

Smart Metering Implementation Programme: A call for evidence on data access and privacy (August 2011).

Local Processing – a Solution to the Smart Meter Data Privacy Problem

A consultation response from Acute Technology Limited, October 2011

[REDACTED]

[REDACTED]

For more details of the technology described here, please visit:

<http://projecthydra.info/local-processing/>

Introduction

Acute Technology Limited is pleased to submit this response to DECC's Data Access and Privacy consultation.

This submission explains role that "local processing" can perform in solving the data privacy problems that have been recognised by the smart meter programme.

If not properly addressed the privacy issues will remain as legal and political threats to the whole smart meter programme:

- ◆ We note that the Dutch courts have found that the proposed Dutch smart meter programme violated the European Convention on Human Rights, as it mandated the export of an unnecessary level of privacy-compromising data. The same EU laws pertain in the UK, and so the UK smart meter programme is vulnerable to the same legal challenge.
- ◆ Furthermore, the UK smart meter programme remains critically vulnerable to an orchestrated campaign targeting privacy deficiencies. You do not have to have a long memory to recall the damage done by irrational press campaigns targeting pay-as-you-drive road pricing ("spy-in-the-sky") and pay-as-you-throw rubbish collection ("chip and bin"). In Holland smart meters are referred to as "espionage meters", and in a sense they are.

Fortunately we have developed a "local processing" technology as an alternative privacy model for the smart meter programme, and have working prototypes. This approach:

- ◆ does not require the export of privacy-compromising raw consumption data from every home to a huge central database;
- ◆ instead, it processes the raw consumption data within the home and exports only the processed data, dramatically enhancing privacy;
- ◆ intrinsically provides firewalls between different "authorised parties", so that each authorised party receives only the information that it is entitled to receive, so relieving the DCC from the onerous task of controlling access to the huge central database;
- ◆ does not compromise functionality or increase the risk of fraud;
- ◆ still allows consumers to share consumption data with third parties, but only involving a process "informed consent";
- ◆ dramatically reduces the volume of data that needs to be transferred through the WAN communications channels and stored centrally;
- ◆ makes the system more distributed and thus reduces the risk that a security compromise at a single point could give an attacker control of the entire smart meter system and all customers' data.

How have we cut the Gordian knot of smart meter privacy? By recognising that smart meters which are truly smart can do all the necessary data processing within the home. And by identifying an existing set of technologies and standards that provide an elegant, secure and extensible platform to perform this processing.

We believe our proposal represents a true example of "privacy by design" (as opposed to "privacy by regulation"), and commend it to the smart meter programme and to stakeholders.


How it Works

The whole smart meter data model can be turned on its head when you realise that the consumption data can be processed locally on in the home just as easily as on remote servers.

If an algorithm is applied to data then the location of the calculation does not affect the result: you get the same answer whether the processing is performed in the home or on a remote server.

However the privacy characteristics of local processing are completely different: if personal data stays in the home then the privacy problems vanish.


The inspiration comes from the smart card industry, which has developed a technology that allow multiple computer programs, called applets, to run on smart cards. This is shown in the figure below.



The characteristics of this technology include:

- ◆ Complete separation of the code and data of each applet, behind internal firewalls.
- ◆ Independent secure communications channels between each applet and its off-card server, with each channel protected by its own cryptographic key.
- ◆ Applets are written in the Java programming language, which allows an applet to be written once then run on any compliant smart card.
- ◆ Applets can be securely, dynamically and remotely managed by “over-the-air” provisioning. New applets can be installed, and old applets deleted up upgraded.
- ◆ Well tested and certified cryptography is used, suitable for protecting financial transactions. In particular, “digital signatures” are used that prove that messages received from an applet must have really originated from within that applet, and could not have been fabricated elsewhere.
- ◆ This technology is used on smart cards, mobile phone SIM cards, and increasingly on “secure elements” in smart phones that protect data associated with NFC applications.
- ◆ Covered by international standards published by ETSI.

This smart card technology can be reused in a smart meter or communications hub, with applets processing the raw energy consumption data:



The architecture is the same. This time each applet processes the raw energy consumption data to derive some “digested” result, from which the privacy-damaging information has been removed.

In the figure above:

- ◆ A monthly bill is calculated and the result (say £124.53) is sent to the energy retailer. The retailer gets all that he needs to bill the customer, and no more.
- ◆ Information about the distribution network, such as voltage levels, is sent to the distribution network operator, who receives the information he needs for the smooth running of the network, but no more.
- ◆ Time of use profile information is generated by the settlement applet, and sent to the settlement authority, who can use this to improve the settlement process. The raw data is not required, and is not transmitted.
- ◆ An applet running on behalf of the customer can provide valuable feedback to the customer, but only if the customer wants to share her data this way.

Any algorithm that could be executed on a server with the raw consumption data can equally well be executed by an applet running locally on the smart meter. As well as keeping private consumption data within the meter, the distributed local processing also reduces the costs associated with exporting the raw data and processing it centrally.

Acute Technology and our Project Hydra partners have implemented the technology needed to perform local processing. The next section describes a time-of-use tariff demonstration.

Privacy by Design: Time-of-Use Tariff Example

To show the effect of local processing as a privacy-enhancing technology, we set up a web-based demonstration that applied a notional time-of-use tariff to data collected by a smart meter.

We divided the work into two parts.

In the first part data is processed remotely from the meter, at a server, in the conventional way. The screen shot below shows this.

The screenshot shows a web browser window titled "Project Hydra Billing demo - Windows Internet Explorer" with the URL "http://localhost/hydrabilling/". The page features the "PROJECT+HYDRA" logo and a header image of a smart meter. The main section is titled "Time-of-Use Tariff (Server-side calculation)".

Below the title, there is a form with the following fields:

- Select a meter:** A dropdown menu showing "ELON0212306".
- Select Dates to calculate the bill:**
 - Start Date:** A date picker showing "30/08/2011".
 - End Date:** A date picker showing "11/09/2011".
- Calculate:** A button.
- Calculated Bill:** A text box showing "£ 6.29".
- Status:** A text box showing "NOERROR".

At the bottom left, there is a small copyright notice: "(c) Brunel University all rights reserved".

On the right side of the form, there is a table titled "Daily Calculations:" with the following data:

Time interval	Whr	Cost (£)
1	124	0.0124000
2	125	0.0125000
3	124	0.0124000
4	125	0.0125000
5	124	0.0124000
6	125	0.0125000
7	124	0.0186000
8	125	0.0187500
9	125	0.0187500
10	124	0.0186000
11	125	0.0187500
12	124	0.0186000
13	125	0.0187500
14	125	0.0125000
15	124	0.0124000

At the bottom right of the table, there is a link labeled "Next >>".

It works like this:

- ◆ The meter collects 30-minute consumption data and sends this to a server.
- ◆ A user enters the start and end date of a billing period in the web application.
- ◆ The web application queries the server for the raw consumption data and applies a tariff.
- ◆ The price of the energy (£6.29) is displayed on the web page, along with the raw consumption data.

In the second part data is processed locally on the meter, using a time-of-use billing applet described above. The screen shot below shows this.

The screenshot shows a web browser window with the URL `http://localhost:54837/Default.aspx`. The page title is "Project+Hydra" and there is a "Logout" link in the top right. Below the header is a banner image of a smart meter. The main content area is titled "Time-of-Use Tariff (Local calculation at meter)". It contains a form with the following fields:

- "Select a meter:" with a dropdown menu showing "ELON0212306".
- "Select Dates to calculate the bill:" with two date pickers:
 - "Start Date:" set to "30/08/2011".
 - "End Date:" set to "11/09/2011".
- A "Calculate" button.
- "Calculated Bill:" showing "£ 6.29".
- "Status:" showing "NOERROR".

It works like this:

- ◆ The meter collects 30-minute consumption data but does not export this to a server.
- ◆ A user enters the start and end date of a billing period in the web application.
- ◆ The web application sends these dates to the billing applet running on the meter.
- ◆ The applet queries the meter for the raw consumption data and applies the tariff.
- ◆ The applet returns the result of the calculation (£6.29) to the web application, which displays it. Note that the table containing the raw consumption data is not displayed in this screen shot – as it is simply not available outside the house.
- ◆ The price of the energy is identical in both cases, since in each case the same tariff calculation is applied to the same consumption data: the only difference is where the calculation is performed.

To reiterate: any algorithm that could be executed on a server with the raw consumption data can equally well be executed by an applet running locally on the smart meter.

If DECC is serious about privacy then it is incumbent upon you to specify a true privacy-by-design solution when it is shown to be available.

Consultation Response: Answers to Questions

In the remainder of this response we provide answers to specific questions posed in the consultation.

Question 1: Please submit any further evidence, such as surveys or consumer research, regarding privacy issues and smart metering. In particular is there evidence available about the extent of any potential consumer concerns about the availability of daily versus half-hourly data?

The UK public might not be expressing widespread concern about the privacy aspects of smart metering yet, but this might yet come. You do not have to have a long memory to recall the damage done by irrational press campaigns targeting pay-as-you-drive road pricing (“spy-in-the-sky”) and pay-as-you-throw rubbish collection (“chip and bin”).

Here are two international examples.

The first image below is from the (successful) Dutch campaign against smart meters. Smart meters are described as “espionage meters”.



The second image is from a persuasive video from the USA: http://youtu.be/8JNFr_j6kdl

DECC should act to ensure such campaigns do not take off in the UK. These can be completely circumvented if you can demonstrate to the public that their data will not leave the house (without their informed consent). This can be done by specifying the privacy-by-design technology described in this response.

Question 2: To what extent would different rules for access to data between suppliers and third parties be expected to impact on the development of an energy services market (in terms of product and tariff innovation and / or entry to the energy market by third parties)? What are the particular data uses to which these concerns apply?

The consultation says (paragraph 16) about access to half-hourly data via DCC:

“Depending on the approach taken to the definition of regulated duties, and the exercise of choice by consumers, suppliers could have readier access to this granular data than other energy service companies.”

We suggest that with the local processing approach described in this document:

- ◆ The supplier can be provided with the information he needs to bill the customer without needing access to the half-hourly data.
- ◆ Consumers can make the half-hourly data available to third-parties if they so choose.

A wide range of innovative services can be developed without placing the supplier in a privileged position, by cleanly separating the data needed for the regulated duty (which can be locally-processed billing data) from the data needed to provide an innovative data analysis service. In each case the DCC need only administer the two secure communications channels to the two different recipients.

Question 5: Should theft management be considered a regulated duty for which suppliers should have access to a certain level of smart metering data? What level of data would be required and how would this be used to manage theft? Please provide practical examples.

Question 6: Does data need to be collected from all customers all of the time, for theft management, or could there be a trigger for accessing more detailed data (for example where theft is suspected)?

We do not have a view on how theft might be detected, beyond monitoring tamper alerts from within the meter itself. However, if consumption data can be analysed to detect theft then:

- ◆ It must be possible to do that by locally processing the data in the meter, rather than exporting the data to the supplier's server.
- ◆ The ability to monitor more fine-grained consumption data (i.e. data points more frequent than every 30 minutes) might be able to reveal theft more readily, and this is better done by local processing within a meter than on a server.

Question 7: What level of take-up of time-of-use tariffs could be expected under different scenarios for access to data? What information is needed to design time of use tariffs? In particular would sample or anonymised data be sufficient?

We have demonstrated in the example on pages 5 and 6 that a time of use tariff can be applied by locally processing consumption data within the smart meter, and that it is not necessary to export consumption data the supplier's server to do this.

By the same token, it must be the case that half-hourly consumption data can be analysed

locally in order to design time of use tariffs, rather than exporting the data to a server.

Usage patterns could be captured and analysed but this should be done by statistical sampling, and by an impartial third-party who would anonymise data, rather than by a supplier. If a supplier claims to need detailed half-hourly data from identified customers then the suspicion must be that they want to customise a time-of-use tariff for that customer in a manner that benefits the supplier rather than the customer.

It seems probable that energy suppliers will need proper motivation to promote time-of-use tariffs. In particular, work that they do in reducing peak energy consumption must be reflected in the settlement process, so that the hard work done by energy company A benefits company A and not all of the other free-loading energy companies. See our answer to Question 8 for how this might be achieved.

Question 8: Do you agree that individual half-hourly data is not currently required for suppliers to meet their obligations in relation to settlement? Over what timescale are any changes to settlement likely to take place and what might the implications be in terms of data requirements?

It is likely that the settlement process needs to be refined in order to provide incentives for time-of-use tariffs, and the attendant peak demand reductions. Otherwise if one supplier moved their customers to time-of-use tariffs and reduced peak demand then they would share this benefit (during the settlement process) with all the other suppliers who did not.

We believe that a “settlement applet” can be sent to the smart meter system to locally process half-hourly data by applying a settlement algorithm. Figures of merit can be sent to the settlement authority, perhaps monthly, that accurately reflects the real consumption profile at that meter. The entire raw data set is not required.

Data transfer and processing costs can be dramatically reduced and almost no personal information is exported from the home. The settlement bodies would only require the data in aggregated form, not the raw data.

As an example, the algorithm would multiply actual half-hourly consumption by a weighting curve and accumulate the result. At the end of the month the total would be a measure of how that household’s consumption varied from a national average, and each figure could be allocated to the correct energy retailer. A single figure might be sufficient for an accurate settlement calculation.

Other algorithms could of course be designed. For example, perhaps a single 30-minute reading could be selected at random once every month by the settlement applet. When processed statistically this may be sufficient to provide the necessary accuracy for the settlement process without disclosing the full privacy-sensitive data set.

With the local processing technology proposed here:

- ◆ A much-improved settlement process could begin immediately.
- ◆ There would be almost no increase in data traffic.
- ◆ Privacy-sensitive data need not be exported from the home.

Question 9: How far would aggregated or sample data provide suppliers' with what they need in the area of wholesale hedging? Please provide examples of how the data would be used and where possible quantify potential benefits and costs.

We do not have a view on how consumption data can be analysed for purposes of wholesale hedging. However, if this is possible then it must be possible to do that by locally processing the data in the meter, rather than exporting the data to the supplier's server.

It is easy to envisage some local processing algorithm that could be used to provide the necessary data to the energy suppliers without divulging all of the half-hourly data.

For example, a single half-hourly reading could be selected at random once a month by a "hedging applet" and sent to the energy supplier. When averaged over their millions of customers this would build a statistically reliable picture of aggregated consumption patterns, with almost no loss of privacy. As an additional protection, perhaps the random measurement could be sent to the settlement authority, or the DCC, who would perform the aggregation.

Question 10: What level of data would be required and how would this be used to manage debt? Please provide practical examples.

Question 11: How would suppliers envisage using daily data to support debt management and what evidence do they have to support claims of additional savings that could be achieved with access to daily data as opposed to less frequent data?

It is not clear if energy suppliers should have a right (or a duty) to analyse data looking for customers with a debt problem. The act of identifying such customers is in itself intrusive of privacy, and raises ethical questions that should be managed by the regulator.

We do not have a view on how consumption data could be analysed for purposes of managing debt. However, if this is possible then it must be possible to do that by locally processing the data by an applet in the meter, rather than exporting the data to the supplier's server.

Question 12: How could smart metering data be used to identify and protect vulnerable consumers? Should such activity be considered a regulated duty and are any licence changes needed to create particular duties on suppliers in this area?

It is not clear if energy suppliers should have a right (or a duty) to analyse data looking for "vulnerable" customers. The act of identifying such customers is in itself intrusive of privacy, and raises ethical questions that should be managed by the regulator.

We do not have a view on how consumption data could be analysed for purposes of identifying and managing vulnerable customers. However, if this is possible then it must be possible to do that by locally processing the data by an applet in the meter, rather than exporting the data to the supplier's server.

Question 13: Do you consider that use of data by network companies to support them in maintaining an efficient and economic network should be considered a regulated duty?

Question 14: Do you agree with the requirement for such data to be anonymised or aggregated wherever possible, and how should this be monitored?

Question 15: Would suppliers be expected to advise consumers of network company usage of data given network companies do not have a direct relationship with customers?

Much of the data needed to run an efficient smart grid can be obtained at the sub-station level, but some house-by-house data might be useful also. This might include monitoring minimum and maximum voltages and floating neutral fault conditions.

Some house-by-house data might still be useful if anonymised or aggregated, but for some purposes an address would be necessary, and this could inevitably be linked back to a named occupant. So DECC should work on the assumption that network operators would be receiving personal data.

As with the other local processing examples, a key privacy principle is that the recipient of the data should receive the minimum required to perform the operation.

So we propose a “DNO applet” that could process energy data locally. If necessary, messages could be sent by this applet to the network operator. Only data essential for the correct operation of the grid would be sent (e.g. abnormal voltage levels and fault conditions); this would not include privacy-sensitive half-hourly consumption records.

Question 18: What current and future technical options exist for energy consumption data minimisation / privacy enhancing technologies? How might aggregated or anonymised data be provided in practice? Would this imply additional services to be provided by DCC?

The whole thrust of this document is to make the case for a truly innovative privacy enhancing technology.

This technology is feasible and Acute Technology can show it working on real hardware in a real smart meter. It is based on the re-use of smart card technologies deployed in billions of smart cards and SIM cards, and it is available now.

The technology could be located in the smart meter, but is probably best located within the communications hub. Thus it can be deployed without needing changes to the meters, rather an extension to the communications hub specification.

We suggest to DECC that if you are really interested in optimising the privacy of customers in the smart meter programme then it is incumbent upon you to adopt this privacy-enhancing technology rather than a business-as-usual approach which you attempt to provide privacy by regulation.

The consultation asks “Would this imply additional services to be provided by DCC?” The beauty of the Java Card applet approach is that it provides a general-purpose computing platform on the communications hub. New functionality can be dynamically deployed over the life-time of the communications hub. So yes, new services can be deployed to the home. These could be services operated by the DCC, but alternatively the DCC could simply operate as a conduit for messages that travel between the applets and their

corresponding servers.

We would welcome the opportunity to discuss this technology in depth with the DECC, and demonstrate our current implementations.

Question 19: What parts of the privacy policy framework do you think should be delivered by regulation and why?

Question 20: What is the most effective way to set out any sector specific protections around privacy (e.g. licence conditions or other alternatives)?

To the greatest possible extent, the privacy policy should be delivered by real “privacy-by-design” rather than by “privacy-by-regulation”. Privacy-by-regulation is necessary only when privacy-by-design has failed to prevent the generation of the data that must be regulated. Local processing is one example of real privacy-by-design.

The Data Protection Act is notorious for the low penalties for breaches, in some cases. DECC must consult with the Information Commissioner's Office and ask Parliament for sufficiently strong penalties to deter and punish smart meter privacy breaches.

Question 21: What practical options for authentication would provide the right balance between allowing easy access to consumer data in the home while providing the necessary privacy protection? Are there any other issues or options that the programme should be considering in developing the approach in this area?

Pairing is a constantly occurring problem in wireless communications, and there are many solutions, with varying levels of robustness. DECC should remember:

- ◆ Do not attempt to invent a new security technology, but reuse something that has been well tested. Security is difficult and roll-your-own security is inevitably compromised.
- ◆ The wireless HAN will prove an attractive point at which to attack the smart meter system, and so it is appropriate to invest in good security technology. Acute Technology advocates the use of “secure elements” (such as those used in smart cards and SIM cards) as a root of trust within the home.

In a system that includes a smart meter communications network, a consumer Internet connection, and smart card technology capable of cryptographic functionality, one viable approach to pair a consumer device with a smart meter system with a reasonable level of security might be:

- ◆ The customer signs up to a third-party data processing website, provides a meter ID, and selects a user-name and password.
- ◆ A pairing/authentication applet is deployed to the smart meter system.
- ◆ The customer logs onto the data processing website with their user-name and password with a request to initialise their new consumer gateway.
- ◆ A message is sent through the smart meter WAN from the website to the pairing/authentication applet, asking the SMHAN to admit the new consumer gateway to the SMHAN.

- ◆ The applet also generates a one-time PIN code which is displayed on the in-house display.
- ◆ The customer types the PIN code into the web browser where it is sent to the website.
- ◆ If the PIN is correct the website sends a command to the consumer gateway, through the consumer's Internet, asking the consumer gateway to pair with the SMHAN.

Other variations of this approach are possible. Balances between ease of use and security can be adjusted to suit the customer's appetite for risk and security. But having the in-house display provide a PIN is simple, cheap, and goes a considerable way to providing security: it is similar to the security devices that banks provide for online banking.

Note that this approach can be used with other consumer devices that might want to connect to the SMHAN. These could include load control devices and devices to support value-added services, such as telehealth.

Question 22: Are there other issues that need to be considered to make using the HAN a viable route for access to data in the home, from either a process or consumer perspective?

Data access via the HAN, if it requires the user to purchase a new consumer HAN gateway device, will raise considerable barriers which will prevent its take-up. If the energy supplier has to provide the gateway to the consumer then this destroys the economics of the smart meter programme.

The consumer HAN gateway, as described in the Industry's Draft Technical Specifications, is ill-defined. At best it sounds like a ZigBee to broadband gateway. As such it could only be used by customers who have working broadband already. It is likely that such a unit would require a separate mains power supply (which must be accounted for in the smart meter power budget) and messy wiring. In reality such a system is likely to be used a few times by the customer and then disconnected.

Every effort must be made to use the WAN as the route to move energy data out of the house and back to the customer, as this involves much lower costs and much lower adoption barriers. Once the data is out of the house it can be processed and sent back to the customer in many different ways: by a website, by email, to a smart phone, or by paper.

Question 23: What sort of arrangements would provide an appropriate balance between providing ease of access for consumers seeking to sign up to new services and adequate protection for consumers' data when accessed via DCC?

Do you have any suggestions for alternative approaches?

Customers' access to data through the WAN is by far the preferred approach, as explained in our answer to Question 22.

A "data access applet" could be installed in the smart meter system, designed to export data to a destination selected by the customer. This could be the energy supplier, a third-party energy services company, or a web browser.

As with other applets, the data would be encrypted using a key only shared with the

intended recipient (in this case the customer or her third-party energy services supplier).

Here is one example of how energy consumption data could be exported from a meter, and processed securely within a web browser under the control of the customer:

- ◆ The customer signs up to a third-party data processing website, provides a meter ID, and selects a user-name and password.
- ◆ A data access applet is deployed to the smart meter system.
- ◆ The customer logs onto the data processing website with their user-name and password with a request to process the smart meter data.
- ◆ A message is sent through the WAN from the website to the data access applet, requesting the data.
- ◆ The applet fetches the data, encrypts it with a one-time cryptographic key and sends it to the data processing website.
- ◆ The applet also generates a one-time PIN code which is displayed on the in-house display.
- ◆ The customer types the PIN code into the web browser where it is sent to the website and used to decrypt the energy data.
- ◆ The website processes the data and displays it to the customer.
- ◆ When the customer logs off the data is deleted from the website's server.

Many variations on this theme are possible. Balances between ease of use and security can be adjusted to suit the customer's appetite for risk and security. But having the in-house display provide a PIN is simple, cheap, and goes a considerable way to providing security: it is similar to the security devices that banks provide for online banking.

Question 25: Do you have any suggestions as to how the Foundation Stage can be used to further learn about our approach to data access and privacy?

Acute Technology and or Project Hydra partners have been adapting existing smart card technologies for use in smart meters.

In this consultation response we have limited ourselves to pointing out how valuable these technologies can be to deliver true "privacy-by-design", through local processing of energy data.

In our response to the other consultation we have described the benefits of reusing smart card technology to deliver good security.

We have smart card technology, described here, ready and working, and connected to smart meters. We would like to work with DECC, STEG and utilities to demonstrate, test and further develop the concepts we have described here.