Experiences and Challenges in using constrained Smart Objects

Introduction

This paper describes the experiences gained from the design and operation of a wireless IP-based system for use in the home and commercial building environments. Initially developed for lighting applications, we and our partners are now expanding the range of devices and applications which can be used in the system to include common home automation functions, energy monitoring and security. The paper describes how the devices are commissioned into the network, the security procedures used and then goes on to look at some wider questions associated with highly constrained smart objects, particularly commissioning and the provisioning of security material

Lighting system experience

Our experience with the Internet of Things (IoT) is in the field of lighting systems, specifically internet-enabled light bulbs, both in terms of Compact Fluorescent Lamps (CFLs) and LED Solid State Lighting (SSL). The deployment model for such bulbs is as direct replacements for existing incandescent bulbs or non-internet enabled CFLs in the existing light infrastructure in both commercial and domestic environments

The system operates in the 2.45GHz band (IEEE 802.15.4) and at the moment supports 250 nodes arranged as a mesh-under network structure; currently the devices in the network are predominantly light bulbs but other devices such as switches, occupancy sensors and energy monitoring devices will be added over time. The communication protocols between devices consist of the mesh network layer with a 6LoWPAN adaptation layer providing compression and fragmentation for a standard IPv6 implementation. UDP is used to carry a proprietary binary-encoded application layer command protocol. An aspect of the system is that it must be able to operate with or without connection to the internet; when the internet is not present, hand-held remote controls are used to operate individual lights or groups, turning them on or off or changing their brightness.

The security aspects of the system relate to ensuring that only nodes intended to be members of the system are allowed to join and participate at the link and network layer, and that security relationships can also be established at the application layer, both in terms of who is allowed to access certain applications within the system (for example only authorised personnel can operate parts of the access control in a building) and the privileges that certain users or clients have when controlling the system (for instance all users can see the temperature from thermostats around a building but only facilities management staff can change the setpoint to raise or lower the ambient temperature)

The network layer is secured using a common network key distributed to all nodes after they have successfully joined the network. At present the security is provided by encrypting all traffic in the

network using the CCM* cipher suite defined in IEEE 802.15.4, utilising a hardware AES engine built into the JN5148 System on Chip device used in all the nodes. The key size is 128 bits.

There are two methods used to join the network depending on whether there is a border router available or if only handheld remotes are present. The border router provides a bridge between the 802.15.4 network and a wired Ethernet or Wifi port depending on configuration to connect into the local home network and thence out to the internet. In the case where a border router is available, authentication is provided by a process in the border router (although it could be moved to a separate authentication server somewhere else in the network).

To join the network a device will establish a link to a node already on the network; in the case where several 802.15.4 networks are in range it is important to make sure that the joining device chooses a responding network node that is a member of the intended network. This is controlled by the border router; a button press or change to a control webpage allows the border router to propagate an indication to all nodes in the network to respond to join requests. This state times out after a short period on the nodes ensuring that there is a defined window when the network will respond to join attempts. During this window the join responses contain a marker indicating that the node which generated it will accept new joining devices which allows the joining device to find the appropriate network and establish an unsecured link to one of nodes.

The authentication process uses a system of pre-shared commissioning keys to establish that the device requesting to join the network is authorised. Before the device is introduced to the network its address and commissioning key are added to a whitelist of devices which are allowed to join the system, by typing in an identifier printed on the device into a webpage served from the authentication server (border router currently), scanning a barcode or similar method. The device also contains the commissioning key so now both ends share the same secret.

The network node with which the joining device established a link sends to the authentication server information to identify the joining device; if the device is in the whitelist the authentication server replies directly to the joining device with the network key encrypted with the joining device commissioning key. The joining device receives the encrypted network key over the clear link it established with the network, decrypts the network key and begins communicating with the network securely.

Questions and Challenges

In this section we list some of the challenges we have identified with constrained devices and questions that are still unresolved about the security procedures needed by such systems

Provision of keys in low cost devices

This is a problem particularly in systems where the user doesn't have an authentication server or does not have access to it. This is the case for devices such as smart meters, where introduction of devices to the network is controlled by the utility company, and the use of certificates and PKI to allow a security

relationship to be established. For relatively complex, expensive devices such as smart meters it may be acceptable to carry the overhead of storing and exchanging certificates, but on highly constrained, low cost devices this is not the case. Certificates are large objects (several kbyte) which is a significant proportion of the storage available; also the cost of buying certificates from a Certificate Authority may not be a burden on devices costing tens of dollars but for a lightbulb will be a significant addition. There is also the question of whether CAs could issue the required number of certificates for the billions of mass market devices which the IoT will eventually become. It would be useful to look at ways of establishing security relationships between simple devices which do not have any means of entering information other than at manufacturing time or through an insecure link, but are still lightweight enough to have a small or (more hopefully) almost zero impact on code size and data storage

Appropriateness

In our system we are using 128-bit keys mainly because we have access to a hardware AES engine which supports this key size. However if it can be shown that this level of security is in excess of what is really needed in a home or commercial building it may be possible to reduce the size of keys; this may then have a knock-on effect in the silicon implementation by reducing the number of gates required by the security engine and hence the overall size and cost of the device. It would be of benefit if work can be done on identifying the likely threats and methods of attack, and the consequences of such attacks to come up with recommendations on the appropriate strength of security, rather than one strong enough for banking-level transactions

Privacy

Work has been done in other parts of the IETF on aspects of privacy related to the movement of laptops from place to place, the implication being that if you know where a laptop is in the world you have a good idea that its owner is also likely to be there, and so you can tell if they are at home or not. This traceability comes from the uniqueness built into the laptop at the MAC level. A similar problem comes about in the devices that are deployed in the IoT; they all have a unique MAC address and this is used in many cases to generate the IP address as a combination of MAC address and network prefix. By monitoring information flows from a particular IP prefix it would be possible to deduce how much activity there is in a home, in itself a useful piece of information which may indicate if the home is occupied or not. With further correlation it may be possible to identify types of device or even individual devices, such as would allow energy use or medical information flows to be extracted. Governments particularly in Europe are keen that such inferred information is not able to leak out and it is likely that legislation will be enacted to require levels of opaqueness in information flows from the home. Will the levels of security we currently have access to be adequate, or do we need to ensure that techniques to anonymise data further during its transport are in place, particularly to break the link between the IP prefix and a location, and the MAC address and a particular device?

Ease of commissioning

This point really is a continuation from the section on provision of keys in low-cost devices; the IoT will not be deployed unless the user experience involved in installing the devices is similar to that of the non-internet enabled devices. It must be simple enough to install an internet controlled lighting system

such that the average person capable of installing a light bulb can perform the same task and not have to worry about the need to enter security parameters by hand or decide which network to connect to. Secure methods of providing this information between the network and the device being introduced need to be seamless so that the user is unaware of what is going on under the hood. From the manufacturer's standpoint, the IoT of simple low-cost objects will fail if there is any need for support to the user; their margins will not stand any level of returns or added cost such as help desks. As always, the simplest, most cost sensitive (and constrained) devices will be the most difficult to make easy to install, since they cannot stand the cost of any form of user interface – even a button. The availability or secure methods for provisioning such a device with little or no user intervention will be paramount

Sleeping and Energy Harvesting Devices

Many, if not the majority, of devices which will make up the IoT will be either battery powered or will harvest energy from their surroundings, in the form of light, vibration, heat and so on. Battery powered devices must attempt to maintain the life of the power source for as long as possible commensurate with performing their task, in order to reduce or eliminate the maintenance cost involved in replacing batteries. In order to extend their battery life such devices power down into a sleep mode for relatively long periods, typically with an active duty cycle of only 1-2%. This means that when updates to security material are made (eg expiring keys, addition to a group) they may be asleep when the information is sent. In order for the information to reach the sleeping node it will need to be buffered somewhere in the system. Mechanisms need to be provided to allow the node to be notified that a message is pending and where it is stored before it can be retrieved during one of the nodes active periods. Transfer of security information when commissioning such a device may not be such a problem since this operation will happen only rarely in its lifetime and we can probably afford to allow the device to maintain power until the commissioning process is complete before entering the operational low duty cycle mode

The case of energy harvesting nodes is often more severe, particularly those which harvest energy from a discrete event such as a switch press. The amount of energy which can be harvested even from continuous processes such as machine vibration or solar cells may place a limit on how often the device is able to operate, forcing it into the same low duty cycle regime as described for a battery powered device. Again communication is only possible for short periods until the energy store of the device has been replenished by the harvester; however for some applications where there is sufficient energy available to transmit and receive a few packets before the store is depleted it may be possible to use the same mechanisms as for the battery powered devices to buffer messages elsewhere in the network and retrieve them. The restriction may be that the buffering node should only be one radio hop away from the EH node in order that the device does not have to maintain its receiver powered up for long periods of time while the information is being retrieved

For the most restricted forms of energy harvesting nodes which derive their energy from discrete events such as button presses or switch operations, it may only be possible to send packets of restricted size, and only a small number with each operation. During operation this may be acceptable since amount of information a light switch needs to provide is rather small; however it is a different story during commissioning. In many cases these devices are so restricted that they are capable only of operating as

transmitters, with no ability to receive information. This means that there is no easy way to get information to them other than through providing external power by plugging them into a programming station or providing a small power source such as a coin cell which would give a limited ability to operate as a receiver. Both cases have their disadvantages; ease of commissioning is compromised by the need to have a programming tool and the cost of the providing the power source which is only used a few times may not be economic when compared against the total cost of the device. It may be necessary at the end of the day to accept that the EH device cannot be a true member of the IoT but must act through a dedicated proxy which maintains security relationships with the rest of the network

Conclusion

In order to be successful and pervasive, the IoT needs to be able to support highly constrained devices which are cheap to manufacture and simple enough for the average skilled person to install without extra hardware or training other than provided in simple instruction leaflets. The methods used to provision security material during commissioning of these devices needs to be lightweight and not use large amounts of storage on the devices. Some of the devices which may be introduced into the IoT may be so restricted in their capabilities that it is impossible for them to be true members, but will need proxies to allow them to participate