

The pattern can be used in applications where users share and publish personal data and contents, but can control their visibility using privacy settings. This includes but is not limited to social networking sites.

Problem

Users do not investigate policies and preferences due to the effort required, and cannot inherently comprehend the consequences of settings otherwise. The poor understanding of these can lead to undesirable disclosures.

Forces and Concerns

- Users want to be able to quickly investigate how much or little information they can comfortably provide while still enjoying the service
- Users want to be guided as to what preferences achieve better privacy
- Controllers want users to configure preferences in ways they actually intend, therefore not processing data without

informed consent

- Controllers also want users to understand the limits of the settings through understanding the policies

Privacy settings and the actual effect of these settings on shared content and data is often not obvious for the user. Not having the active settings constantly in mind might lead to nonoptimal privacy experiences when the perceived privacy settings differ from the actual settings.

Goal

Users receive direct visual cues on the consequences of their privacy settings currently in effect. In order to be more clear about their privacy settings.

Solution

Present the user with standardized color visual cues to help guide them in selecting privacy friendly settings, and in understanding the policies around those settings.

Implementation

The results of privacy settings such as visibility are divided into different levels. A distinct color is assigned to each of these levels. Every time the user is performing an action where privacy settings come into play, the color is used as an indication of the privacy settings currently in effect. The choice of colors should take into account prevalent color meanings, like usage of the color red for warning situations. If privacy settings cannot be grouped into distinct levels, a gradient between different colors could also be used.

same treatment may be applied to policies, or explanations of settings. User rights and affordances may be presented differently from what the controller may do with their data. Aspects which could be perceived to have the greatest impact on privacy should stand out most. Explanations of who has responsibility or accountability, contact details, etc. can also be given a distinct color. Finally purposes and means for processing should be clearly visible.

The results of privacy settings such as visibility are divided into different levels. A distinct color is assigned to each of these levels. Every time the user is performing an action where privacy settings come into play, the color is used as an indication of the privacy settings currently in effect. The choice of colors should take into account prevalent color meanings, like usage of the color red for warning situations. If privacy settings cannot be grouped into distinct levels, a gradient between different colors could also be used.

Constraints and Consequences

C Users will directly see the outcome of their privacy settings. The danger of unwanted actions is decreased, as users will permanently receive visual cues. On the other hand a reduction of complex settings to a few colors may lead to an oversimplification which would render the whole pattern useless. Visual cues must be integrated into the site design but must still be placed prominently enough to be noticeable. Cultural aspects for the different meanings of colors should be taken into account. The same color may not be recognized as a warning label in different cultures.

Motivating Scenario

Scenario Alice uses a social network and shares personal stories only with her friends while she shares mundane content publicly. Hence she always has to change the privacy settings of her posts in order to adjust the visibility of the posts. One day she forgets to change the setting and does not realize that she actually shared a precarious story with her boss. ■

Know Uses and Related Work

A color coding similar to traffic lights is implemented in many modern web browsers for HTTPS connections. A green background indicates a valid certificate while a red background and a warning label shows that there are problems when validating a certificate. Facebook Privacy Watcher [http://www.daniel-

> puscher.de/fpw/] enhances the Facebook website by color-coding shared content and indicating its visibility. Posts with green background are public, yellow indicates visibility for friends only and red content is only visible to the user. Blue background is used for custom audiences such as groups.

Categories

- Distraction
- Visualize
- User-Interface
- Inform
- Explain
- Control

Related Patterns

Ⓟ privacy-aware-network-client

Ⓟ privacy-icons

Ⓟ layered-policy-design

Ⓟ Impactful Information and Feedback

Ⓟ Informed Secure Passwords

Ⓟ Privacy Aware Wording

Ⓟ Awareness Feed

Ⓟ Icons for Privacy Policies

- Trust Evaluation of Services Sides

Supporting Patterns

- privacy-aware-network-client

Sources

- `https://privacypatterns.org/patterns/Privacy-color-coding`
- `https://privacypatterns.eu/#/patterns/privacy-color-coding/0-0-0-0-0-0-0-0-0-0-0-0-1-0-0-0`
- `http://privacypatterns.wu.ac.at:8080/catalog/`

2.37 Appropriate Privacy Icons

Summary

Context

Controllers offering services (or products) to users have various policies regarding privacy. These typically exist within one document catering to legal evaluation, and thus one which is quite long and complex. Users are often encouraged to read such a policy, though as users are exposed to many of these, they mostly do not. As a countermeasure to this, controllers partition their policies, provide simplified versions, or bring relevant aspects to user attention when needed. One method of simplification is the use of privacy icons. This approach has its own issues for controllers to consider.

Problem

Privacy icons are easily misunderstood, as they are oversimplified concepts using imagery shared with numerous other concepts. Even when fully grasped, important information may be overlooked when finer details play a role.

Forces and Concerns

- Users do not want to regularly read long and complex policies
- Users want to understand what risks their data undergoes by using certain features of the service
- Controllers want users to actually take note of the relevant policies rather than process their data without informed consent
- Controllers want to save space so that they can have more appealing interfaces

Goal

G

Solution

Introduce the user to a consistent set of icons, carefully grouped and not excessive, and explain their meaning. Explanations should be short and concise, and these paired with the icons should be put through user tests. Users should be able to understand the icons when shown them in context.

While these icons should be able to stand alone, it is still important that a user has access to clarification. As such provide a mechanism, such as an on hover tooltip, which further explains what the icon attempts to convey. The icon should also be machine readable.

Implementation

When selecting appropriate icons for conveying information, take the following into account:

- primarily prevent misunderstanding,
- use icons users are familiar with,
- do not reassign meaning to familiar icons, and
- keep icon style and design consistent

Perform tests with actual users to determine whether there is any room for misunderstanding and adjust accordingly with further tests. If a concept cannot be reliably conveyed through an icon, then it must not be primarily provided as one.

Regardless of whether an icon perfectly conveys a policy, always allow users to investigate further. This can be achieved through hover, click or tap mechanisms. A tooltip, for example can provide a short explanation, but the full policy being depicted should also be available. As such, a context menu may also be appropriate, especially on single tap for mobile users.

Constraints and Consequences

C Informed users are able to make informed decisions which lead to a more responsible handling of private information. Since icons are an integral part of any kind of [interface], it is important that they convey the right information. Furthermore, users are only able to use [a service] to its full [extent] when they trust it. This effort towards transparency will assist in creating that trust.

Motivating Scenario

Scenario ■

Know Uses and Related Work

Categories

- Privacy-Policy
- Inform
- Explain

Related Patterns

- P Impactful Information and Feedback
- P Informed Secure Passwords
- P Layered Policy Design
- P Privacy Aware Wording
- P Privacy-Aware Network Client
- P Awareness Feed
- P Trust Evaluation of Services Sides
- P Icons for Privacy Policies

Supporting Patterns

P

Sources

- https://privacypatterns.org/patterns/Appropriate-Privacy-Icons

2.38 User Data Confinement Pattern

Summary

Avoid the central collection of personal data by shifting some amount of the processing of personal data to the user-trusted environments (e.g. their own devices). Allow users to control the exact data that shares with service providers

Context

This pattern may be used whenever the collection of personal data with one specific and legitimate purpose still pose a relevant level of threat to the users' privacy

Problem

> The engineering process is biased to develop system-centric architectures where the data is collected and processed in single central entities, forcing users to trust them and share potentially sensible personal data.

Goal

G Avoid the need for trust in service providers and the collection of personal data.

Solution

> The solution is to shift the trust relationship, meaning that instead of having the customer trust the service provide to protect its personal data, the service provider now has to trust the customers' processing.
>
> In the smart meter example, the smart meter would receive the monthly tariff and calculate the customer's bill which will be then sent to the energy provider where it will be processed. The main benefit is that at no moment the personal data has left the users trusted environment.

Constraints and Consequences

C Depending on the type of processing (e.g calculate the bill for the monthly energy consumption or the age from the birth date) the service provider will require some guarantees from the processor (the end user). This may involve the usage of Trusted Platform Modules or cryptographic algorithms (e.g. ABC4Trust).

Motivating Scenario

Scenario The smart grid is a domain with a clear example: having smart meters delivering hourly customers' energy consumption to the energy provider poses a serious threat to the customers' privacy. If the only purpose of collecting these data is to bill the customer, why cannot this calculation be done by the customer based on pre-established tariffs?

Similar examples in other domains are "pay as your drive" insurance policies where the insurance price is calculated based on the drivers behavior or electronic toll pricing.

Know Uses and Related Work

Smart meter, Privacy-enhanced attribute based credentials, pay as your drive insurances, electronic toll pricing.

Categories

- Data-Minimization
- Separate
- Isolate

Related Patterns

P

Supporting Patterns

P

Sources

- https://privacypatterns.org/patterns/User-data-confinemen
- https://privacypatterns.eu/#/patterns/user-data-confineme
  0-0-2-0-0-0-0-0-1-0-0-0-0-0-0-1

2.39 Icons for Privacy Policies

Summary

Context

Services (and products) which users use usually handle user data in ways which justify the use of a privacy policy. These documents are however made for legal purposes first and foremost and thus must cover a lot of detail rigorously. Users however are exposed to many of these documents when they seek to delve into the practices of each of the services they use. Controllers of these services realize the difficulty apparent in understanding full policy documents, but need users to understand risks if their processing consent is to be valid. Some approaches used to simplify policy are the layering of detail levels, general summarization, and contextual explanations. However, even these are subject to shortcomings.

Problem

Users struggle to understand privacy policies, even when reduced to a reasonable length. This discourages them from putting in the effort required to understand risks to their data, and invalidates consent.

Forces and Concerns

- Users want to understand the risks to their data in using a service, but do not want to read long or overly complex policies
- Many users want to be able to decide for themselves which policies apply to them, but do not want to read complex or time consuming summaries just to identify them
- Controllers need users to be informed before their data may be processed
- Controllers do not want to inconvenience users by making them read the privacy policy document intended for lawyers

Goal

G

Solution

Use privacy icons to aid in describing, grouping, and distinguishing the various policies in a privacy policy document. The icons should not allow for misinterpretation, which shall require user testing. Using consistent icons in a standardized way will promote understandability.

Implementation

In this pattern, privacy icons should not be used in place of the full policy document, but used to augment it. They should be shown in a manner which explains the policy explanation, excerpt, summary, or full detail as appropriate. Examples of usage include describing kinds of data processed, means, purposes, and legitimate interests or other justifications.

This usage should aid users in determining not only whether to further explore a policy, but also a rough idea of what each policy entails. More than one icon may be used per policy to achieve this, so long as the content becomes less complex.

Icons should be consistent and yet distinguishable from one another. This may justify a certain size limit. The icons should also be self-explanatory. These aspects need verification from a representative sample of the user population.

Furthermore, using standardized icons aids in both understanding and in promoting further use, but should not conflict with the norm. Doing otherwise may confuse users. If Icons are used in the same way on many of the applications or websites the user visits, it will be easy for the user to learn their purpose and to accept them as assistance. When users are aware of the icons from other purposes it will be also become more easy for them to create a mental model which supports them when reading a policy.

Constraints and Consequences

C Without dedicating too much effort, a user may quickly determine the potential risks of processing under a given policy. The user will be able to also quickly locate the other relevant policies both when first using a service and when revisiting policy.

C When the icons are sufficiently standardized, or at least for the subset which are, the user will not first need to familiarize themselves with explanations. Where not the case, education can assist in changing this if the icons are indeed widely used and consistent.

C The use of this measure will make policy more transparent, which will enhance the level of trust placed by users. Users which provided an invalid form of consent due to lack of policy understanding may then choose to retract it, or modify permitted usage.

Motivating Scenario

Scenario Alice buys a fitness tracker and she is aware that the device collects her location, and sends it to a central web service in order to provide her with her fitness statistics (her fitness routes, the time spent...). [She immediately consents to this even though it asks to first read a privacy policy.] The device controller [consequently] aggregates this data and provides a business analytics service to third parties.

Alice is totally unaware of this secondary use of her data and may not agree to it. But accessing this policy involves accessing a website and going through a lengthy and legally oriented document.

Comparatively, the tracker could have provided a short policy summary on the packaging using icons to convey more information with less space. Alice would have noticed an icon she recognized to convey third party sharing. Curious of whom this third party might be, and what extra risks she might be taking, she searches the online policy and finds it to be a company she does not trust. As a result she would not have consented, and potentially not purchased the device.

See also the Privacy Icons entry at Ideas for a Better Internet (kind of a pattern repository by the Berkman [Klein] Center for Internet and Society in Harvard).

Know Uses and Related Work

Categories

- Inform
- Explain

Related Patterns

- Impactful Information and Feedback
- Informed Secure Passwords
- Layered Policy Design
- Privacy Aware Wording
- Privacy-Aware Network Client
- Awareness Feed
- Privacy Color Coding
- Trust Evaluation of Services Sides
- Appropriate Privacy Icons

Supporting Patterns

P

Sources

- https://privacypatterns.org/patterns/Icons-for-Privacy-Policies
- http://privacypatterns.wu.ac.at:8080/catalog/

2.40 Obtaining Explicit Consent

Summary

Context

In order to offer (or products) to users (data subjects), controllers often need to collect (process) user data. Sometimes this is sensitive, identifying, or just metadata or other information which may be correlated to become more invasive. This nonetheless enables them to offer competitive features and functionality.

However, controllers are required to obtain unambiguous consent from their users in order to process their personal data in any way. Depending on the legal jurisdiction, there are additional considerations to take into account depending on the type of data in question. Typically, sensitive data requires especially rigorous care.

Problem

Controllers which aim to make use of user data, especially that which can be used to identify the user or sensitive aspects about the user, may not do so without a legally binding and sound acquisition of the user's consent.

Forces and Concerns

- Users want to use services without having to invest an inordinate amount of effort into discovering privacy risks
- Controllers need to be sure that users do not consent out of impatience or intimidation
- Users do not want to consent many times to the same service under the same privacy policy for each and every purpose
- Controllers need to be able to prove that users consented

Goal

G

Solution

Provide a clear and concise notification of all pertinent information the service could derive provided it had all the data it asks for. Indicate what this means for features and functionality. Then ask the user whether this tradeoff is something they consent to. If true, digitally signify and timestamp their response, or use Contractual Consent.

Implementation

The controller must ensure each user's sufficient understanding of the potential consequences. Otherwise the consent might not be informed. They must verify their users' willingness despite those consequences to provide their data for the specific purposes they need. If they do not, the consent might not be freely given.

Ensuring that users do not consent based on time constraints, or the intimidation of the information provided, may require testing with a sample. If the sample is representative, it will give the controller a defense against any claims of coercion.

The mechanism used for users to signify their consent should be clear. For example, if it is a button, it could read "I consent."

Constraints and Consequences

Benefits

Controllers can derive clearer potential consequences when the data collected is the same for every consenting user. Users therefore can look over these risks and spend less time making a valid decision. This reduces the chances of users consenting without informing themselves due to the difficult or verbose content presented.

Liabilities

Users do not however want to consent to all purposes necessarily since they might not all be compatible with what they feel comfortable sharing. In these cases users can be presented with a type of Selective Disclosure.

Motivating Scenario

Scenario

Know Uses and Related Work

Categories

- Control
- Consent

Related Patterns

- Lawful Consent
- Obtaining Explicit Consent
- Sign an Agreement to Solve Lack of Trust on the Use of Private Data Context
- Obtaining Explicit Consent
- Informed Consent for Web-based Transactions

Supporting Patterns

Sources

- https://privacypatterns.org/patterns/Obtaining-Explicit-Consent

2.41 Privacy Mirrors

Summary

Context

Controllers which provide services (or products) to users have various policies, Controllers process a lot of personal data within the services (or products) which users use. These users should however be made to understand the risks involved in all the processing. Typically, users need to be encouraged to review what data a service uses, and whether they consent to this. When provided with lock icons for certificates, or privacy coordination through color, users often still overlook warnings. Users want information to be streamlined, quick, and easy to digest in order to benefit from a service without delay.

This pattern is focused on the socio-technical domain, as opposed to purely technical, and as such considers a number of factors that do not play into a developer's perspective.

Problem

Users are frequently unaware of the personal data which a system processes and may use to draw conclusions from. Due to this, they either accept their data's undefined usage, or limit their disclosure, potentially more than needed, which could result in a poorer user experience.

Forces and Concerns

- Controllers want users to take note of important notifications, particularly if it prevents users from over or under sharing
- Many users are complacent with not knowing about what a service does with their data, maintaining a disinterest in detail
- Users may experience notification fatigue, if they are frequently provided with warnings and access notifications
- Users perceive information and notification appropriateness differently, depending on various contexts

- An approach to this transparency which is both noticeable and yet unobtrusive is needed. One which is passively assertive. A purely technical solution is not appropriate, it must also consider physical and social contexts
- This is because personal information has varying levels of sensitivity to users depending upon these contexts (for e.g. a closed room, or a close friend, are contexts in which some personal information may be considered less vulnerable).

Goal

G

Solution

Provide a framework for socio-technical systems which allow users to consider their privacy in context, and make decisions to cater for their personal needs.

This pattern encourages methods, mechanisms, and interfaces which reflect the history, flow, state, and nature of processed personal data which may otherwise have been hidden.

Structure

Privacy Mirrors focus on 5 characteristics:

- history, of data flows;
- feedback, regarding the state of their physical, social, and technical environment;
- awareness, enabled by the feedback;
- accountability, from these; and
- change, enacted by the users.

History

Logging is possible in technical systems from as little to as much granularity as desired. Whatever is kept needs to have been done so bearing in mind the social implications, i.e. the contexts, which may be important to a user's privacy. Who was involved in the processing, where was it processed, when, how, why, etc. are relevant.

The past must be summarized in a way which is easy to understand, but still detailed enough to identify less obvious risks. What information is relevant to different users as opposed to that which is seen as unnecessary noise? What isn't socially acceptable to record? How long must this be kept, should it deteriorate or simply vanish?

Feedback Logging is of little (or detrimental) use if not transparent and appropriately accessible. There needs to be a way to disseminate this history, state, and flow information to the users without inducing notification fatigue and without exposing information which is not contextually acceptable. Visual cues may be less distracting than the use of other senses, and capable of conveying much more information. However, some contexts may call for more distracting notification. A user should be able to choose whether, how much, and by what means they are notified - as some people have higher tolerances, or different tolerances, for distraction than others.

A distinction is suggested between notifications which require 'glancing', 'looking', or 'interacting'. Examples of these in an Android system are toast notifications (ambient display), heads up notifications (in the status bar), and pop up notifications, respectively. Ideally, each level will be available to cater to a user's personal preference. Information about a certain context should by default be found in the location where users would naturally look for it.

In feedback, how should different senses be addressed, to what level, where should it be shown, for how long, etc. These are aspects which should ideally be user configurable, with reasonable defaults. This should cater to what each user determines is important.

Awareness This concept includes the user's knowledge about how they feature in the system, how others feature with regards to the user's personal data, as well as what capabilities and constraints entities are given. The level of information and notification to convey depends on the user, as some will want more detail than others - meeting this balance will make the user more comfortable with their involvement.

This awareness can be divided amongst the three domains:

- Social: Notable usage patterns on access, being able to correlate this with others and encourage better decisions
- Technical: Understanding the limitations of the system, and the capabilities if used correctly, to use the system more effectively. Users should understand the flow, state, and history of their personal data in the system
- Physical: Having regard for the repercussions of their physical state, including location, being perceivable by the system

In maintaining awareness, one difficulty is in adequately informing users of the flows, history, and states of more complex systems. Meeting the balance between overwhelming users and underwhelming them can be difficult.

Accountability

In a ubiquitous system, interpersonal information is something which should ideally be traceable to show who can access, and has accessed, what. In order for social governance to take place, people should be held accountable for what they do. When personal data is accessed, it should be clear who did so, and when - to both the person concerned and the one doing the accessing. Other matters such as how it was accessed, where from, or why, are also subject to the social norms and contexts placed on these aspects by those concerned. This ’you-know-that-I-know-that-you-know’ effect controls the (mis)use of shared personal information. Another balance to make is how much accountability is necessary. Too much exposure of usage may create tension, while too little may do the same. Being able to reliably link usage to an individual is also a matter to consider.

Change

The social norms brought into the foreground, or created through the previous steps will bring about changes in usage. Being able to anticipate repercussions for actions, will cause users to think more carefully about what they reveal and what they access, as well as how (often), when, why, etc. If a user can determine that certain changes to information flow will be overall beneficial, the user may decide to act on those changes. This applies to both the user's needs, and that of the those around them.

Understanding the resulting effect of sharing or controlling information will help users find a level which suits them. Although some may retreat in-wards upon realizing the consequences of over-sharing, or over-share when the consequences of doing so are not immediately apparent. This is why meeting the balance is important. There is also the matter of outliers, where some users will not be as comfortable or as uncomfortable as the majority. Therefore, having the means to share less or more than others is important.

Implementation

Implementing this pattern is a matter of providing logging, reporting, and other informational access and notifications on user-selected/filtered, appropriately defaulted, relevant usage data. The data provided to the notified users should not intrude on other users' sensitive information, apart from the activities which involve the notified users.

The right balance needs to be met, both in the selected defaults and in the minimum and maximum levels available to the users in settings. The settings should be found easily, as should additional information. Balanced defaults can be determined from identified norms among users, while minimums should cater to the least interested and maximums to those most interested in their data's usage.

Users should ideally be able to choose the medium for the notifications, and for information retrieval, which best suits them.

Candidates include email, push notifications on mobile devices, or simply from the same interface as the rest of the system.

An effort can be made to slowly introduce users to this system. For example, starting out more privately and then gradually revealing future information so that users have time to adjust their usage. This way users are less likely to portray an undesirable usage pattern.

These correlations may also be stored in a secure way, so that they cannot be viewed arbitrarily by backend users. If users are given assurances about what information can be seen by who, including backend users, they will be more willing to make use of the information and notification system.

Constraints and Consequences

Advantages:

- Users are less likely to portray a negative usage pattern if they are aware of the correlation of their actions to it. This results in a more positive user experience once adjustments have been made. In some cases, this can increase productivity, and or efficiency in using the system.

Disadvantages:

- Initial discovery of the way they appear from the outside may lead users to retreat into themselves and disclose little to no information, cease using the system, and or call for their usage to be erased. This can be mitigated by slowly introducing users to the system without immediately providing intricate usage history
- Even if introduced slowly, users may be dissatisfied with their usage pattern as it appears to the system (and any authorised backend users). These correlations should ideally be difficult (or impossible) to retrieve if not by the user in question

Motivating Scenario

Scenario A Groupware Calendar System (GCS), 'Augur', Tullio, J., Goeckes, J., Mynatt, E.D, and Nguyen, D.H. Augmenting Shared Personal Calendars. Submitted to UIST'02 Paris, France.

History: This example logs all access to the shared calendars by the group members, and if the calendar is public, especially the users which are not expected to do so. The sharing of these calendars produces social norm information, that is, social trends from how members use the information, spread it, or dismiss it. This includes the usual when/how/what/etc. information around access. As the members are presented this analysis, they are given the ability to react to it, and adjust the sharing or details within their calendars to their privacy needs.

Feedback: Augur informs users who accessed the calendar, when, where, and what in particular was seen. This is especially useful for shared calendars since the feedback mechanism allows users of the calendar to adjust what they add to it, or who is permitted access the information.

While a GCS notification, or information display could reside anywhere, the 'native habitat' for calendar related information is in the calendar application, a mail application, or if the user chooses and if necessary, the area where time sensitive notifications usually appear.

Awareness: Users are given information at the chosen detail and notification levels in order to feel comfortable with the system. Since monitoring can negatively impact the users, the level of this is also configurable.

Accountability: Social norms around mutual understanding of what will/has been accessed in this example will affect calendar viewing and sharing. Prying into other's affairs without reasonable explanation could have social (or other) consequences. This includes a distinction between occasional viewing and constant checks. This may result in less information being shared, different access control settings, or an inquiry into the usage, which may address underlying issues.

Change: Being able to see cause and effect around different personal data sharing, hiding, or specific information flows, will likely bring about changes in how users use the shared calendar system. Important information may be shared while leaving less important details out, increasing efficiency.

Know Uses and Related Work

WebAware provided a view of page accesses, this was extended to a Web Server Log Mirror (WSLM). This was initially shown at http://www.smartmoney.com/marketmap/, but is no longer available.

Categories

- Inform
- Provide

Related Patterns

P Minimal Information Asymmetry

P Impactful Information and Feedback

P Privacy Dashboard

P Personal Data Table

P Appropriate Privacy Feedback

Supporting Patterns

P

Sources

- https://privacypatterns.org/patterns/Privacy-Mirrors

2.42 Appropriate Privacy Feedback

Summary

Context

Users are frequently unaware or unsure about what personal data systems collect and otherwise process. When systems fade into the background users are less likely to take notice and adjust what information is collected. Data controllers who provide services (or products) to such users realize that consent is not valid without users first being sufficiently informed. They aim to do so in a manner which is appropriate for the service.

The controller may have relied on op-out mechanics, but now realizes that within the European General Data Protection Regulation (recital 32) ’silence, pre-ticked boxes or inactivity’ no longer constitute consent. Unnecessarily disruptive notice is also not permitted.

The controller may already consider Fair Information Practices, and have an accessible privacy policy. They may also implement Respecting Social Organizationsand Building Trust and Credibility. However, their service is not immediately obvious to the user when in use.

Problem

Many systems are designed to be seamless or ubiquitous. However, this can make personal data risks less apparent to the user.

As a result users may overlook services without fulling understanding the privacy risks involved. Potentially, these users may realize consequences long after, or worse, not realize them at all.

Forces and Concerns

- Controllers want systems to do their tasks in the background without bothering the user, but need the user’s informed consent
- Controllers often do not want to process data which users feel uncomfortable about, but uninformed users may pro-

vide it

- Users want to get the benefits of a service without having to interact with it, and may not do so at all if they do not have to
- There are users who would avoid these services if they were aware of the privacy risks

Goal

G

Solution

Visible feedback loops, which capture the user's attention, are needed to help ensure that users understand what data is being collected, who can see that data, and how might it be used.

Implementation

Notification should occur before access where possible, and during or shortly after access if earlier notification is not appropriate. In most cases this means preventing a user's use of a service before allowing the core functionality of the service to run at all. Where some features with variable privacy implications are not essential to the service, they may be provided as optional, defaulting to being disabled.

Users should be informed appropriately, providing both concise and understandable explanations of the personal data acquired, and warnings of the risks involved. The service should make a best effort to ensure that the user understands the implications of consent before commencing or resuming functionality. An effort should also be made to make these notifications non-invasive. Using Ambient or Asynchronous Notice is one way to achieve this.

Where users choose to be notified less immediately or less often, and after being warned of the risks involved, then the service may store logs of its privacy affecting activities. The user should then be able to retrieve these logs, in a human readable form,

> at will. As only the user should be able to access these, unless said user provides informed consent otherwise, it should also be secured. Use state of the art means of encryption to do this. If this functionality cannot be done in this manner, due to technical constraints for example, then do not provide logging functionality.

Constraints and Consequences

C The user will be informed before using a service, which will cause the user to be more careful according to their personal privacy preference. Those who find the service too invasive will not use it, or provide feedback towards its improvement. The service will not be liable for user activities where it has informed them of the risks those activities involve.

Constraints

Preventing functionality until consent is acquired lessens the feasibility of various services. However, doing otherwise presents risks of high financial and good-will damages.

Motivating Scenario

> **Scenario** When you share some content on Facebook, it sometimes asks you to review your fundamental privacy settings. In the short tour given, you can see what data is accessed by other users or by third party applications. ■

Know Uses and Related Work

Categories

- Inform
- Notify

Related Patterns

P Awareness Feed

Ⓟ Privacy Awareness panel

Ⓟ Who’s listening

Ⓟ Trust Evaluation of Services Slides

Ⓟ Who’s listening

Ⓟ Increasing Awareness of Information Aggregation

Ⓟ Reasonable Level of Control

Ⓟ Privacy Mirrors

Ⓟ Privacy Dashboard

Ⓟ Ambient

Ⓟ Asynchronous Notice

Supporting Patterns

Ⓟ

Sources

- https://privacypatterns.org/patterns/Appropriate-Privacy-Feedback

2.43 Impactful Information and Feedback

Summary

Context

Users are frequently in a rush to use services at the same pace as their own ever quickening lifestyles. Such value for time can leave them unaware of the potential for mistakes, such as in automatic media sharing, or the careless disclosure of information in their contributions. These mistakes may disclose personally identifiable information, or otherwise undesirable associations. Sometimes whether the information is appropriate is dependent on the audience, or some other contextual element. Controllers who provide services to these users do not fare well when these instances occur, as they provide the means for it to happen. As such, they tend to want to be proactive in handling such issues.

Problem

A lack of user awareness in the moment can lead to regretted disclosure, whether this disclosure is manually or automatically performed.

Forces and Concerns

- Users want to use a service in an immediate and streamlined fashion, but in doing so expose themselves to mistakes
- Many users disclose unintentional information during the use of a service, especially when participating without caution
- Controllers do not want users to use a service in a way which fosters regret
- Controllers want users to learn to use a service responsibly without having to make mistakes

Goal

G

Solution

Use contextual privacy warnings, through analytical measures and historical queues to provide relevant information and suggestions regarding pending disclosures.

Implementation

Prior to submissions taking place, and provided that the user has consented to contextual privacy warnings, analyze the content of the disclosure using natural language processing. This may also entail additional metadata, such as factors pertaining to the expected audience, social comparisons against similar users or ones which the user in question has connected with. All users from which the analysis is derived should also first have provided their explicit, informed, and freely-given consent.

Search for strings which are likely to heighten the sensitivity of the content, and evaluate this against the expected audience. Where users disregard warnings, take note for improvement of future predictions through a feedback loop. Additionally, allow users to signify that despite ignoring the warning, they later regretted the post (or detect deletions which imply this) to distinguish false positives from inaccurate warnings.

One way to increase user understanding of the risks involved is to demonstrate by example a disclosure which matches the approximated sensitivity or contextual appropriateness of their content. This example will need to be one which they could usually view on their own, so as not to inadvertently cross the boundary of another user's privacy. This approach is also susceptible to inaccuracies, and would also need to be improved overall by the userbase.

The learning algorithm may at first be trained using text mining from logs of users who have opted-in to the, at first, experimental feature. While assumptions may be made, possibly inaccurately, users could also give feedback about regretted submissions or contextual appropriateness. Which type of learning is chosen is dependent on what information the controller has at their disposal at the time. If starting fresh, the implementation will likely be less sophisticated. While if available, solutions can be as complex as a reinforced classification learning algorithm.

Constraints and Consequences

By applying this pattern, users who choose to partake will have a better realization of what might happen when they disclose certain content. This can apply to any information they put online, and may show who will be able to see what. This can be both beneficial and disadvantageous, as this means users will be more cautious and less likely to contribute. They may also have worries about the trustworthiness of the learning algorithm which may access their content before they themselves have seen it fit for publication.

Motivating Scenario

Scenario Systems can reduce user uncertainty about factors important to disclosure choices. For example, systems may be able to estimate the audience for a particular disclosure at decision-time, thereby reducing uncertainty and influencing user choices. Systems could use social comparison, such as decisions made by friends or other users in similar context, to reduce uncertainty about relevant norms for disclosure. Finally, tools for viewing photo “disclosures” in ways similar to how others will view these photos could help users understand the content and appearance of their disclosures. ■

Know Uses and Related Work

Categories

- Inform
- Explain

Related Patterns

Awareness Feed

Privacy Mirrors

Ⓟ Unusual Activities

Ⓟ Preventing Mistakes or Reducing Their Impact

Ⓟ Privacy Awareness Panel

Ⓟ Informed Credential Selection

Ⓟ Privacy Dashboard

Ⓟ Appropriate Privacy Icons

Ⓟ Icons for privacy Policies

Ⓟ Privacy Color Coding

Ⓟ Privacy Aware Wording

Ⓟ Layered Policy Design

Ⓟ Privacy-aware Network Client

Ⓟ Appropriate Privacy Feedback

Ⓟ . • Increase Awareness of Information Aggregation

Supporting Patterns

P

Sources

- https://privacypatterns.org/patterns/Impactful-Information-and-Feedback

2.44 Decoupling (content) and Location Information Visibility

Summary

Context

Users often share content in socially oriented services on the Internet. The applications used for uploading this content may attach location information. Controllers can use or publicize this information, allowing others to use it. Sufficient correlations can infringe upon the user's privacy expectations.

The organization in question (likely the controller) does not wish to undermine these expectations, and seeks to enable the user to assign contextually specific privacy settings

Problem

Concerns about disclosing location information conflict with the appeal of location information for [content] organization.

Forces and Concerns

- Location is highly indicative of life patterns and significant contexts of the users' daily lives
- Location data is increasingly available in various consumer devices
- Users want streamlined processes for organizing their content in socially oriented services. That simplicity could be achieved by using location automatically
- Users do not want to set privacy requirements every time they generate content, nor to sweepingly deny all sharing if they intend to use the service

Goal

G

Solution

Allow users to retroactively decide upon the disclosure of location information with respect to the context of the particular content. Record location information by default, but do not automatically share it.

Structure

Give users an interface or control to configure an access policy regarding the privacy of location information. That is, a place where users may, granularly or in bulk, define who may access location information of their content.

Implementation

A basic solution could feature an interface or control for selecting the allowed users from all the types of users of the socially oriented service (e.g. built-in or user-defined groups, individuals, or anonymous users). This control could apply to individual content, or to multiple selections, or groups.

Prior to this grant of additional consent the content itself, or versions containing location, might only be available by unpublished Private Link. The protection of the content itself is however not the focus of this pattern.

If a user chooses, certain individuals or groups may have default access to the attached location information. Default access like this, however, invalidates the following approach.

Removing Controller Trust Requirement

An extended solution may aim to be further privacy preserving.

The service may accept ciphertext as the location coupled with the content. When (and only if) the user chooses to make that specific location accessible, their client-side device decrypts the location and provides the service with the plaintext location.

The user may choose how granular the location is before the service receives it. In this way, by default the controller only needs Lawful Consent to store the content itself. This solution, like many, is dependent on the trustworthiness of the client's own device.

Constraints and Consequences

C

Benefits

- Users can define in one place, or where contextually relevant, the granular privacy settings for the location information of their content
- Users do not need to consider settings when generating content, only later when sharing them, or if they choose, automatically with select individuals or groups

Liabilities

- If users do not configure the policy, then the default configuration shares nothing and the service is not being used
- Users could require fine-grained location configuration, such as how specific the location is per content. This could be addressed with additional settings
- This pattern assumes the controller is trustworthy, as all location information attached to the content is still given to the service. Alternatively, the service could endeavor to by default also restrict its own access to the information (e.g. client-side decryption)

Constraints

By applying this pattern the controller prevents location access by default, and thus risks a low location sharing rate. This is due to the tendency of users to leave settings in the default state. However, depending on the effort, the controller may encourage positive public image, and raise adoption overall.

Motivating Scenario

Scenario

- Flicker (basic implementation)
- Twitter (basic implementation)
- Facebook (basic implementation)

Know Uses and Related Work

Categories

- Control
- Retract

Related Patterns

- Support selective Disclosure
- Lawful Consent
- Private Link
- Discouraging Blanket Strategies
- Negotiation of Privacy Policies
- Reasonable Level of Control
- Buddy List

Supporting Patterns

Sources

- https://privacypatterns.org/patterns/Decoupling-[content]-and-location-information-visibility

2.45 Platform for Privacy Preferences

Summary

Context

Users are frequently intimidated or discouraged by the size and complexity of legal texts. Privacy policies are an example of such texts, which are in the user's best interest to understand. As these policies are also written for the sake of legal compliance, balancing or reconciling comprehensiveness with comprehensibility is non-trivial. Different users will have varying thresholds to the amount of detail they will readily look through. The controller in this case wants to make their privacy policy more accessible to their users.

Problem

Users regularly do not read privacy policies, as they are too verbose, complex, and repetitive amongst the sites they visit.

Forces and Concerns

- Users typically do not want to read walls of texts, often needing to be persuaded to inform themselves
- Controllers want to ensure that users are not surprised and or upset about what is done with their data
- A number of users want to really understand what risks they are taking regarding their privacy
- Controllers want to be legally compliant, and minimize the costs involved in catering to data protection

Goal

G

Solution

Controllers may use the P3P standardization of terms and data elements to construct their privacy policies, allowing users to instead immediately see the policy distinctions which matter before using the service. The policies they share with other controllers the user is subject to will already have been reviewed, or are separated such that minimal time is spent reviewing policy.

Rationale

By removing redundancies, there is far less to read. By standardizing, comprehension is strengthened.

Implementation

The controller must publish the P3P syntax files and policy reference file to their live site. The files may be generated by automated tools. It is encouraged that the policy reference file be published in the well-known location, /w3c/p3p.xml. A link tag or HTTP Headers may also be used. The policies used may also cover the entire site, or specific areas.

Further information is available at https://www.w3.org/TR/P3P/

Structure

Constraints and Consequences

C Users will be able to construct preferences for a privacy standard (risk appetite) which they personally can accept. This template will allow them to quickly review the privacy policy of the controller while avoiding repetition, and understanding distinctions. They may additionally choose to have site-specific preferences which point out what is most relevant to them.

Constraints

The human readable privacy policy should be compatible with what can also be expressed using the P3P standardization. While extensions can be made to the specification, there is a limit to this. Careful consideration will need to be used when constructing the policy to ensure full coverage. This may require additional explanation beyond what the P3P specification can provide, which needs to be clearly indicated and explained to users.

Motivating Scenario

Scenario The following example is taken from the P3P1.0 specification: Claudia has decided to check out a store called

CatalogExample, located at http://www.catalog.example.com/. Let us assume that CatalogExample has placed P3P policies on all their pages, and that Claudia is using a Web browser with P3P built in.

Claudia types the address for CatalogExample into her Web browser. Her browser is able to automatically fetch the P3P policy for that page. The policy states that the only data the site collects on its home page is the data found in standard HTTP access logs. Now Claudia's Web browser checks this policy against the preferences Claudia has given it. Is this policy acceptable to her, or should she be notified? Let's assume that Claudia has told her browser that this is acceptable. In this case, the homepage is displayed normally, with no pop-up messages appearing. Perhaps her browser displays a small icon somewhere along the edge of its window to tell her that a privacy policy was given by the site, and that it matched her preferences.

Next, Claudia clicks on a link to the site's online catalog. The catalog section of the site has some more complex software behind it. This software uses cookies to implement a "shopping cart" feature. Since more information is being gathered in this section of the Web site, the Web server provides a separate P3P policy to cover this section of the site. Again, let's assume that this policy matches Claudia's preferences, so she gets no pop-up messages. Claudia continues and selects a few items she wishes to purchase. Then she proceeds to the checkout page.

he checkout page of CatalogExample requires some additional information: Claudia's name, address, credit card number, and telephone number. Another P3P policy is available that describes the data that is collected here and states that her data will be used only for completing the current transaction, her order.

Claudia's browser examines this P3P policy. Imagine that Claudia has told her browser that she wants to be warned whenever a site asks for her telephone number. In this case, the browser will pop up a message saying that this Web site is asking for her telephone number, and explaining the contents of the P3P statement. Claudia can then decide if this is acceptable to her. If it is acceptable, she can continue with her order; otherwise she can cancel the transaction.

Alternatively, Claudia could have told her browser that she wanted to be warned only if a site is asking for her telephone number and was going to give it to third parties and/or use it for uses other than completing the current transaction. In that case, she would have received no prompts from her browser at all, and she could proceed with completing her order.

Know Uses and Related Work

Categories

- Inform
- Provide

Related Patterns

Ⓟ Dynamic Privacy Policy Display

Ⓟ Privacy Policy Display

Ⓟ Policy Matching Display

Ⓟ Privacy-Aware Network Client

Supporting Patterns

Ⓟ

Sources

- https://privacypatterns.org/patterns/Platform-for-Privacy-Preferences

2.46 Selective Access Control

Summary

Context

Users enjoy social reaction when posting content in socially oriented services on the Internet. Though sometimes the reactions are not as ideal. Some content is inappropriate for some audiences, and some users would rather keep some content mostly private. While users are capable of sharing content privately, perhaps through Private Link, they may wish to have better control over whom they share with in their service of choice. The controller providing this service may too want its users to share more specifically.

Problem

Users want to control the visibility of the content being shared, because it may not currently be appropriate for all users.

Forces and Concerns

- Users aim to share content aimed at different kinds of users because they have varying social proximities (friends, family, colleagues, etc.).
- Users want to share content to certain other users based on the content's nature for that user specifically
- Users could have trouble over-sharing, dealing with content aimed at the wrong audience, or under-share as a precaution

Goal

G

Solution

> Provide users with the option to define the audience of their contributions by specifying the access rules to their [content].

Implementation

Implement visual controls to help users to define access control rules when they create or modify content.

These rules could be defined based on users, groups of users, or based on context-aiding attributes like age or location. For groups, it should be possible to directly define who may view the post being published (e.g. a post with personal data aimed only at a group of close friends). Contextually, it may be possible to define an attribute constraint based on whom in general the post is intended for (e.g. a post aimed at people in a specific town or region).

Constraints and Consequences

Benefits

Users have the possibility to control access as they want in every submission. It allows configuration based on kinds of users or the content's context.

Liability

Users could find granular configurations time consuming or tedious, so a default configuration for new submissions would be helpful.

Motivating Scenario

Scenario

- Facebook
- YouTube

■

Know Uses and Related Work

Categories

- Control
- Choose

Related Patterns

P Selective Access Control

P Private Control

P Support Selective Disclosure

P Reasonable Level of Control

P Discouraging Blanket Strategies

P Decoupling [content] and location information visibility

P Lawful Consent

Supporting Patterns

P

Sources

- https://privacypatterns.org/patterns/Selective-Access-Control
- http://privacypatterns.wu.ac.at:8080/catalog/

2.47 Pay Back

Summary

Context

In services where users may contribute content, or provide the system with account or profile information, the information is only valuable if relevant and accurate. For controllers providing this service (or product), worthless information does not typically generate income or future participation. Without consistent usage, a service becomes less popular and eventually may run at loss. This is particularly true in socially oriented services. To keep the service working, it is crucial that its users maintain content. Users however, might not feel inclined to do so. Keeping content up to date, or adding it in the first place, requires effort, and in some cases an acceptance of privacy risk.

Problem

Users do not necessarily want to provide and maintain content, they need a motivation to do so. Without this, a service will not flourish.

Not all users will be equally motivated, so the service may not receive contribution at the level required to keep the service competitive. Furthermore, some users might not contribute at all. Thus, it is difficult to maintain Reciprocity.

Forces and Concerns

- A service that depends on information flow requires a continuous feed of user activity
- If users are not motivated they likely will not continue to contribute content
- Some users do not require much motivation, and the use of the service could be enough for them to contribute. But this alone is insufficient for most services to run

Goal

G

Solution

Provide users with different kinds of benefits when they contribute or maintain content for the service and make sure they do so consensually.

Implementation

Depending on the kind of service that is provided, different benefits could be considered: virtual or real currency, use of services, social benefits, and so on.

When using virtual or real currency, the controller should first define how much in value users would receive depending on the contributions. In the case of virtual currency, the places where the currency could be used should be defined.

Regarding use of service, some criteria could be considered non-exhaustively: feedback on content, frequency of contributions, the use of service for a minimum duration, access to a service earlier than others, or use of special features within the service.

When users reach a limit, they could additionally assist with virtual or physical events for learning, meeting people, etc. In virtual scenarios, users could receive attention (feedback) from one another.

Constraints and Consequences

Benefits

More users will be motivated to provide information, so the service could continue to be competitive.

It could help to maintain Reciprocity.

Liabilities

It could be necessary to monitor the quality of the contributions before giving the user benefits.

Consent will not be genuine if users are coerced into providing their personal data. An example of this is the sunk cost fallacy. As the user builds emotional investment, the controller has more power over them. A service which was once available freely, or anonymously, can push users into accepting terms they do not truly consent to. When using this pattern it is important to make sure that Lawful Consent is also used.

YouTube financial retribution. Dropbox increasing storage programs. Local guides for Google Maps. Likes, comments, followers in Facebook, Instagram.

Motivating Scenario

Scenario ■

Know Uses and Related Work

Categories

- Control
- Choose

Related Patterns

P Reciprocity

P Incentivized Participation

P Lawful Consent

Supporting Patterns

P

Sources

- https://privacypatterns.org/patterns/Pay-Back

2.48 Privacy Dashboard

Summary

A single point of access to monitor and control large quantities and various types of personal information.

Context

A service (or product) which processes personal data of users may make that data accessible to them. This is often the case whether conforming to laws about self-determination and notice, or merely wanting to provide an additional privacy consideration for the sake of users. The controller concerned wants to open up about the data they have processed, and to improve the ease of use for configuring privacy settings.

This pattern applies specially to data controllers which hold large quantities of personal data, of different types, for different purposes, and share it with different third parties.

Problem

A system should succinctly and effectively communicate the kind and extent of potentially disparate data that has been processed.

Users may not remember or realize what data a particular service or company has collected, and thus can't be confident that a service isn't collecting too much data. Users who aren't regularly and consistently made aware of what data a service has collected may be surprised or upset when they hear about the service's data collection practices in some other context. Without visibility into the actual data collected, users may not fully understand the abstract description of what types of data are collected; simultaneously, users may easily be overwhelmed by access to raw data without a good understanding of what that data means.

Forces and Concerns

- Controllers want to provide users with sufficient information to determine how it is used, and to prevent regrettable sharing decisions
- Controllers want to prevent both over and under-sharing, so as to provide users with the best experience possible
- Users often do not realize the privacy risks in providing their personal data
- Users do not want to be subjected to too many or overly detailed notifications, as they will quickly make a habit of overlooking them

Data controllers may hold loads of personal data of different types. Service provision often involves using those data for different purposes and sharing them with several third parties. Users may not be aware of all the data being collected, created, maintained, processed and shared by the service provider or third parties. Moreover, the access to this data is scattered through different interfaces, which pose further difficulties for data subjects to manage their personal data.

Goal

G Allow users to monitor their personal data at a glance, and easily control them and the associated permissions.

Solution

Provide successive summaries of collected or otherwise processed personal data for a particular user, representing this data in a meaningful way. This can be through demonstrative examples, predictive models, visualizations, or statistics.

Where users have choices for deletion or correction of stored data, a dashboard view of collected data is an appropriate place for these controls (which users may be inspired to use on realizing the extent of their collected data).

Structure

A variation of the privacy dashboard, Privacy Mirrors, focuses on history, feedback, awareness, accountability, and change.

Implementation

Implementing this pattern is a matter of providing logging, reporting, and other informational access and notifications on user-selected/filtered, appropriately defaulted, relevant usage data.

Aspects which the controller wishes to inform their users about may include the collection and aggregation of their data, particularly personal data which:

- changes over time,
- is [processed] in ways that might be unexpected,
- is invisible or easily forgotten, or
- is subject to correction and deletion by users.

Provide the user with an easy-to-access single view that summarizes the different types of personal data held by the data controller at a glance, together with user interface controls to control that data and the associated permissions (i.e. amend them, erase them, modify the purposes for which they can be used, or the parties with which it can be shared, when applicable).

Constraints and Consequences

Constraints

As in other access mechanisms, showing a user's data back to them can create new privacy problems. Implementers should be careful not to provide access to sensitive data on the dashboard to people other than the [user]. For example, showing the search history associated with a particular cookie to any user browsing with that cookie can reveal the browsing history of one family member to another that uses the same computer.

Also, associating all usage information with a particular account or identity (in order to show a complete dashboard) may encourage designers to associate data that would otherwise not be attached to the user account at all. Designers should balance the access value against the potential [considerations within] Deidentification.

The data controller must provide a user interface (it is difficult to apply this pattern to, e.g. surveillance devices) which is capable of authenticating the data subjects whose data is managed. Fine-grained control of personal information is not directly provided by the dashboard.

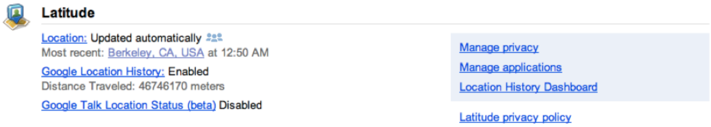

Latitude

Location: Updated automatically
Most recent: Berkeley, CA, USA at 12:50 AM

Google Location History: Enabled
Distance Traveled: 46746170 meters

Google Talk Location Status (beta) Disabled

Manage privacy
Manage applications
Location History Dashboard

Latitude privacy policy

Figure 2.5: Google Privacy Dashboard

Motivating Scenario

Scenario The Google Dashboard shows a summary of the content stored and/or shared by many (but not all) of Google's services (Latitude, Google's location sharing service, is shown above). For each service, a summary (with counts) of each type of data is listed, and in some cases an example of the most recent such item is described. An icon signifies which pieces of data are public. Links are also provided in two categories: to actions that can be taken to change or delete data, and to privacy policy / help pages.

Google Accounts: About the Dashboard

See Also Dashboards are a widely-used pattern in other data-intensive activities for providing a summary of key or actionable metrics.

A Web-oriented corporation provides different services that span personal communication (e-mail, instant messaging), hosting and publishing (blogs, photo galleries, videos), cloud-based content storage and management (of documents, pictures, personal agenda, etc.) and Web search. They make their business from leasing the "screen real estate" on their sites and elsewhere, for external advertisements. They track the user's browsing both on their own websites and elsewhere, capturing the user interactions in order to serve the best possible advertisements (i.e. more targeted, and thus more suitable as well as more effective). All these user activities allow the corporation to amass a large quantity of user data, which would be simply unmanageable by the data subjects if they needed to do it on a per-item basis. The typical users are not even aware that the corporation holds all those data about them.

■

Know Uses and Related Work

- Google Privacy Dashboard now integrated in the broader [My Account] (https://myaccount.google.com/) which manages personal data held and processed by their different services (Google Search, Drive, Blogger, Google+, Android, etc.)
- Microsoft Account which manages personal data held by their services including Bing, Outlook.com, Skype, etc
- Meet-U: The key points of TAC that affect the user's privacy the most, are displayed on one screen. Hence, the gathering and processing of data are addressed and summarized briefly. The long version of the TAC is linked. The user has to agree on that before continuing with the application

See also Privacy Dashboard by Nick Doty at privacypatterns.org for an alternative presentation of this pattern.

Categories

- Transparency
- Access
- Location
- Inform
- Provide
- Control

Related Patterns

P Awareness Feed

Privacy Mirrors

Appropriate Privacy Feedback

Supporting Patterns

Sources

- https://privacypatterns.org/patterns/Privacy-dashboard
- https://privacypatterns.eu/#/patterns/privacy-dashboard/0-0-0-0-0-1-1-0-0-1-0-1-0-0-0-0
- http://privacypatterns.wu.ac.at:8080/catalog/

2.49 Preventing Mistakes or Reducing their Impact

Summary

Context

Numerous services (or products) are designed with the purpose of sharing amongst the public or a specific subset of users. In content sharing implementations, it is commonplace to streamline disclosure so that users do not need to publish manually. Content they generate is often automatically shared with the controller, even if not immediately made available to other users or the public. This of course requires the prior consent of users, though it is also possible for users to forget about that consent, or change their mind. If the distinction lies in a simple setting, it may not be apparent to the user that it is still in effect.

Problem

Immediate and automatic content publication without notification or confirmation of consent leads to unintentional disclosure and may invalidate prior consent.

Forces and Concerns

- Users of the service want to share content with others, but not all of the content they generate is fit for sharing
- Most users do not want to manually upload content case by case, sometimes long after creation
- Controllers want to make it easy for users to contribute content
- Controllers do not want users to disclose content which they regret disclosing and potentially ruins the user's experience

Goal

G

Solution

Use contextual measures to predict whether content should be processed, re-establishing consent, to prevent accidental disclosure.

Implementation

Through the study of patterns in disclosure behavior, systems may be able to helpfully warn users when disclosing following potentially significant change in context, perhaps reducing potential for mistakes. [These] privacy decisions are often correlated with the context of capture and the [content] as indicated [by the user. It] could be feasible to use these patterns for prediction or recommendation of privacy settings. In addition, providing an optional "staging area" before disclosure actually takes place and an easy way to review recent disclosures may reduce the immediate consequences of quickly regretted or accidental disclosure decisions.

Constraints and Consequences

C Clearing up mistakenly shared data adds additional overhead, especially if the service does not offer simple modification or removal of information. As sharing more than actually intended may result in potential damage for users, they will benefit from services which reduce these risks.

Motivating Scenario

Scenario Through the study of [trends] in disclosure behavior, systems may be able to helpfully warn users when disclosing following potentially significant change in context, perhaps reducing potential for mistakes. As [Ahern et al.] found that privacy decisions are often correlated with the context of capture and the content of the photo as indicated by user-specified tags, it could be feasible to use these patterns for prediction or recommendation of privacy settings. In addition, providing an optional "staging area" before disclosure actually takes place and an easy way to review recent disclosures may reduce the immediate consequences of quickly regretted or accidental disclosure decisions.
■

Know Uses and Related Work

Categories

- Inform
- Notify

Related Patterns

P Impactful Information and Feedback

P [Informed] Credential Selection

P Asynchronous Notice

P Ambient Notice

Supporting Patterns

P

Sources

- https://privacypatterns.org/patterns/Preventing-Mistakes-or-Reducing-Their-Impact

2.50 Obligation Management

Summary

The pattern allows obligations relating to data sharing, storing and processing to be transferred and managed when the data is shared between multiple parties.

Context

The developer aims to make sure that multiple parties are aware of and comply with required user/organizational policies as personal and sensitive data are successively shared between a series of parties who store or process that data.

Problem

Data may be accessed or handled by multiple parties that share data with an organisation in ways that may not be approved by the data subject.

Goal

G

Solution

Service providers use an obligation management system. Obligation management handles information lifecycle management based on individual preferences and organizational policies. The obligation management system manipulates data over time, ensuring data minimization, deletion and notifications to data subjects.

Constraints and Consequences

C Benefits: privacy preferences and policies are communicated and adhered to among organisations sharing data. Liabilities: additional effort to set obligations.

Motivating Scenario

> **Scenario** A service provider subcontracts services, but requires that the data to be deleted after a certain time and that the service provider requires to be notified if there is further subcontracting.
> ■

Know Uses and Related Work

> Pretschner et al (2009) provide a framework for evaluating whether a supplier is meeting customer data protection obligations in distributed systems. Researchers at IBM propose Enterprise Privacy Authorization Language (EPAL) (2004) to govern data handling practices according to fine-grained access control. Casassa Mont (2004) discusses various important aspects and technical approaches to deal with privacy obligations. Pretschner, A., Schtz, F., Schaefer, C., and Walter, T.: Policy Evolution in Distributed Usage Control. Electron. Notes Theor. Comput. Sci. 244, 2009 IBM, The Enterprise Privacy Authorization Language (EPAL), EPAL specification, http://www.zurich.ibm.com/security/enterprise-privacy/epal/, 2004 Mont, M. C., Dealing with Privacy Obligations, Important Aspects and Technical Approaches, TrustBus, 2004.

Categories

- Enforce
- Uphold

Related Patterns

P

Supporting Patterns

P

Sources

- `https://privacypatterns.org/patterns/Obligation-managemen`
- `https://privacypatterns.eu/#/patterns/obligation-manageme`
  `0-0-0-0-0-0-0-0-0-0-0-0-0-0-0-0`

2.51 Informed Credential Selection

Summary

Context

Controllers offering services (or products) often provide a means to authenticate users in order to permit them access. This access can be to a secure function, such as fulfilling a transaction. Since this action may have difficult to reverse consequences, the controller needs to be certain of the user's identity and informed consent. In order for consent to be valid, the controller must ensure that it is informed, as well as freely given, specific, and unambiguous. In order to determine identity, however, personally identifying information is needed. Some methods of authentication are also more invasive than others, allowing users to provide more information than necessary.

Problem

Credentials which users supply may be more invasive than necessary, this is a kind of consent which legally must be informed.

Forces and Concerns

- Users want to authenticate so that they know only they can obtain access
- Users do not want to provide more information than they feel comfortable or than is necessary
- Controllers want to prevent unauthorized access to user actions, as this can seriously affect their experience
- Controllers do not want to process user data for which they do not have valid consent

Goal

G

Solution

> Allow granular credential selection which explains to users the various ways in which personal data can be used, including

who may access it, and how it may be used to derive further information.

Implementation

Present the user with a selection mechanism that shows the user what possible choices are available and then show a summary page that contains the data to be sent. The explanation of consequences must be shown as the user investigates the available credentials. It should be clear to the user which information authenticates them with the least privacy impact.

One mental model for this could be the use of a credit card for identification. See the HCI Pattern Collection for further information on this example.

Independent of a mental model, the credential selection UI should contain two steps, namely, selection and summary. During the first step, all graphical elements of the selection mechanism should be based on the mental model. Thus, if working with the card based metaphor this should be apparent from the UI. During the second step, the invoked mental model is not as important as the key issue is to clearly convey which selected data and which meta-data is being sent.

Constraints and Consequences

C Allows a user to identify themselves in a granular way, controlling how much information they reveal by doing so.

C This approach should be used to make it easy for users to select the appropriate credentials. It also should inform them about which (personal) data and meta-data the recipient of the information will have after the transaction.

Motivating Scenario

Scenario Jiang et al. (2010). "A Classified Credential Selection Scheme with Disclosure-minimizing Privacy". International Journal of Digital Content Technology and its Applications, 4

(9), December 2010. 201 - 211. ■

Know Uses and Related Work

Categories

- Inform
- Explain

Related Patterns

P Informed Secure Passwords

P Preventing Mistakes or Reducing Their Impact

P Unusual Activities

Supporting Patterns

P

Sources

- https://privacypatterns.org/patterns/Informed-Credential-Selection
- http://privacypatterns.wu.ac.at:8080/catalog/

2.52 Anonymous Reputation-based Blacklisting

Summary

Get rid of troublemakers without even knowing who they are.

Context

A service provider provides a service to users who access anonymously, and who may make bad use of the service.

Problem

Anonymity is a desirable property from the perspective of privacy. However, anonymity may foster misbehaviour, as users lack any fear of retribution.

A service provider can assign a reputation score to its users, based on their interactions with the service. Those who misbehave earn a bad reputation, and they are eventually added to a black list and banned from using the service anymore. However, these scoring systems traditionally require the user identity to be disclosed and linked to their reputation score, hence they conflict with anonymity. This has made, for instance, Wikipedia administrators to take the decision to ban edition requests coming from the TOR network, as they cannot properly identify users who misbehave.

A Trusted Third Party (TTP) might be introduced in between the user and the service provider. The TTP can receive reputation scores from the service provider so as to enforce reputation-based access policies, while keeping the identity hidden from the service provider. However, this would require the user to trust the TTP not to be a whistle-blower indeed.

How can we make users accountable for their actions while keeping them anonymous?

Goal

A service provider wants to prevent users who misbehave from accessing the service anymore, without gaining access to their identity.

Solution

First, the service provider provides their users with credentials for anonymous authentication.

Then, every time an authenticated user holds a session at the service, the service provider assigns and records a reputation value for that session, depending on the user behaviour during the session. Note that these reputation values can only be linked to a specific session, but not to a specific user (as they have authenticated anonymously).

When the user comes back and starts a new session at the service, the service provider challenges the user to prove in zero-knowledge that he is not linked to any of the offending sessions (those that have a negative reputation associated). Zero-knowledge proofs allow the user to prove this, without revealing their identity to the service provider. Different, alternative proofs have been proposed, e.g. prove that the user is not linked to any of the sessions in a set of session IDs, prove that the last K sessions of the user have good reputation, etc.

In practice, more complex blacklisting rules can be applied as well. For instance, several reputation scores can be assigned to the same session, each regarding different facets of the user behaviour. Then, the blacklisting thresholds may take the form of a Boolean combination or a lineal combination over individual session and facet reputation values.

A service provider wants to prevent users who misbehave from accessing the service anymore, without gaining access to their identity.

Constraints and Consequences

C Different implementations may only be practical for services with a reduce number of users, require intense computations, limit the scope of the reputation to a constrained time frame, be vulnerable to Sybil attacks, etc. Nonetheless, protocols are being improved to overcome these and other issues. See the cited sources below for the specific discussion.

Motivating Scenario

Scenario A wiki allows any visitor to modify its contents, even without having been authenticated. Some malicious visitors may vandalize the contents. This fact is signalled by the wiki administrators. If a visitor coming from the same IP address keeps vandalizing the site, they will earn a bad reputation, and their IP will be banned from modifying the contents anymore. However, users accessing through a Tor anonymity network proxy cannot be identified from their IPs, and thus their reputation cannot be tracked. ■

Know Uses and Related Work

- Au, M. H., Kapadia, A., & Susilo, W. (2012). BLACR: TTP-free blacklistable anonymous credentials with reputation
- Au, M. H., & Kapadia, A. (2012, October). PERM: Practical reputation-based blacklisting without TTPs. In Proceedings of the 2012 ACM conference on Computer and communications security (pp. 929-940). ACM.

Categories

- Separate
- Hide
- Restrict

Related Patterns

P

Supporting Patterns

P onion-routing

P anonymity-set

Sources

- https://privacypatterns.org/patterns/Anonymous-reputation-base
- https://privacypatterns.eu/#/patterns/anonymous-reputation-bas 0-1-0-5-2-0-2-3-1-0-0-1-0-1-0-0

2.53 Negotiation of Privacy Policy

Summary

Context

Often when users find a service (or product) they would like to use, and begin signing-up, they are immediately exposed to assumptions which may not hold for them. As users have differing privacy priorities, a controller cannot guess as to what settings best accommodate them. Since these preferences may be intricate, users cannot be expected to specify them in detail all at once or before using the service.

Problem

Users have sometimes wildly different priorities regarding their privacy, though a controller does not know these details when a user first joins a service. There is a temptation to provide these users the settings the average user uses.

Forces and Concerns

- Users are different and do not all fall under one universal setting without some being unsatisfied
- The controller wants to cater to user individuality
- Getting users to specify all of their individual tastes before using a service will make some users abandon the process. Some settings may be missed, and many users will be upset

Goal

G

Solution

As users begin to use a service, determine their individual privacy sensitivities by allowing them to opt-in/opt-out of account details, targeted services, and telemetry. When a user's preference is not known, assume the most privacy-preserving settings. It should always take more effort to over-share than to under-share.

Implementation

Unauthenticated users should enjoy the most privacy-preserving defaults. When a user joins the service, they may be presented with [excerpts or summaries of] a privacy policy, which they can use to inform their choices. Using simple, recognizable controls, users can be asked to opt-in (for explained benefits) or opt-out (at explained costs) before any of their data is used. They can then be asked for additional consents further down the line as they become contextually relevant.

In this way, only the needed consent is asked for as the controller's understanding of the user's preferences improves. This can allow the service to determine which solicitations users are individually likely to consider, and which ones will only waste their time or upset them.

Constraints and Consequences

C Private defaults will often not be the appropriate settings for a user, as most users may be less privacy-concerned. The additional effort taken to share more, with users or the controller, will reduce the valuable data collected. However, providing users with invasive defaults would risk public outrage by the vocal few, who may affect opinions holistically.

Motivating Scenario

Scenario ■

Know Uses and Related Work

Categories

- Control
- Choose

Related Patterns

P Support Selective Disclosure

P Discouraging blanket Strategies

P Decoupling[content] and location information visibility

P Negotiation of Privacy Policy

P Reasonable Level of Control

P Enable/Disable Functions

P Lawful Consent

P Ambient/Asynchronous Notice

P Preventing Mistakes or reducing their impact

P Awareness Feed

P Privacy Dashboard

Supporting Patterns

P

Sources

- https://privacypatterns.org/patterns/Negotiation-of-Privacy-Policy

2.54 Reasonable Level of Controls

Summary

Context

Users have certain expectations about what level of privacy they can expect in certain contexts. In general, they are given the means to provide themselves with as much or little shielding from intrusions as they need. This expectation carries over to usage of services (or products) offered by a Controller. Users expect that they can have an impact on what about them is known to a service, or others that use the service.

Problem

Users expect to be afforded sufficient self-determination over what information about them is collected or otherwise processed. The level of information and control desired, however, varies from person to person, as does the negative response when expectations are not met.

Forces and Concerns

- Users want to share and be shared with, but have varying limits on what they feel comfortable sharing
- They have their own conceptions on what is worth withholding, and different regards for information sensitivities
- Not all users trust a service to handle their information with the same care they feel is due
- Many users want others to be able to know certain things about them on request, sometimes even in real-time

Goal

G

Solution

> Allow users to selectively and granularly provide information to a service, or its users, and have select information available to

user-defined or predetermined groups.

Structure

Users should be able to push their chosen information to (or have it pulled by) those they grant access. Using push mechanisms, users will have the greatest level of control due to the fact that they can decide the privacy level of their data case by case.

Pull mechanisms are less granular, as granting access to a group or individual continues until that access is denied. Within this time frame, the sensitivity of the data may fluctuate. However, the user should have the ability to retract access at will, and thus, can manage their own identified risks.

Users should also be made aware of the potential risks of oversharing and increased sensitivity of data over time. This creates a complementing relationship between many Inform patterns, including Ambient/Asynchronous Notice, Preventing mistakes or reducing their impact, as well as Awareness Feed, Privacy Dashboard and their compounded patterns.

Additionally, Blur Personal Data and Partial Identification patterns could be used inside the implementation.

implementation

When users are pushing their information to a service, design the user interface such that where appropriate, controls define the access, granularity, completeness, accuracy, etc. of the information being shared.

Elsewhere, ensure that any required fields are truly required, and that the completeness needed for those fields be indicated. When there are automatic suggestions, let users redefine or remove the information before it is collected by the service. These automatic suggestions should also not take place without consent.

Where information is provided on a continual basis to those granted access, provide the user with the necessary tools. They should be able to indicate who falls within a group, and what exactly that group can access, for how long, at what granularity, how far back they can look, and so forth.

Constraints and Consequences

C

Motivating Scenario

Scenario

- Google Maps (simple implementation)
- Facebook (simple implementation)

Know Uses and Related Work

Categories

- Control
- Update

Related Patterns

P Selective Access Control

P Negotiation of Privacy Policy

P Decoupling [content] and location information visibility

P The Negotiation of Privacy Policy

P Lawful Consent

P Masquerade

P Support Selective Disclosure

P Discouraging blanket strategies

P Private link

P Active broadcast of presence

P Selective Access Control

P Ambient/Asynchronous Notice

P Preventing mistakes or reducing their impact

P Awareness Feed (and components)

P Privacy Dashboard (and components)

P Appropriate Privacy Feedback

P Limited Data Retention

P Fair Information Practices

P Privacy Sensitive Architecture

Supporting Patterns

P

Sources

- https://privacypatterns.org/patterns/Reasonable-Level-of-Control

2.55 Masquerade

Summary

Context

Users are frequently monitored for various reasons by a service (or product), for instance to associate them with shared activity. Monitoring is sometimes needed to allow users to know certain attributes about one another which can assist them in communicating or otherwise participating. This monitoring is sometimes apparent to the user, opted-in, or unavoidable. This may cause some users distress, or affect their actions for better or for worse. Many working environments additionally feature productivity tracking software or the ability to Gaze Over the Shoulder. This of course allows any altered activity to have an effect on work performance, or its perception. Mandatory tracking is commonly undesirable for users, and in these cases can negatively affect user experience.

Problem

Users act differently under active supervision, and this may negatively impact their content generation.

Forces and Concerns

- Controllers may require monitoring for the functioning of the service or depend on it as a business model
- Users have an interest in restricting the amount to which they are monitored
- Every user could require a different level of identifiability depending on the context
- It would be necessary to at least have Partially Identification of the user when implementing Reciprocity

Goal

G

Solution

Allow users to select their desired identifiability for the context in question. They may reveal some subset of the interaction or account attributes and filter out the rest.

Implementation

For implementing this pattern, a configuration interface will be required. Two approaches could be considered: levels of publicity or publicity profiles.

In levels of publicity, all possibly revealed information could be arranged on a scale depending on how identifying each kind of information is alone or when shown together. A visual element could be used to select a specific publicity level. When the users select one level, all information with the same or smaller publicity level will be revealed. This is taken into account when measuring where upon the scale a piece of information falls.

In publicity profiles, all possibly revealed information could be depicted using visual elements and the users have to select each kind of information that they want to reveal. Furthermore, depending on the kind of information, the users could define different granularity for each one (E.g. regarding location it is possible to define the country, region, city, department and so on).

Reciprocity could implemented by connecting privacy levels with permissions for interaction.

Constraints and Consequences

Benefits

Since users can explicitly control how much personal information they provide to other users, they [no] longer have to fear that their personal information [is being] misused by [other users]. This provides them with an environment that is as private as the situation [demands]. [Users] can decide to discuss private matters without the possibility of being monitored by other users by simply changing [their] privacy profile [or privacy level].

Liabilities

Anonymous interaction with the system may lower the inhibition threshold for destructive or forbidden behavior. The users do not have to fear that destructive activities are associated with their identity. Thus, [the controller] should provide only limited functionality for anonymous users (e.g. only read access or only moderated postings to a discussion board).

Motivating Scenario

Scenario

- Video systems: NYNEXPortholes (Lee, Girgensohn, and Schlueter 1997);
- TUKAN (Schummer and Haake 2001);
- Anonymous access in web-based communities.

■

Know Uses and Related Work

Categories

- Control
- Retract

Related Patterns

- Ⓟ Reasonable Level of Control
- Ⓟ Private link
- Ⓟ Active broadcast of presence
- Ⓟ Support Selective Disclosure
- Ⓟ Buddy List
- Ⓟ Reciprocity

Supporting Patterns

P

Sources

- https://privacypatterns.org/patterns/Masquerade

2.56 Buddy List

Summary

Context

Users frequently interact upon various media, forums, and communication channels. There are however far more users on these channels than most would be comfortable wading through. As controllers for such channels, many services wish to aid their users in finding familiar and comfortable interactions. Users may also seek to participate outside their immediate circles, but may aim not to stray too far.

Problem

When many users are able to interact in the interaction space, it is hard to maintain an overview of relevant interaction partners since the number of users exceeds the number of relevant contacts for a specific user. User lists grow very large and it is hard to find people who the local user knows. On the other hand, the local user is [more interested in close contacts].

A service aims to provide users with shortcuts to interaction with users who they are most likely to interact with within a particular context (close contacts within social circles).

Forces and Concerns

- Large socially oriented or interaction-oriented mediums often hold more participants overall than any one user can manage
- Users want to interact in a way which is familiar and comfortable, most likely with people they know
- Users want to get to using the service without being blockaded by walls of text, but the also do not want to be blindsided about policy
- Some users aim to make new interactions with people bordering their friend circles, or sharing connections

Goal

G

Solution

Allow users to find and assign others to a user-maintained directory of social circles and contexts to interact with. This is optionally only visible to the users themselves.

Implementation

Users should be able to view the Buddy List on demand, either during a search operation or persistently. They should be able to add or remove users from the relevant list with minimal effort.

The list may be seen as a set of user objects. This buddy list has the possibility of adding or removing user objects. In the first case, whenever the local users interact with another user, they can add the other user to their buddy list. To reach this goal, in the user interface, the local users can select the representation of the another user and execute a command for adding (e.g. a menu item associated to the user object). For removing users, when the buddy list is shown, the local users can select the representation of the another user and execute a command for removing (e.g. a menu item associated to the user object).

Extending Functionality

The Buddy List may fuse with other common interaction idioms to constitute a more comprehensive approach to the problem, making it more than an idiom.

- The list may extend to the full User Gallery during a search operation, listing 'buddies' distinctively before the rest of the userbase;
- Common connections or nearby outliers can be suggested to the user, both during search and while viewing the list itself;
- The list may indicate the activity or status of each user, as a User List, additionally doing so where consented for users outside the list;

- Users who also list the local user in their Buddy List may be indicated, perhaps even when not explicitly in the local user's list; and
- Users may choose to block other users from seeing them.

Constraints and Consequences

C

Benefits

Connecting the means for adding users to the buddy list with the user's representation (or the interface elements that are used to interact with the other user) makes the process of adding a user to the buddy list intuitive and reminds a user to consider adding the user.

using the Buddy List to make connections about the user, the service can recommend relevant contact suggestions.

Liabilities

If users only consider buddy lists for maintaining contacts to other users, they will hardly find new users in the system. Thus you should ensure that users can also browse other users who are not on their buddy list (e.g. by providing a User Gallery).

The service can trivially derive the social structure of its userbase which may put trust at jeopardy.

Motivating Scenario

Scenario

- Instant Messaging Systems
- Email address books and mailing lists
- Reddit Subreddits
- Facebook Friends
- LinkedIn Connections

Know Uses and Related Work

Categories

- Control
- Choose

Related Patterns

- P Masquerade
- P Reciprocity
- P Incentivized Participation
- P Lawful Consent

Supporting Patterns

- P

Sources

- https://privacypatterns.org/patterns/Buddy-List

2.57 Privacy Awareness Panel

Summary

Context

Numerous services (and products) make an impact on user privacy in ways which are not immediately apparent to the user. Unaware and thus uninformed users are likely to make regrettable decisions in the services they use. Certain kinds of information, especially when combined or viewed over time by others, can reveal details about the user they did not intend. The consent for these disclosures cannot be valid if they do not understand the risks inherent in doing so. Controllers of such personal data therefore seek to minimize these risks.

Problem

Users do not anticipate the pitfalls of disclosure. They may be under the false impression that their activities are inherently anonymous.

This can manifest in the use of online services where a user shares information with an unknown audience using a pseudonym. Entities within can potentially discover detail the user does not intend, especially if the user loses track of who knows or has access to what. Providing publication history, or reusing aliases in various services, for example can have unintended consequences.

Furthermore, the controller themselves typically has more capability for identifying the user. If users do not know any better, they might behave or contribute in a manner which assumes they cannot be identified.

Forces and Concerns

- Users sometimes want to use services without being identified, but do not know how to maintain their pseudonymity
- Users want to understand what using a service might reveal about them to various parties
- Controllers want to protect users from unknowingly mak-

ing disclosures which are invasive

- Controllers do not want to process any personal data without informed consent

Goal

G

Solution

Provide the user with reminders on who can see the content they have or will disclose, what is done with it, why, and how it might become identifying.

Structure

First, it should be made clear to users which persons will be able to access their contributions. Second, users should know that [controllers] get additional information about them for instance their IP addresses, browser versions, location information etc. and thus that they are not completely anonymous [within] the [service].

Implementation

The potential consequences of content disclosure may depend on the service in question, and should be investigated in a general sense.

The user does not need to be shown every potential consequence, but rather must be aware of the need to consider their submission before disclosure. This may require access to an illustrative example to assist in conveying the risks in an accessible manner.

Prior to disclosure, controllers should primarily indicate the access capabilities of different types of users and entities. For example, those on a Buddy List, or unauthenticated users. Entities include themselves, their processors, and any third parties. Wherever this might entail personal data, purposes and means are also required before informed consent to the submission. If the user is already aware of these, reminders need not be as frequent and prominent.

Constraints and Consequences

C Improved awareness of users about who exactly will be [and has been] able to see the [content they disclose] will hopefully make them consider [disclosure] more carefully.

Motivating Scenario

Scenario In a forum setting, a Privacy Awareness Panel may include login and account information, any personalizations, as well as information relating to their browser, session, IP, or other metadata which can uniquely identify them to a degree. It could also show post and user interaction history, and what, if any, of this information is more widely available or public. The panel should be easily located and known about by users, for instance introduced on first use of the forum. Unauthenticated users should also have access to this panel, though there would be less information on these users. ■

Know Uses and Related Work

Categories

- Inform
- Provide

Related Patterns

P Awareness Feed

P Who’s Listening

P Impactful Information and Feedback

P Appropriate Privacy Feedback

P Increasing Awareness of Information aggregation

Supporting Patterns

P

Sources

- https://privacypatterns.org/patterns/Privacy-Awareness-Panel

2.58 Lawful Consent

Summary

This pattern covers in detail the legal and social obligations surrounding a data subject's consent to processing of their data in specific circumstances. Every use of the subject's personal data should be covered by an explicit agreement in which the data subject was made aware of the implications of their consent.

Context

Where data controllers (e.g. organisations) aim to provide a service (or product) to users, there may be opportunities to reuse data, gather feedback, or make use of user data to further their system's value. Many controllers seek to continually collect and utilise this data, often in ways which warrant privacy concerns. For any data processing (including collection), controllers should first obtain consent from the users in question.

There are social norms surrounding the use of personal data which need to be adhered to if an controller wishes to avoid scrutiny. Users do not inherently trust controllers to handle their personal data with care for privacy. Without clearly defined boundaries, these users may have justifiable concerns about what is learned about them, and how this information may be used. Additionally, various jurisdictions supply varying compliance requirements, and these controllers need to cater to every market they provide to.

Doing otherwise, possibly by disinterest or negligence, may have financial consequences in addition to potential public outcry. Despite this, controllers regularly consider the impact that their decisions may have on competitive edge and resulting profits. The link between better decision making, possibly less sharing, and reduced monetary gains sways some controllers into unlawful forms of consent.

Concerned controllers aim to promote trust in any number of ways, potentially including an Awareness Feed and or Privacy Dashboard to properly inform their users. The controller in this context may wish to adhere to the corresponding laws for their users, or above that, genuinely value their users' rights to self-determination.

Problem

A controller aims to maximise the value of their services by gathering as much sharing and participation as possible, potentially seeing user consent as a barrier to functionality and efficiency. They may inadvertently subvert notions of consent by unnecessarily bundling together desirable services with needs for personal information, or downplaying the significance of the data involved. They undermine self-determination at the risk of losing trust from their users, and attracting legal investigations which may rule their practices unlawful.

Forces and Concerns

- Controllers want to encourage participation, and thus may be less concerned with investigating or revealing tradeoffs
- Controllers may be tempted to bundle various services under a single broad consent request, pressuring users into agreements they might not otherwise accept
- Users often want to make use of new and exciting features, and therefore easily overlook downplayed privacy risks
- • Some users avoid certain services as they realise the potential privacy risks are not being acknowledged

Goal

G

Solution

A user should be given every opportunity to assess their sharing choices prior to making their consent. The controller should aid the user in comprehending the tradeoffs apparent in using each of their services, without over-burdening the user. These consented services should be purposed-separated, so that users may make use of functionality without first granting unnecessary consent.

Rationale

Controllers need to ensure that anything they do with a user's sensitive or potentially identifying data is legal. This pertains to lawfully obtained consent, for purposes which are clear and lawful in their own right. Additionally, anything they do should be resistant to backlash from users.

Implementation

Separate Purposes

Services should be separated into distinct processes for which distinct consent is acquired. Each purpose requires its own consent. These permissions need to be given subsequent to ascertaining sufficient awareness in the user about the consequences of that consent.

Given Consent

The users should not be pressured into providing consent. Instead, the benefits may be presented along with the trade-offs so that the users may make an informed decision. Some users are not necessarily capable of making these decisions themselves (e.g. children) and thus provisions need to be made to cater to this. The provided information should not be misleading, as coerced consent is not a valid form of permission. One way to present policies in an accessible manner is through comparative examples (e.g. in addition to further detail, what is unique about our privacy policy?).

Providing too much information may also intimidate users into making uninformed decisions, and thus awareness must be garnered in a way which is broadly accessible (see Awareness Feed). Opportunities for further reading should be available, though should not be necessary to understand the trade-offs involved.

Personalized Negotiation

In more personal services (i.e. one-on-one), personal privacy policies may undergo a formal negotiation. As opposed to user preferences (both at sign-up and through appropriate defaults), understanding a user's personal privacy requirements may benefit from the facilitation of a human representative. This, however, suffers from it's own drawbacks where the representative may misunderstand the user's requirements. Even in interpersonal exchanges, controllers should err on the side of caution. Where available, explicit signing of an agreement aids in proving consent.

Constraints and Consequences

C With the ability to choose exactly what tradeoffs are agreeable to them, users will be more content, and trusting of the system. They may as a result use more services, and participate more than they otherwise would. Being aware of what information is actually needed to perform certain functionality may also prevent its use, but rightfully so as to prevent backlash. The need for certain information for some services will bring inappropriate business processes to the foreground to be rectified, or otherwise questioned. This will likely bring the controller towards better practices, and may affect others as well. Once the public sees the controller's willingness to cooperate, trust will grow even further.

C Overall adoption will grow for controllers who are shown to be trustworthy and upfront about their data processing practices. This may very well offset the costs involved in maintaining transparency.

Constraints

Allowing informed and specific consent prevents controllers from soliciting misplaced consent, which greatly reduces the adoption of invasive services. These are often the most profitable services.

Motivating Scenario

Scenario ■

Know Uses and Related Work

Categories

- Control
- Consent

Related Patterns

- P Informed Consent for Web-based Transactions
- P Obtaining Explicit Consent
- P Sign an Agreement to Solve Lack of Trust on the Use of Private Data Context
- P Decoupling [content] and location information visibility
- P Discouraging blanket strategies
- P Single Point of Contact
- P Buddy List
- P Reasonable Level of Control
- P Outsourcing [with consent]
- P Negotiation of Privacy Policy
- P Private link

P Active broadcast of presence

P Pay Back

P Reciprocity

P Selective Access Control

P Enable/Disable Functions

P Incentivized Participation

P Support Selective Disclosure

Supporting Patterns

P

Sources

- https://privacypatterns.org/patterns/Lawful-Consent

2.59 Privacy Aware Wording

Summary

Context

Users are exposed to many privacy policies and notifications which seek to inform them of various issues. The controllers who provide these explanations require that users fully understand the circumstances around the use of their data. Specifically, the purposes for which and means by which their personal data is collected or otherwise processed. There is much information however, and so users are likely to overlook important details.

Problem

Information the controller conveys to the user is frequently overlooked due to length and complexity of both the content and the vocabulary within, which compromises validity of consent.

Users should clearly understand the content of and terms used within privacy and security software. The terms are usually formulated on an expert-basis and therefore often difficult to understand for the average user.

Forces and Concerns

- Users do not want to read complex and long policies
- Users still want to understand what risks they might be taking with their data by using the service (or product)s
- Controllers want to ensure that users understand risks
- Controllers need consent given by users to be informed

Goal

G

Solution

> Construct privacy related information using easily parsed and low difficultly vocabulary, with short concise sentences and

enough flow to persuade the user to process it.

Implementation

Users should not need to be familiar with the subject matter. They should also not be given unnecessary detail at the highest level of abstraction. Consider combining techniques from other patterns such as Layered Policy Design.

Before using the terms, one should be sure that they are clear and understandable for the target-users. Therefore, it is recommended to either refer to standardized terms [or] to conduct user tests on the understandability of [utilized] terms and phrases. These tests do not have to be extensive. Asking only few representative users from the target-group about their understanding of the terms should suffice.

Constraints and Consequences

C

Motivating Scenario

Scenario Referring to the user as the data subject or otherwise introducing terms to the user may reduce reading comprehension. Instead of focusing on legally accurate terms, the information should make sense to the user. It should not be provide a false interpretation, however. The PrimeLife example features a mock corporation which summarises information according to 'what', 'how', and 'who'. ■

Know Uses and Related Work

Categories

- Inform
- Explain

Related Patterns

P Awareness Feed

- P Appropriate Privacy Icons
- P Icons for Privacy Policies
- P Privacy Labels
- P Privacy Color Coding
- P Abridged Terms and Conditions
- P Privacy Policy Display
- P Layered Policy Design
- P Privacy-Aware Network Client
- P Impactful Information and Feedback
- P Dynamic Privacy Policy Display
- P Policy Matching Display
- P Privacy Policy Display

Supporting Patterns

- P

Sources

- https://privacypatterns.org/patterns/Privacy-Aware-Wording
- http://privacypatterns.wu.ac.at:8080/catalog/

2.60 Sticky Policies

Summary

Machine-readable policies are sticked to data to define allowed usage and obligations as it travels across multiple parties, enabling users to improve control over their personal information.

Context

Multiple parties are aware of and act according to a certain policy when privacy-sensitive data is passed along the multiple successive parties storing, processing and sharing that data.

Problem

Data may be accessed or handled by multiple parties that share data with an organisation in ways that may not be approved by the data subject.

Goal

G The goal of the pattern is to enable users to allow users to control access to their personal information.

Solution

Service providers use an obligation management system. Obligation management handles information lifecycle management based on individual preferences and organisational policies. The obligation management system manipulates data over time, ensuring data minimisation, deletion and notification to data subjects.

Constraints and Consequences

C Benifits:Policies can be propagated throughout the cloud to trusted organisations, strong enforcement of the policies, traceability. Liabilities: Scalability: policies increase size of data. Practicality may not be compatible with existing systems. It may be difficult to update the policy after sharing of the data and existence of multiple copies of data.

It requires ensuring data is handled according to policy e.g. using auditing.

Motivating Scenario

Scenario When data is shared by an organisation they can use privacy preserving policy to enforce respecting user privacy by third party organisations that use, process and store such data. For example, a hospital may share data with third party organisations requiring adhering to specific privacy policies associated with the data. ■

Know Uses and Related Work

Examples of policy specification languages include EPAL, OASIS XACML and W3C P3P. Tracing of services can use Identifier-Based Encryption and trusted technologies. An alternative approach using Merkle hash tree has been proposed by Pohls (2008). A Platform for Enterprise Privacy Practices (E-P3P) (2003) distinguishes the enterprise-specific deployment policy from the privacy policy and facilitates the privacy-enabled management and exchange of customer data. References: Pearson, S., Sander, T. and Sharma, R., Privacy Management for Global Organisations, Data Privacy Management and Autonomous Spontaneous Security, LNCS 5939, Springer, pp. 9-17., 2009 Phols, H.G., Verifiable and Revocable Expression of Consent to Processing of Aggregated Personal Data. ICICS, pp.279-293, 2008 Karjoth, G., Schunter, M., & Waidner, M., Platform for enterprise privacy practices: Privacy-enabled management of customer data. In Privacy Enhancing Technologies, pp. 69-84, Springer Berlin Heidelberg, 2003.

Categories

- Privacy-Policy
- Enforce
- Uphold

Related Patterns

P

Supporting Patterns

P

Sources

- https://privacypatterns.org/patterns/Sticky-policy
- https://privacypatterns.eu/#/patterns/sticky-policies/0-0-0-0-0-0-0-0-0-0-0-0-0-0-0-0
- http://privacypatterns.wu.ac.at:8080/catalog/

2.61 Personal Data Table

Summary

Context

Controllers which maintain software systems that process user data, especially identifying or sensitive data, are subject to various laws. In the case of personal data, transparency about processing is particularly important. Users (the data subjects) also care to know about what data is used, and what might be done with that data, at various degrees. Users do not often want to be constantly notified or reminded, as many of them would rather spend their time actually using the system. Some users, however, care about more intricate detail, and are entitled to it. Nonetheless, if verbose information is provided, it should be sensible.

Problem

The controller wants to be upfront about what they know and can do with personal data which might be of importance to those users. They only want users to know about data and risks pertaining to them specifically.

Forces and Concerns

- The controller wants to show the actual data they process, as well as what they do with it, as opposed to just describing policy
- Users want full transparency, with detailed explanation as well as easily and quickly understood overviews
- Controllers do not want this transparency to ruin trust, but to strengthen it
- The controller wants to keep the data on their servers, while still allowing users to automatically view their own data

Goal

G

Solution

Keep track of the processing that occurs on personal data so that users can view the activities associated with their data and review their preferences in a tabular environment.

Structure

Structure

Which information A table that shows the overview. The overview could show: Which data Why collected How used/for which purpose collected Who has access to the data Who the user authorized for access Which consent the user has given for specific data To which parties the data is disclosed Who has seen the data Whether the data can be hidden Whether the data can be removed How long the data is stored How datasets are combined to create richer (privacy sensitive) information. Note that this may violate local laws and regulations With which other information the data is combined

Where in the application flow Options are (not mutually exclusive): At the service's help section At the service's privacy section Through a separate menu item At a myData section of the service

Amount of information A table can show a lot of information or can be adjustable by the user to tweak which information to show, and which values (e.g. which range). From the table links to applicable other pages/screens can be given, to allow a user to easily change privacy settings (or possibly delete data) or visit websites of data buyers. A way to present more detail than visible at the overview table is to apply the Overview beside detail user interface pattern (Laakso 2003).

Implementation

Provide users with access to an interface which displays their data in useful dataset views, and give them the option for raw information. See the following table for an example.

|Type of Data|Data|Date Recorded|Accessed by| |–|–|–|–| |data

> type a|data itself|date a|person one| |data type b|data itself|date b|person one, person two|
>
> To be really transparent, also show things like how and why data was used, who of your organization has access to the user's personal data, what was downloaded or sent to a specific third party, and when all these events happened. The table can present all the data at once, or order it in categories, that may be further detailed when the user selects a category.

Constraints and Consequences

Benefits: - Actual data: Users can see the actual personal information you have, real-time*. - Details: A table can show all the personal information at once, in a structured way. - Details: Users may see errors in their data and ask for rectification, thereby improving the data quality. - Security and Availability: Data can remain stored in a secure place and still be available to your users whenever they want to see it. - Usability: Users get a better understanding of the personal data your systems holds and how you handle their data. Users may even decide to better control access to their data (not part of this pattern), increasing their own privacy and limiting the risks of privacy incidents, caused by e.g. an external attack on the system. - Trust: Providing transparency in a user-friendly manner increases the trust that users have of you as an organization. - Automation: A table is relatively easy to implement and automatically generate, compared to for example graphic data visualisations.

Liabilities: - Actual data: Everything that happens to all user data must be logged. This may impose a privacy problem of its own. - Details: Users may be overwhelmed by the amount of data you have and decide to stop using your system - Security and Availability: Some users may want to have the data on their own systems, for example to run their own analyses. This pattern does not make that possible, but the functionality could be easily provided. Trust: Users may decide to delete some of their data or otherwise restrict access to their data in a way that decreases the amount of data available for your systems. Or users may even decide that the privacy infringement is too large and stop using your system all together. *Providing real-time insight in all personal data that a system contains is not common practice; currently people usually have to put in a formal request (e.g. by email) and wait for a couple of weeks until they receive a reply with zip-file attached.

Motivating Scenario

Scenario ■

Know Uses and Related Work

> Figure 1 shows the actual design of the personal data table pattern implementation for a Quantified Self data store, the Nutritional Research Cohort (NRC). The NRC is a cohort of researchers in nutrition and health sciences who gather self-assessment data on their lifestyle and their health. NRC gives access to information on personal health trajectory, and the effects of diet on personal health. For each column, a mouse overlay details the meaning of the column name. This solution implements an overview of which data is collected, whether data is private or shared with others, for which purpose the data is used, which external parties requested the data, and who downloaded the data and when. This overview is shown on a special page in a myData section of the NRC application.

Categories

- Transparency
- Access
- Inform
- Provide

Related Patterns

P Minimal Information Asymmetry

P Privacy Mirrors

Supporting Patterns

P

Sources

- https://privacypatterns.org/patterns/Personal-Data-Table

2.62 Informed Consent for Web-based Transactions

Summary

Context

User data is frequently collected for various purposes. Sometimes this data is personal, personally identifying, or otherwise sensitive. The data may serve to improve a service (or product) offered by a controller, or to provide relevant suggestions or advertisements to users. This is particularly prevalent on the web, as many websites derive most of their income from this data. Where income is instead in the form of purchase, user data is nonetheless needed to provide billing or shipping information. This includes auditing, logging, or other non-repudiation purposes to facilitate transactions.

Problem

Before collecting data, controllers must make sure users provide informed consent.

Controllers utilize persistent local or server-side storage to process potentially identifiable or sensitive information about users in order to perform a transaction. However, users are often resistant to disclosing personal information because they are uncertain if it will be used without their consent or against their interests.

Controllers need to be able to inform their users about these purposes and means before the user consents.

Forces and Concerns

- Users want to visit websites and make use of the services (or products) offered, but do not want their privacy to be undermined
- Users want to have control over their personal information
- Controllers may need to process data to conduct business, and may in some cases deny service to those who withhold their data
- Controllers may profit from additional user data, and users may too enjoy enriched services

- They wish to protect their users from privacy violations, and protect themselves from responsibility, but also aim to secure a viable business model

Goal

G

Solution

Provide the user with clear and concise information regarding what may be learned from their data, and how that data can be used to offer or improve the service. Then acquire their explicit, freely-given consent.

Structure

To the extent possible given the limits imposed by web technology, provide the user with the six elements of informed consent: Disclosure [of purpose specification and limitation,] Agreement [and disagreement capabilities,] Comprehension [through easily understandable, comprehensive and concise explanations,] Voluntariness [showing that consent is freely-given,] Competence [to make reasonable legally binding decisions, and] Minimal Distraction [which may otherwise aggravate the user].

Implementation

Human Computer Interaction concepts expressed in the work of Fischer-Hübner et al. (2010) allow implementing this pattern in various ways:

- Just-In-Time-Click-Through Agreements (JITCTAs), i.e. click-trough agreements that instead of providing a large list of service terms confirm the user's understanding or consent on an "as-needed basis". The information shown in JITCTAs includes what data is requested, the controller's identity and the purpose of processing.
- Selection via cascading context menus, where users have to choose more consciously the menu options of data to be released. This option is intended for simple data request

forms with not many fields to be filled.

- Drag-and-Drop Agreements (DADAs), which also requires user to make more conscious drag and drop actions for consenting to data disclosures. The user has to choose an icon that represents some kind of personal data and drag and drop it to an icon representing the controller.

Constraints and Consequences

Benefits

- Helps to reduce information asymmetry between the user and the [controller].
- Empowers users to make informed decision that do not conflict with their tolerance for private information disclosure.
- Provides a basis for trust between the consumer and website owner by establishing an expectation of practice by the website. Consider the risk of lost trust for ecommerce, medical and financial companies such as eBay, Amazon, Bank of America, ehealthinsurance.com, etc..
- This pattern can be applied to many other systems that interact with the user and external systems such as email and location aware devices (e.g. cellphones, PDAs).

Liabilities

- This pattern cannot provide any assurance that a website will comply with the informed consent model.
- Privacy policies are generally known to be confusing for the user to read and fully understand.
- The website may not wish to disclose their ability to track users without their knowledge.
- The website may not have the infrastructure to offer and support each of the solution elements for every user. For example, the ability for users to opt-out of the agreement.
- If the distraction due to implementing this pattern is sufficiently great, the user may simply cancel the transaction altogether.
- Information provided to gain consent is necessarily a) limited and b) manipulated by the site to obtain consent – this implies that the actual consequences of the revelation of personal information may remain unknown to the user.

Motivating Scenario

Scenario ■

Know Uses and Related Work

- Yahoo! Registration Form
- Intuit Registration Form
- Google Registration Form

Categories

- Inform
- Control
- Consent

Related Patterns

P Informed Consent for Web-based Transactions

P Lawful Consent

P Obtaining Explicit Consent

P Privacy Policy Display

Supporting Patterns

P

Sources

- https://privacypatterns.org/patterns/Informed-Consent-for-Web-based-Transactions

2.63 Added-noise measurement obfuscation

Summary

Add some noise to service operation measurements, but make it cancel itself in the long-term.

Context

A service provider gets continuous measurements of a service attribute linked to a service individual.

Problem

> The provision of a service may require repeated, detailed measurements of a service attribute linked to a data subject to e.g. properly bill them for the service usage, or adapt the service according to the demand load. However, these measurements may reveal further information (e.g. personal habits, etc.) when repeated over time.

Goal

G A service provider can get reliable measurements of service attributes to fulfil its operating requirements; however, no additional personal information can be inferred from the aggregation of several measurements coming from the same user.

Solution

> A noise value is added to the true, measured value before it is transmitted to the service provider, so as to obfuscate it. The noise abides by a previously known distribution, so that the best estimation for the result of adding several measurements can be computed, while an adversary would not be able to infer the real value of any individual measurement. Note that the noise needs not be either additive or Gaussian. In fact, these may not be useful for privacy-oriented obfuscation. Scaling noise and additive Laplacian noise have proved more useful for privacy preservation.

Constraints and Consequences

The pattern applies to any scenario where the use of a resource over time is being monitored (e.g. smart grid, cloud computing). The device providing the measurement must be trustworthy, in order to ensure that it abides by the established noise pattern.

Some information is lost due to the noise added. This loss of information may prevent the information from being exploited for other purposes. This is partly an intended consequence (e.g. avoid discovering user habits), but it may also preclude other legitimate uses. In order for information to be useful after noise addition, the number of data points over which measurements are aggregated (i.e. the size of the aggregated user base) needs to be high; otherwise, either the confidence interval would be too broad or differential privacy could not be effectively achieved.

Constraints

The pattern applies to any scenario where the use of a resource over time is being monitored (e.g. smart grid, cloud computing). The device providing the measurement must be trustworthy, in order to ensure that it abides by the established noise pattern.

Some information is lost due to the noise added. This loss of information may prevent the information from being exploited for other purposes. This is partly an intended consequence (e.g. avoid discovering user habits), but it may also preclude other legitimate uses. In order for information to be useful after noise addition, the number of data points over which measurements are aggregated (i.e. the size of the aggregated user base) needs to be high; otherwise, either the confidence interval would be too broad or differential privacy could not be effectively achieved.

Motivating Scenario

Scenario An electric utility operates a smart grid network with smart meters that provide measurements of the instantaneous power consumption of each user. The utility employs that information to both adapt the power distribution in a dynamic fashion,

> according to user demand at each moment, and bill the each client periodically, according to his aggregated consumption over the billing period. However, this information can also be exploited to infer sensitive user information (e.g. at what time he or she leaves and comes back to home, etc.). ■

Know Uses and Related Work

- Bohli, J.-M.; Sorge, C.; Ugus, O., "A Privacy Model for Smart Metering," Communications Workshops (ICC), 2010 IEEE International Conference on , vol., no., pp.1,5, 23-27 May 2010
- Xuebin Ren; Xinyu Yang; Jie Lin; Qingyu Yang; Wei Yu, "On Scaling Perturbation Based Privacy-Preserving Schemes in Smart Metering Systems," Computer Communications and Networks (ICCCN), 2013 22nd International Conference on , vol., no., pp.1,7, July 30 2013-Aug. 2 2013
- Mivule, K. (2013). Utilizing noise addition for data privacy, an overview. arXiv preprint arXiv:1309.3958

Categories

- Minimize
- Hide
- Obfuscate

Related Patterns

P aggregation-gateway

P trustworthy-privacy-plug-in

Supporting Patterns

P

Sources

- `https://privacypatterns.org/patterns/Added-noise-measurem`

- https://privacypatterns.eu/#/patterns/added-noise-measurement-0-0-0-2-0-1-0-3-3-0-0-1-1-0-0-0

2.64 Information Aggregation Awareness

Summary

Context

Controllers process mass amounts of user data in order to provide enhanced services (or products). Aggregating this data with other sources unlocks new insights which could not be determined alone. This kind of aggregation is distinct from the notion of abstracting information away from personal data, effectively making it less sensitive. Instead this may turn seemingly harmless data into identifying, intrusive, or inferred information, some of which not even the user is aware of. This makes aggregation very useful for marketing, as well as other more usability-centric features, but places a heavier burden on users to disclose with care.

Problem

Poor awareness of data aggregation capabilities can lead to unintentionally revealing information being disclosed. Processing this personal data goes against the principles of data protection.

Forces and Concerns

- Users do not want to inadvertently disclose information which may become sensitive or identifying
- Users are less familiar with the risks of information only becoming invasive sometimes long after disclosure
- Controllers do not want users to unwittingly consent to disclosures they later regret due to poor awareness
- Controllers want users to be cognizant of the sensitivity and contextual applicability of their disclosures and how these may be changed by aggregation

Goal

G

Solution

Provide users with knowledge of data aggregation's ability to reveal undesirable information to prevent them from over sharing. Take users through a hypothetical example to aid in conveying this.

Implementation

Prior to allowing users to submit content to the controller or other parties, provide them with a warning about data aggregation. This warning is only necessary where aggregation is applied. As such if it is determined after collection that data should be aggregated, this warning would be given prior to obtaining consent for that further processing.

The warning must make it clear to the user that content they disclose may be more sensitive that it first appears. The context in which they provide it may be subject to changes, and these potential contexts should be provided to the user, or else consented to as they become applicable. The user should not have to deal with broad or otherwise unclear usage of their data.

At the same time, the user should not be exposed to deep, complex, and lengthy detail unless they choose to review it further. Instead, concise and clear explanations should be used. One approach to this is to provide a hypothetical example in which a controller reveals surprising characteristics about a user from combinations of data, which alone are less informative.

Consider the use of user tests to determine the level of clarity an explanation or example provides. It is important that if a user chooses to accept the risks (and benefits) of aggregation, then they do so knowingly. It is also important not to force aggregation onto users if they choose not to consent. This may prevent the user from gaining a feature, but should not lock them out of functionality which does not require it.

Constraints and Consequences

C If users understand the power of data aggregation better, they are better able to put any new data they're about to share in perspective to all the data they've already shared, and may consider the total picture this creates of them more carefully. But this also means that it becomes harder for organizations to create accurate profiles of people and may result in improper labeling based on the little data that is known.

Motivating Scenario

Scenario ■

Know Uses and Related Work

CryptPad Provides a thorough and clear explanation of their Data Aggregation usage which is linked to from the 'What is CryptPad' page in every instance. Towards the end of the blog post they include graphs to show how useful the data can be, but they also explain what they access, can (but do not) access, and what they cannot access. While this example explains aggregation well, and features a concise summary at the beginning, it could still be better highlighted before a user's first use of the service.

Categories

- Inform
- Provide

Related Patterns

P Awareness Feed

P Privacy Awareness Panel

P Appropriate Privacy Feedback

P Impactful Information and Feedback

Supporting Patterns

P

Sources

- https://privacypatterns.org/patterns/Increasing-Awareness-of-Information-Aggregation

2.65 Attribute Based Credentials

Summary

Attribute Based Credentials (ABC) are a form of authentication mechanism that allows to flexibly and selectively authenticate different attributes about an entity without revealing additional information about the entity (zero-knowledge property).

Context

ABC can be used in a variety of systems including Internet and smart cards.

Problem

Authentication of attributes classically requires full and unique authentication of an entity. For example, attributes (like age) could be put into a certificate together with name of the user, email address, public key, and other data about that entity. To corroborate an attribute (for example, that the user is an adult) the certificate has to be presented and all information have to be revealed. This is not considered a privacy-preserving solution.

Goal

G To allow a user to selectively prove specific attributes like age > 18 to a verifying party without revealing any additional information.

Solution

There are multiple schemes to realize ABCs and implementations are also available. They typically all include a managing entity that entitles issuers to issue credentials to entities that could then act as provers of certain facts about the credentials towards verifiers.

A formal model can be found here. [http://sit.sit.fraunhofer.de/smv/pattern-models/ABC-pattern-model.pdf]

Constraints and Consequences

C ABC schemes require substantial compute power or optimization, so implementation may not be straightforward. Some projects like IRMA developed at Radboud University Nijmegen have shown that even resource restricted devices like smartcards can implement ABCs.

Motivating Scenario

Scenario You want to issue an ID card that holds a users birthdate bd and can be used to prove that the card holder is old enough to view age-restricted movies in a cinema. Depending on the rating of the movie (minimum age x), the card holder can run a proof that:
"today - bd > x"
Multiple uses of the card at the same cinema should not be linkable.

■

Know Uses and Related Work

The most popular implementations include:

- IBM's IDEMIX developed as part of the PRIME/PRIMELIFE project
- Microsoft's U-Prove
- Radboud University Nijmegen's IRMA project

Categories

- Anonymity
- Authentication
- Zero-Knowledge
- Minimize
- Hide
- Restrict

Related Patterns

P

Supporting Patterns

P

Sources

- https://privacypatterns.org/patterns/attribute-based-cred
- https://privacypatterns.eu/#/patterns/attribute-based-cre 0-0-0-0-0-0-0-3-0-0-0-0-0-0-0-0
- http://privacypatterns.wu.ac.at:8080/catalog/

2.66 Trustworthy Privacy Plug-in

Summary

Aggregate usage records at the user side in a trustworthy manner.

Context

A service provider gets continuous measurements of a service attribute linked to a service individual. Applicable service tariffs may vary over time.

Problem

The provision of a service may require repeated, detailed measurements of a service attribute linked to a data subject to e.g. properly bill them for the service usage. However, these measurements may reveal further information (e.g. personal habits, etc.) when repeated over time.

Goal

G

Solution

Host a Privacy Plugin at a consumer-trusted device, in between the metering and the billing systems and the service provider in charge of billing for the service usage. This privacy plugin, under the consumer's control, computes the aggregated invoice and sends it to the service provider (or to its billing subsystem), which does not need any fine-grained consumption records anymore. Cryptographic techniques (homomorphic commitments, zero-knowledge proofs of knowledge, digital signatures) are used to ensure trustworthiness of the generated invoices without requiring tamper-proof hardware.

Constraints and Consequences

C The service provider does not need anymore to access detailed consumption data in order to issue reliable bills.

Motivating Scenario

Scenario An electric utility operates a smart grid network with smart meters that provide measurements of the instantaneous power consumption of each user. Depending on the power demand, dynamic tariffs are applied. The utility employs that information to bill each client periodically, according to his aggregated consumption over the billing period and the respective tariffs at each moment. However, this information can also be exploited to infer sensitive user information (e.g. at what time he or she leaves and comes back to home, etc.) ■

Know Uses and Related Work

- Alfredo Rial and George Danezis. 2011. Privacy-preserving smart metering. In Proceedings of the 10th annual ACM workshop on Privacy in the electronic society (WPES ’11). ACM, New York, NY, USA, 49-60.
- Rial, A., & Danezis, G. (2011, October). Privacy-preserving smart metering. In Proceedings of the 10th annual ACM workshop on Privacy in the electronic society (pp. 49-60). ACM.

Categories

- Aggregate
- Hide
- Restrict

Related Patterns

aggregation-gateway

added-noise-measurement-obfuscation

Supporting Patterns

user-data-confinement-pattern

Sources

- `https://privacypatterns.org/patterns/Trustworthy-privacy-plugi`
- `https://privacypatterns.eu/#/patterns/trustworthy-privacy-plug`
  `0-0-0-1-1-0-1-1-0-0-0-0-1-1-0-0`

2.67 (Support) Selective Disclosure

Summary

Many services (or products) require the collection of a fixed, often large, amount of personal data before users can use them. Many users, instead, want to freely choose what information they share. This pattern recommends that services Support Selective Disclosure, tailoring functionality to work with the level of data the user feels comfortable sharing.

Context

Controllers aim to design services to be both maintainable and extensible, though as a result blanket strategies are used to simplify designs. Users are individuals and do not always respond the same way to different approaches. Restricting user choice on processing displeases users, and bundling purposes for that processing conflicts with international law. Users want a service which works without the data they do not want to provide, even so far as effectively anonymous usage.

Problem

Controllers typically want to collect data by default, and tend to limit the diversity of their services, and the choices they provide, to encourage that. This goes against the best interests of the users, who have varying data collection tolerances.

The underlying issues are discussed in more detail below.

An All or Nothing Mindset

Controllers are tempted to see consent as all-encompassing, see held personal data as data available for use, and the lack of that data as a barrier to service. This mindset reduces adoption of the offering and may introduce a lack of trust.

The Temptation to Share by Default

User information is frequently acquired before users are given the opportunity to decide whether to share. An example of this is in cookie policies, where the whole site is loaded before the user is shown the policy. From this loading, metadata is often generated even if the user chooses to leave the site.

This problem is also present when users register for or acquire a service, as unnecessary information is often requested as part of the process. In the case of account registration users are often provided with inappropriate default settings. They are typically sent additional offers by default as well. The negative implications of these defaults are also not necessarily reversible, as the Internet is notorious for its inability to forget.

Data Gluttony

Services tend to collect a surplus of information, especially in contexts where monitoring is integral to the system, such as in productivity tracking. This unnecessary level of detail results in negative experience factors for the tracked individuals (for e.g. increased levels of anxiety) which in a work environment may affect their actual productiveness.

Forces and Concerns

- Controllers want their system to be applicable to as many potential users as possible, but do not want this to heavily inflate costs or jeopardise profits
- Users want to be able to use a system anonymously, or with as little leakage of their personal information as needed to perform functionality
- Controllers do not want users to be capable of malicious activity without consequence
- Many controllers want to benefit from the data they collect from their users, but do not want to violate their users' trust by using their information excessively or otherwise inappropriately

Goal

G

Solution

Determine what information is integral to the functioning of the system. If functionality may be sustained with less, it should be an option for the user, even if doing so comes with reduced usability. Additionally, provide anonymous functionality only where it cannot jeopardise the service. Lower levels of anonymity may be provided in relation to various capabilities for abuse.

Rationale

The key to a successful solution is meeting the correct balance between how little the system can work with and what is feasible in performance. This will affect the applicability of the system; a system which optionally functions with less will attract far more users. This increases the popularity of the system and therefore offsets a certain amount of additional implementation costs.

Implementation

Anonymous Usage

At one extent it may be possible to benefit from the system anonymously, though whether this is feasible will depend on the level of abuse anonymous usage might attract. Alternatively, this can be approached from the perspective of revocable privacy. That is, tentative or eroding anonymity. If this would result in an unsustainable business model, however, a re-balance of usability may be sufficient.

It is important to note that while anonymous usage might not translate into direct profit, additional contributors and positive public perception may increase overall user activity. Furthermore, there are payment methods which support [some level of] anonymity if necessary.

Assumption of Modesty

Where users choose to register, it should not be assumed that they wish to use all of the system's services. Short of explicitly opting for 'best experience settings' (with sufficient explanation; not the default option), user preferences should default to the most privacy-friendly configuration.

There exist numerous strategies for streamlining the preference-tailoring process, including gradual elevation as they begin to use services (see Awareness Feed and Lawful Consent).

The Right to Reconsider

User decisions should be amendable. For example, an agreement to share activity with another user may not carry over to all future usage. A user may decide to share something once in a while, or share regularly, but not always. The system should be able to account for this behaviour if it aims to prevent mistaken actions.

Look Before You Leap

In situations where there are requirements for personal data, particularly when strict, users should be aware of this prior to their consent. These services should also not be coupled with other services holding lower requirements unless it would be infeasible not to. Where users are required to use the system, no unnecessary information should be used. In a productivity tracking example, this may mean that users are only identified when their productivity falls, or perhaps if they opt to receive credit for their work.

Constraints and Consequences

C

Benefits

Due to increased control over their data, users may be able to share pieces of information which they otherwise wouldn't due to it otherwise being coupled with what they perceive to be more sensitive.

Users will be less likely to mistakenly release personal information to the public, since they would perhaps be able to set their own defaults, or by default stay private. To a further extent, users may be capable of participating or benefiting from a system anonymously. Where this is the case, the activity levels of the system will benefit, and users who stayed anonymous due to mixed feelings about the system may decide to register and authenticate later, once trust has been built.

Liabilities

The system's complexity will increase by a certain degree, as not only will each user need to have their preferences set, stored, and adhered to, but also services will need to account for variable inputs. As such, flaws in the system will be felt with greater effect.

Providing anonymity for some contexts may result in increased undesired behaviour, depending on the level of anonymity provided. Anonymising a service often requires additional processing power, especially in the case of revocable privacy.

Constraints

By separating functionality according to purpose and personal data needed, as well as providing variations where feasible, the system will be more complex. Services will need to be designed while taking into consideration the potential for limited access to data.

Improvements to results may therefore be limited as well. However, the controller may be able to gauge adoption in data-rich services while they are investing in them. The same holds for determining how valuable non-invasive alternatives are, as users will express their [in]tolerance for invasiveness through their actions.

Motivating Scenario

Scenario TUKAN; a collaborative software development environment which introduces Modes of Collaboration: lightweight contexts which filter collaboration possibilities according to user privacy preferences.

Anonymous access; to a degree, there exist many examples of websites which allow access to content without a need to identify users. Especially in cases where usage analytics are kept to a minimum, or tracking is disabled completely, users may use a service without a need to be monitored.

■

Know Uses and Related Work

Categories

- Control
- Choose

Related Patterns

- Masquerade
- Support Selective Disclosure
- Negotiation of Privacy Policy
- Buddy List
- Discouraging blanket strategies
- Decoupling [content] and location information visibility
- Reasonable Level of Control
- Selective Access Control
- Enable/Disable Functions
- Private link
- Lawful Consent

Supporting Patterns

P

Sources

- https://privacypatterns.org/patterns/Support-Selective-Disclosure

2.68 Private Link

Summary

Context

The controller provides a service which hosts resources, potentially constituting personal data. When users want to share (and enable re-sharing of) these resources, they may wish to do so privately using existing communication mechanisms. This is particularly relevant when users are sharing with contacts who would rather not, or cannot, simply authenticate.

Problem

Users want to share a private resource with unauthenticated users in a way that respects the sensitivity of that resource. The solution must not allow users to access resources that weren't intended to be shared, nor publicize the location of the intended resource to unintended recipients.

Forces and Concerns

- The controller should keep the links confidential in order to honor the user's privacy expectations
- The link should not be guessable (e.g. by convention or brute force) to prevent unintended recipients from accessing unlisted links
- The user should be able to limit the access to a resource by version or time restriction
- The recipient should be able to forward the link to another trusted recipient, so long as the link is valid
- The recipient should be able to access the link again at a later date, unless the resource is version or time restricted

Note that the URL will be retained in recipients' browser history and could easily be inadvertently shared with others. Services should help users understand these limitations.

Goal

G

Solution

Provide the user a private link or unguessable URL for a particular resource, such as a set of their personal information (e.g. their current location, an album of photos). Anyone who has the link may access the information, but the link is not posted publicly or guessable by an unintended recipient. The user can share the private link with friends, family or other trusted contacts who can in turn forward the link to others who will be able to access it, without any account authentication or access control lists.

Services may allow users to revoke existing private links or change the URL to effectively re-set who can access the resource. Additionally, users may set a time-limit for the resource's validity, or have it invalidated upon modification.

Implementation

The controller allows their users' online resources to be shared by publishing an unlisted URL with a complex, long, and randomly generated string. This can be part of a query string as opposed to an on disk location. In this case, the preprocessor intercepts the query and redirects the user to the correct resource. This may be an actual file on disk (probably not served by direct link), generated on the fly, or extracted from a database or compressed file. The preprocessor can verify validity dynamically before serving the resource.

The situation in which the user has a direct link to the resource location is not ideal, however, as it will need to change in the event of a time or version restriction since access to the file is not controlled by the preprocessor.

Constraints and Consequences

C

Motivating Scenario

Scenario

Know Uses and Related Work

> 1. Flickr "Guest Pass"
> 2. Google "anyone with the link" sharing
> 3. Tripit "Get a link"
> 4. Dropbox "Share Link"

Categories

- Distribution
- Media
- Access
- Control
- Choose

Related Patterns

Private link

Active broadcast of presence

Support Selective Disclosure

Masquerade

Decoupling content and location information visibility

Selective Access Control

Reasonable Level of Control

Lawful Consent

Supporting Patterns

P

Sources

- https://privacypatterns.org/patterns/Private-link

2.69 Anonymity Set

Summary

Context

This pattern is applicable in a messaging scenario, where an attacker can track routing information. Another possible scenario would be the storage of personal information in a database.

Problem

> In a system with different users we have the problem that we can often distinguish between them. This enables location tracking, analyzing the behavior of the users or other privacy-infringing practices.

Goal

G The goal of this pattern is to aggregate different entities into a set, such that distinguishing between them becomes infeasible.

Solution

> There are multiple ways to apply this pattern. One possibility is, to strip away any distinguishing features from the entities. If we do not have enough entities, such that the anonymity set would be too small, then we could even insert fake identities.

Constraints and Consequences

C One factor to keep in mind is that this pattern is useless if there are not many entities, such that the set of probable suspects is too small. What "too small" means depends on the exact scenario. Another factor is a possible loss of functionality.

Motivating Scenario

Scenario Assuming that there are two companies, one is a treatment clinic for cancer and the other one a laboratory for research. The Clinic releases its Protected Health Information (PHI) about cancer victims to the laboratory. The PHI's consists of the patients' name, birth date, sex, zip code and diagnostics record. The clinic releases the datasets without the name of the patients, to protect their privacy. A malicious worker at the laboratory for research wants to make use of these information and recovers the names of the patients. The worker goes to the city council of a certain area to get a voter list from them. The two lists are matched for age, sex and location. The worker finds the name and address information from the voter registration data and the health information from the patient health data. ■

Know Uses and Related Work

Anonymity sets are in use in various routing obfuscation mechanisms like Onion Routing. Hordes is a multicast-based protocol that makes use of multicast routing like point-to-multipoint delivery, so that anonymity is provided. Mix Zone is a location-aware application that anonymizes user identity by limiting the positions where users can be located.

Categories

- Hide
- Aggregate
- Anonymity
- Mix-Networks
- Obfuscation
- Mix

Related Patterns

P

Supporting Patterns

P

Sources

- `https://privacypatterns.org/patterns/Anonymity-set`
- `https://privacypatterns.eu/#/patterns/anonymity-set/0-0-0-0-0-0-0-0-0-0-0-0-1-0-0-0`

2.70 Active Broadcast of Presence

Summary

Context

Controllers provide an interface for acquiring information about the user. When one such user wants to share or broadcast their information, such as location or other presence data, that user may want to constrain the information. In this way, they may wish to prioritize data that is contextually relevant, or avoid a full stream of data which may be either noisy or intrusive. The controller wants the user to be able to provide this data at will, to maximize the applicability of their services. However, they do not want the user to regret providing too much data, nor to bother the user with constant requests.

Problem

A service aims to acquire or broadcast a user's real-time data, particularly presence or location information, to a platform (e.g. social network). They wish to do so without revealing sensitive data (e.g. private locations, histories, or health information) nor overwhelming recipients with noisy data or users with constant requests.

Forces and Concerns

- The controller wants to use the user's current data to provide more relevant information to the users of their service, but without violating the user's privacy
- The user wants to participate in the service and provide useful information, but not all information, as they consider some aspects more sensitive than others
- Users who intend to use the service do not want to have the service flooded with irrelevant data

Goal

G

Solution

Allow the user to actively choose when to share information, whether to broadcast it, and when not to. Assume that sharing settings do not apply holistically to all situations and seek clarification when in doubt.

Structure

The service may present distinct contexts in which to honor explicit settings, but in absence of this context assume that further consent is required. The user may choose not be be asked again, but must make this decision explicit.

Implementation

In addition to privacy settings with appropriate defaults, allow the user the option to be asked again, every time the context changes.

By default, users should actively choose to broadcast rather than the service deciding based on general settings which may not apply to the present context. Various contexts may be provided distinct settings.

In these situations users need only be reminded prior to setting the values themselves. After this, they may choose to be notified about broadcasting, but not about sharing with the service itself. In this way, the user may decide later.

Constraints and Consequences

C

Motivating Scenario

Scenario ■

Know Uses and Related Work

- Foursquare check-in model prior to Pilgrim
- Google services

Categories

- Location

- Mobile
- Control
- Update

Related Patterns

P Active broadcast of presence

P Reasonable Level of Control

P Masquerade

P Private link

P Lawful Consent

Supporting Patterns

P

Sources

- https://privacypatterns.org/patterns/Active-broadcast-of-presence
- http://privacypatterns.wu.ac.at:8080/catalog/

2.71 Unusual Activities

Summary

Context

Services (or products), particularly over the Internet, tend to use username and password based authentication. This security mechanism proves most convenient for users, as it is commonplace and simple compared to the more secure alternatives. It is also subject to common shortcomings, however. Passwords become less secure the longer they remain unchanged, are often vulnerable to brute force, snooping, and phishing attacks, and cannot be proven to be held solely by the user.

This complicates the certainty of the authentication, and thereby the authenticity of any decision made by the user, including consent. Controllers may also derive additional factors, however, such as device or access specific information. If location is provided, for example, it may hint at unlikely account activity.

Problem

Username and password authentication alone has varying reliability for proving decisions taken by a user, especially when concerning more sensitive actions. Controllers need to enhance their certainty that any consent provided is legitimate.

Forces and Concerns

- Users want to be able to authenticate easily and quickly, but also do not want controllers to accept decisions made by intruders
- Users want to know that their password is compromised, so that they can change it, especially if they use derivatives elsewhere
- Controllers want to protect user accounts from unauthorized access
- Controllers do not want to allow actions which the user did not truly consent to

A balance should be made between the insecurity of username and password authentication and the inconvenience of multi-factor authentication. If measures affect usability or privacy too greatly, users will stop using the system. While the rate of false positives must not be too high, they are far preferable to undetected intrusions.

- In the provided example, Facebook makes use of its resource of friendship and photos. Their decision is based on the assumption that it is very unlikely for a hacker to recognize the friends. Actually the assumption may not hold true in some scenarios, because many of the photos are public and can be viewed under another account, or can be identified with the help from a large-scale tagged photo collection and machine learning
- Persuading the user into carrying a hardware token everywhere only for occasional multi-factor authentication may be difficult, but it might worth the effort for financial services

Goal

G

Solution

Analyze the available information for which there is consent to establish an access norm. Test this against future access to identify unusual activities. When this occurs, alert the user and use multi-factor authentication while re-establishing certainty. The authenticated user should be able to review and take further action.

Implementation

Typically, a sign-in to a website is in the form of an HTTP request, which contains many customized settings of the browser, including the type of the browser and operating system as well as the architecture (User-Agentheader), the Cookie (Cookie header), language preferences (Accept-Languages header). Apart from these, the website can get the IP address of the user, which may be mapped to a certain country/area through GeoIP. [These] can be used to tell if a browser is ’new’ to the website. The website can have its rules to determine if an access is ’suspicious’, for example, an access from a new country / browser / operating system is considered suspicious.

By running native code, the application can [consensually] collect some [device identifiers], including the operating system environment settings (e.g. the list of running processes), the hardware parameters (such as the ID of the CPU), and device UUIDs (provided by mobile operating systems like iOS). By completing a network request, the service also retrieves the IP address of the [device]. [These] can be used to tell if a [device] is ’new’ to the service. The service can have its rules to determine if a sign-in is ’suspicious’, for example, an access from a new country / [device] / operating system is considered suspicious.

- Require Multi-factor Authentication
  In case of a suspicious [activity], multi-factor authentication may be a way to let the legitimate user in. The service can request [further authentication], such as:
- A software token Examples include Google Authenticator which runs on mobile phones and implements RFC6238 TOTP security tokens.
- A hardware token (disconnected) Examples include a token issued by a bank which displays digits, which is similar to a software token.
- A hardware token (connected) The token may exchange a longer secondary password than the previous one, which means it’s safer.
- Personal data like date of birth, [or civil identification]. Obviously not a good choice here because it cannot be changed.
- An one-time password (OTP) sent to the registered E-mail address / mobile phone Depending [on] the type of the service, [the user may use] the same password for the E-mail address, or [may lose their mobile phone].
  Using multi-factor authentication only in case of suspi-

cious [activity] is more convenient [than] using it all the time, but is less secure.

- Notify Account Holders of Unusual Activities
  When a suspicious sign-in is detected, it may be a sign that the password has already been leaked. Depending on the type of the service, it can notify the user about the suspicious sign-in through E-mail, telephone, or other means.

  Here the immediate notification can also be used in the multi-factor authentication. For services that can be logged on from multiple devices at the same time, the user should be able to check the existence of other sessions, and review recent [activity].

Constraints and Consequences

C Users will be able to use an easier, more familiar method of authentication in most scenarios, only having to resort to multi-factor authentication when there is potential cause for concern.

Constraints

This pattern has some limitations. For example, it relies on accurate identification of suspicious [activity] based on meta information, where the meta information including the IP address can be spoofed by an experienced attacker. If the fallback multi-factor authentication only happens occasionally to the legitimate account owner, they may be unprepared to [handle] such authentication, leading to [decreased] usability.

Motivating Scenario

Scenario 1. Gmail

- Gmail displays information about other sessions (if any) in the footer, linking to a page named "Activity on this account" which lists other sessions and recent activities to the Gmail account. The user has the option to sign out other sessions
- In case of annoying false positives, the user may

choose to disable the alert for unusual activity. The disable takes about a week, "to make sure the bad guys aren't the ones who turned off your alerts."

2. Facebook
    - When Facebook detects an unusual sign-in, it shows 'social authentication' that displays a few pictures of the user's friends and asks the user to name the person in those photos.
3. Dropbox
    - The 'Security' tab of the 'Settings' of the Dropbox website displays all web browser sessions logged in to the account, and enables the user to log out one or more of them. The name of the browser, operating system, and the IP address and corresponding country are displayed to help the user make a choice.
    - It also displays all devices that are linked to the account, and allows the user to unlink one or more of them.

Know Uses and Related Work

Categories

- Notice
- Authentication
- Inform
- Notify

Related Patterns

Data Breach Notification

Informed Secure Passwords

Impactful Information and Feedback

Supporting Patterns

P

Sources

- https://privacypatterns.org/patterns/Unusual-activities

2.72 Strip Metadata

Summary

Metadata that is not needed and poses a potential threat to privacy should be hidden.

Context

This pattern is applicable in a system in which metadata is shared, published or sent.

Problem

There are multiple types of metadata. There is user-generated metadata data like exif-data. Exif is a format for storing metadata in pictures. There is also metadata which exists to ensure the functionality of some services like headers in email or http, or timestamps in files. Often the user is not aware of this additional data that is attached to the content. When publishing data, this could lead to a potential loss of privacy.

Goal

G The possibly identifying information must not be accessible after publication.

Solution

Erase metadata which is not needed for the functionality.

Constraints and Consequences

C Private information will be protected by stripping metadata with sensitive content. The data without the according metadata uses less space and is thus easier to store or transmit.

C Another consequence is, that the process of removing metadata is not reversible. When additional services require information, Metadata can be mandatory. This could lead to a loss of functionality. Geolocations can help placing pictures on a map. Another example would be when accessing a website with a mobile device and stripping device information, the webserver cannot provide an optimized version for mobile devices of the website, decreasing user experience.

Motivating Scenario

Scenario Alice is a food blogger and she takes a picture of her meal. She uploads the photo on her blog. Assuming that Mallory, a malicous reader of Alice's blog wants to know from where the picture was taken. So she looks at the metadata and can tell by looking at the coordinates, the exact location. ■

Know Uses and Related Work

Anonymous Type I Remailer forward emails by modifing the message header and removing sender related information. Flickr.com give users the option to hide Exif data from public disclosure. The Anonymizer is a well-known tool for anonymous web interaction. For example by using a proxy between a request sender and a recipient to strip header information like $HTTP_{U}SER_{A}GENT$ in packet headers because they contain metadata about packet senders.

Categories

- Hide

Related Patterns

P

Supporting Patterns

P

Sources

- `https://privacypatterns.eu/#/patterns/strip-metadata/0-0-0-0-0-0-0-0-0-4-0-0-0-0-0-0`

2.73 Identity Federation Do Not Track Pattern

Summary

All information has been extracted from http://blog.beejones.net/the-identity-federation-do-not-track-pattern

The Do Not Track Pattern makes sure that neither the Identity Provider nor the Identity Broker can learn the relationship between the user and the Service Providers the user us.

Context

This pattern is focused on identity federation models.

Problem

When an identity system provides identifying information about a user and passes this to a third party service, different parties can do correlation and derive additional information.

Goal

G Avoid the correlation of end user and service provider data.

Solution

Include an orchestrator component, that must act in behalf and be controlled by the user. The orchestrator makes sure that the identity broker can't correlate the original request from the service provider with the assertions that are returned from the identity provider. The correlation can only be done within the orchestrator but that's no issue because this acts on behalf of the user, possibly on the device of the user.

Constraints and Consequences

C In practice, the orchestrator could run in the browser of the user as a javascript program or as an App on his device.

Motivating Scenario

Scenario ■

Know Uses and Related Work

> Identity federations and ecosystems.

Categories

Related Patterns

P

Supporting Patterns

P

Sources

- https://privacypatterns.eu/#/patterns/identity-federation
  0-0-0-0-0-0-0-0-0-0-0-0-0-0-0-0

2.74 Dynamic Location Granularity

Summary

When locating users, let them blend in with the crowd.

Number of people who share location at a time is an essential factor in determining the sensitivity of user's location.

Context

A service that exploits user location information.

Problem

Location-based services (LBS) rely on location information for e.g. service delivery or customization. However, this kind of information may be exploited to infer other, sensitive personal data (e.g. habits, presence at home, etc.) A further problem appears when the anonymity of the users needs to be preserved. By accessing detailed enough location information, it is easy to single out the data subjects and effectively deanonymize them. Some minimization solutions for location data consist in blurring this data by reducing its accuracy and precision, so that data is only provided only in a coarse-grained form. However, these do not take into account the deanonymization problems (i.e. they are focusing on l-diversity, rather than on k-anonymity).

Goal

G Make location information achieve k-anonymity while keeping it useful.

Solution

Adjust the granularity of location information, according to the amount of people that share the same location at the same moment. For instance, a granularity of 100 m at a sports stadium in the outskirts may provide anonymity among a group of 100000 people at the main tournament finals, while the same granularity in the early hours of the next morning may not be enough to provide anonymity even among 10 people.

Constraints and Consequences

C The service consumer cannot rely on having the same granularity throughout time. The simple adjusting propose is not resistant to correlation along time.

Motivating Scenario

Scenario A location service provides coarse-grained location measurements of individuals, with a 5-km uncertainty. Individuals are only identified by a service-internal, random-generated pseudonym. A consumer of this service gets to know that a specific individual goes every day in the morning from neighbourhood A to district B, at the other side of the city. Neighbourhood A has broad streets and detached houses, which effectively yields a very low population density. It happens that John Doe is effectively the only person that lives in A and works at B, thus his identity has been deanonymized. ■

Know Uses and Related Work

Gedik, B., & Liu, L. (2005, June). Location privacy in mobile systems: A personalized anonymization model. In Distributed Computing Systems, 2005. ICDCS 2005. Proceedings. 25th IEEE International Conference on (pp. 620-629). IEEE.

See also the Location Granularity pattern by Nick Doty at privacypatterns.org

Categories

- Minimize
- Aggregate

Related Patterns

P

Supporting Patterns

P pseudonymous-identity

Sources

- https://privacypatterns.eu/#/patterns/dynamic-location-granula
  0-0-0-0-0-0-0-0-0-0-0-0-0-0-0-0
- http://privacypatterns.wu.ac.at:8080/catalog/

3. Design Space

3.1 How to Use Design Patterns

Bibliography

[1] Giuseppe D’ Acquisto et al. *Privacy by design in big data*. Technical report. European Union Agency For Network and Information Security, 2015, pages 1–80 (cited on page 16).

[2] Ian F. Akyildiz and Ismail H. Kasimoglu. “Wireless sensor and actor networks: research challenges”. In: *Ad Hoc Networks* 2.4 (2004), pages 351–367. ISSN: 15708705 (cited on page 12).

[3] Altimetergroup.com. *Consumer Perceptions of Privacy in the Internet of Things: What Brands Can Learn from a Concerned Citizenry*. Technical report. Altimeter Group, June 2015 (cited on page 16).

[4] Article 29 Data Protection Working Party. *Opinion 8/2014 on the on Recent Developments on the Internet of Things*. Technical report. 2014, pages 1–24 (cited on page 16).

[5] Kevin Ashton. “That ’Internet of Things’ Thing In the real world, things matter more than ideas”. In: *RFID Journal* (June 2009) (cited on page 10).

[6] Atmel Team. *A look back at the history of the Internet of Things | Atmel Bits and Pieces*. 2015. URL: http://blog.atmel.com/2015/04/09/a-look-back-at-the-history-of-the-internet-of-things/ (cited on page 10).

[7] F Bushmann, Regine Meunier, and Hans Rohnert. “Pattern-oriented software architecture: A system of patterns”. In: *John Wiley&Sons* 1 (1996), page 476. ISSN: 0007-1250. DOI: 10.1192/bjp.108.452.101 (cited on page 17).

[8] Casaleggio Associati. *The Evolution of Internet of Things*. Technical report. Casaleggio Associati, Feb. 2011 (cited on page 10).

[9] Ann Cavoukian and Jeff Jonas. *Privacy by Design in the Age of Big Data*. Technical report. Information and Privacy Commissioner, Ontario, Canada, 2012, pages 1–17 (cited on page 16).

[10] George Danezis et al. *Privacy and Data Protection by Design - from policy to engineering*. Technical report. European Union Agency for Network and Information Security (ENISA), 2014, pages 1–79 (cited on page 16).

[11] Data Council. *No Title*. Technical report. London: Direct Marketing Association (UK), 2014, pages 1–81. URL: http://dma.org.uk/guide/data-guide (cited on page 16).

[12] European Commission. *Internet of Things in 2020 Road Map For The Future*. Technical report. Working Group RFID of the ETP EPOSS, May 2008 (cited on page 10).

[13] European Communities. *Internet of Things IoT Governance, Privacy and Security Issues*. Technical report. European Commission Information Society and Media, 2015 (cited on page 16).

[14] European Union Agency for Fundamental Rights. *Handbook on European data protection law*. Luxembourg, 2014, pages 1–203 (cited on page 16).

[15] Federal Trade Commission. *Internet of Things: Privacy and Security in a Connected World*. FTC Staff Report. Federal Trade Commission, Jan. 2015 (cited on page 16).

[16] Patrick Guillemin and Peter Friess. *Internet of Things Strategic Research Roadmap*. Technical report. The Cluster of European Research Projects, Sept. 2009 (cited on page 11).

[17] Dominique Guinard. “Towards the web of things: Web mashups for embedded devices”. In: *In MEM 2009 in Proceedings of WWW 2009. ACM*. 2009 (cited on page 10).

[18] Hewlett Packard Enterprise Development LP. *Securing the Internet of Things*. Technical report. 2015, pages 1–16 (cited on page 16).

[19] IEEE. *IEEE INTERNET OF THINGS JOURNAL - Home*. 2015. URL: `http://iot-journal.weebly.com/` (cited on page 11).

[20] Information Commissioner’s Office. *The Guide to Data Protection*. Technical report, pages 1–131. URL: `https://ico.org.uk/media/for-organisations/guide-to-data-protection-2-4.pdf` (cited on page 16).

[21] International Telecommunication Union. *Internet of Things Global Standards Initiative*. 2015. URL: `http://www.itu.int/en/ITU-T/gsi/iot/Pages/default.aspx` (cited on page 11).

[22] Jaewoo Kim et al. “M2M Service Platforms: Survey, Issues, and Enabling Technologies”. In: *IEEE Communications Surveys & Tutorials* 16.1 (2014), pages 61–76. ISSN: 1553-877X. DOI: `10.1109/SURV.2013.100713.00203`. URL: `http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=6644332` (cited on page 11).

[23] Rob Kitchin. *Getting smarter about smart cities: Improving data privacy and data security*. Technical report. Data Protection Unit, Department of the Taoiseach, Dublin, Ireland, 2016, pages 1–83 (cited on page 16).

[24] Ronald Leenes, Jan Schallaböck, and Marit Hansen. *PRIME white paper v2*. Technical report. 2007, pages 1–22 (cited on page 16).

[25] Tan Lu and Wang Neng. “Future internet: The Internet of Things”. In: *3rd International Conference on Advanced Computer Theory and Engineering(ICACTE).* Volume 5. Aug. 2010, pages V5–376—-V5–380. DOI: `10.1109/ICACTE.2010.5579543` (cited on page 11).

[26] Knud Lasse Lueth. *Why it is called Internet of Things: Definition, history, disambiguation.* 2014. URL: `https://iot-analytics.com/internet-of-things-definition/` (cited on pages 10–12).

[27] Roberto Minerva, Abyi Biru, and Domenico Rotondi. *Towards a definition of the Internet of Things (IoT).* Technical report. IEEE Internet Initiative, 2015. URL: `http://iot.ieee.org/` (cited on page 13).

[28] Natalia Olifer and Victor Olifer. *Computer Networks: Principles, Technologies and Protocols for Network Design.* John Wiley & Sons, 2005. URL: `http://au.wiley.com/WileyCDA/WileyTitle/productCd-EHEP000983.html` (cited on page 9).

[29] Postscapes.com. *Internet of Things History | Background and Timeline of the Topic.* 2016. URL: `https://postscapes.com/internet-of-things-history/` (cited on page 10).

[30] Gil Press. *A Very Short History Of The Internet Of Things - Forbes.* 2014. URL: `http://www.forbes.com/sites/gilpress/2014/06/18/a-very-short-history-of-the-internet-of-things/3/%7B%5C#%7D2506b07543c5` (cited on page 10).

[31] Rodrigo Roman, Jianying Zhou, and Javier Lopez. “On the features and challenges of security and privacy in distributed internet of things”. In: *Computer Networks* 57.10 (2013), pages 2266–2279 (cited on page 15).

[32] Sasha Romanosky et al. “Privacy patterns for online interactions”. In: *Proceedings of the 2006 conference on Pattern languages of programs - PLoP ’06* (2006), page 1. DOI: `10.1145/1415472.1415486` (cited on page 17).

[33] Trustworthy Computing Next. *Protecting Data and Privacy in the Cloud.* Technical report. Microsoft Corporation, 2014 (cited on page 16).

[34] UK Government Chief Scientific Adviser. *The Internet of Things: making the most of the Second Digital Revolution*. Technical report. London: The Government Office for Science, 2014, pages 1–40 (cited on page 16).

[35] University of Amsterdam and the Massachusetts Institute of Technology. *Privacy Bridges: EU and US Privacy Experts in Search of Transatlantic Privacy Solutions*. Technical report. 2015, pages 1–52 (cited on page 16).

[36] Yang Wang and Alfred Kobsa. "Privacy-Enhancing Technologies". English. In: *Handbook of Research on Social and Organizational Liabilities in Information Security*. Edited by Manish Gupta and Raj Sharman. IGI Global, Jan. 2009, pages 203–227. ISBN: 9781605661322. DOI: `10.4018/978-1-60566-132-2`. URL: `http://www.igi-global.com/chapter/privacy-enhancing-technologies/21343` (cited on page 16).

[37] John Woulds. *A Practical Guide to the Data Protection Act*. Technical report. The Constitution Unit, 2004, pages 1–32 (cited on page 16).