Building Stakeholder Trust in Internet of Things (IoT) Data Services using Information Service Level Agreements (SLAs)

C. Peoples, M. Abu-Tair, B. Wang, K. Rabbani, P. Morrow,
J. Rafferty, A. Moore, and S. McClean
Ulster University, School of Computing
Jordanstown Campus
Newtonabbey, UK
{c.peoples; m.abu-tair; b.wang1; rabbani-k; pj.morrow;
j.rafferty; aa.moore; si.mcclean}@ulster.ac.uk

M. Fisher
BT
Adastral Park
Martlesham
Ipswich, Suffolk, UK
mike.fisher@bt.com

Abstract-A Service Level Agreement (SLA) defines a contract between network service providers and consumers, specifying the terms of a service which providers will make available and the conditions which consumers will accept. To date, SLAs have been specified using basic terms, such as availability and network performance, with a consumer being compensated in the event that the service provided does not meet the terms agreed. Given changes in the ways which network services are now made available, however, SLA terms are changing to capture both the differences in service provision and, additionally, in the responsibilities of the parties involved. It is this aspect of information SLAs which we respond to in this work, and we propose a SLA model which accommodates the requirements of these new relationships. We also propose a set of metrics, a selection of which are presented in this paper to demonstrate our concept, and recommend that a selection can be adapted by consumers. Finally, due to the intricate relationships between data consumers and data providers in the IoT environment and the fact that metric adaptation may lead to SLA violation, we discuss SLA conflict resolution through prioritizing non-functional metrics on a per-customer basis.

Keywords—Conflict Resolution, Consumer Trust, Internet of Things (IoT), Policy-based Network Management (PBNM), Service Conflict, Service Level Agreement (SLA).

I. INTRODUCTION

In the IoT environment, a vast amount of personal data can be collected from citizens, who may not have control over its collection, availability and use. The IoT continues to evolve in an ad hoc, non-standardized and unregulated manner and, due to this, in addition to data subjects generally being unhappy about the volume of data collected and how their data is or may potentially be used, there has been, to date, a general lack of trust in data collected using IoT technologies [1]. Exacerbating this problem is the fact that it is difficult for citizens to enforce a right to remove their personal data on demand. It is in response to such concerns that the General Data Protection Regulation (GPDR) will enforce that IoT service providers set in place actionable solutions on behalf of their customers. These solutions will be enforceable using Service Level Agreements (SLAs).

A Service Level Agreement is part of a contract between the provider of a service and its end-user or customer that defines which services the provider will offer and the level of performance it must meet. The customer will agree to this level of service as being acceptable. In addition, a SLA will incorporate any remedies or penalties should the agreed-upon levels not be achieved. When the service level which was agreed upon in the SLA is unfulfilled, the customer is credited for lost service, measured by the duration of downtime. The SLA captures both what the service provider will offer and what the customer will accept. A service provider is not required in all scenarios; two parties, which might involve a data provider and a data consumer, will have some agreement which governs their expectations and obligations associated with exchanging information between them. In this paper, however, we are taking the viewpoint of the service provider and investigating what it can do to increase trust in the services it is offering to its customers by, for example, managing SLAs with original sources of information which it is using to derive services for customers. One example of a particular kind of SLA is measured according to system uptime, an aspect which is generally specified as a service being available and accessible, on average, for example, 99% of the time. SLAs, to date, are measured by provider performance, and not for the behavior of customers in their use of the service. Using traditional SLAs, a customer could purchase storage space made available by a service provider to host their personal data; the SLA for this service will specify any potential risk associated with the platform's availability and the customer's compensation in the event that the platform's availability fails to be fulfilled. Data uploaded to the IoT platform and made publically available has, so far, been done on the free will of humans engaging in these environments - they have not been committed through contract in their IoT activity; there are, for example, no penalties if a customer who is storing data on the service provider's platform suspends delivery of a dataset without warning.

The IoT environment, however, is one in which services can be made available using the datasets provided by customers, hosted on the public platform. This can occur in the instance that the customer permits public access to their data, which may be anonymized or not. In this way, service consumers implicitly become service providers for other service consumers. This is where relationships in the IoT become more difficult to manage, and also govern using SLAs – the scenario described above would require two SLAs if managed via an intermediary service provider, as opposed to directly between the data provider and

978-1-5386-4980-0/19/$31.00 ©2019 IEEE

the data consumer. In the absence of SLAs which act as contracts to define and agree the behavior of all participants, customers who subscribe to IoT data use are vulnerable to the publicly-available data becoming unavailable, being provided in an inconsistent manner, or not being kept up-to-date. SLAs can support customers in their vulnerable positions where they may come to rely on these datasets, and provide protection for situations where promised, but not guaranteed, service levels fail to be achieved. As a SLA defines, "*obligations that each actor must meet*" [3], one could have a provider reimburse a customer if they are unable to access a dataset provided by another customer to which they have subscribed; it would be expected that the data provider would be penalized by the service provider for failing to fulfil their obligation of providing data throughout the period to which they had agreed in their SLA. This is a complex management challenge.

In addition, more complex relationships in the IoT are also possible; for example, if the Service Provider takes a more active role in curating, processing, and selling access to data, then data may be provided, implicitly or explicitly, by an individual as a consequence of using some service or IoT device and the service provider could then seek to exploit this data itself or by selling services to third parties. This sort of behavior should be explicit in the SLA and clear to all participants.

Therefore, given changes in the ways which network services are now made available in the IoT - "*The relationship will be between people-people, people-things, and things-things*" [4] - terms specified in SLAs are similarly changing to capture the differences in party responsibilities. It is this aspect of SLAs which we respond to in this work. We propose that SLAs are made available separately to both customers who are consuming data from the IoT and customers who are providing data to the IoT. We make recommendations with regard to a selection of the metrics considered to be necessary in next generation SLAs to support the functional and non-functional requirements of customers, given the role played by citizens in providing data and the dependence of other citizens on it. (The list of SLA metrics presented in this paper is not exhaustive, and includes a selection to demonstrate our concept.) We also describe the ways in which conflicts can occur as a result of the metrics chosen to capture the functional and non-functional requirements in SLAs supporting both data providers and data consumers, and the ways in which conflicts may be resolved through SLA attribute prioritization. In our definition, it is our objective that a service provider needs to be able to manage the risk it takes on by having a set of SLAs with both providers and consumers of data where the ability to deliver against some SLAs depends on relationships with third parties, which are governed by different contracts and SLAs, particularly where some of these may be changed dynamically as a result of right of erasure or withdrawal of consent, for example.

The remainder of the paper is organized as follows: In Section II, a literature review is carried out of the key aspects which are pertinent to the SLA metric definition presented in this paper. This takes into account the state-of-the-art in SLA definition, customer acceptance of risk in IoT services, the capture of functional and non-functional requirements in IoT SLAs, and the prioritization of SLA metrics during periods of policy conflict. Our research proposal is presented in Section III, and continues in Section IV, where we describe the approach proposed to manage conflicts which might occur between SLA contracts upon the renegotiation of their metrics. The paper concludes and discusses future work in Section V.

II. LITERATURE REVEIW

Concepts presented in this section are reviewed from the perspective of their contribution to SLA definition for the IoT, in preparation for our SLA definition proposal in Section III.

A. State-of-the-Art in IoT SLAs

When we consider SLAs for online services in general, we might think about performance, reliability, availability, security, and so on. Within the context of a traditional SLA, in the sense of one defined prior to the IoT, services are offered according to *delivery* (customer commit date, service credits for delivery, application of service credits), *availability* (measurement of downtime, service credits for downtime), and *network performance*. In terms of *availability*, a SLA can be offered for an uptime between 97% and 99.99%. The price paid for the service will depend on the uptime package requested, with any deviation below the agreed uptime repaid in service credits.

SLAs have been typically commonly defined using a standard language, with a well-recognized vocabulary. This is evident in the SLA defined above, with system operation and performance measured using the common metrics of *availability* and *network performance*. A standard definition with easily recognizable terms is important, allowing the SLAs from one service provider to be directly compared with those from others. Similar terms are observed in SLAs for IoT platforms: The Microsoft Azure IoT Hub, as one example, is a platform which can be used to, "*securely connect, monitor and manage billions of devices to develop Internet of Things (IoT) applications*" [5]. Azure services are provided with a SLA, and customers therefore pay for an agreed level of service. Service levels for the Microsoft Azure Hub are characterized according to: *Deployment minutes*, *Maximum available minutes*, *Downtime*, and *Monthly uptime percentage*. The SLA defines the credit percentage given in the event that platform availability falls below 99.9% uptime (< 99.9% uptime is compensated with 10% service credit, < 99% uptime is compensated with 25% service credit). The Amazon AWS platform offers a similar service (< 99.99% uptime is compensated with 10% service credit, and <99% uptime is compensated with 30% service credit) [6]. Terminology used to describe an Amazon AWS SLA includes: *Monthly Uptime Percentage*, *Availability Zone*, and *Unavailable/Unavailability*. The metrics in this instance are influenced by the operational design of the Amazon platform, with organization of its physical and virtual resources according to regions, with the understanding that it may be due to region unavailability that the SLA is unable to be fulfilled.

While similar terms are therefore observed in SLAs for both IoT and pre-IoT networks in the sense of platform availability, there are typically additional attributes considered for the IoT due to the service levels and configuration choices on offer. Google Cloud [7], as one example, offers services which are influenced by aspects including:

Sustained use discounts: Automatically up to 30% off workloads that run for a significant portion of the billing month.

978-1-5386-4980-0/19/$31.00 ©2019 IEEE

Pre-emptible VM instances: Up to 80% off workloads that can be interrupted.

Committed use discounts: Purchase resource requirements for a committed usage term of 1 year or 3 years.

Each service offered will have an associated SLA which defines the expectations that customers can have in relation to its operation. While traditional services, together with their SLAs, were relatively generic and intuitive to understand, IoT services can be more technical, and may not be obvious to the general public. It might be expected, for example, that a general customer will not appreciate what a "*Pre-emptible VM instance*" is, nor why it is important in supporting their service provision. Nonetheless, it is possible to appreciate the evolution of SLA definition in terms of the range and technicality of metrics: *"Although the engagement costing model takes into consideration many delivery factors such as the number of the managed servers and the type of managing tools, it assumes the use of a standard set of service level agreements (SLAs). This … limits its capability to handle non-standard SLAs that are frequently asked by the customers to fit their special needs"* [8]. It is in this vein that the SLA proposal is presented in this paper.

B. Risk Acceptance in IoT SLAs

In [2], it is noted that, "*A lack of trust and understanding among users could become a barrier to the continued development of innovative services and applications, and the benefits for consumers that they bring*". However, with the evolution of SLAs to bolster service levels across the IoT, it is interesting that the authors in [8] explore the likelihood that customers will accept variable levels of risk from their SLA. The purpose of their study is to assist customers with evaluating service providers such that they can select the one which most suitably meets their needs while being aware of the limitations associated with each. Risks are captured in their study in terms of six service areas, which include: *availability*, *reliability*, *number of instance types*, *response time*, and *data storage*; these are prioritized on a per customer basis. The authors advise that the estimated service level achievable, given the real-time network condition, should be reinforced clearly to customers before they sign the SLA contract. This information is calculated using historic performance data for the previous six months, and the rise and fall of each of the performance metrics over that period of time. It is based on the priority of each performance metric that the customer can determine suitability of the service based on their needs. Taking the findings presented in this work into account, the SLA definition presented in this paper allows customers to accept a degree of risk in association with the network performance before their SLA is considered to have been violated, and is subsequently terminated.

C. Functional & Non-functional Requirements in IoT SLAs

We distinguish between functional and non-functional requirements as functional being *what* the system should achieve, and non-functional being *how* the system should achieve it. Functional requirements are commonly not observed in traditional SLAs - as described in Section II.A, the attributes in non-IoT SLAs typically relate to the level of performance achieved, such as the platform's uptime, and are therefore largely non-functional, performance-based. There is an opportunity, however, for the inclusion of functional requirements in IoT SLAs, which may help to increase customer trust in its operation and management. Kounelis et al. (2014) express the opinion that trust in IoT services is dependent on the users' ability to negotiate the terms of their SLA [9]. This is somewhat agreed with by the authors in [10], who have the opinion that, "*individuals should have a basic right to opt-out, delete, or mask their information from systems in the IoT, …*" However, published in 2012, it was noted that, "*The typical model being that the user relies on the Service Provider to deploy data/services in a cloud environment according to specific requirements*" [11]. In our opinion, the inclusion of functional requirements in next generation IoT SLAs will help to control the type of services being offered. This takes into account the ability to access a particular dataset, to control that a specific dataset is forwarded to a particular device endpoint, to view a historical dataset from a particular device, or to aggregate statistics for a particular data source. Furthermore, given the dynamic and evolving nature of the IoT, we recommend that these functional requirements are adaptable on customer demand, in recognition of the fact that there will be charges associated with some SLA changes and ability to cancel without penalty if the other party introduces a change. This is in addition to the fact that some parts of a SLA may need to be changeable where required by law or regulation. It is therefore also from this perspective that the SLA proposal is presented in this paper.

III. Research Proposal: SLA Terms for the IoT

In this section, we define the terminology used in our proposal for SLAs which are specific to the IoT. This terminology takes into account our perception of the SLA model in terms of the significant parties involved, and the functional and non-functional service requirements necessary to support their needs.

A. Components Contributing to IoT SLAs

The SLA model proposed exists between the *Customer* and *Service Provider*. We consider customers to include both *Data Providers* and *Data Consumers*: *Data Providers* consume storage resources and *Data Consumers* consume data resources. The *IoT Data* and *IoT Infrastructure/IoT Hub* are also considered to be components contributing to the SLA in the sense that each must be available to support the *Data Provider* and *Data Consumer* in fulfilling their role. Relationships

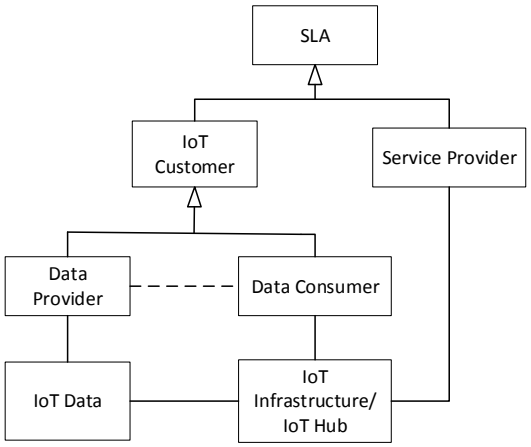

Figure 1 Components Contributing to the IoT SLA

978-1-5386-4980-0/19/$31.00 ©2019 IEEE

between these terms are presented in Figure 1, and are described in more detail below:

IoT Infrastructure: The IoT Infrastructure describes the network and storage resource capacities which are deployed, operated and managed by the Service Provider.

IoT Hub: The IoT Hub refers to the platform on which data collected from IoT sensors reside. We make an assumption in this business model that the IoT Hub is distributed on an international basis, and also that IoT Hub data may be replicated internationally. The IoT Hub is operated by the Service Provider.

IoT Data: IoT Data refers to metrics produced from IoT sensors which is disseminated for storage on and distribution by the Hub. IoT data will be characterized according to attributes which include those presented in Table 1 and Table 2.

IoT Customer: An IoT Customer can access datasets stored in the IoT hub, and/or upload data for storage in the IoT hub. IoT Customers include both Data Consumers and Data Providers.

Data Consumer: A Data Consumer accesses data stored on the IoT Hub. The Data Consumer is bound by a contract.

Data Provider: A Data Provider contributes data from one or more sensors to the IoT Hub. Their participation in the IoT scenario is bound by a contract.

Service Provider: A Service Provider provides the contract to Data Providers and Consumers, agreeing to separate contracts with each Data Provider and Data Consumer.

Together, these terms refer to the IoT environment for which the SLA structure is presented in this paper. A Data Provider and a Data Consumer, both of which are IoT Customers, have a relationship with one another, given that a Data Consumer will be able to access the IoT Data which has been made available by the Data Provider (assuming that the Data Consumer has the appropriate access rights) and which is hosted on the IoT Infrastructure/IoT Hub. The Service Provider makes the IoT Infrastructure/IoT Hub available for IoT Customers, and both parties agree to (providing and accepting) a certain quality of service, which is captured in the SLA. It is not required that each of the human roles represented in Figure 1 always exists; for example, a Service Provider is not always required, as in instances where we have a direct relationship between a Data Provider and a Data Consumer. Within the context of this model, we consider a SLA as being strictly between 2 parties. Metrics used within the SLA Contract to support the agreed roles and responsibilities of the SLA parties are presented in Section III.B (attributes supporting functional requirements) and Section III.C (attributes supporting non-functional requirements).

B. SLA Attributes Supporting Functional Requirements

A selection of attributes which will be agreed in the Data Consumer and/or Data Provider's SLA are presented in Table 1; these attributes will work to fulfil the achievement of the functional requirements of these parties. One particularly important aspect which should be included in a SLA is the unique identifier for the dataset to which the Data Consumer has subscribed and which the Data Provider is supplying (`Dataset ID`). This will result in the Consumer agreeing to be able to access a particular dataset for an agreed minimum period of time

Table 1 A Selection of SLA Attributes Supporting Customer Functional Requirements*

Data Attribute	Attribute Description	Data Consumer / Data Provider SLA Attribute	Is Attribute Value Renegotiable?
`Dataset ID`	Unique dataset identifier	Consumer and Provider	x
`Minimum Duration Data Collection`	Minimum duration of time the data will be collected	Provider	✓
`Minimum Duration Data Available`	Minimum duration of time the data will be available	Provider	✓
`Anonymization Required`	Ability to anonymize the dataset	Provider	✓
`Interruptible`	Can delivery of data to the Hub be interrupted?	Provider	✓
***This list is not exhaustive**			

(`Minimum Duration Data Available`). Another requirement relates to the Data Provider ensuring that the dataset which is being contributed or has been contributed to the Hub having been collected for an agreed duration of time (`Minimum Duration Data Collection`). This helps to ensure that the minimum size of dataset provided is sufficient for the Data Consumer's needs. `Anonymization` refers to the ability of a Data Provider, or *obligation* if the data provider is not the data subject but is providing data from third parties under the terms of a separate SLA, to remove any data attributes which might be used to identify them personally, such as their location and date of birth. We make an assumption that data will be anonymized by default. (Note that anonymization and access control are not synonymous – data could be anonymous and publically available.) We also consider the ability for a Data Provider to agree to the delivery of their data to the Hub to be interrupted, which will help to avoid SLA violation that can occur when the IoT Infrastructure is constrained (`Interruptible`).

All attributes in Table 1, apart from `Dataset ID`, could be configured to be renegotiable on-demand by either Data Provider or Data Consumer customers, dependent on the contract, once the SLA moves into the active state (assuming that the SLA will progress through a life cycle, such as one presented in [12]). It is a result of the ability to re-negotiate these terms which help to support the achievement of a customer's functional requirements, that conflicts and SLA violations can occur. This is discussed in more detail in Section IV.

We do not consider this to be an exhaustive list of metrics to support IoT SLA requirements, but include only a selection to demonstrate the way which problems with IoT SLAs can evolve and how our proposal can help to resolve them while fulfilling customer service levels. These can be described as attributes supporting functional requirements because they aid the function of uploading and downloading data to and from the

978-1-5386-4980-0/19/$31.00 ©2019 IEEE

Hub, for example, and define the way that data should be presented, both of which are aspects of *what* the system can do.

C. SLA Attributes Supporting Non-Functional Requirements

A selection of the non-functional attributes used within customer SLAs are presented in Table 2. As one non-functional requirement, we consider a customer's ability to renegotiate the terms of their SLA (`Renegotiable`) once it has become active, to support the desire for increased flexibility of SLA configuration. This could be used to influence their data confidentiality, with customers initially allowing all data collected to be publically available, and later adapting this condition of their SLA so that it becomes anonymized. Using the `Renegotiable` attribute, data consumers can also renegotiate the level of service required, and accept a lower level of service when the network is constrained and compromised, without their SLA being violated. This is achievable using the `Acceptable Risk` metric, which allows a customer to indicate that they can tolerate a dataset becoming unavailable, as one example of a lower quality service. This metric can be quantified according to a value of true or false; a value of true indicates that a data consumer is prepared to accept risk in the sense of data becoming unavailable if the provider changes its mind.

Each of the non-functional requirements presented in Table 2 is renegotiable: The renegotiable capacity of the SLA terms introduces a layer of complexity to the management of SLAs due to the knock-on effects which each renegotiation can have. A Data Provider who once made their data non-anonymized and publicly available, for example, can create a problem if a Data Consumer who was accessing the dataset has captured in their SLA that they are unable to cope with any risk with regard to its public availability. Ability to renegotiate the terms of the SLA metrics therefore increases the management challenge through the need to deal with real-time conflict between SLA policies.

Table 2 A Selection of SLA Attributes Supporting Customer Non-functional Requirements*

Data Attribute	Attribute Description	Data Consumer / Data Provider SLA Attribute	Is Attribute Value Renegotiable?
`Renegotiable`	Do data consumer and provider wish to have option to renegotiate their SLA once Active?	Consumer and Provider	✓
`Accept Anonymised Data`	Will anonymized datasets be acceptable to the data consumer?	Consumer	✓
`Acceptable Risk`	Can data consumer accept risk with regard to dataset availability?	Consumer	✓
***This list is not exhaustive**			

IV. Managing SLA Conflicts between Service Providers and Data Providers/Consumers

Given the design approach which we are proposing for IoT SLAs, with the addition of metrics supporting functional and non-functional requirements, and ability for customers to re-negotiate the terms of their SLA on-demand, we recognize that conflicts become possible between SLAs: changes to the SLA terms of one customer may cause the SLA of another customer to be violated. In this section, we therefore discuss the dependencies between several attributes which are used in the proposed SLA definition to demonstrate the reasoning behind our concept, together with consideration of the ways in which SLAs may be violated as a consequence of them being modified on-demand. The section concludes by considering the prioritization of SLA attributes such that SLA violations are avoided and customer service can continue unaffected.

A. Conflict Potential between IoT SLA Requirements

Within the context of our SLA proposal, we recommend that customers can adapt certain aspects of their SLA on-demand. This is in line with research, such as [13] and [14], which advocates that flexibility can improve trust levels. Additionally, in line with the General Data Protection Regulation (GDPR) [15], it is necessary to provide ability for individuals to make changes to previous agreements and consents. A selection of the reconfigurable metrics which we propose for inclusion in IoT SLAs are presented in Table 3. We also acknowledge that changing the values of these metrics can lead to SLA conflicts, and therefore describe in Table 3 the ways in which a change to one SLA may lead to violation of another.

In recognition of the SLA conflicts defined in Table 3, we therefore propose that prioritization of the non-functional requirements should also feature as a capability of IoT SLA management, as described in Section IV.B. Attributes leading to policy conflict will be monitored throughout the lifetime of the consumer's subscription. These are difficult to predict as they

Table 3 Dependency Conflicts with SLA Configurations

Data Attribute	Attribute Change	Dependency Conflict
Consumer: `MinPeriod DataAvailable` Provider: `MinDuration DataAvailable`	Reduction in `MinDuration DataAvailable` to be less than `MinPeriod DataAvailable`	The dataset will not be available for the duration of time which the consumer has subscribed to.
`DatasetPublic Availability`	From available to unavailable	The dataset will not be accessible for the duration of time which the consumer has subscribed to.
Consumer: `AcceptableRisk`	From being able to cope with risk to being unable to cope with risk	The consumer has indicated that they can tolerate the dataset becoming unavailable, but then become unable to cope with any availability change.

978-1-5386-4980-0/19/$31.00 ©2019 IEEE

depend on the change of a single attribute value at the customer's discretion. Prompt action needs therefore to be taken once violation has occurred to satisfy consumer requirements.

B. *Prioritising Non-functional Requirements in the IoT*

Priority in relation to non-functional SLA requirements lies in the customer's ability to renegotiate the terms of their SLA. Customers are required to indicate in their SLA their ability to cope with risk (`Acceptable Risk` in Table 2): risk in this context refers to the fact of a dataset which a customer has subscribed to becoming unavailable during the period which the Data Provider indicated that it would be available for (`Minimum Duration Data Available` in Table 1), a situation which might occur due to the Data Provider suspending their collection of the data or changing the anonymity of the data from being public (individual identifiable) to private (anonymized) and thereby restricting its access. Customers are also required in their SLA to indicate their desire (potential) to re-negotiate the terms of their SLA once it moves into the active state. Both aspects will have a bearing on the SLA pricing, with it perhaps being an attractive proposition to encourage customers to accept both risk and re-negotiation of their SLA terms in order to tolerate constrained network conditions and unpredictable customer behavior – maximizing the opportunity that a SLA is not violated will help to satisfy customer requirements.

As an example of how the attributes defined in Table 1 and Table 2 are used to influence a SLA agreed between a Data Consumer and Service Provider, consider the following: in the event that Customer 1 indicates that he is unable to tolerate any risk with regard to dataset availability, Customer 1 will not be offered the dataset of Customer 2, if Customer 2 is able to re-negotiate the terms of his SLA to make his originally public data now private - this action would cause a SLA violation for Customer 1. Prioritization in this case is placed on Customer 1's ability to cope with risk, which then allows Customer 2 to adapt the privacy of his data as desired, without negatively impacting Customer 1, who is unable to tolerate this event.

V. Conclusion & Further Work

Earlier in this paper, it was stated that, "*A standard definition with easily recognizable terms is important, allowing the SLAs from one service provider to be directly compared with those from another*". We then presented a selection of the attributes which we consider could be used to support next generation IoT SLA operation, fulfilling customers' functional and non-functional requirements while avoiding SLA violation and subsequent SLA termination. It is obvious from the limited amount of detail provided in relation to our SLA approach that this indeed is *not* a standard definition, particularly when compared with the SLAs on offer from other service providers which include metrics such as availability. It is our opinion that a standard definition which captures the range of customer configuration requirements and service choices possible is difficult to achieve in practice; furthermore, we do not wish to suggest that the services made available by one provider can be compared directly with those from another (although, in line with GDPR, we recognize that some parts should be comparable, particularly those relating to personal data where informed consent is a requirement of use). We believe however, that the way in which this service is being constructed in our proposal fulfils the requirement of "*easily recognizable terms*" to support the general public for whom this mechanism is essentially designed. As some services are inherently technical, we recognize in other SLAs that the terminology used may not be understandable for a general user. Continued lack of transparency in IoT service provision will allow the low level of customer trust in IoT systems observed to date to continue.

Further work is now continuing to expand the set of terms used to provide configurable SLA options to customers and mechanisms to accommodate policy conflicts. We will propose these to support the overall objective of increasing trust, and therefore ideally participation and usage, of future IoT systems.

Acknowledgement

This research work was conducted under the BT Ireland Innovation Centre (BTIIC) project and was funded by Invest Northern Ireland and BT

References

[1] Cisco, "*Cisco Survey Reveals Divide between IoT Value and Trust*," Dec. 2017; Available: https://newsroom.cisco.com/.

[2] Communications Consumer Panel, "*Online Personal Data – Communications Consumer Panel*"; Available: https://www.communicationsconsumerpanel.org.uk/.

[3] G. Gaillard, D. Barthel, F. Theoreyre, and F. Valois, "*SLA Specification for IoT Operation – The WSN-SLA Framework,*" [Research Report] RR-8567, INRIA, 2014, pp. 71.

[4] J. Morgan, "*A Simple Explanation of 'The Internet of Things',*" Forbes, May 2014; Available: https://www.forbes.com/.

[5] Microsoft Azure, "*Service-Level Agreements*"; Available: https://azure.microsoft.com/.

[6] AWS, "*AWS IoT*"; Available: https://aws.amazon.com/iot/.

[7] Google Cloud, "GCP *Pricing* | *Google Cloud*"; Available: https://cloud.google.com/pricing/.

[8] B. Yadranjiaghdam, K. Hotwani, and N. Tabrizi, "*A Risk Evaluation Framework for Service Level Agreements*," in Proc. of IEEE Int. Conf. on Computer and Information Technology, Dec. 2016, pp. 681-685.

[9] I. Kounelis, G. Baldini, R. Neisse, et al., "*Building Trust in the Human-Internet of Things Relationship*," IEEE Technology and Society Magazine, 2014, pp. 73-80; DOI: 10.1109/MTS.2014.2364020.

[10] F. Berman and V. G. Cerf, "*Social and Ethical Behaviour in the Internet of Things*," Communications of the ACM, Vol. 60, No. 2, Feb. 2017, pp. 6-7; DOI: 10.1145/3036698.

[11] T. Kirkham, K. Djemame, M. Kiran, et al., "*Risk Based SLA Management in Clouds: A Legal Perspective*," in Proc. of 7th Int. Conf. for Internet Technology and Secured Transaction, Dec. 2012; INSPEC Accession Number: 13370312.

[12] IBM Knowledge Center, "*Service Level Agreement Lifecycle*"; Available: https://www.ibm.com/.

[13] I. Kounelis, G. Baldini, R. Neisse, et al., "*Building Trust in the Human-Internet of Things Relationship*," IEEE Technology and Society Magazine, 2014, pp. 73-80.

[14] F. Berman and V. G. Cerf, "*Social and Ethical Behaviour in the Internet of Things*," Comm. of the ACM, Vol. 60, No. 2, Feb. 2017, pp. 6-7.

[15] EU GDPR Portal, "*GDPR Glossary of Terms*"; Available: https://www.eugdpr.org/.

978-1-5386-4980-0/19/$31.00 ©2019 IEEE