ICME: An Informed Consent Management Engine for Conformance in Smart Building Environments

Chehara Pathmabandu
chehara.pathmabandu@monash.edu
Faculty of Info. Tech., Monash University
Australia

John Grundy
john.grundy@monash.edu
Faculty of Info. Tech., Monash University
Australia

Mohan Baruwal Chhetri
Mohan.BaruwalChhetri@data61.csiro.au
CSIRO's Data61
Australia

Zubair Baig
zubair.baig@deakin.edu.au
School of Info. Tech., Deakin University
Australia

ABSTRACT

Smart buildings can reveal highly sensitive insights about their inhabitants and expose them to new privacy threats and vulnerabilities. Yet, convenience overrides privacy concerns and most people remain ignorant about this issue. We propose a novel *Informed Consent Management Engine* (ICME) that aims to: (a) increase users' awareness about privacy issues and data collection practices in their smart building environments, (b) provide fine-grained visibility into privacy conformance and infringement by these devices, (c) recommend and visualise corrective user actions through "digital nudging", and (d) support the monitoring and management of personal data disclosure in a shared space. We present a reference architecture for ICME that can be used by software engineers to implement diverse end-user consent management solutions for smart buildings. We also provide a proof-of-concept prototype to demonstrate how the ICME approach works in a shared smart workplace. Demo: https://youtu.be/5y6CdyWAdgY

CCS CONCEPTS

• **Security and privacy → Privacy protections**; **Human and societal aspects of security and privacy**; • **Social and professional topics → Privacy policies**.

KEYWORDS

IoT, Smart Buildings, Smart office, Informed consent, Privacy Preservation, awareness, Privacy Rights, Privacy policies, Compliance

ACM Reference Format:
Chehara Pathmabandu, John Grundy, Mohan Baruwal Chhetri, and Zubair Baig. 2021. ICME: An Informed Consent Management Engine for Conformance in Smart Building Environments. In *Proceedings of the 29th ACM Joint European Software Engineering Conference and Symposium on the Foundations of Software Engineering (ESEC/FSE '21), August 23–28, 2021, Athens, Greece.* ACM, New York, NY, USA, 5 pages. https://doi.org/10.1145/3468264.3473118

ESEC/FSE '21, August 23–28, 2021, Athens, Greece

ACM ISBN 978-1-4503-8562-6/21/08…$15.00
https://doi.org/10.1145/3468264.3473118

1 INTRODUCTION

A smart building relies upon human-centric processes to automatically control the building's operations, including access, HVAC, lighting, security and energy management, with intelligence and communication capabilities facilitating operations [5][15]. IoT is the principal empowering technology that transforms ordinary buildings into "smart" buildings manipulating hundred to millions of sensing and actuating devices to offer a self-sustaining ecosystem. IoT devices installed in smart buildings are capable of capturing a huge amount of data, including data directly related to an individual inhabitant. The primary intent of capturing user behaviour, occupation, or choices is to promptly cater to end-user needs without the need for interference or identification of its users. However, such data collection and usage can possibly infringe the privacy of the inhabitants, and smart building sensors can permeate various private aspects of individuals' lives without their prior consent [9]. To improve the privacy of smart building occupants, it is essential to clearly define (a) how personally identifiable information (PII) is collected, used, shared and disclosed, (b) how permission is sought and granted for the accumulated data to be shared and disclosed to third-parties, and (c) the liabilities related to any breaches of private information [16].

With the initiation of GDPR law [6], and other similar regulations (e.g., California Consumer Privacy Act - CCPA, Personal Information Protection and Electronic Documents Act - PIPEDA), people are becoming more aware of their privacy rights [2] [12]. Yet, despite the increased awareness, users rarely make active efforts to protect their personal information as their privacy-related behaviours are influenced by the increased functionality and convenience offered by IoT devices [3]. This discrepancy between the consumers' stated concerns and their actual behaviour is referred to as the *privacy paradox* [14]. To address this problem, we have previously proposed a model for managing *informed consent* in smart spaces [11]. As shown in Fig. 1, privacy conformance checking is done by intercepting the device API calls and checking the list of permissions requested by a device (and granted by a user) against the device privacy policies. When an end-user triggers an action, the consent model checks the permissions granted by the user and the event's payload data against the device's privacy policy

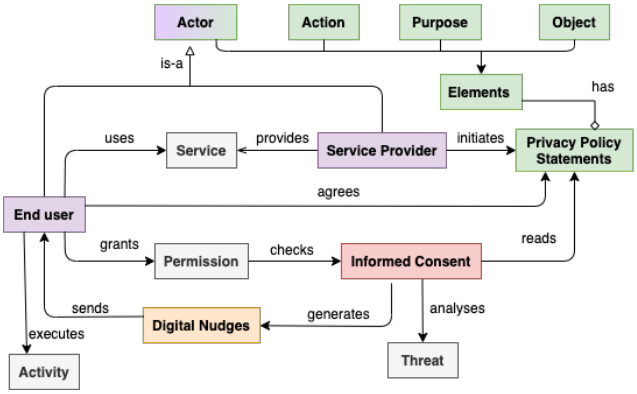

Figure 1: Informed consent model

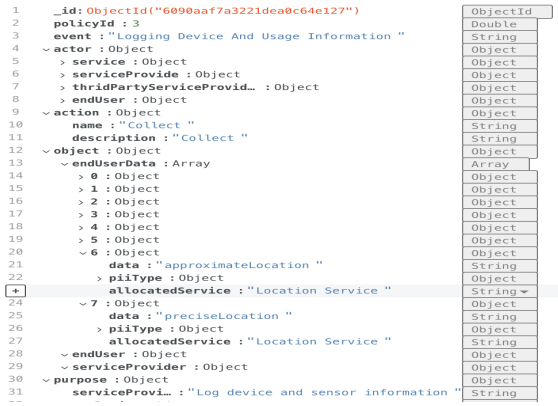

Figure 2: Example privacy policy statement

statements. If the consent engine detects a policy infringement, the end-user is made aware of it through a *nudge*. Following the nudge, the user can decide if they want to take appropriate actions against the infringement, such as not using the device, refusing permission for data capture, or leaving the smart space.

In this paper, we present a reference architecture for our informed consent management engine (ICME), which can be used by software engineers to implement consent management solutions for smart buildings. We also provide a proof-of-concept prototype to demonstrate how it works in a shared smart workplace.

2 APPROACH

A smart building can house many *organizations* that are spread across multiple *floors* and occupy many individual and shared *rooms* that can be grouped into different *zones*. Each of these "smart spaces" are instrumented with different types of devices that are shared and used by multiple users with different levels of accessibility, authority, and privacy preferences. The nudging mechanism used and the nudging content presented to the inhabitants of these smart spaces may differ significantly depending upon their level of accessibility and authority. A central concern in managing such *system of systems* is the complex relationship between the smart spaces, shared devices, and their users. Below, we outline the five key phases through which ICME manages the privacy of the inhabitants in such shared smart spaces.

Phase 1: Extract privacy policies by applying textual patterns. For each device installed in the shared smart space, its privacy policies are extracted based on a pre-defined textual pattern, converted into an appropriate machine-readable format (e.g. JSON [1], XML), and stored in the *privacy policy document database* (PDD). As shown in Fig. 2, each privacy policy statement has the following elements: Actor, Action, Object(s), Purpose(s) [10], e.g. *the service provider collects end-user usage frequency for service provisioning/analytics*. A key-value data format is suitable for handling unstructured, dynamic schema in creating documents, and each privacy policy is represented as an object of objects, which is more expressive and powerful than a traditional row/column model.

Phase 2: Maintain a list of risky permissions for smart devices.

A *PII bank* is created to maintain a list of sensitive and non-sensitive user information that may be collected and transferred when users interact with devices in the shared smart space. Each PII type is linked to a risk level based on its sensitivity, e.g. *a user's precise location is highly sensitive compared to an approximate location*. Each time a new device is installed in the shared smart space, all permissions (mandatory and optional) defined by the device's manufacturer are checked to gain more visibility into the personal data that the vendor can collect about the device user at both the software and hardware level. Following this, possible dangerous permissions that have the risk of revealing sensitive information about the users and could be used for profiling, tracking, advertising and identification [4] are identified, labelled and listed. Each policy statement extracted in Phase 1 includes one or more such PII types under the object element.

Phase 3: Check for privacy breaches or potential risks when users interact with devices.
Every time a user interacts with a smart device, the payload data for the user executed action is captured and pre-processed into a lightweight data-interchange format that is machine-readable. The payload is then assessed against the privacy policy statements for that device stored in the PDD. More specifically, the PII types under the object element of the privacy policy are compared against the parameters collected in the payload. Two types of incidents are flagged – direct breach of the privacy policy and potential disclosure of high-risk PII (even if no direct infringement is identified).

Phase 4: Track and log events associated with privacy breaches or potential risks.
Once the payload has been assessed against the device privacy policy statements, the user-triggered action (event, action, time, current status) is tracked and logged for all incidents that are flagged in Phase 3. These incidents and past incidents are analysed to nudge the users to change their behaviour.

Phase 5: Recommend corrective actions using nudges.
Corrective actions are recommended to the user via nudges to control and mitigate privacy issues that have occurred and/or may transpire in future. Nudging is a promising approach at influencing people's judgment, choice or behaviour in a range of domains, including cybersecurity, in a desirable way [7]. In our approach, a

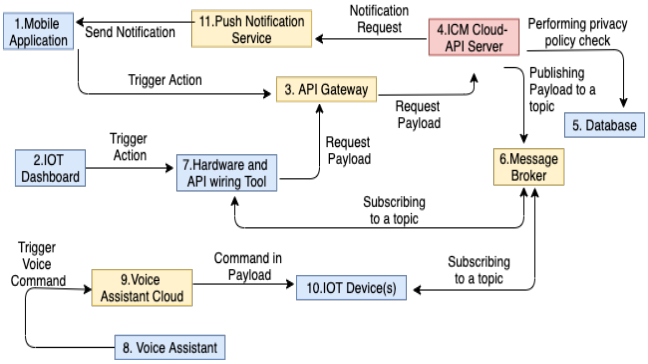

Figure 3: Reference architecture (showing data flows)

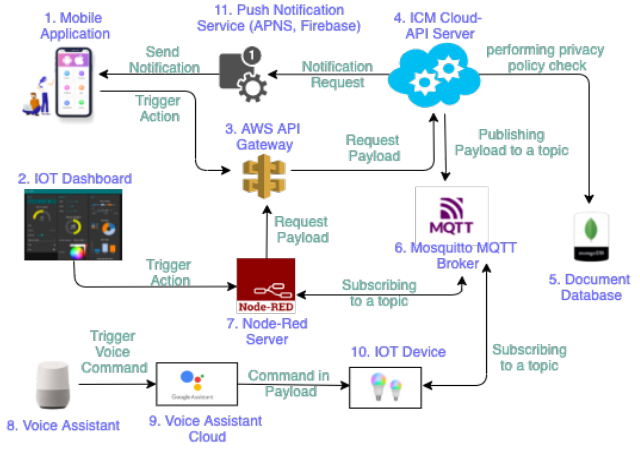

Figure 4: Implemented solution overview

nudge consists of three elements: (1) a summary of the risky actions, (2) redirect instructions to configure privacy settings, and (3) an option to decline and allow data to be shared. The nudging technique is used for positive reinforcement, highlighting what privacy settings should be practised by the smart building inhabitants to secure and protect their PII. In designing the nudge content, we have considered the concepts of protection motivation theory (PMT) [17][8], herd behaviour [18], social norms [13], privacy paradox [19] and privacy calculus theories [13] to enhance user awareness, share good practices, and to motivate users to review and possibly adjust their privacy settings.

3 REFERENCE ARCHITECTURE

Fig. 3 presents the reference architecture for our proposed approach and shows the data flow between the different components. The privacy policies for each IoT device in the shared smart space **(Phase 1)**, the list of pre-defined permissions **(Phase 2)**, the events associated with privacy breaches or potential risks **(Phase 4)**, and the nudge details **(Phase 5)** are stored in the database (5). A user can interact with the IoT Device (10) in three different ways; via a mobile application (1), via the IoT dashboard (2) and by using a voice assistant (8). These act as consumers of the ICME API deployed on the ICM Cloud-API server (4). If we consider the scenario of the user interacting with an IoT device via her mobile application, then the following data flow occurs. The user request from the application (1) is forwarded to the ICM Cloud-API server (4) via the API Gateway (3) by making Rest API calls. On the ICM Cloud-API server, each variable of the request payload is assessed against the privacy policies in the PDD (5) **(Phase 3)**. Simultaneously, the API publishes the request payload to a dedicated topic on the Message Broker (6), which routes it to the appropriate IoT device subscribed to the same topic (10) after identifying the end-point. In the event of a privacy breach or a potential disclosure of high-risk PII, the event is logged in the database **(Phase 4)**. Following this, the ICM API server also sends a nudge recommending corrective actions to the mobile application using the Push Notification Service (11) **(Phase 5)**. Finally, the device's current status is reflected on the IoT dashboard (2) by forwarding the message via a hardware and API writing tool (7).

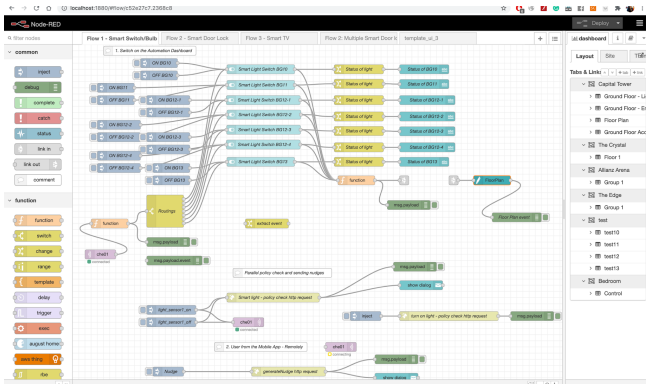

Figure 5: Data flow definition for smart lights in the HumaniSE Lab.

4 IMPLEMENTATION

Fig. 4 shows one specific implementation of the reference architecture shown in Fig. 3. We used Node-RED[1] to implement the IoT Dashboard (2) and for wiring together the hardware and the ICM APIs (7). Similarly, we used the AWS API Gateway[2] (3) to act as the "front door" for our ICM Cloud API Server, which was deployed on Amazon EC2 (4). The privacy policy database was implemented using MongoDB Atlas[3] (5). We used MQTT[4] as the Message Broker (6). For the purposes of the prototype demonstration, we created "mock-up" interfaces for the different IoT devices, including Smart Door Lock, Smart TV and Smart Light Switch, by writing "device" interfaces that mimic their behaviours (10). For the push notification service (11) we used Firebase[5]. Finally, we developed a prototype mobile application for displaying the nudges received by the user. The code for ICME is available on Github[6].

[1] https://nodered.org/
[2] https://aws.amazon.com/api-gateway/
[3] https://www.mongodb.com/cloud/atlas2
[4] https://mqtt.org/
[5] https://firebase.google.com/
[6] https://github.com/chehara/ICM_repo

Figure 6: Node-Red Dashboard for the HumaniSe Lab (with floor map integration)

4.1 Designing Data Flows for the Smart Office Space

For our prototype implementation, we have simulated the setup of the HumaniSE lab[7]. As shown in Fig. 6, the lab includes private rooms, meeting rooms, open plan working spaces and shared areas. For the purpose of demonstration, we have configured the lab with seven smart light interfaces, four smart door lock interfaces, two smart TV interfaces, and five smart temperature sensor interfaces. We used Node-RED to implement smart space automation operations by wiring together these different devices, the ICM APIs and the Nudge Notification Service. Node-RED is a programming tool that allows wiring together hardware devices, APIs and online services and supports the remote management of IoT devices in a shared space. Node-RED is specifically designed to be used with MQTT. To provide remote access and control, we have implemented an MQTT client and server network node to communicate with an 'MQTT broker node'. We have also used many JavaScript 'nodes' to process the messages received by the nodes and to create virtual wires between flows to make them reusable. Fig. 5 shows the data flow definition for the smart lights in the HumaniSE Lab. It also shows multiple tabs for the data flows for the smart TVs and the smart door locks. We can define such data flows for any type of smart space setup, including hierarchical ones as discussed in Section 2.

4.2 Dashboard to Interact with the Smart Space

We have used Node-RED to implement the dashboard, which provides users with a real-time view of the different IoT devices in the shared space. Fig. 6 shows the dashboard for the HumaniSE lab with the floor map integration. Depending upon the assigned authority level, users will have access and control to different devices in the lab via the dashboard. The privacy nudges are also sent and displayed on the dashboard in addition to being sent to the user's mobile application.

[7] https://www.monash.edu/it/humanise-lab/home

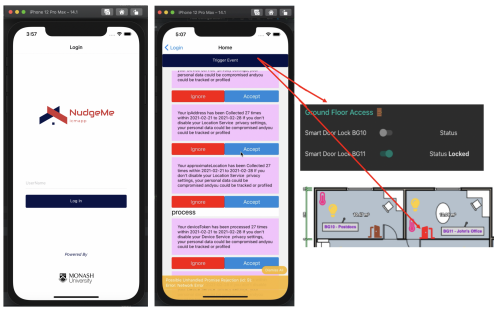

Figure 7: Nudge visualisation and device control via mobile app

4.3 Mobile App to Interact with the Smart Space and for Nudge Visualisation

In addition to the Node-RED dashboard, we have also developed a prototype mobile application that can be used to interact with the devices and to visualise and respond to digital nudges about privacy threats and settings (refer to Fig. 7). The app displays all recent nudges received by the user. Different types of nudges are delivered to the user based on how many times the user has declined the recommended corrective actions for the same privacy infringement or potential risk. This encourages the use of empirically validated techniques rather than relying on traditional measures that might not achieve the desired behaviour change. We have considered nine different nudge types in our prototype implementation. The following example details how the same information can be presented differently depending upon the selected nudge type.

Type 4: Awareness and threat appraisal – "*Your <PII Type> has been shared <n> times during the last <x> days. If you do not disable your <PII Type> sharing privacy settings, your personal data could be compromised, and you can be tracked or profiled*".

Type 5: Awareness and herd behaviour – "*Your <PII Type> has been shared <n> times during the last <x> days. <Percentage> of your colleagues do not share <PII Type> with others.*"

Users can control the smart devices they have access to via the mobile app, similar to the dashboard. Once a user decides to change the privacy settings, the system will automatically update their preferences. Finally, all generated nudges with the user action are recorded in the database to measure the nudge effectiveness against user behaviour change in the future.

5 CONCLUSION

We presented ICME, an informed consent management engine for shared smart spaces. We introduced a reference architecture for ICME followed by concrete implementation. We then demonstrated its practical feasibility and how it works by simulating a smart shared workplace. Our proposed ICME approach can be used by software developers to implement diverse end-user consent management solutions for a variety of shared smart spaces.

ACKNOWLEDGEMENTS

Pathmabandu is supported by a CSIRO Data61 PhD scholarship. Grundy is supported by ARC Laureate Fellowship FL190100035.

REFERENCES

[1] 2001. JSON. https://www.json.org/json-en.html.

[2] 2020. California Consumer Privacy Act (CCPA) | State of California - Department of Justice - Office of the Attorney General. https://oag.ca.gov/privacy/ccpa.

[3] Noura Aleisa, Karen Renaud, and Ivano Bongiovanni. 2020. The privacy paradox applies to IoT devices too: A Saudi Arabian study. *Computers and Security* 96 (2020), 101897. http://www.sciencedirect.com/science/article/pii/S0167404820301711

[4] E. Alepis and C. Patsakis. 2017. Monkey Says, Monkey Does: Security and Privacy on Voice Assistants. *IEEE Access* 5 (2017), 17841–17851.

[5] Alex H Buckman, Martin Mayfield, and Stephen BM Beck. 2014. What is a smart building? *Smart and Sustainable Built Environment* (2014).

[6] GDPR. 2018. General Data Protection Regulation (GDPR) – Official Legal Text. https://gdpr-info.eu/.

[7] Pelle Guldborg Hansen. 2016. The Definition of Nudge and Libertarian Paternalism: Does the Hand Fit the Glove? *European Journal of Risk Regulation* 7, 1 (2016), 155–174. https://doi.org/10.1017/S1867299X00005468

[8] Daniel Kahneman. 2012. *Thinking, fast and slow.*

[9] Milan Markovic, Waqar Asif, David Corsar, Naomi Jacobs, Peter Edwards, Muttukrishnan Rajarajan, and Caitlin Cottrill. 2018. Towards automated privacy risk assessments in IoT systems. 15–18.

[10] Nazila Mohammadi, Jens Leicht, Ludger Goeke, and Maritta Heisel. 2020. Assisted Generation of Privacy Policies using Textual Patterns. 347–358.

[11] Chehara Pathmabandu, John Grundy, Mohan Baruwal Chhetri, and Zubair Baig. 2020. An Informed Consent Model for Managing the Privacy Paradox in Smart Buildings *(ASE '20)*. Association for Computing Machinery, New York, NY, USA, 19–26. https://doi.org/10.1145/3417113.3422180

[12] PIPEDA. 2019. The Personal Information Protection and Electronic Documents Act (PIPEDA) - Office of the Privacy Commissioner of Canada. https://www.priv.gc.ca/en/privacy-topics/privacy-laws-in-canada/the-personal-information-protection-and-electronic-documents-act-pipeda/.

[13] Maija Elina Poikela. 2020. *Theoretical Background to Location Privacy.* Springer International Publishing, Cham, 13–32.

[14] Karen Renaud, Ivano Bongiovanni, and Noura Aleisa. [n.d.]. The privacy paradox: we claim we care about our data, so why don't our actions match? https://theconversation.com/the-privacy-paradox-we-claim-we-care-about-our-data-so-why-dont-our-actions-match-143354.

[15] Nico Saputro, Ali Yurekli, Kemal Akkaya, and Selcuk Uluagac. 2016. *Privacy Preservation for IoT Used in Smart Buildings.* 129–160. https://doi.org/10.1201/b19516-10

[16] Amit Kumar Tyagi, G. Rekha, and N. Sreenath. 2020. Beyond the Hype: Internet of Things Concepts, Security and Privacy Concerns. In *Advances in Decision Sciences, Image Processing, Security and Computer Vision*, Suresh Chandra Satapathy, K. Srujan Raju, K. Shyamala, D. Rama Krishna, and Margarita N. Favorskaya (Eds.). Springer International Publishing, Cham, 393–407.

[17] René van Bavel and Nuria RodrÃguez-Priego. 2016. Nudging Online Security Behaviour with Warning Messages: Results from an Online Experiment.

[18] Ali Vedadi and Merrill Warkentin. 2020. "Can Secure Behaviors Be Contagious? A Two-Stage Investigation of the Influence of Herd Behavior on Security Decisions ". *Journal of the Association for Information Systems* (01 2020), 428–459. https://doi.org/10.17705/1jais.00607

[19] M. Williams, J. R. C. Nurse, and S. Creese. 2016. The Perfect Storm: The Privacy Paradox and the Internet-of-Things. In *2016 11th International Conference on Availability, Reliability and Security (ARES).* 644–652. https://doi.org/10.1109/ARES.2016.25