

Demonstrably doing accountability in the Internet of Things

Lachlan Urquhart*, Tom Lodge[†] and Andy Crabtree[‡]

ABSTRACT

This article explores the importance of accountability to data protection (DP), and how it can be built into the Internet of Things (IoT). The need to build accountability into the IoT is motivated by the opaque nature of distributed data flows, inadequate consent mechanisms and lack of interfaces enabling end-user control over the behaviours of Internet-enabled devices. The lack of accountability precludes meaningful engagement by end users with their personal data and poses a key challenge to creating user trust in the IoT and the reciprocal development of the digital economy. The European Union General Data Protection Regulation 2016 (EU GDPR) seeks to remedy this particular problem by mandating that a rapidly developing technological ecosystem be made accountable. In doing so, it foregrounds new responsibilities for data controllers, including DP by design and default, and new data subject rights such as the right to data portability. While GDPR is ‘technologically neutral’, it is nevertheless anticipated that realizing the vision will turn upon effective technological development. Accordingly, this article examines the notion of accountability, how it has been translated into systems design recommendations for the IoT and how the IoT Databox puts key DP principles into practice.

KEYWORDS: Internet of Things, privacy engineering, edge computing, personal information management systems, accountability, data protection

INTRODUCTION

The ‘connected home’ currently sits at the ‘peak of inflated expectations’ in Gartner’s often-cited hype cycle, and the Internet of Things (IoT) is a key driver of the hype.¹ A cursory glance at the consumer IoT market reveals swathes of household goods

* Lachlan Urquhart, Research Fellow in Information Technology Law, Horizon Digital Economy Research, School of Computer Science, University of Nottingham, Wollaton Road, Nottingham NG8 1BB, UK. E-mail: lachlan.urquhart@gmail.com

[†] Tom Lodge, Research Fellow, Mixed Reality Lab, School of Computer Science, University of Nottingham, Wollaton Road, Nottingham NG8 1BB, UK. E-mail: tom.lodge@nottingham.ac.uk

[‡] Andy Crabtree, Professor of Computer Science, Mixed Reality Lab, School of Computer Science, University of Nottingham, Wollaton Road, Nottingham NG8 1BB, UK. E-mail: andy.crabtree@nottingham.ac.uk

1 Kasey Panetta, *Top Trends in the Gartner Hype Cycle for Emerging Technologies 2017* (Gartner 2017) <<http://www.gartner.com/smarterwithgartner/top-trends-in-the-gartner-hype-cycle-for-emerging-technologies-2017>> accessed 21 June 2018.

with the prefix ‘smart’ or ‘intelligent’ on offer, spanning white good to fixtures and fittings embedded in the fabric of the home.² The promise of the IoT is greater convenience, security, safety, efficiency and comfort in a user’s everyday life. While the necessity of many IoT products and services may be questionable,³ anticipated growth in the sector is vast: major IT firms like Cisco, Ericsson, General Electric and Accenture all predict billions of networked devices in the coming years.⁴ The IoT essentially trades on data, both actively and passively, with inputs ranging from explicit spoken voice commands to sensed data inputs implicated in such things as movement or temperature monitoring. The IoT also aligns with other trends in computing, particularly big data, cloud computing and machine learning, with personal data collected by IoT devices typically being distributed to the cloud for processing and analytics.

Accompanying the diversity of IoT devices and services are concerns centring on privacy and trust. When sensing occurs in the home, for example, patterns of behaviour can be detected and inferences made about inhabitants’ lifestyles. Depending who is making these inferences, and who they share the data with, privacy harms can emerge. As Nissenbaum argues, inappropriate flows of information between contexts can cause harm to an individual’s sense of privacy.⁵ The nascent nature of the industry means there is a lack of harmonized standards for building IoT devices in ways that sufficiently foreground and anticipate data protection (DP) concerns.⁶ Building trustworthy relationships with consumers in the new IoT infrastructure is critical, and not least because an increasing awareness that IoT devices leak data or can easily be hacked and implicated in widespread distributed denial of service attacks contributes to a diminishing sense of trust in the emerging infrastructure.⁷

Against this background we elaborate key challenges posed by the IoT from a regulatory perspective and how these practically occasion the need for accountability. These include challenges posed by devices that lack or only provide partial user interfaces and compliant consent mechanisms; the opacity of data flows to end users and the spectrum of General Data Protection Regulation (GDPR) control rights; machine to machine (M2M) communications and the legitimacy of access; and cloud storage and international data transfer safeguards. We move on to explore various aspects of the Accountability Principle, first its history in DP governance and

2 ‘IoT List’ <<http://iotlist.co>> accessed 21 June 2018.

3 ‘The Internet of Useless Things’ <<http://www.internetofuselessthings.io>> accessed 21 June 2018.

4 Gartner, ‘Gartner Says 8.4 Billion Connected “Things” Will Be in Use in 2017, Up 31 Percent from 2016’ *Gartner Newsroom* (7 Feb 2017) <<https://www.gartner.com/newsroom/id/3598917>> accessed 21 June 2018.

5 Helen Nissenbaum, *Privacy in Context: Technology, Policy, and the Integrity of Social Life* (Stanford University Press 2009).

6 Ian Brown, *Regulation and the Internet of Things* (International Telecommunications Union 2015) <https://www.itu.int/en/ITU-D/Conferences/GSR/Documents/GSR2015/Discussion_papers_and_Presentations/GSR_DiscussionPaper_IoT.pdf> accessed on 21 June 2018; Karen Rose, *Internet of Things: An Overview* (Internet Society 2015) <<https://www.sfbayisoc.org/wp-content/uploads/2015/12/ISOC-IOT-FEB-2016-Rose.pdf>> accessed 21 June 2018.

7 World Economic Forum, *Rethinking Personal Data: A New Lens for Strengthening Trust* (World Economic Forum 2014) <http://www3.weforum.org/docs/WEF_RethinkingPersonalData_ANewLens_Report_2014.pdf> accessed 21 June 2018.

then how it is presented in Article 5(2) of GDPR.⁸ This exploration involves questioning the nature of the account to be provided, how it is to be provided and to whom. We situate Article 5(2) within the wider context of GDPR, turning to various requirements of Article 24 as interpreted in GDPR recitals, and other related articles, to map how they intersect with accountability. The requirements of GDPR pose distinct challenges to the development of technological systems and we subsequently turn to consider the recommendations of the Article 29 Working Party, and how they envisage GDPR playing out in the IoT, as a preface to presenting the IoT Databox. We conclude by mapping how the IoT Databox addresses the different accountability requirements of GDPR.

THE PRACTICAL NEED TO BUILD ACCOUNTABILITY INTO THE IOT

From May 2018, GDPR will be enforced across all European Union member states (Regulation 2016/679). It will also affect data controllers outside Europe if they target goods and services to, or otherwise monitor, European Union (EU) citizens.⁹ Seeking to bring DP laws into the 21st century, GDPR replaces the pre-Internet Data Protection Directive 1995.¹⁰ The IoT sector is heavily driven by personal data, meaning it is critical that IoT developers negotiate their relationship with the new user rights and controller responsibilities mandated by GDPR. This includes a raft of fresh legal rules governing the processing of personal data, along with extension of the rights provided to data subjects and the responsibilities incumbent on data controllers, all of which are impacted by the underlying technological infrastructure.

Lack of or partial user interfaces and consent

The design of IoT devices is heterogeneous. Unlike mobile phones, where users can develop mental models about how these devices work,¹¹ ‘interfaces’ to the IoT vary immensely. Many IoT devices do not have screens and communication with users and rely instead on lights or sounds or haptic feedback; text notifications to mobile phones may also be leveraged in the absence of direct device feedback occasioned by the desire to create aesthetically pleasing devices, which may in turn result in opacity about device functionality. This diversity makes it hard for users to understand what personal information is being collected and how it is being used. From a regulatory perspective, this shapes the nature of consent mechanisms. Consent is one legal basis for processing personal data. Consent follows a notice and choice model, meaning it should be informed, unambiguous, freely given and specific to a particular process, and enable a clear indication of the data subject’s will.¹² Data subjects need to affirm

8 European Union, ‘Regulation 2016/679 General Data Protection Regulation’ (2016) 59 OJ 4. Hereinafter ‘GDPR’.

9 art 3(2) GDPR.

10 European Union, ‘Directive 95/46/EC of the European Parliament and of the Council of 24 October 1995 on the Protection of Individuals with Regard to the Processing of Personal Data and on the Free Movement of Such Data’ (1995) 281 OJ L 31, 50 (hereinafter ‘DPD’).

11 Martina Ziefle and Susanne Bay, ‘Mental Models of a Cellular Phone Menu’ in Stephen Brewseter and Mark Dunlop (eds), *Mobile Human Computer Interaction* (Springer 2014) 25–37.

12 art 2(h) DPD; art 4(11) GDPR.

their choice, and if the type of data being processed is within special categories of personal data (eg health, gender, race or biometric information) explicit consent is needed. Such consent cannot be obtained through pre-ticked boxes, silence or inactivity by the subject.¹³ The dominant web-based model takes advantage of the affordances of mobile devices, using screens to display privacy policies, and terms and condition contracts containing large blocks of text. Extensive research shows users do not read this text, as it would take an incredibly long time to do so hence they often agree in any case.¹⁴ Even if they did read it, they cannot renegotiate as it is a form contract and may not understand it due to complex literacy requirements.¹⁵ This situation is not ideal and challenges the notion of legally compliant consent. The heterogeneity of IoT devices could be good or bad for consent processes. On the one hand, consent could be frustrated by devices which, by design, ambiently collect data and have interfaces that lack affordances for communicating clear information. This could be particularly challenging for homes, where children and adults cohabit, as GDPR introduces stricter requirements about delivering clear, concise, comprehensible information to children about data processing.¹⁶ However, on the other hand, the IoT poses an opportunity to redesign how consent is done with users. Taking advantage of new interaction methods may provide for the ‘ongoing’ negotiation of terms of consent.¹⁷

Opacity of data flows to end users and control

IoT devices and the digital ecosystems they feed into are largely opaque in how they handle data. Insofar as end users may struggle to understand how their devices work, given the lack of effective interfaces, this may in turn lead to lack of legibility in how data is being processed, why, by whom, where it is being stored, for how long, etc.¹⁸ This has the knock-on effect of making it hard for users exercise their legal rights and to control use of their information. While no hierarchical framing of rights is encoded in GDPR, a spectrum of various control rights enabling data subjects to escalate action from controllers is nevertheless discernible and underpin accountability in GDPR:

1. Article 15 the ‘right to access’, or the right to discover what data is held by the controller about the data subject.

13 Recital 32 GDPR.

14 Aleecia McDonald and Lorrie Cranor, ‘The Cost of Reading Privacy Policies’ (2008) 4 I/S 543.

15 Ewa Luger, Stuart Moran and Tom Rodden, ‘Consent for All’, Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, Paris, 2013, 2687–96.

16 art 12 GDPR.

17 Ewa Luger and Tom Rodden, ‘The Value of Consent: Discussions with Designers of Ubiquitous Computing Systems’, Proceedings of the International Conference on Pervasive Computing and Communication Workshops, New York, USA, 2014, 388–93; Lachlan Urquhart and Tom Rodden, ‘New Directions in Information Technology Law: Learning from Human–Computer Interaction’ (2017) 31 IRLCT 150.

18 Peter Tolmie and others, ‘“This Has to Be the Cats”—Personal Data Legibility in Networked Sensing Systems’, Proceedings of the 19th ACM Conference on Computer-Supported Cooperative Work & Social Computing, San Francisco, CA, 2016, 490–501.

2. Article 16 the 'right to rectification', or the right to correct erroneous data held by the controller.
3. Article 21 the 'right to object', or the right to object to the processing of data by the controller.
4. Article 18 the 'right to restriction of processing', or the right to require the controller to restrict processing of data.
5. Article 20 the 'right to data portability', or the right to have a controller provide data to the data subject in a commonly used, machine readable format to take to another controller.
6. Article 17 the 'right to erasure', or the right to have data deleted by controller and for the data subject to thereby be forgotten.

Each of these control rights occasions practical challenges of implementation. If we take data portability, for example, how can data from sensors be moved between IoT service providers in a usable way?¹⁹ Equally challenging and key to control is the need to surface and make visible what information is being processed in the first place.

M2M communications and access

The connected home consists of a network of connected devices, many of which may interact with one another. We already see this commercially, with home management system like 'Works with Nest' or Apple's 'HomeKit' linking together manufacture devices and third-party offerings. However, and again due to the paucity of interfaces to the IoT, the lack of human oversight in M2M communications makes it hard for users to know what is being shared between devices, and if this is contextually appropriate or not. A good example is sensitive personal data collected by a smart mirror detecting someone's skin condition or smart bathroom scales sensing rapid weight loss over time indicating health conditions.²⁰ Ideally, to respect the agency of users and build their trust, this should not be shared with a health insurance mandated wearable health tracker, unless the user wants it to. Similarly, access by an Amazon Dash inspired replenishment button, perhaps sponsored by a pharmaceutical firm pushing a new skincare range, should have human oversight too. The challenge here is balancing the movement of personal data, utility from devices, business models and ensuring legitimate access to data by different devices and services. By limiting the role of users in the loop, it becomes harder to know if appropriate access is being given (or not) by devices. Linking data sets without adequate access management could also have impacts for data controllers, who need to ensure compliance with DP rules, and users, who may suffer information privacy harms through unexpected data sharing.

19 Lachlan Urquhart, Neelima Sailaja and Derek McAuley, 'Realising the Right to Data Portability for the Domestic Internet of Things' (2017) 22 *Pers Ubiqu Comput* 317.

20 Nokia Health—<<https://health.nokia.com/eu/en/>> accessed 21 June 2018.

Cloud storage and international data transfer

The nature of remote, cloud-based data storage utilized by most IoT devices is also problematic under GDPR. Services using IoT sensor data often store collected data in servers located outside of the EU. This enables businesses to create large data sets, used in training of machine learning algorithms and finding patterns that can be used either in service delivery, or creation of new services. Managing big data sets raises challenges addressing the velocity, variety, veracity and volume of data.²¹ From the perspective of ensuring GDPR compliance, users will struggle to know where their data is, or how they can access and control it when its storage location is likely unknown or geographically distant. Again, oversight over what it is being used for becomes difficult and from a legal perspective, issues of jurisdiction and applicable law in contract clauses can come to the fore. From a DP stance, adequate protection of data when it leaves the EU is difficult, and measures to guarantee protection, like Privacy Shield (which replaced the former Safe Harbor agreement) or model contract clauses all have their flaws.²² Furthermore, as mentioned above, Article 3(2) expands the reach of GDPR for controllers outside of the EU monitoring or targeting goods and services towards EU citizens. Cloud providers may not be able to ignore the importance of GDPR in compliance. The alternative of local data storage, keeping information proximate to end users is preferable for ensuring their control over how it is processed, and ensuring more user centric, ethical IoT applications can emerge in the future. The IoT ecosystem, by design, is opaque, and its actions often invisible to end users. In contravention of DP law principles, interactions are being designed that provide little information about how devices function, what data is collected and what trade-offs consumers are making in order to receive relevant services. This is not sustainable, and risks growth of the sector. It is for these reasons that we argue that accountability needs to be built into the IoT. But what exactly do we mean?

ACCOUNTABILITY?

We are of the view that the answer to many of the regulatory issues surfaced by IoT is to build accountability into products and services, by design. Increased dialogue between data controllers and data subjects is needed so that citizens can exercise better control over how their personal data is exploited in the digital economy. Due to GDPR, interest in accountability as a governance mechanism is growing. However, it remains a difficult concept to succinctly pin down. The accountability principle is only substantively mentioned 'once' in GDPR, in Article 5(2), excluding reference in recitals 61 and 85 in the context of data breaches. However, its implications quickly spiral when read in the wider context of GDPR, in conjunction with Article 24, various recitals, and other relevant articles.

21 ICO, *Big Data, Artificial Intelligence, Machine Learning and Data Protection* (Information Commissioner's Office 2017) <<https://ico.org.uk/media/for-organisations/documents/2013559/big-data-ai-ml-and-data-protection.pdf>> accessed 21 June 2018.

22 See European Data Protection Supervisor, *Opinion 4/2016 on the EU-US Privacy Shield Draft Adequacy Decision* (EDPS 2016) and art 29 Working Party, *WP238 Opinion 01/2016 on the EU-US Privacy Shield Draft Adequacy Decision* (art 29 Working Party 2016).

Historically, there has been a strong relationship between accountability and DP compliance. In this context, accountability has traditionally been invoked as a mechanism for implementing DP principles.²³ As Aldahoff and others point out:

... even in instruments where accountability is not called out as a separate data protection principle, many of its substantive provisions were in fact designed to enable accountability.²⁴

The Article 29 Working Party has argued that accountability obliges data controllers to put in place effective policies and mechanisms to ensure compliance with DP rules.²⁵ This view is endorsed by Aldahoff and others who underscore the importance of making data processing entities answerable—of ‘calling them to account’—for the implementation of appropriate safeguards.²⁶ The European Data Protection Supervisor (EDPS) emphasizes that accountability is not a prescriptive bureaucratic measure merely concerned with validation, but is about proactive leadership to foster a broad culture of accountability.²⁷ The introduction of GDPR puts measures in place that further develop this culture of accountability.

Adopting a similar framing of the accountability principle created 37 years ago in the Organization for Economic Cooperation and Development (OECD) Guidelines on the Protection of Privacy and Trans-Border Flows of Personal Data 1980 (paragraph 14), GDPR states:

The controller shall be responsible and be able to demonstrate compliance with, paragraph 1 (‘accountability’). (Article 5(2))

This means the controller is responsible for processing personal data in compliance with principles found in GDPR, which are themselves similar to OECD good DP governance principles (paragraphs 7–13). Article 5(1) GDPR includes: (i) lawfulness, fairness and transparency; (ii) purpose limitation; (iii) data minimization; (iv) accuracy; (v) storage limitation; (vi) integrity and confidentiality. Where OECD and GDPR differ is in the explicit requirement for ‘demonstration’ of compliance with the different principles. Accordingly, there is a two-part responsibility on data controllers: firstly, to put the necessary measures in place to comply with Article 5(1), and secondly, to find ways to demonstrate they have complied. This could be viewed as firstly a ‘substantive compliance with principles’ requirement, and secondly as a ‘procedural demonstration of compliance to relevant stakeholders’ requirement. We shall revisit these distinctive aspects of accountability in due course. First we

23 Charles Raab, ‘The Meaning of “Accountability” in the Information Privacy Context’ in Daniel Guagnin and others (eds), *Managing Privacy through Accountability* (Palgrave Macmillan 2012) 15–32.

24 Joseph Alhadeff, Brendan Van Alsenoy and Jos Dumortier, ‘The Accountability Principle in Data Protection Regulation: Origin, Development and Future Directions’ in Guagnin and others, *ibid.*, 49–82.

25 art 29 Working Party, *WP173 Opinion 3/2010 on the Principle of Accountability*.

26 Alhadeff, Van Alsenoy and Dumortier (n 24).

27 EDPS website <https://edps.europa.eu/data-protection/our-role-advisor_en> accessed 21 June 2018.

wish to consider what nature an account needs to take and to whom accountability should be demonstrated.²⁸

The nature of an account and to whom accountability must be demonstrated

The current approach in GDPR of not explicitly defining what accountability requires of data controllers is intentional. This again follows OECD 1980 guidelines, which as Alhadeff and others state:

... do not prescribe to whom the controller should be accountable (the 'accountee'), nor what this relationship should look like.²⁹

In their 2010 Opinion on the Principle of Accountability, the Article 29 Working Party (A29 WP) suggested that putting an explicit accountability principle into GDPR would enable case-by-case analysis of appropriate measures, and be preferable to predefining requirements due to this approach being more flexible and scalable. Seven years on, if we look to the most recent A29 WP guidance on Data Protection Impact Assessments,³⁰ it retains a non-prescriptive stance about measures needed for accountability, beyond publishing Data Protection Impact Assessments (DPIAs) and the obligation for record keeping. Lack of detailed prescriptive guidance around such a central concept is consistent with original OECD practice, and keeps accountability sufficiently flexible as a notion. Despite the virtues of flexibility, a sticking point for accountability in practice is the form a 'demonstrable' account needs to take.

In seeking to answer this, Raab argues that giving an account is akin to 'telling a story' and can be seen to operate at three sequential levels.³¹ At its most simple, accountability merely obliges an organization to report back on its actions. The next level requires mechanisms for that story to be questioned, and for data subjects to offer their own. The third level puts sanctions in place for when an account is poor, either due to inaction or lack of a proper story being offered in the first place. Whilst this provides some abstract navigational aid, it does not pin down the precise dimensions of a good 'account'. A series of European projects including Galway, Paris and Madrid have been grappling with the nature of accountability.³² The Paris project document elaborates elements organizations need to put in place to demonstrate

28 Colin Bennett, 'International Privacy Standards: Can Accountability Ever Be Adequate?'(2010) 106 Privacy L & Bus Intl 21.

29 Joseph Alhadeff, Brendan Van Alsenoy and Jos Dumortier (n 24).

30 art 29 Working Party, WP248 *Guidelines on Data Protection Impact Assessment (DPIA) and Determining Whether Processing Is "Likely to Result in a High Risk" for the Purposes of Regulation 2016/679* (art 29 Working Party 2017).

31 Raab (n 23).

32 Martin Abrams, *Data Protection Accountability: The Essential Element* (Centre for Information Policy Leadership 2009) <https://www.hunton.com/files/webupload/CIPL_Galway_Accountability_Paper.pdf> accessed on 21 June 2018; Martin Abrams, *Demonstrating and Measuring Accountability* (Centre for Information Policy Leadership 2010) <http://informationaccountability.org/wp-content/uploads/CIPL_Accountability_Phase_II_Paris_Project-2.pdf> accessed on 21 June 2018; Martin Abrams, *Implementing Accountability in the Marketplace* (Centre for Information Policy Leadership 2011) <http://informationaccountability.org/wp-content/uploads/Centre_Accountability_Phase_III_White_Paper.pdf> accessed on 21 June 2018.

accountability.³³ These largely consist of organizational measures, such as establishing policies based on relevant law, setting up internal bodies to enforce these, providing staff training on information privacy, analysing risks on a regular basis, setting up mechanisms to respond to customer complaints and providing appropriate redress mechanisms. This sits against the wider work of the Galway project,³⁴ which states accountability in general requires organizational buy-in, particularly putting in place internal standards that correlate with external requirements; access to resources to support compliance with policies (training, etc); and internal oversight mechanisms to ensure adherence, coupled with approaches for appropriate sanctions and rule enforcement.

Examining guidance offered by the EDPS,³⁵ and UK Information Commissioner Office,³⁶ we also find a range of new measures in GDPR to assist with accountability requirements. We cluster these in terms of ‘technical’ or ‘organizational’ measures:

Technical measures

Data protection by design and default; including use of anonymization, pseudonymization and end-to-end encryption; IT security risk management.

Organizational measures

Assigning DP officers (DPOs); prior consultations; certification schemes; DPIAs; transparent policies; documentation and record keeping on processing for organizations with over 250 staff;³⁷ internal compliance and audits for effectiveness of approaches; training.

GDPR thus puts in place a raft of new organizational and technical ‘responsibilities’ for controllers. Executing these responsibilities is not, as the EDPS puts it, simply a ‘box ticking exercise’.³⁸ The challenge lies in implementing these organizational and technical measures as a basis for demonstrating compliance. Thus, it is only through the work of ‘doing’ compliance that accountability comes to life. As Ihde reminds us ‘Left on a shelf, the Swiss army knife or the cell phone “does” nothing’.³⁹ The same can be said for the measures mandated by GDPR. It is only when they are built into the everyday practice that complex negotiations between controller and

33 Abrams, *Demonstrating and Measuring Accountability*, *ibid.*

34 *ibid.*

35 Giovanni Buttarelli, *The Accountability Principle in the New GDPR*, Speech at the European Court of Justice, Luxembourg, 30 September 2016 <https://edps.europa.eu/data-protection/our-work/publications/speeches/accountability-principle-new-gdpr-0_fr> accessed on 21 June 2018.

36 ICO, *Guide to the General Data Protection Regulation (GDPR)* (Information Commissioner’s Office 2018) <<https://ico.org.uk/for-organisations/guide-to-the-general-data-protection-regulation-gdpr>> accessed 21 June 2018.

37 The nature of what should be recorded is laid out in art 30 GDPR. See art 30(5) on conditions when organisations smaller than 250 persons also need to keep records, for example, if they are handling special categories of personal data, information relating to criminal convictions, etc. See also recitals 13, 39 and 82 for more detail on reporting.

38 Buttarelli (n 35).

39 Don Ihde, ‘Smart? Amsterdam Urinals and Autonomic Computing’ in Mirielle Hildebrandt and Antoinette Rouvroy (eds), *Law Human Agency and Autonomic Computing: Philosophy of Law Meets the Philosophy of Computing* (Routledge 2011) ch 1.

user will emerge, and we can understand what an ‘account’ may demonstrably look like.

It is equally difficult to succinctly pin down to whom accountability should be demonstrated. The Madrid Resolution attempts to set up international standards on accountability and states that demonstrations should be to supervisory authorities and data subjects.⁴⁰ However, GDPR is not framed as narrowly. Whilst data subjects and supervisory authorities are clear stakeholders, Article 5(2) is not limited to them, and it is artificial to read Article 5 in isolation of the rest of GDPR, which places many other responsibilities on data controllers. Article 24 specifically focuses on the nature of their wider responsibilities:

Taking into account the nature, scope, context and purposes of processing as well as the risks of varying likelihood and severity for the rights and freedoms of natural persons, the controller shall implement appropriate technical and organisational measures to ensure and to be able to demonstrate that processing is performed in accordance with this Regulation. (Article 24(1))

Article 24 surfaces concepts that need to be read in conjunction with Article 5(2), to situate the full extent of data controller responsibilities in GDPR. We focus on the four key issues emphasized above.

Nature, scope, context and purposes of processing, and risks of varying likelihood and severity. As these two elements are linked, we consider them side by side. In determining the ‘nature, scope, context and purposes of processing’, recital 76 GDPR states ‘objective risk assessment’ is necessary to establish the level of risk attendant to data processing, for example, if it is risk or high risk. Recital 75 provides examples of particular kinds of risk occasioned by data processing, including when it results in discrimination, financial loss, identity theft or fraud, damage to reputation and reversal of pseudonyms, to name a few. Whilst Article 24 requires assessment of risk, in general, it does not call for a DPIA in all cases. However, the focus on risk analysis clearly links to Article 35 which requires a DPIA for processing ‘likely to result in high risks’. The nature of ‘high risk’ is explored in depth A29 WP DPIA guidance,⁴¹ which provides nine examples of high risk processing including processing of data concerning vulnerable data subjects, combining datasets, innovative use of data for new technological/organizational solutions or preventing data subjects accessing a service. Determining the need for DPIAs, and differentiating the distinctions between risk assessment in Articles 24 and 35 is a complex exercise. The nine A29 WP examples are quite broad and many IoT applications will likely require a DPIA. This is not necessarily a bad thing, as DPIAs are an important accountability mechanism providing for ‘building and demonstrating compliance’. Nonetheless, it is uncertain why Article 24 does not just state DPIAs are always necessary, as this seems to be the practical implication of A29 WP guidance.

40 Abrams, *Implementing Accountability in the Marketplace* (n 32).

41 art 29 Working Party (n 30).

Implementation of appropriate technical and organizational measures. The language of ‘technical and organizational’ measures to demonstrate compliance in Article 24 closely aligns with Article 25 requirements on ‘DP by design and default’ (DPbD). DPbD obliges data controllers to safeguard the freedoms and rights of individuals at the time of the determination of the means for processing ‘and’ at the time of the processing itself. This may require minimizing the processing of personal data, pseudonymizing personal data as soon as possible, enabling transparency with regard to the functions and processing of personal data, and allowing the data subject to monitor data processing.⁴² In addition, by default, technical and organizational measures should be taken to ensure that:

1. Only personal data which are necessary for each specific purpose of processing are processed.
2. The amount of personal data collected, the extent of their processing, the period of their storage and their accessibility is controlled.
3. Personal data are not made accessible without the individual’s intervention to an indefinite number of natural persons.

In accordance with Article 24, taking into account the nature, scope, context, purposes and risks of processing, DPbD shall reflect the ‘state of the art’.⁴³ This includes putting appropriate ‘technological’ measures in place to demonstrate accountability and achieve compliance. Recital 63 of GDPR, for example, states that in regard to data subject rights of access:

... where possible, the controller should be able to provide remote access to a secure system which would provide the data subject with direct access to his or her personal data.

We need to acknowledge, then, that GDPR invokes a turn to the systems design community to engage with DP challenges, though we acknowledge the nature of design’s new, explicit role in DP regulation remains unsettled.

Processing is performed in accordance with this regulation. This requirement brings us full circle back to Article 5(2), that the controller shall demonstrate compliance with accountability, which in turn pulls on other GDPR provisions. The ‘lawful, fair and transparent’ principle in Article 5(1), for example, requires a turn to Chapter II GDPR articles on lawful processing (Article 6) and consent (Article 7), to name but two. When Article 5(2) is read in context of GDPR as a whole, we also need to examine the nature of data controller responsibility documented in Article 24. Upon doing this, the ‘breadth’ of responsibilities implicated by the accountability principle become apparent. We believe it requires measures for compliance and subsequent demonstration with the entire GDPR. It is hard to isolate provisions of GDPR, as they often connect to and explicitly call on other provisions. This is clear with

42 Recital 78 GDPR.

43 art 25(1) GDPR.

accountability, which starts as a narrow principle and grows in scope hugely as we dig deeper. Nevertheless, some elements of GDPR align more naturally with the principle. Two examples are transparency (Article 12) and record keeping (Article 30). Thus, accountability turns on the ability to question accounts provided by data controllers around their data handling practices. This requires that record keeping about data processing is in place to demonstrate that compliance with GDPR has been considered and acted upon. Similarly, transparency is intrinsically linked to accountability. Transparency mainly focuses on communication by requiring that processing information be provided in clear, concise language which data subjects (and the public at large) can easily comprehend. As framed in GDPR, transparency is less about providing mechanisms to hold controllers to account. Instead it intersects with accountability by dictating the ‘nature of account giving’.

ACCOUNTABILITY REQUIREMENTS

Translation of the complex provisions of GDPR into more accessible principles is needed if IoT developers are to build accountability into the IoT. We thus propose seven actionable accountability requirements, which seek to address key challenges occasioned by the IoT (as outlined above). These requirements highlight manifold ‘clusters’ of GDPR obligations. This clustering is not exhaustive, but given the broad nature of accountability, we think it provides a useful starting point for considering the nature of an account and substantive elements of GDPR data controllers need to comply with to demonstrate accountability, as outlined below and summarized in Table 1.

Requirement 1: limiting initial data collection

GDPR retains the classic DP principles in Article 5(1) of ‘purpose limitation’, ‘data minimization’ and ‘storage limitation’. According to GDPR, personal data should only be collected for ‘specified, explicit and legitimate purposes’ and not processed in a manner incompatible with those initial purposes.⁴⁴ Only what is ‘adequate, relevant and necessary’ for those initial purposes should be processed.⁴⁵ Furthermore, the data should not be kept in a manner which identifies subjects (ie not anonymized) longer than necessary for these purposes.⁴⁶ Strict oversight over what is being collected, why and how it is managed is necessary.

Requirement 2: limitations on international data transfer

GDPR provides strict requirements on when personal data can be sent outside Europe. Article 44 states data should only be transferred to third countries on basis of various conditions. Article 45 states transfers can occur to countries deemed to provide adequate protection by the European Commission, including Uruguay,

44 art 5(1)(b) GDPR.

45 art 5(1)(c) GDPR.

46 art 5(1)(e) GDPR, with the exception of longer storage for archiving in the public interest, scientific or historical research or statistical purposes.

Table 1: Accountability Requirements in GDPR

Accountability Requirement	Source in GDPR
1. Limiting initial data collection	Purpose limitation Article 5(1b); data minimization Article 5(1c); storage limitation Article 5(1e)
2. Restrictions on international data transfer	Data sent outside Europe on basis of adequacy decision Articles 44 and 45; binding corporate rules Article 47; appropriate safeguards Article 46
3. Responding to the spectrum of control rights	Right to access Article 15; to rectification Article 16; to object Article 21; to restrict Article 18; to portability Article 20; to erasure Article 17; information supply chain (passing down requests for rectification, erasure, restriction) Article 19
4. Guaranteeing greater transparency rights	Transparency of information Article 12; rights to provision of information Articles 13 and 14; algorithmic profiling Article 22; record keeping Article 30
5. Ensuring lawfulness of processing	Legality based on specific grounds (Article 5(1a) and Article 6, eg performance of contract legitimate interest); consent requirements Article 4 (11), Article 7, Article 8 and Article 9
6. Protecting data storage and security	Accuracy of data Article 5(1d); integrity and confidentiality Article 5(1f); breach notification to authorities Article 33 and to data subject Article 34; security of processing Article 32
7. Articulating and responding to processing responsibilities	Articulating responsibilities: Data Protection Impact Assessments Article 35; certifications including seals, marks and certification bodies Articles 42 and 43; new codes of conduct Articles 40 and 41 Responding to responsibilities: DPO Articles 37 and 39; DPbD Article 25

Israel or New Zealand.⁴⁷ Other grounds mandate that appropriate safeguards be put in place (Article 46), such as use of standard DP contract clauses or binding corporate rules that govern data handling in an organization (Article 47). The Privacy Shield 2016 agreement now covers data transfers to USA.⁴⁸ It requires companies apply the principles of notice and choice, and accountability for onward travel. Minimal oversight is provided by the US Department of Commerce.⁴⁹

47 List of Third Countries with Adequacy Decisions <https://ec.europa.eu/info/law/law-topic/data-protection/data-transfers-outside-eu/adequacy-protection-personal-data-non-eu-countries_en> accessed 21 June 2018.

48 It replaces the Safe Harbor Agreement, which was deemed inadequate due to the Schrems decision and Snowden revelations about mass surveillance programmes.

49 Fact Sheet on Privacy Shield <<https://www.commerce.gov/news/fact-sheets/2016/02/fact-sheet-over-view-eu-us-privacy-shield-framework>> accessed 21 June 2018.

Requirement 3: responding to the spectrum of control rights

GDPR provides a spectrum of new control rights around data processing, as described above in Articles 15–21. We frame these as rights users can ‘escalate’ as needed from access (Article 15) to rectification (Article 16), objection (Article 21), restriction (Article 18), portability (Article 20) and ultimately erasure and the ‘right to be forgotten’ (Article 17).

Requirement 4: guaranteeing greater transparency rights

GDPR provides for increased transparency in the relationship between data controller and data subject. Information about processing, particularly concerning data subject rights, is to be provided in:

... concise, transparent, intelligible and easily accessible form, using clear and plain language, in particular for any information addressed specifically to a child ... the information shall be provided in writing, or by other means, including, where appropriate, by electronic means. (Article 12)

Controllers need to furnish the data subject with their identity, the purpose of and legal basis for processing, recipients of their data and so forth (Article 13). They also need to maintain records of processing under their control, including the actors involved, the nature of processing, type of data collected, security measures taken and so on (Article 30). The infamous Article 22 also tackles accountability in algorithms and profiling. It provides a right for data subjects not to be subject to decisions based solely on automated processing where the result has significant legal effects (eg refusal of credit) or similar (eg prejudice from algorithmic profiling).⁵⁰ Measures to protect data subjects should be implemented, at minimum, by providing human oversight over such decisions and enabling subjects to voice concerns and contest outcomes.⁵¹ This assumes that the actions and concomitant reasoning of algorithms can be made accountable, a challenge in itself, particularly for machine learning algorithms deployed in conjunction with IoT devices.

Requirement 5: ensuring lawfulness of processing

Consent is the most discussed grounds for lawful processing of personal data. As discussed in detail above, GDPR provides various requirements for consent mechanisms (see Articles 4, 7, 8 and 9), which are problematic for the IoT. However, consent is not the only basis for lawful processing. Article 6 includes other grounds, namely the legitimate interests of the data controller, the necessity of processing for performance of a contract the subject is party to, or for controller to satisfy a legal obligation they are subject to. Nonetheless, and insofar as the IoT finds its way at

50 Unless provisions in art 22(2) apply, for example, automated processing is by virtue of a contract, authorized by law or by explicit consent of subject.

51 Literature is emerging on the existence of a ‘right to explanation’ and its utility see Lilian Edwards and Michael Veale, ‘Slave to the Algorithm? Why a ‘Right to an Explanation’ is Probably not the Remedy You Are Looking for’ (2017) 16 *Duke L Tech Rev* 1.

scale into consumer goods, consent will remain an important ingredient in ensuring the lawfulness of processing.

Requirement 6: protecting data storage and security

Numerous security and storage requirements exist in GDPR. Accuracy of data is key and appropriate security should be provided, particularly against unauthorized or unlawful processing and against accidental loss, destruction or damage, using appropriate technical or organizational measures (Article 5). This is accompanied by Article 32 requirements to put in place appropriate technical and organizational measures for general security of processing drawing on pseudonymization and encryption, regular security testing, ensuring resilience of services and timely restoration of data after an incident.⁵² GDPR also has strict breach notification provisions around information required and the time frame for reporting, within 72 hours to authorities (Article 33). For data subjects, what is reported and when is more contingent on severity of breach (Article 34).

Requirement 7: articulating and responding to processing responsibilities

GDPR encourages the adoption of mechanisms for data controllers to articulate their responsibilities. Data protection impact assessments have a key role to play in mapping risks, forecasting their likelihood of occurrence, considering appropriate safeguards, implementing these and making this process of reflection public (Article 35). In highlighting compliance with GDPR, an increased role is envisaged for certification processes, using seals and marks (Articles 42 and 43). Similarly, it is envisaged that new industry codes of conduct will emerge (Articles 40 and 41). In responding to established responsibilities, GDPR guides action by controllers. For organizations of a certain size, an appointed DPO will play a key internal oversight and guidance role (Articles 37 and 39). More generally, the turn to technical measures, encapsulated in Article 25 DPbD is 'critical' for the IoT.

BUILDING ACCOUNTABILITY INTO THE IOT

Article 25 introduces DPbD into law. It presupposes not only that technical but 'also' technological measures will be put in place to enable demonstrations of compliance with the principle of accountability. However, as GDPR is 'technologically neutral', it offers no insight to 'systems designers' as to how to build accountability into the technological ecosystem generally or into the IoT specifically. Here we consider how the IoT Databox model⁵³ enables a technological response to the two-part compliance and demonstration requirements of the GDPR Accountability Principle and meets the accountability requirements detailed in Table 1.

52 art 33 GDPR stipulates, 'Taking into account the state of the art, the costs of implementation and the nature, scope, context and purposes of processing as well as the risk of varying likelihood and severity for the rights and freedoms of natural persons', which echoes the DPbD provision in art 25.

53 Alex Hern, "Roomba Maker May Share Maps of Users" Homes with Google, Amazon or Apple' *The Guardian* (London 25 July 2017) <<https://www.theguardian.com/technology/2017/jul/25/roomba-maker-could-share-maps-users-homes-google-amazon-apple-irobot-robot-vacuum>> accessed on 21 June 2018.

Requirement 1: limiting initial data collection

Limiting data collection in the IoT is challenging insofar devices are intentionally designed to collect extensive information in order to provide contextually aware services (such as fine-grained heating management or home security). Limiting initial data collection also sits uneasily with commercial (cloud-based) models underpinning IoT technologies, which seek to repurpose data. Nonetheless, only collecting what is functionally necessary for an IoT device, application or service to operate, that is, for specific purposes, is clearly mandated by GDPR. The IoT Databox model limits data collection through a number of architectural design choices, which are reflected in Figure 1. The IoT Databox model sits on a networked minicomputer that can be situated at the edge of the network (rather than in the cloud) in the user's home and implements the local control recommendation advocated by A29 WP to 'enforce' transparency and control.⁵⁴ It exploits a Security by Design approach at the outset, locally storing data in manifold containers or 'data stores' (rather than one container or store) to minimize the potential attack surface and security problems associated with general purpose operating systems.⁵⁵ The model posits a 'user' (by or about whom data is created), 'data sources' (eg connected devices or online accounts, which generate or contain data about the user), 'data stores' (which collate the data produced by data sources and can be accessed via an Application Programming Interface or API), a 'dashboard' (which allows the user to manage third party access to their data) and 'data processors' (external machines exploited by data controllers who wish to make use of the user's data in some way). Data processing is done by 'apps' which (like data stores) run within isolated containers on-the-box and interact with data stores to perform a specified (purposeful) task. Apps are obtained by users from an 'app store' and app development is supported by a bespoke software development kit or 'SDK'. Apps may query data stores, write to a connected device's store to perform actuation, and/or write to a communications data store if they send the results of data processing to processing entities operating on the controller's behalf. Note, only 'the results' of data processing are distributed by the IoT Databox; the raw data stays on-the-box. The IoT Databox model thus represents a distinctive approach to personal data processing, limiting the amount of data leaving connected devices and aggregating data at the nearest point of data collection (ie, at the edge of the network) to enable 'data minimization'.

The IoT Databox also limits data collection by enabling users to exercise choice, and in the fine-grained ways advocated by A29 WP:

Device manufacturers must provide granular choices when granting access to applications. The granularity should not only concern the category of collected data, but also the time and frequency at which data are captured. (WP223)⁵⁶

54 art 29 Working Party, WP223 *Opinion 8/2014 on Recent Developments on the Internet of Things* (art 29 Working Party 2014).

55 Anil Mudhavapeddy and David Scott, 'Unikernels: The Rise of the Virtual Library Operating System' (2014) 57 *Commun ACM* 61.

56 art 29 WP (n 54).




Figure 2: Enabling granular choice.

The architecture of the IoT Databox negates many—but by no means all—questions to do with international data transfer, in locating data on-the-box, that is, on a physical device situated in the data subject's home. While the question of where the data is stored potentially arises if the data subject runs a virtual backup of the data held on the physical box but in such circumstances, and unlike current cloud-based PDMSS, the decision of where to put the data is ultimately up to the user; the IoT Databox enables data subjects to control where their data is stored. However, the question of where the data is stored also arises with respect to the 'results' of third-party processing run on-the-box: the outcomes of analytics 'can' travel. The manifest again plays an important role here in providing an account to the user not only of what will be done with their data but by whom, including (in the condensed layer) other recipients of the data and, where applicable, relevant adequacy decisions or safeguards. The manifest also makes it accountable (in the short layer) whether or not data will be taken off the box, and whether or not online access is provided to the data subject if so.

That some degree of risk is occasioned by apps that take data off the box is also clearly flagged to the data subject by an app's risk rating. While the IoT Databox cannot prevent data being taken off the box, it can incentivize a reduction in data transfer through the risk rating mechanism, with apps that take data off the box and pass it on to other recipients, particularly those located outside the EU, and/or which do not provide online access being 'severely' rated. An app's risk rating is not only displayed in the manifest but on the app store (Figure 3) to motivate and drive the development of low-risk and even no-risk apps that do not export data, that provide users with granular choice over data sampling and reporting frequency, and that provide online access if apps take data off the box. Low-risk apps approved by the platform display a 'verified' status, and apps that do not take data off the box are also badged with a check mark or tick. Further demonstrations of compliance may be achieved in the future through the use of machine-readable add-ons to data transfers such as blockchain, smart contracts or sticky policies (see, for example, Pearson and Casassa-Mont on sticky policies⁵⁷ or Christidis and Devetsikiotis on smart contracts).⁵⁸

Requirement 3: responding to the spectrum of control rights

In enforcing local control, the IoT Databox negates the spectrum of control rights to some extent, for insofar as data stays on-the-box there is no need for access, rectification, objection, restriction, portability or erasure. External support for the spectrum of rights will still be required in cases where the results of local processing leave the box, though the raw data is retained on-the-box and the specific processing operations performed on the raw data are logged for audit. External support for the spectrum of rights will also be required where IoT devices first export data to the cloud, but that is beyond the IoT Databox's remit. However, insofar as APIs make data

57 Siani Pearson and Marco Casassa-Mont, 'Sticky Policies: An Approach for Managing Privacy across Multiple Parties' (2011) 44 *Computer* 60.

58 Konstantinos Christidis and Michael Devetsikiotis, 'Blockchains and Smart Contracts for the Internet of Things' (2016) 4 *IEEE Access* 2292.

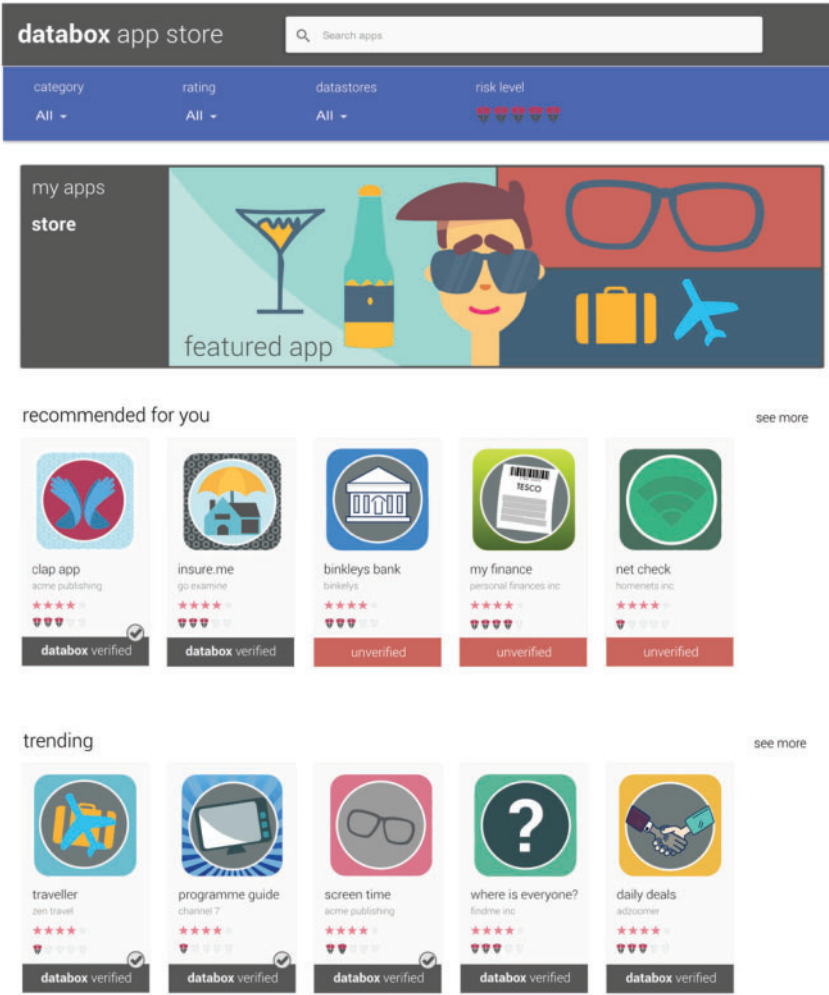


Figure 3: At-a-glance risk (shield) and user (star) ratings.

available to the data subject or data is provided in a common machine-readable format, as per the right to data portability, then data subjects may store such data on the box and reuse it for other purposes.

There is a particular aspect of the right to portability that is problematic, namely it does not cover statistical inferences which are common to IoT data processing. The spectrum of control rights does not necessarily prevent potential ‘harms’ that stem from lack of control over inferences, as opposed to raw data, then.⁵⁹ The IoT Databox seeks to address this situation as by incentivizing local processing. This not includes making potential risk accountable to data subjects but also to ‘app developers’. The IoT Databox model includes a bespoke software development kit or SDK,

59 Urquhart, Sailaja and McAuley (n 19).

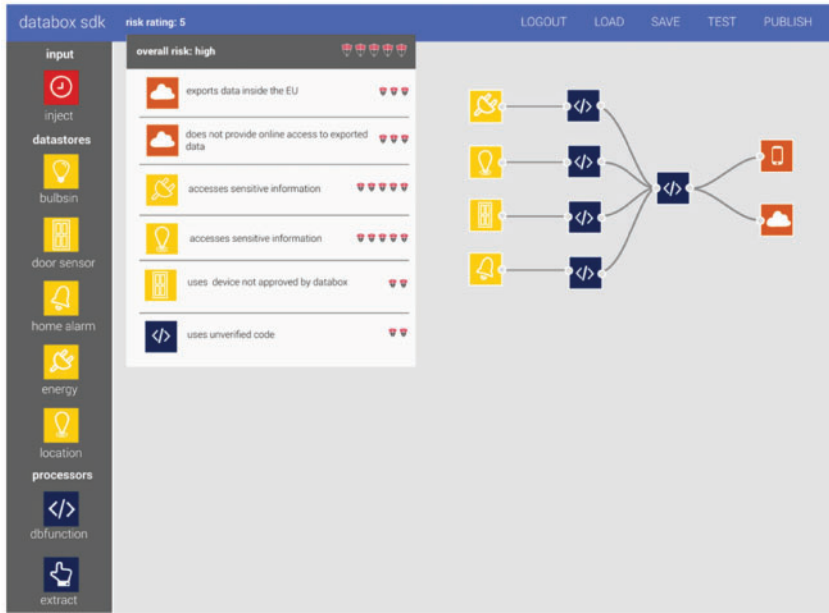


Figure 4: The SDK: risk-rating apps during their development.

which surfaces the potential risks created by an app and makes them accountable to developers (Figure 4).

Our interest in risk goes beyond the requirements of GDPR to also include technological and social risk (eg the risks occasioned by actuating physical infrastructure such as windows and doors or of building apps that exploit sensing to monitor environmental conditions but unintentionally reveal insights into users' everyday lives). Each 'node' in the SDK (Figure 4) allows developers to rapidly construct apps by connecting data, processors and outputs together, has a predefined spectrum of risk attached to it from 0 to 5. The final risk rating assigned to an app sits within this spectrum and is determined by how the nodes are configured (eg the hardware they work with, the data they process, what they do with the data, etc). The final risk rating is encoded in the app and made available to users on the app store and in the app manifest as outlined above but the value here lies in sensitizing developers to risk 'in the course of' app construction, thereby reducing the overhead of building apps that will be severely rated (and potentially unpopular) and promoting the development of low-risk apps that are socially, technologically and legally acceptable. The spectrum of control rights is further enabled by IoT Databox Dashboard's notification function, which allows the results of processing to be 'previewed' prior to distribution (should an app take data off the box), and the Dashboard also allows users to exercise the ultimate sanction and immediately revoke access by terminating data processing entirely.

Requirement 4: guaranteeing greater transparency rights

GDPR establishes a mandate for opening up the opaque IoT and providing more transparent information to data subjects. Transparency is key to enabling control

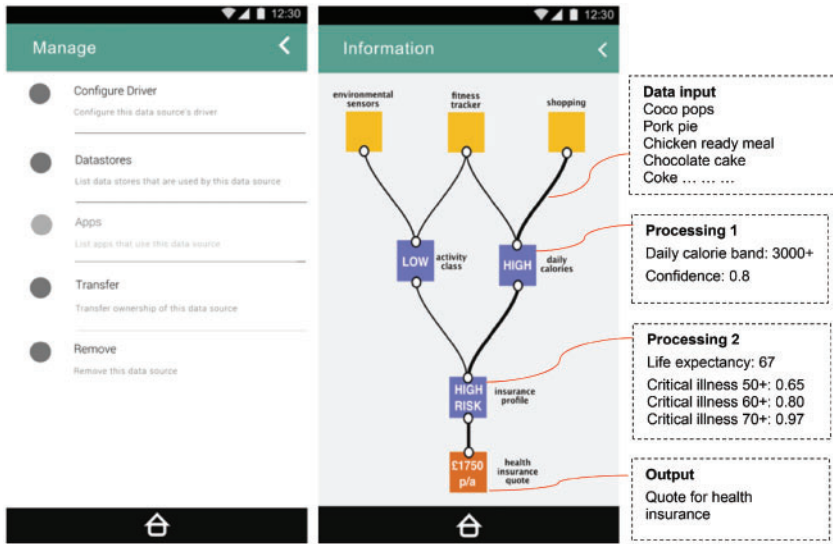


Figure 5: (a) The IoT Databox dashboard (b) inspectable data processing.

rights insofar as data subjects cannot action them without knowing who controllers are, and what they are doing to their data. The IoT Databox takes significant steps to increase transparency in surfacing M2M interactions and the social actors on whose behalf they operate. The manifest clearly plays an important role in this respect, with the multi-layered notice approach scaffolding information depending on intended audience, providing legal and technical information as well as ‘user-friendly’ accounts of apps and their data processing operations. The app store also enhances transparency in providing social feedback through a commonly understood ‘rating’ mechanism.

Further to this, the IoT Databox Dashboard provides the data subject with a range of functions that are designed to make data processing transparent. In addition to enabling users to download and rate apps, the Dashboard (Figure 5a) allows users to manage data sources and device drivers to enable data sources to write to data stores; to manage data stores, including sharing, clearing or deleting stores; to manage notifications, including previewing the results of data processing prior to distribution; and to audit data processing operations, including viewing all accesses to data stores and all data transactions, and enabling data processing to quickly put on hold or terminated. The Dashboard also allows users to inspect ‘how’ data flows through an app and how some action or decision is ‘arrived at’ which speaks to the transparency and accountability in algorithms requirement as detailed in Article 22 GDPR. For example, Figure 5b shows how a health insurance quote is arrived at by processing environmental sensors in the home and fitness tracker data to classify level of activity and processing shopping data to identify calorie band. These results are further processed to predict life expectancy and risk of critical illness. The final results are then used to output a quote. Each of the input and processor nodes is inspectable to reveal the data and processing results. Importantly, the ‘revelation’ is entirely local as processing happens on-the-box, thus ensuring the data subject’s privacy.

Requirement 5: ensuring lawfulness of processing

It is a condition of GDPR that the lawful grounds upon which processing stand be made accountable to data subjects. While a number of legitimate grounds are possible in law, the IoT Databox provides for consent, and consent alone, as the legal basis for data processing and for apps to operate on the platform. The manifest ‘coupled with the underlying architecture’ plays a key role in turning the various permissions implicated in consenting to data processing, particularly the data sources used and the frequency of sampling and reporting, into locally ‘enforceable’ data processing policies. This is done by the IoT Databox’s arbiter component, which issues data store access tokens to apps and checks their data processing permissions. As noted above, the data subject may withdraw consent at any time and terminate processing. Obviously IoT devices which first export data to the cloud and subsequently make it available to data subjects via APIs render data subjects reliant on third-party terms and conditions, including other grounds for processing.

Requirement 6: protecting data storage and security

Given the poor standard of security that currently affects the IoT, there is a lot of work to be done in safeguarding data and satisfying quick data breach notifications. GDPR clearly pushes towards technical approaches to doing this.⁶⁰ A combination of IoT Databox approaches assist with this requirement. First, the IoT Databox distributes data across data stores or ‘containers’ in more technical language. Containers are manifold, with each unique data source having not only a container of its own but a potential array of containers associated with it, which are produced as the result of data processing (eg an app may process data from temperature, humidity and air quality sensors and create a new data store that holds data on environmental conditions in the home). This containerized approach exploits a ‘honeycomb’ rather than ‘honey pot’ approach (unlike centralized cloud servers) forcing an attacker to ‘hack’ each container to access data, which is itself encrypted at rest. This approach does not prevent attacks on online data sources (ie IoT devices which first export data to the cloud), but it does make attacking data on-the-box extremely challenging. Furthermore, the arbiter component (in line with consent permissions) regulates and manages which apps get to run on-the-box, what data are accessed and what processing operations are allowed to run on the data. The IoT Databox also logs all data processing operations, including data export, for audit. These logs may be analysed to identify potential data breaches, though challenges exist in working out how to best achieve this (eg it may be possible to detect some breaches automatically, whereas others may require manual identification by an expert).

Requirement 7: articulating and responding to processing responsibilities

GDPR encourages the adoption of mechanisms for data controllers to articulate their responsibilities, which is key to increasing public trust in the IoT. The IoT Databox SDK provides one such mechanism, enabling a relatively lightweight technological approach towards articulating compliance challenges ‘to developers’ working on a controller or associated data processor’s behalf. This recognizes that IoT developers,

60 art 32 GDPR.

especially those in SMEs and start-ups, often lack the organizational resources needed to understand and respond to these challenges,⁶¹ a situation more broadly confounded by the knowledge and skills deficit confronting organizational actors (eg lawyers lacking technical knowledge and technologists lacking legal knowledge). As one developer puts it, by way of example,

Unfortunately the European Union's new GDPR, introduced on 25th May 2018, creates uncertainty and risk that I can't justify taking. GDPR threatens [developers] with fines of 4% of turnover or €20 million (whichever is higher) if they do not jump through a number of ambiguously-defined hoops. The law, combined with parasitic no-win-no-fee legal firms, puts [developers] at risk of vindictive reporting. Young websites and non-profits cannot afford legal teams. Therefore the risk posed by GDPR is unacceptably high. Perversely, this new EU law hurts small and ethical startups, but helps reinforce the dominance of Facebook, Google and Twitter, who are able to prepare and defend themselves using established legal teams and cash reserves, and who now face less competition from startups. The EU Cookie Law, EU VAT regulation and now the EU GDPR are all examples of poorly-implemented laws that add complexity and unintended side-effects for businesses within the EU.⁶²

The SDK plays a key role in helping developers understand key accountability requirements of GDPR and enables them to respond effectively in articulating a range of potential risks, social and technological as well as legal in the course of data processing app development. The app store in turn conveys the risks occasioned by an app to the public along with an app's verified status and accredits apps that do not take data off-the-box, thus certifying that the 'highest' degree of responsibility is exercised in data processing by app developers and the parties on whose behalf those apps operate.

DEMONSTRABLY DOING ACCOUNTABILITY

As noted in introducing this article, the need to build accountability into the IoT is motivated by the opaque nature of distributed data flows, inadequate consent mechanisms and lack of interfaces enabling end-user control over the behaviours of Internet-enabled devices. This lack of accountability precludes meaningful engagement by end users with their personal data and undermines both trust and the ongoing development of the digital economy. GDPR seeks to remedy the problem by mandating that a rapidly developing technological ecosystem be made accountable. In doing so, it foregrounds new responsibilities for data controllers and anticipates that these will implicate the technological ecosystem itself. Just how these responsibilities, which emphasize putting in place compliance measures and demonstrating accountability, might be exercised in the technological ecosystem is problematic,

61 Lachlan Urquhart (in print), 'White Noise from the White Goods? Conceptual and Empirical Perspectives on Ambient Domestic Computing' in Lilian Edwards, Burkhard Schafer and Edina Harbinja (eds), *Future Law: Emerging Technology, Ethics and Regulation* (Edinburgh University Press forthcoming) <<https://arxiv.org/pdf/1801.07185.pdf>> accessed 21 June 2018.

62 Streetlend is no more <<https://www.streetlend.com>> accessed 21 June 2018.

Table 2: How the IoT Databox Meets Accountability Requirements

Accountability Requirement	IoT Databox Compliance Measure
1. Limiting initial data collection	<ul style="list-style-type: none"> • The IoT Databox architecture situates data processing at the edge of the network in the data subject's environment, enables the data subject to control external access to data via app manifests that provide granular choice encoded as enforceable data processing policies on-the-box, and constrains data distribution to the results of processing.
2. Limitations on international data transfer	<ul style="list-style-type: none"> • Insofar as data is not taken off-the-box by an app, the requirement is negated. • Insofar as data is taken off-the-box by the user backing up data in the cloud, the user retains control over where the data is placed for storage. • Insofar as data is taken off-the-box by app, the manifest provides the data subject with the information required by Article 13 GDPR in the condensed layer of the manifest's multilayered notice. • The app manifest also flags data transfer as a risk in the short layer, which increases in direct relation to access, further recipients and their location, adequacy decisions and safeguards.
3. Responding to the spectrum of control rights	<ul style="list-style-type: none"> • Insofar as data is not taken off-the-box by an app, the requirement is negated. • Insofar as data is taken off-the-box by an app, the spectrum of rights is outside the control of the IoT Databox, though the raw data is retained on-the-box and the specific processing operations performed on the raw data are logged for audit. • Insofar as data is taken off-the-box by an app and online access is provided, any data made available to the data subject (either via APIs or in common machine-readable formats) may be imported into the IoT Databox for other processing as per the right to data portability. • As the right to data portability does not include statistical inferences derived from data processing, the IoT Databox seeks to incentivize the construction of apps that do not take data off-the-box through risk-rating mechanisms in the SDK, the app store and the app manifest, and through accreditation. • The IoT Databox allows data subjects to preview the results of data processing prior to distribution, to

(Continued)

Table 2: (continued)

Accountability Requirement	IoT Databox Compliance Measure
	terminate data processing at any time and to revoke access.
4. Guaranteeing greater transparency rights	<ul style="list-style-type: none">• The app manifest surfaces M2M interactions and the social actors on whose behalf they operate and processing takes place.• Manifests exploit a multilayered approach to make data processing accountable in legal, technical and ‘user-friendly’ terms.• The dashboard enables the data subject to view data processing in real time, including notifications of processing operations, previews of the results of processing and inspection of how results were arrived at as per Article 22.• The app store exploits a commonly understood social ‘rating’ mechanism to enhance transparency.
5. Ensuring lawfulness of processing	<ul style="list-style-type: none">• Data processing on the IoT Databox operates on the grounds of consent only.• Consent is provided for by the app manifest, with specific permissions to access data sources and granular choices over data sampling and reporting frequencies being encoded as enforceable processing policies by the IoT Databox arbiter.• Consent may be withdrawn and processing terminated by the data subject at any time.
6. Protecting data storage and security	<ul style="list-style-type: none">• The IoT Databox stores data in a distributed array of containers, which encrypt data at rest.• The arbiter component regulates and manages app access to data stores and data processing operations, based on permissions set by the data subject.• All data processing operations, including data export, are logged for audit.
7. Articulating and responding to processing responsibilities	<ul style="list-style-type: none">• The SDK articulates compliance challenges to IoT developers in the course of app and manifest construction.• The app store articulates risks associated with an app to the public, including legal, technical and social risks, and an app’s verified status.• The app store also accredits apps that do not take data off-the-box, certifying the highest level of responsibility in data processing.

however. This is not only because GDPR is ‘technologically neutral’ but because it does not explicitly define what compliance with the accountability principal should amount to or what an account should look like. In reviewing GDPR, we have therefore sought to identify key components of an account and to subsequently articulate how these might be built into the technological ecosystem with specific respect to the IoT. We have thus explored how the IoT Databox might enable the implementation of compliance measures and demonstration of accountability. Accountability requirements are laid out in [Table 1](#). The specific ways in which the IoT Databox implements measures to comply with them are laid out in [Table 2](#).

The IoT Databox provides a *prima facie* example of what a demonstrably compliant account might ‘look like’ from a technological perspective. It also makes it perspicuous to whom accountability is demonstrated: primarily to data subjects and IoT developers/data controllers, but the demonstration also extends to a range of other parties from technical experts to regulators and supervisory authorities, whose inquiries might be facilitated by the transparent and auditable characteristics of the IoT Databox. The account is articulated through a distinctive software architecture and development environment that implements compliance measures (data limitation, enhancing transparency, articulating responsibilities, etc) and interactional arrangements (app stores, apps, manifests, dashboard, etc) that are designed to demonstrably respond to the accountability requirements of GDPR. It is in ‘doing being accountable’ as a demonstrable feature of constructing, publishing and using IoT Databox apps that compliance is therefore achieved. The demonstration reflexively tackles the various regulatory problems of the IoT, which occasion the need for accountability in the first instance. Thus, the IoT Databox makes an opaque technological infrastructure visible, provides clear consent mechanisms rendering data processing legible and enables data subjects to control the flow of data. It also provides oversight on M2M communications along with what the data is being used for, where and by whom. In making data processing in the IoT accountable, the IoT Databox seeks to go beyond remedying existing problems with the IoT, however. Ultimately it seeks to foster a culture of accountability that results in widespread data minimization, and local processing wherever possible to engender widespread trust in the IoT.

Data protection must move from ‘theory to practice’ . . . accountability based mechanisms have been suggested as a way [to] . . . implement practical tools for effective data protection.⁶³

FUNDING

This research was funded by the Engineering and Physical Sciences Research Council [grant numbers EP/M001636/1, EP/N028260/2, EP/M02315X/1].

63 art 29 Working Party, *WP173 Opinion 3/2010 on the Principle of Accountability* (art 29 Working Party 2010).