Security and Privacy in Smart Grid Demand Response Systems

Andrew Paverd[1], Andrew Martin[1], and Ian Brown[2]

[1] Department of Computer Science, University of Oxford
andrew.paverd@cs.ox.ac.uk andrew.martin@cs.ox.ac.uk
[2] Oxford Internet Institute, University of Oxford
ian.brown@oii.ox.ac.uk

Abstract. Various research efforts have focussed on the security and privacy concerns arising from the introduction of smart energy meters. However, in addition to smart metering, the ultimate vision of the smart grid includes bi-directional communication between consumers and suppliers to facilitate certain types of Demand Response (DR) strategies such as demand bidding (DR-DB). In this work we explore the the security and privacy implications arising from this bi-directional communication. This paper builds on the preliminary work in this field to define a set of security and privacy goals for DR systems and to identify appropriate and realistic adversary models. We use these adversary models to analyse a DR-DB system, based on the Open Automated Demand Response (OpenADR) specifications, in terms of the security and privacy goals. Our analysis shows that whilst the system can achieve the defined security goals, the current system architecture cannot achieve the privacy goals in the presence of honest-but-curious adversaries. To address this issue, we present a preliminary proposal for an enhanced architecture which includes a trusted third party based on approaches and technologies from the field of Trusted Computing.

1 Introduction

It is widely acknowledged that the upgrade to a smart energy grid presents multiple new challenges in terms of security and privacy. There has been extensive research on the security and privacy issues that arise from the Advanced Metering Infrastructure (AMI) in which smart meters record fine-grained energy consumption measurements and send these to the energy supplier or other external entities. In particular, privacy-preserving smart metering has been the subject of numerous research efforts and various privacy-preserving protocols have been proposed [1][2][3][4].

However, whilst the AMI is a critical part of the smart grid infrastructure, it is not the only aspect from which security and privacy concerns arise. In addition to the AMI protocols for measuring energy consumption, the future architecture of the smart grid includes Demand Response (DR) protocols for managing energy consumption. Specific types of DR, such as demand bidding (DR-DB) protocols,

involve bi-directional communication between the consumers and entities such as the energy supplier in order to co-ordinate the consumers' actions towards reducing their consumption during periods of high demand. As a result of this bi-directional communication, these protocols also present various security and privacy challenges that must be addressed.

The Open Automated Demand Response (OpenADR) specification is an example of a data model that can be used in DR communication. OpenADR 1.0 [5] was developed by the Demand Response Research Center operated by Lawrence Berkeley National Laboratory as a means for communicating DR information between energy suppliers, network operators and consumers. This formed the foundation of OpenADR 2.0 [6] which has now been developed by the OpenADR Alliance, an industry coalition that promotes the development and adoption of OpenADR-compliant systems. The alliance claims that over 60 vendors are currently producing OpenADR-compliant systems [7]. Section 2 of this paper provides background information about DR and presents an overview of OpenADR.

Building on preliminary research in the area of security and privacy for DR systems, this paper describes the threats to security and privacy that arise from bi-directional DR communication. Although we use the OpenADR specification as a case study, our analyses can be applied to similar DR systems. Section 3 defines the security and privacy goals that we have identified for a generic DR system. Section 4 presents the possible adversary models and describes their capabilities. Using these adversary models, Section 5 presents an analysis of a DR-DB system, based on the OpenADR specification, in terms of the security and privacy goals. In order to address the identified privacy challenges, this paper presents a preliminary proposal for a technical architecture that enhances consumers' privacy in DR protocols. Using approaches and technologies from the field of Trusted Computing (TC), this architecture is designed to mitigate against the major security and privacy threats that have been identified. An overview of this proposed architecture is presented in Section 6 but the full design and analysis will be carried out as future work. The three main contributions of this paper are therefore: the development of a threat model for DR protocols through the combination of appropriate adversary models and security and privacy goals; the application of this model to a concrete protocol based on the OpenADR specification; and the proposed architecture for mitigating against these threats.

2 Background

This section contextualizes the work by providing background information about demand response (DR) systems in general as well as an overview of OpenADR.

2.1 Demand Response Systems

In the absence of grid-scale storage capacity, electrical energy must be used as it is generated. Electrical energy consumption can be divided into a *base-load* that

remains relatively constant and a *peak-load* that varies with time. As demand for energy increases relative to supply, it is necessary to either increase generation or reduce demand. Although additional peak-load generation capacity might be available, it is often expensive and might not be sufficient to satisfy the full demand. The same objective can be achieved by reducing peak demand through the use of demand response (DR) techniques.

The United States Department of Energy (DoE) defines DR as:

"Changes in electric usage by end-use customers from their normal consumption patterns in response to changes in the price of electricity over time, or to incentive payments designed to induce lower electricity use at times of high wholesale market prices or when system reliability is jeopardized." [8]

DR refers to a set of actions with the aim of dynamically reducing energy demand at specific times and in specific locations in response to a relative shortage in supply. These so called *DR events* could be caused either by an increase in demand or a decrease in generation capacity at a particular time. It is well known that peak-load demand increases at specific times during the day (known as *peak times*) due to human behaviour and energy generation capacity is dimensioned to accommodate these known variations. However, in addition to this variability in demand, we are also faced with increasing variability in supply as we integrate renewable energy sources such as wind and solar power into the grid.

The simplest and most extreme form of DR is forced curtailment or *load-shedding* in which specific consumers are forcibly disconnected in order to reduce demand. However, load-shedding can result in significant productivity and economic losses for the affected areas. There has been significant interest in improving this situation through more participative forms of DR that involve the consumers in the DR activities. As indicated in the US DoE definition [8] and explained in the categorization by Albadi and El-Saadany [9], there are two major categories of participative DR, namely *price-based* and *incentive-based* DR:

Price-Based Demand Response In a price-based system, the DR manager uses time-based pricing in an effort to reduce demand at certain times. If possible, consumers will reduce demand when the price is high in order to minimize their energy bills. This could be implemented in various ways:

- **Time-of-use (ToU) pricing:** The energy price varies predictably according to the time at which it is used.
- **Critical peak pricing:** The energy price is specifically increased for periods of peak demand.
- **Dynamic pricing:** The energy price varies dynamically in time or geographic location depending on the ratio between supply and demand.

All of these approaches require a reliable mechanism for communicating the current price information to the consumers (e.g. in-home displays) as well as the implementation of appropriate billing (e.g. ToU billing using smart meters).

Incentive-Based Demand Response As an alternative to price-based DR, incentive-based schemes provide certain incentives (usually financial in nature) to consumers who participate in DR events.

One type of incentive-based DR is a *demand bidding* (DR-DB) system. A DR-DB system requires a bidding protocol in which the DR manager (e.g. the energy supplier) initiates a DR event and consumers send *bids* indicating the amount by which they are willing to reduce demand at the specified time. These bids might include each consumer's desired incentive price if this is not specified by the DR manager. The DR manager accepts these bids until the DR objective has been met. Although it is not required, it may be desirable to check that consumers with accepted bids actually do reduce or shift their consumption accordingly. An overview of the communication in this type of protocol is shown in Figure 1.

The bi-directional communication in the bidding protocol provides a feedback loop for the DR manager. Without this feedback, the DR manager would be required to predict the effects of specific DR actions on consumers' behaviour. Depending on the dynamics of the system, incorrect predictions could lead to instability in the system characterized by large swings in demand. Instead, the inclusion of the bidding protocol makes this a closed-loop feedback system which can be controlled effectively.

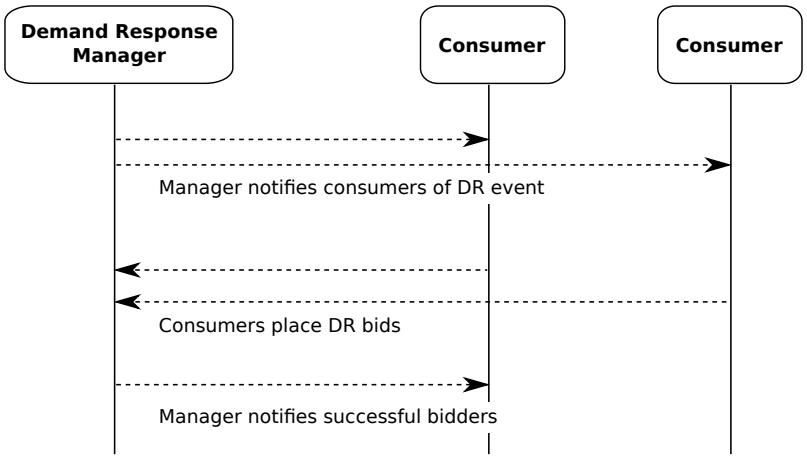

Fig. 1. Bidding process in a generic demand bidding (DR-DB) system.

2.2 OpenADR

The OpenADR specifications describe an open standards-based communications data model to facilitate DR communication between service providers and con-

sumers [5][6]. The specification defines various XML-based messages that can be exchanged over any IP-based network using protocols such as Hypertext Transfer Protocol (HTTP), Simple Object Access Protocol (SOAP) or XML Messaging and Presence Protocol (XMPP) [5][6]. The OpenADR 1.0 specification [5] introduces the concept of the Demand Response Automation Server (DRAS). The role of this component is to automate the communication between various entities in the system. The DRAS augments the generic bidding procedure by serving as an intermediary between the DR manager (usually the energy supplier) and the consumers. The DR manager informs the DRAS of a DR event and the DRAS in turn publishes this information to the consumers. Consumers have the option to set up standing bids with the DRAS so that when a new DR event is announced, they can either place new bids, maintain their standing bids or cancel their standing bids by opting-out of the event. The DRAS forwards the new bids or standing bids to the DR manager who accepts bids until the DR objective is met. These interactions are shown in Figure 2 [5].

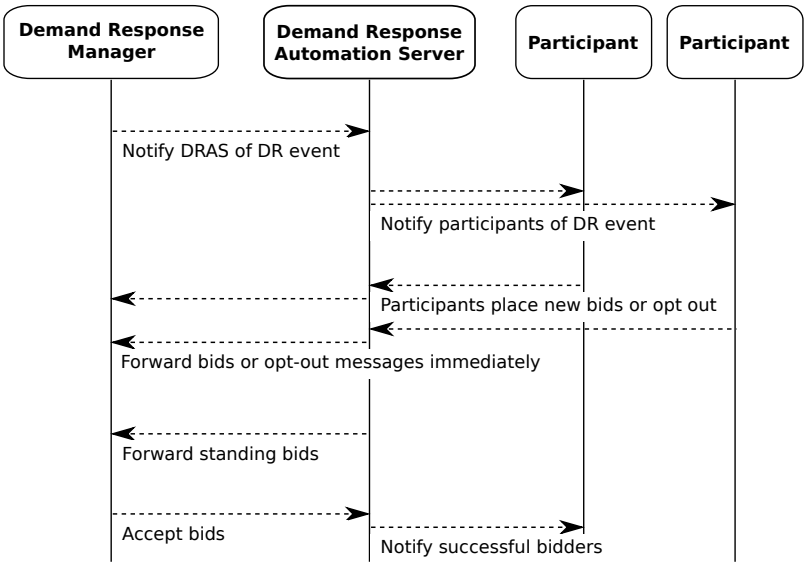

Fig. 2. Bidding process in an OpenADR system [5][10].

The OpenADR 2.0 specification [6], developed by the OpenADR alliance, differs significantly from the OpenADR 1.0 specification. In OpenADR 2.0, there are two types of nodes in the system: nodes that publish or transmit information about events are called *Virtual Top Nodes* (VTNs) and nodes that receive and respond to information are called *Virtual End Nodes* (VENs). The specification intentionally does not define the behaviour on the nodes once a message has been received. Although it is not defined in the specification, the functional role of

the DRAS can therefore be recreated using a specifically-designed VTN. There is no peer-to-peer communication between VENs but a hierarchical structure can be used in which a node receives information as a VEN and retransmits it downwards to subordinate nodes as a VTN. Since the behaviour of the nodes is not specified, it is possible that this hierarchical structure can also be used in the reverse (upwards) direction since a VTN can retransmit information from its subordinate VENs to nodes further up the hierarchy. This also makes it possible for a node to aggregate information from multiple subordinate nodes. In this paper, all references to the DRAS are therefore based on the OpenADR 1.0 specification since this provides a concrete definition of this component's behaviour. In general, our analysis is also applicable to OpenADR 2.0 systems which may or may not include intermediary nodes.

3 Security and Privacy Goals

This section defines a set of security and privacy goals for a DR system. These represent the overall goals for the system rather than the individual security and privacy requirements for specific scenarios or the mechanisms through which these are achieved. These goals are used as a frame of reference for the analysis in Section 5. Although the security and privacy goals are presented separately, it will be shown that there is a strong relationship between them.

3.1 Security Goals

Given the critical nature of the electricity supply infrastructure, the primary security objective is to ensure that only legitimate entities can participate in the DR protocol. This can be defined through the following two goals:

S-1: Consumers must be able to verify the authenticity and integrity of all DR events and bid notifications.

S-2: The DR manager must be able to verify the authenticity and integrity of all DR bids.

Goal **S-1** refers to any DR event in either a price-based or incentive-based approach and also includes the acceptances of bids in an incentive-based approach. This goal means that actions such as setting a higher ToU energy price or requesting bids for energy reduction can only be performed by a legitimate entity since the authenticity of the message must be verifiable. It also means that these messages cannot be modified by an adversary since the integrity must be verifiable. Goal **S-2** is only applicable in the incentive-based approach and enforces the same restrictions as **S-1** on messages sent by the consumers containing bids for energy reduction.

Similar versions of these security goals are present in the OpenADR 2.0 specification which also describes mechanisms for achieving these goals. The specification defines two security levels: standard and high security [6]. All OpenADR-compliant systems must implement at least standard level security in which

Transport Layer Security (TLS) with mutual authentication is used to protect the confidentiality and integrity of the communication and authenticate the communicating entities [6]. Some OpenADR systems implement the high security level in which XML signatures are used in addition to TLS to ensure the integrity and authenticity of the messages [6]. In the OpenADR specification, the confidentiality of the messages is an important concern but in this paper we classify this as a privacy goal as described in the next subsection.

In addition to the above goals, the specific hardware elements used in the smart grid might introduce requirements on the security mechanisms, for example, that any cryptographic operations used in the protocol must be achievable on a smart meter with limited computational capabilities. However, the requirements of the security mechanisms as well as the mechanisms themselves are beyond the scope of this section. It should be noted that the security goals presented in this section are broadly similar to those used in most other protocols for secure communication. From a communication perspective, DR systems do not introduce any new security goals beyond those already in place elsewhere. However, it is precisely because of these security goals that certain conflicts with the privacy goals arise as explained in our analysis. Therefore, it is critical to recognize the existence and impact of these seemingly general security goals.

3.2 Privacy Goals

The privacy goals for the system aim to protect the privacy of individual consumers. Initially, the participants in DR programmes have been large consumers such as industrial sites or building complexes. However, it is anticipated that DR programmes will be extended to all consumers including residential homes. For residential consumers, the protection of personal or private information is an important requirement in the smart grid. This is illustrated by the significant privacy concerns raised in response to the introduction of smart meters [11][12][13] as well as the various research efforts to develop privacy-preserving smart metering protocols [1][2][3][4]. However, despite their importance, these privacy goals are not addressed in the OpenADR specifications. The specifications only call for confidentiality of the communicated messages with respect to an external adversary. Building on the research about privacy-preserving smart metering as well as the preliminary research on privacy in DR systems by Karwe and Strüker [10], we define the following privacy goals:

P-1: Untrusted entities must not be able to link DR bids to individual consumers.

P-2: Untrusted entities must not be able to infer private information about individual consumers from the DR system.

These goals should be interpreted from the perspective of the individual consumer as he or she is the owner of the private information. Goal **P-1** requires that entities that are not trusted by the consumer must not be able to link DR bids to specific consumers since this could reveal private information about the

consumer. If bids were visible to an untrusted entity and could be linked to individual consumers, the untrusted entity would learn information such as the consumers chosen energy supplier and tariff plan. Furthermore, the energy reduction specified in the bid reveals some information about the consumer's total energy consumption. In the same way that frequent energy measurements from smart meters can be used to make inferences about the occupants of resident premises, DR bids could also be used to infer private information. For example, a bid to decrease a large load, equal to that of a plug-in electric vehicle, indicates that the consumer probably owns such a vehicle and would otherwise be recharging it. The ability to link the bids to individual consumers also allows the untrusted entity to build up a profile of the consumer's behaviour. Any deviations from this profile could lead to further inferences about the user's behaviour. Continuing the previous example: if a particular consumer regularly bids to stop recharging an electric vehicle at peak times, any deviation from this pattern could indicate that the electric vehicle and its owner are away from home at that time. Even if an untrusted entity cannot view the individual bids or link them to specific consumers, goal **P-2** aims to ensure that untrusted entities either outside or within the system cannot make inferences such as those described above from the DR system.

4 Adversary Models

This section defines the adversary models used in our analyses in terms of the adversary's capabilities. The main adversary models used in this work are the Dolev-Yao (D-Y) model and the Honest-But-Curious (HBC) model.

4.1 Dolev-Yao Adversary

In the model proposed by Dolev and Yao [14], the adversary has full control of the communication network. The adversary can eavesdrop, intercept, block or modify messages as well as replay old messages or synthesize falsified messages. The adversary is only limited by the constraints of the cryptographic systems. It is assumed that the adversary cannot break cryptographic primitives. This means the adversary can neither read encrypted messages without the correct decryption key, nor forge cryptographic signatures, nor reverse cryptographic hash functions. Although the D-Y model is already considered to be the strongest type of adversary, it is sometimes also assumed that the DY adversary might be able to guess passwords with some defined success probability or recover encryption keys after a defined period of time.

4.2 Honest-But-Curious Adversary

In contrast to the D-Y model, the HBC adversary is more limited in terms of its capabilities. The HBC adversary does not deviate from the defined protocol in terms of sending and receiving messages. This adversary is also limited by the

constraints of cryptographic systems and cannot break cryptographic primitives. However, this adversary aims to learn as much as possible from any messages it can receive. This usually also involves linking messages together or making inferences based on message contents. This model is sometimes referred to as the semi-honest model [15]. The HBC model differs from a passive D-Y adversary. The passive D-Y adversary attempts to avoid detection by not performing any active actions (i.e. by neither modifying messages nor sending falsified messages) but will still attempt to eavesdrop on all messages in the system. In contrast, the HBC adversary does not attempt to eavesdrop on messages for which it is not the intended recipient. Therefore, the HBC model is deliberately more limited than even a passive D-Y adversary so that it can be used to accurately model the behaviour of real entities in our system.

5 Analysis of a Demand Bidding System

This section presents an analysis of a DR-DB system, based on the OpenADR specification, in terms of the security and privacy goals defined in Section 3 and the adversary models described in Section 4. The aim is to provide a realistic representation of the potential adversaries within the system using an appropriate model for each adversary. This representation can then be analysed with respect to the defined security and privacy goals. Figure 3 shows the communication architecture of the system and indicates the potential adversaries we consider in this analysis.

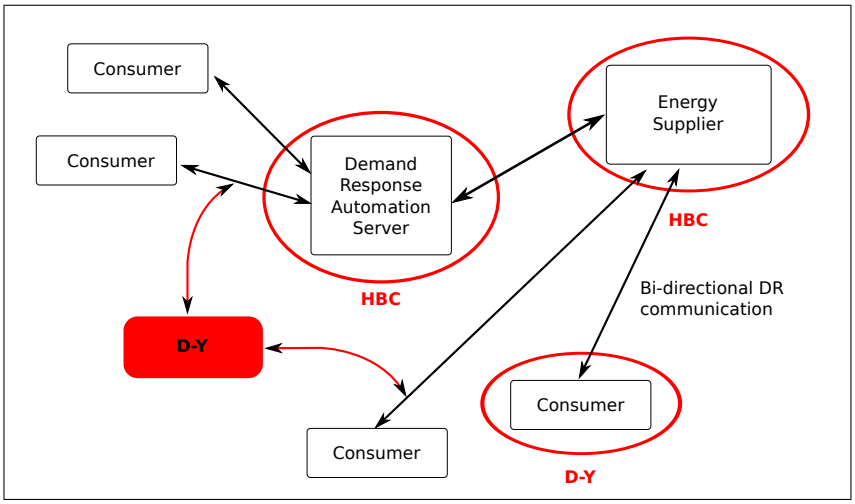

Fig. 3. Adversary model for a demand response system.

5.1 External D-Y Adversary

The most widely used adversary representation is that of an external D-Y adversary who controls the communication network. This adversary is neither authorized to initiate events nor respond to events and so must be prevented from doing so in order to satisfy **S-1** and **S-2**. In the OpenADR specifications, this is achieved through the use of mutually authenticated TLS connections between all nodes and optional XML signatures on messages [6]. The privacy goals **P-1** and **P-2** are also achieved with respect to this adversary because of the confidentiality provided by TLS. The adversary could perform traffic analysis on the encrypted messages but could be prevented from learning any private information by introducing dummy traffic from the consumer at regular intervals. Although it is assumed that the adversary cannot break the underlying cryptographic primitives, the security and privacy of the system are still fully dependent on all secret keys being protected from the adversary. If any of the nodes exhibits endpoint vulnerabilities, it might be possible for the adversary to obtain these keys. Therefore, the protection of these keys is of critical importance. Techniques such as that described in [16] aim to address this challenge taking into account the unique constraints of the smart grid.

5.2 Consumer as a D-Y Adversary

The second possible type of adversary in the system is a dishonest or malicious consumer. This adversary is modelled as a D-Y adversary because he or she might deviate from the defined protocol. In the worst case it can be assumed that this adversary exhibits the same level of control over the network as the external D-Y adversary. This is a realistic assumption because the dishonest consumer might collaborate with the external adversary or the external adversary might also be a consumer in the system. This adversary is stronger than the external D-Y adversary because he or she is also a legitimate agent in the communication protocol and thus has access to a set of cryptographic keys required to respond to DR events. For example, this adversary could represent a dishonest consumer who attempts to claim larger incentives by submitting high bids but does not reduce demand by the bid amount. Assuming that the bids are attributable to the dishonest consumer because of TLS mutual authentication (**S-2**), it should be possible for the supplier to identify and take action against this adversary. A more malicious consumer might try to masquerade as multiple different consumers in order to evade detection. Unless the system has a robust mechanism for distinguishing between different consumers, this attack will succeed. If the false bids are not detected, this type of attack could be used to destabilize the electricity grid through the submission of multiple false bids from a large number of consumers. The privacy goals would still be maintained under the same conditions as for the external D-Y adversary. Since OpenADR does not permit peer-to-peer communication between VENs, the adversary gains no personal information about other consumers by becoming a consumer in the protocol.

5.3 DRAS as an HBC Adversary

The third type of possible adversary is the DRAS as an HBC adversary as described by Karwe and Strüker [10]. In this section we use the term *DRAS* to refer to the functionality of the DRAS node as defined in the OpenADR 1.0 specification [5] or to the equivalent functionality provided by an OpenADR 2.0 VTN. Since this entity is an important part of the infrastructure, it must be assumed to be weaker than a D-Y adversary due to external forces such as regulation, auditing and legal intervention. If this entity had the capabilities of a D-Y adversary, it would have the capability to cause a catastrophic system failure by sending falsified data to the energy supplier. Real-world implementations are therefore designed to minimize the probability of this occurrence and so the most realistic way to model these implementations is to use an HBC rather than a D-Y adversary model. It is therefore assumed that the DRAS will follow the defined protocol and will not violate the security goals (**S-1** and **S-2**). However, since the DRAS acts as an intermediary node in the communication architecture, it already has legitimate access to all the messages passing between consumers and the supplier. Even if it executes the protocol correctly, it could still violate the privacy goals (**P-1** and **P-2**) if it is not trusted by consumers. Using only information it has legitimately obtained, the DRAS could link bids to individual consumers and therefore make inferences about these consumers and their behaviour. Karwe and Strüker [10] propose a solution to this problem by introducing end-to-end encryption between the consumers and the DR manager so that messages cannot be read by the DRAS.

5.4 Supplier as an HBC Adversary

The final type of adversary is the energy supplier as an HBC adversary. As in the previous section, it is assumed that external forces such as regulation limit the capabilities of the supplier. Giving this entity the capabilities of a full D-Y adversary would again result in catastrophic system failure since this entity is the only legitimate initiator of DR events. Therefore, an HBC adversary model must be used to achieve a realistic representation of the system. As above, the security goals are satisfied because the supplier is always assumed to follow the protocol correctly. However, the supplier can violate the privacy goals by linking bids to individual consumers and making inferences based on these bids. This challenge cannot be overcome by anonymizing bids as this would allow the external D-Y adversary or the consumer D-Y adversary to violate the security goals by submitting multiple falsified bids which could not be attributed to specific consumers. Furthermore, one of the functional requirements is that the supplier must be able to link bids to individual consumers in order to allocate the relevant incentives. This means that with the current architecture, neither of the privacy goals can be achieved unless the supplier is trusted. However, in reality energy suppliers are not always trusted by consumers as illustrated by the Dutch case in 2008 [17][13]. This challenge could be addressed through regulation of the energy supplier or through modification of the system architecture as we propose in the next section.

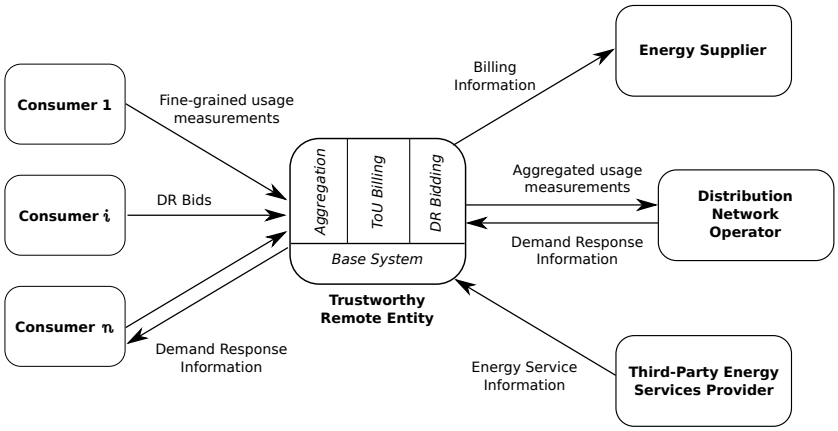

Fig. 4. Enhanced system architecture using a Trustworthy Remote Entity [18].

6 Proposed Architecture

In order to address the privacy challenges identified above, we propose an enhanced system architecture to facilitate communication between consumers and other entities such as energy suppliers or distribution network operators (DNOs). This architecture has been developed as part of our ongoing research into privacy-enhancing technologies for the smart grid [18]. Figure 4 shows our proposed system architecture. The significant innovation of this architecture is the inclusion of a trusted third party called a Trusted Remote Entity (TRE). The TRE is an information processing node situated as an intermediary between the consumers and all external entities. The TRE makes extensive use of Trusted Computing (TC) technologies and approaches. As specified by the Trusted Computing Group (TCG), TC can be used to obtain various security guarantees about computational systems. The TCG-specified Trusted Platform Module (TPM) is a secure cryptographic co-processor that can be used as a root of trust in the system [19]. The *secure boot* procedure ensures that the system will boot into a secure state and the process of *remote attestation* uses the TPM to generate unforgeable proofs of this state which are sent to the relying parties in order to establish trust. The TRE uses these approaches to prove its secure state to all the relying parties. Unlike TC in the PC domain, the TRE avoids the problem of scalability in attestation by running a very small Trusted Computing Base (TCB). A full description of the TRE will be presented in future work.

The fundamental aspect of the TRE is that it is mutually trusted by parties that do not necessarily trust each other. As explained in Section 5, there is evidence that consumers do not necessarily trust the energy supplier to store and perform computations on their private information. Similarly, the energy supplier does not necessarily trust the consumers to calculate their own energy bills honestly. However, in our architecture, both consumers and the energy supplier trust the TRE to perform these operations on their behalf. Critically,

these parties have good grounds for trusting the TRE because of its use of TC approaches and technologies. Whilst the use of TC does not remove all risks from this architecture, it significantly reduces the likelihood of a large class of software-based threats. TC secure boot and remote attestation virtually eliminate the possibility of a remote adversary compromising the software of the TRE without being detected immediately. TC does not mitigate against all hardware-based threats such as eavesdropping on the physical memory bus within the system, however, attacks of this type are significantly more complicated and expensive than software-based attacks and so present a significantly lower risk. In practice, these risks would be mitigated through certification or auditing processes. Just as the hardware of a TPM is trusted based on a certificate from its manufacturer, a similar certificate from the TRE manufacturer could be used to establish trust in the TRE hardware which would in turn support trust in the TRE software.

6.1 Distributed TREs

There will be multiple TREs throughout the network, each serving a group of consumers. The maximum number of consumers per TRE will depend on the computational and network capacities of the TRE but it is expected to be in the order of thousands of consumers per TRE. One of the primary weaknesses of any architecture that includes a trusted third party is that this node could become both a single point of failure as well as the most attractive target for attack. This would also be true of the TRE if it were a single node in the architecture. However, the use of multiple distinct TREs (i.e. with differing cryptographic keys etc.) throughout the network mitigates against this risk. There is still a non-zero probability that a single attack could affect all TREs in the network but this is very similar to an attack affecting all smart meters in the grid. The smart meter attack is arguably more likely since the meters generally do not include the hardware-based security capabilities used in the TRE.

6.2 Smart Grid TRE Functionality

In the smart grid, the TRE provides three main types of functionality: Firstly it aggregates the the high-frequency measurements from smart meters before sending them to the DNO for use in network optimization. Secondly, it performs ToU billing calculations on behalf of the energy supplier. Thirdly, it provides the functionality of a DRAS in an OpenADR 1.0 system or an aggregator in an OpenADR 2.0 architecture. In this role, the TRE does not forward the bids to the energy supplier but instead aggregates the bids so that they cannot be linked to individual consumers. Since the TRE also performs the billing calculations, it can apply the respective incentives to successful bidders without revealing their identities to the supplier. This architecture therefore mitigates against both types of HBC adversaries identified in the previous section. Consumers can use TC remote attestation to verify that a particular TRE is trustworthy. Even if the energy supplier or DNO are untrusted HBC adversaries, the aggregation

of energy measurements and DR bids performed by the TRE prevents these adversaries from learning any private information about consumers.

In OpenADR 2.0, the TRE appears as a VTN for the consumers and as a VEN for the energy supplier. This means that the TRE can be deployed as a plug-in enhancement to the smart grid without requiring any modification of the primary information flows. The only additional communication that would be required are the remote attestation protocols for establishing trust in the TRE. Furthermore, a heterogeneous smart grid architecture could be used in which some users communicate directly with the supplier whilst others communicate via a TRE. In a real-world deployment scenario, the TREs could therefore be deployed gradually without causing major disruptions to the smart grid. The specific TRE deployment scenarios are the subject of future research.

7 Conclusion

Security and privacy concerns arising from the introduction of smart meters have been the subject of various research efforts. However, less attention has been given to the security and privacy of demand response protocols, such as demand bidding (DR-DB), which will be an important part of the future smart grid. This paper builds on preliminary work in this area to define a set of high level security and privacy goals for demand bidding systems highlighting the fact that the bi-directional communication used in these systems poses a risk to consumers' privacy. We identify the appropriate types of adversary models and use these to present an analysis of a DR-DB system based on the OpenADR specifications, in terms of the security and privacy goals. Although this system achieves the security goals, it cannot achieve the defined privacy goals if external entities such as the energy supplier are not trusted by the consumers. In order to address this issue, we have proposed a system architecture to enhance consumers' privacy in the smart grid. The key innovation of this architecture is the inclusion of a Trustworthy Remote Entity (TRE) which uses Trusted Computing (TC) approaches and technologies to establish trust relationships with both the consumers and the external entities. The TRE is mutually trusted by parties that do not necessarily trust each other and the use of TC provides good grounds for this trust. Through the functionality provided by the TRE, the defined security and privacy goals can be achieved whilst maintaining the overall functionality of the demand response system.

8 Acknowledgements

The research described in this paper was conducted as part of the *Future Home Networks and Services* project at the University of Oxford, funded by BT.

References

1. Efthymiou, C., Kalogridis, G.: Smart Grid Privacy via Anonymization of Smart Metering Data. In: Proc. IEEE International Conference on Smart Grid Communications (SmartGridComm). (2010) 238–243
2. Danezis, G., Kohlweiss, M., Rial, A.: Differentially private billing with rebates. In: Proceedings of the 13th international conference on Information hiding. IH'11, Berlin, Heidelberg, Springer-Verlag (2011) 148–162
3. Ács, G., Castelluccia, C.: I have a DREAM!: differentially private smart metering. In: Proceedings of the 13th international conference on Information hiding. IH'11, Berlin, Heidelberg, Springer-Verlag (2011) 118–132
4. Borges, F., Martucci, L.A., Muhlhauser, M.: Analysis of privacy-enhancing protocols based on anonymity networks. In: Proc. IEEE Third International Conference on Smart Grid Communications (SmartGridComm). (November 2012) 378–383
5. Piette, M.A., Ghatikar, G., Kiliccote, S., Koch, E., Hennage, D., Palensky, P., McParland, C.: Open Automated Demand Response Communications Specification (Version 1.0). Technical Report April, California Energy Commission, PIER Program (2009)
6. OpenADR Alliance: OpenADR 2.0b Profile Specification. Technical report (2013)
7. OpenADR Alliance: The OpenADR Primer. Technical report (2012)
8. United States Department of Energy: Benefits of Demand Reponse in Electricity Markets and Recommendations for Achieving Them. Technical Report February (2006)
9. Albadi, M., El-Saadany, E.: A summary of demand response in electricity markets. Electric Power Systems Research **78**(11) (November 2008) 1989–1996
10. Karwe, M., Strüker, J.: Maintaining Privacy in Data Rich Demand Response Applications. In Cuellar, J., ed.: First Open EIT ICT Labs Workshop on Smart Grid Security - SmartGridSec12. Volume 7823 of Lecture Notes in Computer Science. (2013) 85–95
11. Quinn, E.L.: Privacy and the New Energy Infrastructure. SSRN eLibrary (2009)
12. Brown, I.: Britain's Smart Meter Programme: A Case Study in Privacy by Design. International Review of Law, Computers & Technology (February 2013)
13. Cuijpers, C., Koops, B.J.: Smart Metering and Privacy in Europe: Lessons from the Dutch Case. (February 2013)
14. Dolev, D., Yao, A.: On the security of public key protocols. IEEE Transactions on Information Theory **29**(2) (March 1983) 198–208
15. Goldreich, O.: Foundations of Cryptography: Volume 2, Basic Applications. Volume 2. Cambridge University Press (2009)
16. Paverd, A.J., Martin, A.P.: Hardware Security for Device Authentication in the Smart Grid. In Cuellar, J., ed.: First Open EIT ICT Labs Workshop on Smart Grid Security - SmartGridSec12. Volume 7823 of Lecture Notes in Computer Science. (2012) 72–84
17. Cuijpers, C., Koops, P.B.j.: Het wetsvoorstel slimme meters: een privacytoets op basis van art . 8 EVRM Onderzoek in opdracht van de Consumentenbond. Technical report, Universiteit van Tilburg (2008)
18. Paverd, A.J.: Student Research Abstract: Trustworthy Remote Entities in the Smart Grid. In: ACM Symposium on Applied Computing (SAC). (2013)
19. Trusted Computing Group: TPM Main Specifications, Version 1.2, Revision 116, Part 1: Design Principles. Technical report (2011)