MWR InfoSecurity

An Overview of ZigBee Networks

A guide for implementers and security testers

Matt Hillman

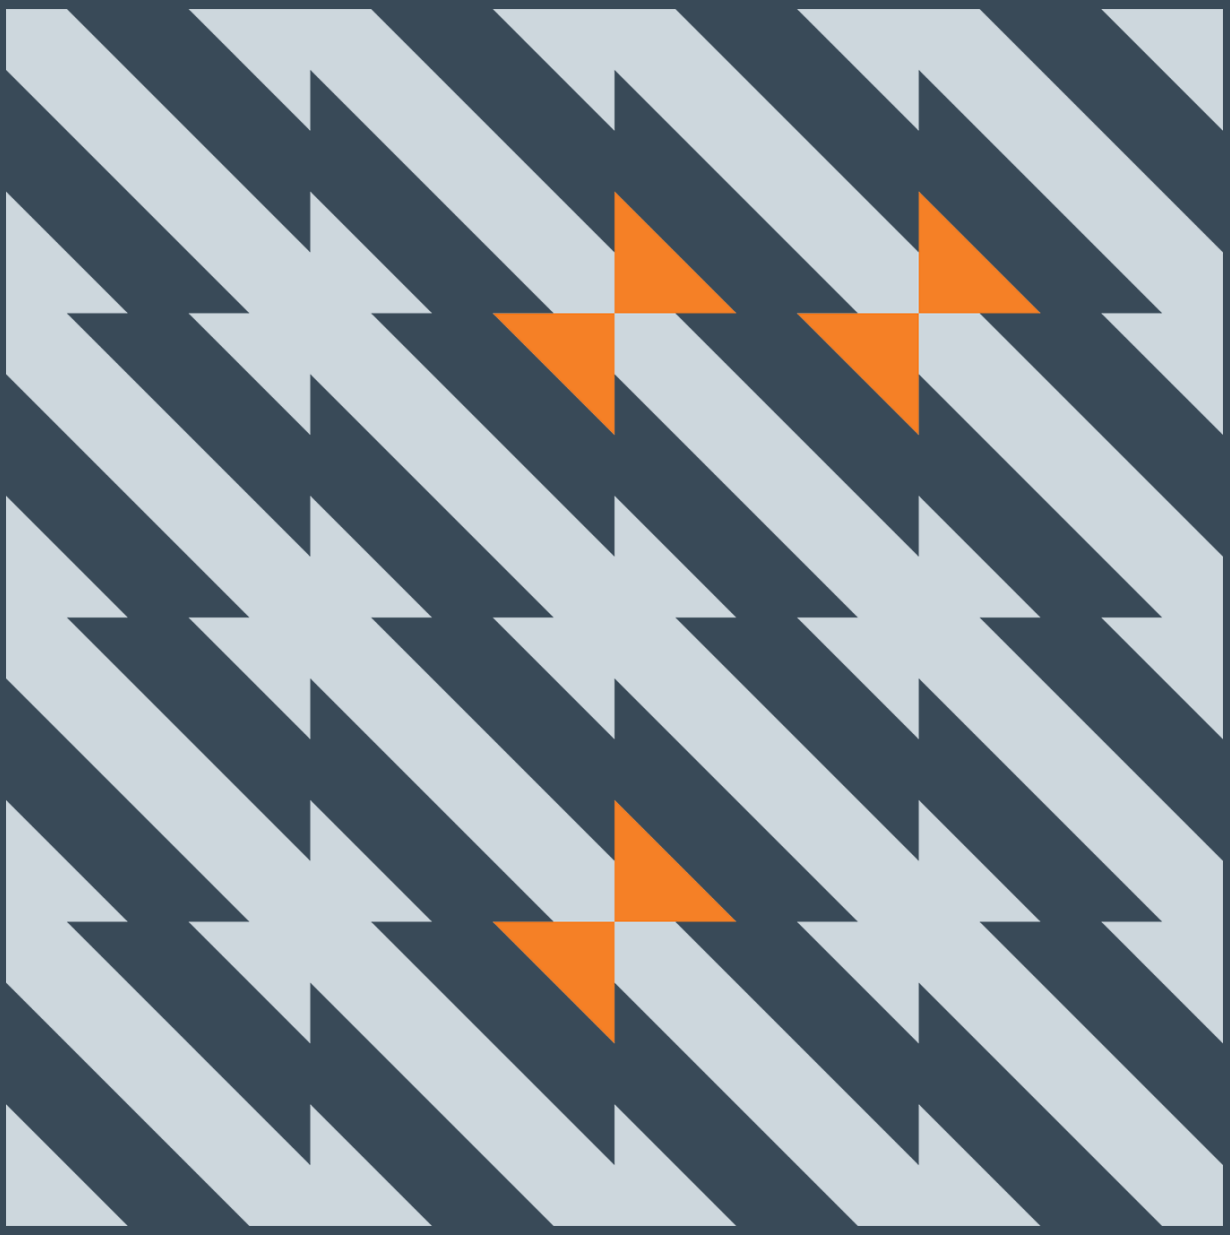

Contents

1. What is ZigBee?

ZigBee is a standard for low-power Wireless Personal Area Networks (WPANs), which is to say wireless networks with a short range, typically 10-100 meters. ZigBee is commonly used for wireless control and monitoring applications such as wireless sensor networks (WSNs), industrial plant monitoring, building control, hospitals, smart metering and home automation. There are actually public profiles defined in the ZigBee specification for many of these use cases.

ZigBee operates in the Industrial, Scientific and Medical (ISM) radio bands and the exact frequency will depend where you are in the world. It can use the 868 MHz band in much of Europe, 915 MHz in the USA and 2.4 GHz in many other locations. The 2.4 GHz band is very common as many of the available chipsets use it. The speeds available depend on which band you are using, but the maximum is 250 Kbps. This is slower than other popular wireless technologies such as WiFi but is also cheaper and lower cost.

1.1 ZigBee Versions

There have been several updates to the original ZigBee 2004 specification, and it is common to see references to both “ZigBee 2006” and “ZigBee PRO”; ZigBee PRO is also sometimes referred to as ZigBee 2007. Among other things, ZigBee PRO allows for more complex routing and dynamic channel switching if interference is detected. While ZigBee PRO is backwards compatible with ZigBee 2006 there are some limitations. Specifically, a ZigBee PRO device on a ZigBee 2006 network must operate as an End Device (more on device types later in this article), and similarly a ZigBee 2006 device on a ZigBee PRO network must also be End Devices.

Being the newer version, some of this article discusses ZigBee PRO concepts, although much of it will apply to both versions. Many people consider pre-ZigBee PRO to be “legacy” but it is important for those concerned with the security of ZigBee networks to understand older versions.

2. How Does ZigBee Operate?

More specifically, ZigBee is built on top of the 802.15.4 specification which defines the Physical (PHY) and Media Access Control (MAC) layers for low-rate WPANs (LR-WPANs). ZigBee adds layers on top of this to add more network and application intelligence. 802.15.4 is the basis for many other industrial wireless protocols as well so understanding it can be very useful to a security consultant.

A ZigBee network allows a set of devices to communicate wirelessly via one of several possible topologies. Packets of data can be sent between nodes, and may be routed by intermediary devices to more distant nodes that would otherwise be out of range. Each device has both a MAC address and a ZigBee network address, and the network as a whole has its own PAN ID shared by all devices. Packets can be protected by encryption but for this to work all nodes will need a key and as we will see later there can be issues around how such keys are deployed to devices.

The following sections provide more detail over how ZigBee actually operates.

2.1 The ZigBee Stack

A simplified view of the ZigBee stack looks like this:

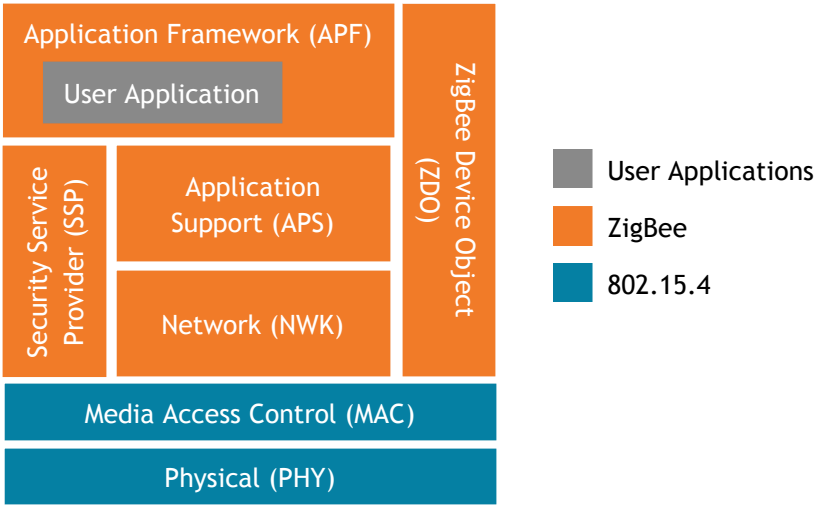

PHY – Defined by 802.15.4 the PHY layer is responsible for the modulation, demodulation and physical transmission of packets over the air and handles various things needed for robust radio transmission in noisy, interference prone environments.

MAC – Also defined by 802.15.4 and similar to MAC layers in other protocols, ZigBee does not actually use all of its features. This layer performs functions such as CSMA/CA to avoid collisions when transmitting frames and defines a frame format with things like MAC addresses etc. The MAC layer also defines network topologies which ZigBee builds upon and enhances at higher levels of the stack.

NWK – A ZigBee layer that builds on 802.15.4 and is one of the more complex ZigBee layers. It provides the ability to discover and join networks and expands on the topologies defined by 802.15.4 at the MAC layer to allow mesh networking, a popular feature of ZigBee. The NWK layer also determines routes through the ZigBee network and supports ZigBee addresses which are different to the MAC addresses present at the MAC layer.

APS – A ZigBee layer implementing features needed by ZigBee applications and acting as an interface to the NWK layer. It performs some filtering of duplicate packets from the NWK layer and maintains a binding table of nodes in the network.

SSP – Provides ZigBee security services to the NWK and APS layers including key establishment and transport, device management and frame protection.

ZDO – Responsible for the overall management of the ZigBee device. The ZDO initialises the APS and NWK layer, allows device discovery, manages binding requests and defines the device mode (coordinator, router or end device).

APF – The APF is an execution environment for ZigBee user applications and facilitates sending and receiving of data by those applications. It also provides an **Endpoint** for each application, with Endpoint 0 being reserved for the ZDO and Endpoint 255 for a broadcast address. Applications themselves implement the function of the ZigBee device (eg. a sensor).

2.2 ZigBee Node Types

There are three node types that a device can act as within a ZigBee network. Whatever node type a device is acting as it can also be doing some useful work such as acting as a sensor. Node types are only relevant to the topology of the ZigBee network and how devices help to route messages. The available node types are described below:

Coordinator – Every ZigBee network must have a single Coordinator. This node is the first node to start up and initialises the rest of the network, selecting the frequency to use, the PAN ID of the network, and allowing other nodes to join the network. It acts as the parent to nodes that connect to the network through it (its children). The Coordinator also often runs other services such as routing and certain security services, although many of these services are options and some can be run on separate dedicated nodes.

Router – Routers are not required in all ZigBee topologies but are still commonly found. They are responsible for relaying messages to other nodes. Nodes can also join the network via a Router, with the Router becoming their parent node; this can include one Router being the parent of another Router.

End Device – An End Device is a simple node that sends and receives messages but performs no other special function in the network. Other nodes cannot join the network through them. End Devices are the only nodes that can sleep according to the ZigBee specification with the parent node (a Router or Coordinator) buffering messages until it wakes up again.

2.3 ZigBee Network Topology

ZigBee networks can have one of three different topologies which affect how messages are routed and which devices talk to which other devices. These topologies are summarised below:

2.3.1 Star

Star topology is the simplest and most limited topology available to ZigBee. Devices all connect to a single Coordinator node and all communication goes via this Coordinator. It is interesting to note that this topology is actually defined by the underlying 802.15.4 specification which ZigBee builds on.

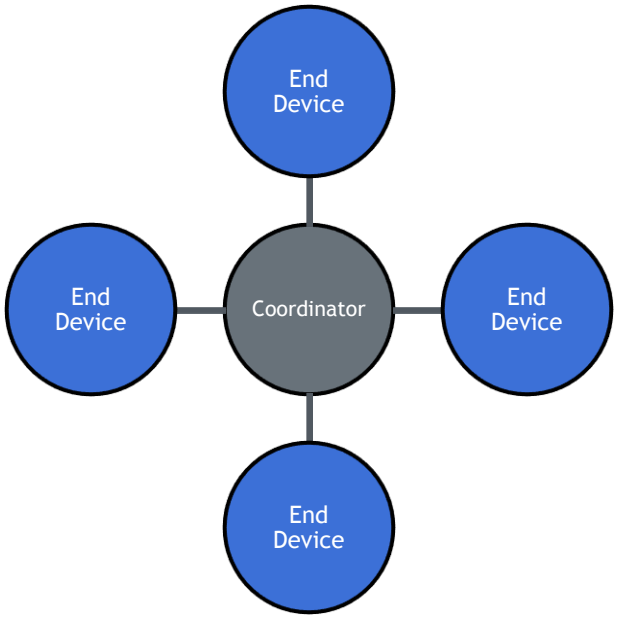

Child nodes can also be Routers, although in this topology they will not perform any routing functionality and essentially act like an End Device.

With Star topology the throughput of the network is limited by the Coordinator and if the Coordinator fails the whole network fails. The range of the network is also limited to the rage of the Coordinator itself.

2.3.2 Tree

With Tree topology the Coordinator forms the root node of a tree of child nodes. End Devices are leaf nodes (although a Router could also be a leaf if no children have joined it yet) and intermediate nodes are Routers. Direct communication can only occur between a child node and its parent, but all nodes can communicate together by messages traversing up the tree to a common ancestor and then down to the target node.

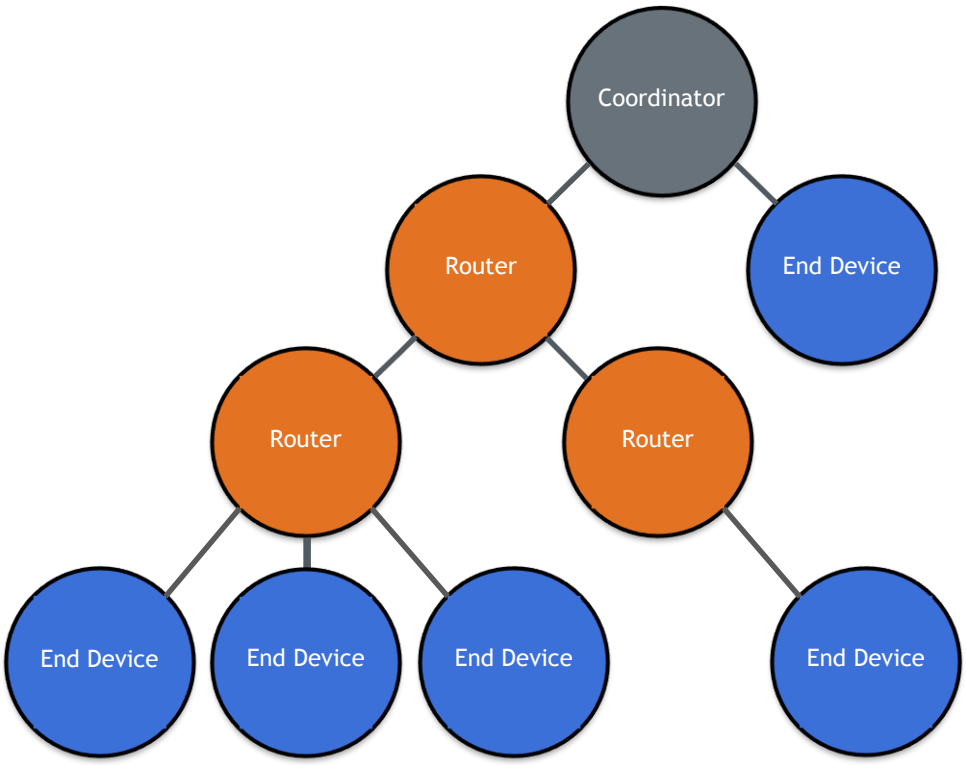

In this topology Routers are able to extend the range of the network beyond that of any single device-to-device link. However, if a Router fails there is no alternative route and portions of the network can become disconnected.

2.3.3 Mesh

Mesh topology is one of the most flexible offered by ZigBee. It is similar to Tree topology but without following the rigid tree structure and a Router can communicate directly with any other Router or the Coordinator if it is in range. This means there can be many different routes through the network to a given node, and ZigBee has a route discovery feature to find the best route to a given node and can therefore be “self-healing”.

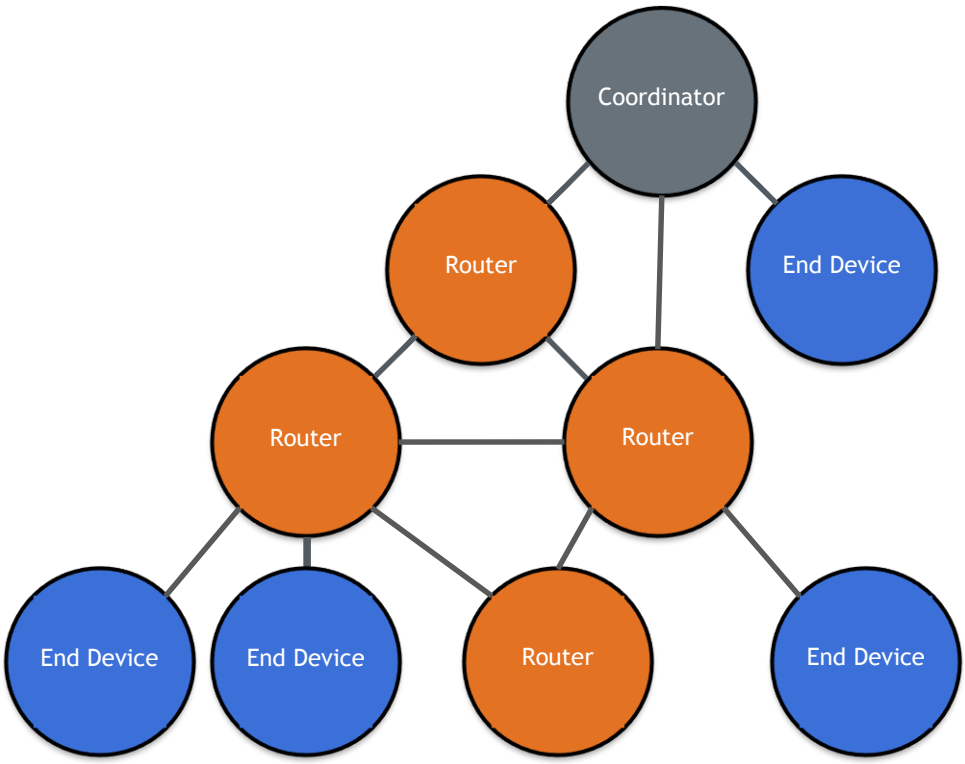

2.3.4 Hardware Device Types

A device can be classified as a Full Function Device (FFD) or a Reduced Function Device (RFD). These classifications relate to the physical capabilities of the hardware and come from the 802.15.4 specification. RFDs are often battery powered and sleep between transmissions to save power, where FFDs are usually not battery powered and do not sleep. In a ZigBee network an RFD must be an End Device and not a Router or Coordinator, as the latter two could not perform their functions correctly if they went to sleep.

2.4 Addressing and Identity in a ZigBee Network

2.4.1 Device Identity

Individual devices in a ZigBee network have two addresses, a MAC address and a Network Address (NwkAddr). The MAC address comes from the underlying 802.15.4 protocol where the NwkAddr is actually part of the ZigBee layer itself. The difference between these addresses is detailed below:

MAC Address – Sometimes referred to as the "extended address" this is a 64-bit address just like the MAC addresses you may be used to in the world of Ethernet. This is meant to be assigned to the device at the time of manufacture and should never change; it should be unique and no other device in the world should ever have the same address. It is an important part of 802.15.4 and is used for low level packet delivery. At the ZigBee layer the MAC address is rarely used, except in certain cases such as binding when the mapping between the MAC address and NwkAddr is needed.

Network Address (NwkAddr) – Also sometimes called the "short address" this is a 16-bit address that is unique only within an individual ZigBee network. The NwkAddr is assigned when a device joins the network and can change if it leaves and re-joins the network. The Coordinator always has the NwkAddr 0x0000. Depending on the version or configuration of the ZigBee stack addresses will either be assigned according to a devices position in a tree topology, or will be assigned randomly by the Coordinator. In the latter case, a Device Announcement will be broadcast at the time of assignment to allow an Address Conflict to be sent by any device that happens to already have the selected address.

Translation between MAC address and NwkAddr is necessary to actually deliver packets correctly and ZigBee provides a mechanism to allow the NwkAddr to be discovered. You can superficially compare this to the way ARP resolves MAC addresses to IP addresses in IP networks.

It is also possible to configure group addresses which allow sending of messages to a specific group of nodes who are subscribed to that address.

2.4.2 Network Identity

There are two identifiers that can be used to identify a ZigBee network, the Personal Area Network Identifier (PAN ID) and the Extended PAN ID (EPID). These are described below:

PAN ID – Part of 802.15.4, the PAN ID is a 16-bit identifier selected at random by the Coordinator when the network starts up and is used by the MAC layer to filter out packets that are not part of the same network.

EPID – A ZigBee concept, the EPID is a 64-bit identifier which can be used for more fine grained uniqueness and identification of ZigBee network. This is not sent in all packets but can be used in some situations such as when resolving PAN ID conflicts.

2.4.3 Application Level Addressing

A single ZigBee node may run one or more applications. In order to direct messages to a specific application **Endpoints** are used which are numbered from 1 to 240. There is also a broadcast Endpoint, 255, which allows a message to be sent to all applications on a given node.

Although not technically an address, messages can also be sent to a specific part of an application by specifying a **Cluster ID**. Clusters provide context or meaning to an application and allow commands and data to be exchanged in a certain standard way. Clusters are divided into **Input/Server** and **Output/Client** clusters. An Input Cluster stores attributes and allows them to be manipulated by incoming messages. An Output Cluster is where messages are sent from to manipulate attributes in an Input Cluster, and receive responses to those messages.

ZigBee applications also run in the context of a particular profile. These define a set of Clusters with particular attributes and commands and allow devices to be compatible out of the box in a specific domain. There are a number of public profiles such as Home Automation, Telecom Applications, Industrial Plant Monitoring and more. These allow devices designed for a specific purpose to all interact together in a meaningful and standard way. Private profiles can also be defined for custom behaviour not present in the public profiles.

2.5 How ZigBee Messages Propagate

ZigBee messages can be sent to a specific node, a group of nodes, or broadcast to (potentially) all nodes; they can even be sent between ZigBee networks with different PAN IDs. An overview of all of these scenarios is given below:

Broadcast – While 802.15.4 has its own broadcast mechanism, ZigBee builds on this. When a broadcast is sent, any Coordinator or Router in range will retransmit it unless it has reached the maximum number of retransmissions. This means a broadcast can only propagate a certain number of hops before it stops being sent. A process of passive acknowledgement is used to provide a level of reliability but without requiring additional messages to be sent. When a device transmits or re-transmits a broadcast it will listen to its neighbours to be sure they also re-transmit the broadcast within a certain period of time. If they do not, it will send it again.

Broadcast messages are identified by the Network Address being set to one of the predefined broadcast addresses. The most common is 0xFFFF which broadcasts to all devices, but the following also exist:

0xFFFF – Broadcast to all devices

0xFFFD – Broadcast to all devices with receiver turned on permanently

0xFFFC – Broadcast to all Routers and Coordinators

0xFFFB – Broadcast to Low Power Routers

Unicast – Unicast messages are directed towards a single node. Obviously not all nodes can communicate directly based on transmission range and network topology so messages must often pass through multiple nodes to reach their final destination. A route discovery algorithm can be used to automatically discover a route. At the Network level a Network ACK is returned to the original node once the messages reaches its destination. At the MAC level a MAC ACK is sent between each hop as the message propagates.

Group Multicast – This is used to send a message simultaneously to a group of nodes. Messages are sent with a group address which is basically a Network Address when configured for group addressing. In fact, Group Multicast messages are sent just like broadcasts and each node checks if it has Endpoints in the group before processing the message.

Bound Transfer – This is when a message is send to Endpoints which the sender has been bound to. More information on this is given in the Service Discovery and Binding section below.

Inter-PAN Transfer – This is when a message is sent to a node in a different network with a different PAN ID. 802.15.4 supports inter-PAN addressing, and while not really a part of ZigBee some devices to support this. Usually these messages are not forwarded or routed through the network and are simply sent directly to an out-of-network device that is in range of the sender. This can allow some data transfer to other devices that may not be compatible with the whole network. Importantly, such transmissions are not encrypted as they are destined to a foreign network outside of the normal security framework. Of course it could be possible for developers to add application level encryption using a shared key or similar.

For messages that need to pass through several nodes to reach their final destination, two address fields are used; “next hop” and “final destination”. The next hop is determined by a routing table maintained by each Router or Coordinator node and is filled in each time one of these nodes forwards a packet on. If no routing table entry is found the route discovery mechanism is used which uses a route discovery broadcast.

2.5.1 Service Discovery and Binding

To allow meaningful communication, depending on the purpose of the ZigBee network, nodes must be able to identify the service they need to interact with on other nodes. There are two main ways to accomplish this, Service Discovery and Binding.

With Service Discover, a node sends a broadcast requesting a certain service and nodes which have that service respond with their address. A node can store this information and communicate with the other nodes services with “direct addressing” using the Network Address, Endpoint and Cluster ID.

Rather than using discovery and direct communication, nodes can also make use of the binding mechanism. When nodes are bound it allows messages to be automatically routed without specifying the destination address and Endpoint. Binding can be one-to-one, one-to-many or many-to-one.

2.6 Encryption, Integrity and Authentication

As with any wireless network where security is a concern encryption plays an important role. There are a good number of security options that can be configured that affect encryption in a ZigBee network, and the way keys are handled and exchanged can be partly defined by the Application Profile being used and whether the network is using Standard Security Mode (called Residential Mode in ZigBee 2006) or High Security mode (called Commercial Mode in ZigBee 2006).

Ultimately, encryption can be applied at three different levels, the MAC layer, Network (NWK) layer and the Application Support (APS) layer. A ZigBee frame will contain fields from all of these layers encapsulated within one another, and each later can encrypt its data payload.

In fact the underlying 802.15.4 specification allows for encryption of frames at the MAC layer but does not specify important things such as key management and authentication schemes which are left to upper layer. ZigBee implements these controls at the NWK and APS layers.

All encryption in a ZigBee network uses AES-128, including at the 802.15.4 level. This allows hardware designed for use with 802.15.4 that contains optimised AES components to be used throughout a ZigBee implementation. AES provides symmetric encryption meaning that both sides must share a key. How nodes obtain these keys is an important security concern. There are three main methods a node can obtain keys:

Pre-installation – Keys are placed on devices out-of-band (eg. physically programmed in before deployment).

Transport - Keys are transported over the network to the device.

Establishment – Through a negotiation process keys are established without ever actually sending them over the network. Three methods of Establishment are:

- Symmetric-Key Key Establishment (SKKE)
- Certificate-Based Key Establishment (CBKE)
- Alpha-Secure Key Establishment (ASKE)

Where keys are not pre-installed there must be a device which can store and distribute or negotiate keys. This is known as the **Trust Center** and is often the Coordinator of the network, although it does not have to be for ZigBee PRO; the Trust Center also authenticates devices onto the network. Where multiple keys are used it can be possible to distribute some keys securely, but there are situations where ZigBee can be configured such that keys will be sent unencrypted over the network which introduces a short period of vulnerability.

Depending on configuration and profile ZigBee can actually use a number of keys. Keys can be used by different layers in the stack but all still have unique purposes. Keys that can be present are:

Link Key – This is uniquely shared between two devices and can be used to encrypt unicast messages between them. If a device shared a Link Key with the Trust Center (known as a Trust Center Link Key) it can be used to encrypt the transfer of the Network Key to a node joining the network. Other Link Keys operate at the APS layer and are known as Application Layer Link Keys.

Network Key – Shared between every device on the network and can be used for NWK layer encryption and protecting broadcast traffic. A set of Network keys are stored on the Trust Center and identified by a key sequence number. They can be pre-installed on devices or transported from the Trust Center. Such a transport will only be encrypted if another key is available to be used as a key-transport key such as the Trust Center Link Key.

Master Key – Used for SKKE Establishment of Link Keys. Usually pre-installed, some mechanisms may use a "key-load key" to help securely transfer a Master Key from the Trust Center to a device.

These keys also provide the basis for authentication through a challenge response mechanism. ZigBee PRO also supports Mutual Symmetric-Key Entity Authentication between any devices as well as simply authentication to the network as a whole.

A frame counter is also used which is incremented with each transmission. This helps to prevent replay attacks and forms part of a nonce value which ensures the freshness of the frame with regards to cryptography. If a device receives a frame with a frame counter lower than the previous value it will reject it. The frame counter is also associated with the current Network Key, and the Trust Center will periodically change the Network Key and reset the frame counter to zero to avoid locking the network up when the counter reaches its maximum value. The Network Key must be changed to reset the frame counter.

While encryption can provide confidentiality of data being transmitted and allow authentication to occur, it can also be used to provide assurances over message integrity. This is achieved by the Message Integrity Code (MIC) also sometimes referred to as the Message Authentication Code (MAC), which is not to be confused with the MAC layer in the protocol stack. The MIC uses AES to effectively sign the message to assure that the contents has not been tampered with.

ZigBee can operate in a number of configurations with respect to encryption level and this can enable or disable encryption and the MIC independently depending in the Security Level; this level also specifies things such as how many bits the MIC contains, where a longer MIC is considered more secure and more difficult for an attacker to forge or guess. The MIC can be 32, 64 or 128 bits.

It is important to remember that any Inter-PAN messages between networks will not be encrypted as they are outside the normal key management processes within a ZigBee network.

Proper key management is important in maintaining the security of a ZigBee network and implementers should be aware that there are situations where keys may be transmitted in the clear which opens a clear window of vulnerability to the network.

3. Conclusion

ZigBee networks can be configured and operate in many different and often subtle ways. Sometimes the exact way a given aspect of the network will operate is based on the manufacturer of the ZigBee chipset. What's

more, by nature ZigBee networks can be highly flexible with devices sleeping and waking up, connecting and disconnecting, altering the layout of the mesh network, switching channel or PAN ID and so on. In order to deploy these networks securely, or to analyse such a network as a security tester or researcher it is important to understand all these core concepts.