Introduction to Z-Wave

An Introductory Guide to Z-Wave Technology

Table of Contents

Z-Wave Overview and Functionality

Z-Wave communicates using wireless technology designed specifically for remote control applications. Z-Wave operates in the sub-gigahertz frequency range, around 900MHz. This band competes with some cordless telephones and other consumer electronics devices, but avoids interference with Wi-Fi and other systems.

Z-Wave Technology Quick Overview

- Uses a **"Network ID"** and a "**Node"** ID (Similar to an IP Address)
- Uses RF technology to transmit between Nodes (Phases do not matter)
- Uses a Mesh Network configuration
- Each A/C Powered node can act as repeaters, for extending the distance (Battery operated nodes do not repeat)
- Must have a **"Primary Controller"** to learn in the modules
- Can have a maximum of 232 devices

Radio Specifications

- Bandwidth: 9,600 bit/s or 40 kbit/s, fully interoperable
- Modulation: GFSK
- Range: 75 feet assuming a non-intrusive environment (non interference), with an **optimum range** of 30 feet.
- Frequency band: uses the 900 MHz ISM band: 908.42MHz (U.S.)
- Power limit: 1mW transmission

Network and Topology

Z-Wave is a low powered mesh networking technology where each node or device on the network is capable of sending and receiving control commands through walls or floors and use intermediate nodes to route around household obstacles or radio dead spots that might occur.

Z-Wave uses a source-routed mesh network topology and has one master (primary) controllers that control routing and security. Devices can communicate to another by using intermediate nodes to actively route around and circumvent household obstacles or radio dead spots that might occur. The following example assumes that other devices exist on the network to create the mesh.

> **Example:** A message from node A to node C can be successfully delivered even if the two nodes are not within range, providing that a third node B can communicate with nodes A and C.
>
> If the preferred route is unavailable, the message originator will attempt other routes until a path is found to the "C" node.

This allows a Z-Wave network to span much farther than the radio range of a single unit; however, with the use of several hops a delay could occur between the control command and the desired result. (Z-Wave, 2011)

For more information on the Z-Wave:

- http://en.wikipedia.org/wiki/Z-Wave
- www.z-wave.com

Z-Wave. (2011, October 11). Retrieved October 31, 2011, from Wikipedia: http://en.wikipedia.org/wiki/Z-Wave

Z-Wave Components and Terminology

Controllers

A controller is defined as a unit that has the ability to compile a routing table of the network and can calculate routes to the different nodes. There are different roles for each controller. Some of the most common are Primary and Secondary roles, also known as static controllers.

- **Primary Controller** is the device that contains a description of the Z-Wave network and controls the outputs. It assigns the "Network or Home ID" and "Node ID" to the Z-Wave node during the enrollment process.
- **Secondary Controller** contains the same "Network ID" as the primary and is required to remain stationary to maintain the routing table.

Notes:

- Any controller can be primary, but only **one** primary controller can exist on a network at a time.
- The primary controller manages the allotment of node IDs and gathers information about which nodes can reach each other. For Example:
    1. Honeywell's L5100 will support up to 40 lights, 3 thermostats and 4 Locks
    2. Honeywell's Tuxedo Touch™ Controller will support 232 devices
- The secondary controllers can obtain the network routing information gathered by the primary controller.

Slave Nodes

Slave nodes are nodes that do not contain routing tables, but may contain a network map. This map contains information about routes to different nodes if assigned to it by the controller.

- **Slave Nodes** has the ability to receive frames and respond to them if necessary
- **Routing Slave** have the ability to host a number of routes for talking to other slaves and controllers
- **Frequently Listening Routing Slave (FLiRS)** is configured to listen to a wake up **beam** during every wake up interval. → See **"Beaming"** for more information.

Notes:

- Any slave node can act as a repeater if the nodes state is set to "listen" mode. However, it is important to note that some Z-Wave manufacturers require software to enable the repeating option in the node
- If the Routing Slave is A/C powered they can be used as repeaters, battery powered devices do not repeat in an effort to control the battery life

Home ID

To separate networks from one another the Z-Wave network uses a unique identifier called the **Home ID**. It refers to the ID that the Primary Controller assigns the node during the inclusion process.

- This is a 32-bit code established by the primary controller
- Additional controllers will be assigned the same Home ID during the inclusion process
- All slave nodes in the network will initially have a Home ID that is set to zero (0)
- Once the slave node contains a Home ID it must be **excluded** before you can assign it to a different network

Node ID

A node is the Z-Wave module itself. A **Node ID** is the identification number or address that each device is assigned during the inclusion process. The logic works very similar to that of an IP Address.

- The primary controller assigns the ID to each node
- There are a total of 232 nodes available on each network
- Important Note: the Primary Controller is considered part of the network and must be subtracted from the overall node count. Therefore, the total numbers of slave nodes available are 231.

In the example below, you can see where the primary controller has a home ID of 16 (0x00001111) and each node has an id of 02 and 03. Note the primary controller will always contain the Node ID of 01.

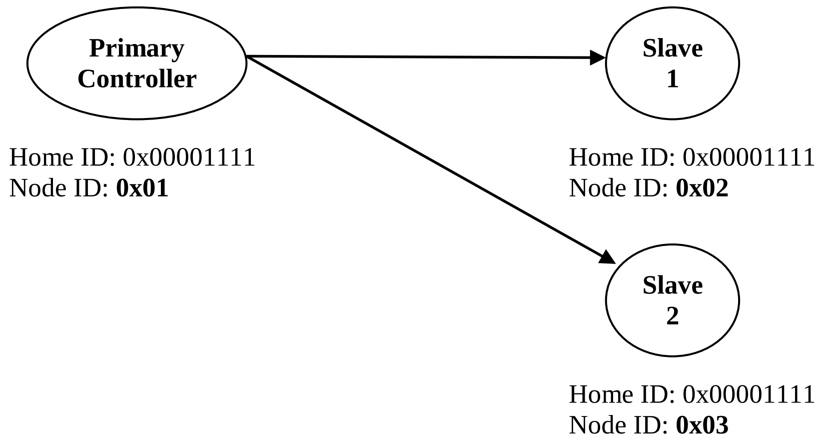

Routing

All controllers have a routing table that enables the controller to calculate the routes in the Z-Wave network. It keeps track of these routes and knows which 'path' to take to communicate with the destination node.

In the following example, the controller knows to reach the second slave node it must pass through the first node so the path would be as follows:

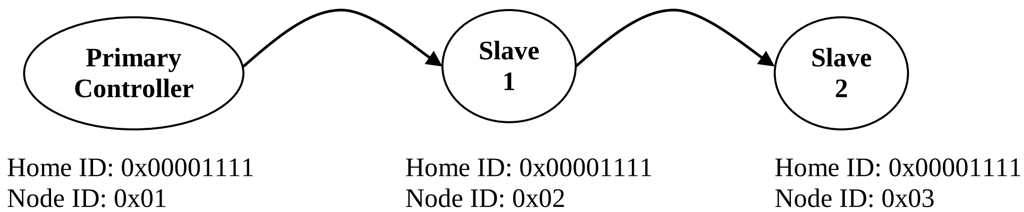

When slave 1 receives the message it will look at the destination node ID then cross-reference that with its own ID. If it does not match it will forward (**repeat**) the signal along. See "**Hopping**" and "**Beaming**" for more information.

Note: Battery powered devices, such as door locks and battery-powered thermostats, will not repeat.

Z-Wave also determines the most efficient path to take to get to the correct node. The example below describes how the Z-Wave mesh routing network decided to get to node 6. In this example we can see how the primary controller decides to use the path of “Slave 1 → Slave 4 → Slave 6” (highlighted in green) versus the path “Slave 2 → Slave 3 → Slave 5 → Slave 6” (highlighted in red). It knows that it has fewer nodes to go through so the message will be delivered more efficiently.

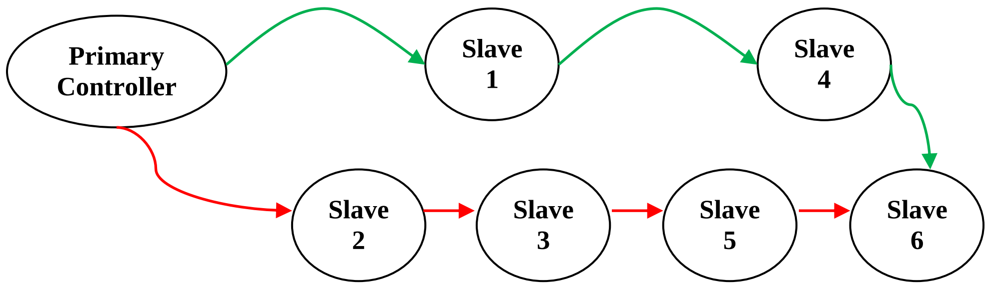

Z-Wave also has an intuitive intelligence where it uses its two-way communication to determine the position of each node, hence may determine the most efficient path is direct communication versus hopping through other nodes.

Beaming

Battery powered nodes (FLiRS) have a battery saver feature in which the node will be awake for 1.5 seconds then sleep for 1.5 seconds. The **Beam** is a carrier that is transmitted for a preset period. The 2-way battery devices operate as follows: the radio turns off (to conserve battery) for 1.5 seconds and then wakes up for 1.5 seconds. If a carrier 'Beam' is detected the battery device will fully wake up, checks into the network with a broadcast to all the surrounding nodes. When the FLiRS device receives this broadcast, it responds with an “I do” and **beams** the command to that particular slave node. If it does not hear the beam command, it falls back to sleep and 1.5 seconds later starts the cycle over.

The key is that the Beam Command is issued from devices that use a permanent power source (like AC). In addition, the repeating device must have recent firmware that have the ability to transmit the Beam Command.

How Does The Beam Command Effect My Installation?

- If there are no battery powered 2-way Z-wave devices in your network then the Beaming Command is not needed.
- If your Z-Wave installation includes a lock, battery powered thermostat, or any other Z-Wave battery powered device, read on.
- Not all repeating devices can issue a beam command. If a controller tries to send a signal to a 2-way battery powered device it will depend on another Z-wave module to repeat the command, that repeating module must support the Beam Command.
- Z-Wave nodes not capable of receiving or sending beams will disregard the data and drop the packet.
- Without a Z-Wave device that supports Beaming, the destination node may not receive the message because it was “Sleeping”

How Does Beaming Differ From Repeating?

A slave node that simply repeats the signal grabs the message, understands that it is not intended for itself and sends the packet a long. A module that support beaming will recognizes that there is a battery powered device nearby and will hold the message in a buffer until the node wakes up with the broadcast request mentioned above.

Note: In a multi-hop situation, only the last node before the destination node needs to be a module that supports beaming.

Hopping

Hopping is the number of times a message can be repeated before the message is dropped. The maximum number of hops a Z-Wave device can perform is five, however the optimum number of hops is two. While designing a system the installer or applications engineer must realize that the more hops a message takes the more susceptible it is to incorrect translations of the message. The image below represents hopping:

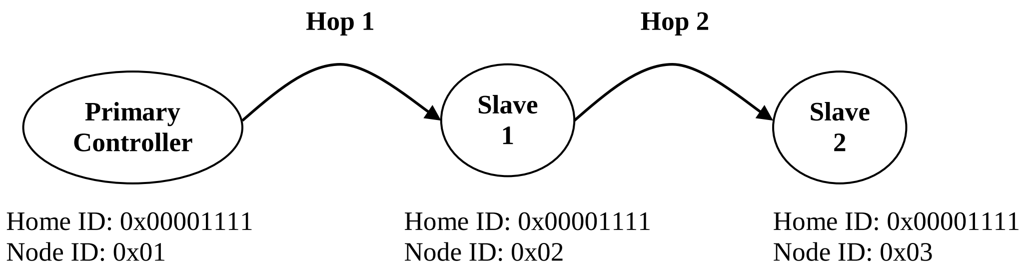

Inclusion

Including a node into a network is the process of assigning the Home ID and Node ID to the slave unit. The inclusion process is initiated on the Primary Controller first then activating the function button or command on the node itself. When an un-initialized node is activated, the Primary Controller will assign it the Home ID and Node ID.

Notes:

- If the node will not include it may have been assigned to a different Home ID.
- A possible solution to this issue is to send it an **Exclusion** command or perhaps factory default the unit.

Exclusion

Excluding a node from the network is the process of removing the assigned Home ID and Node ID. The exclusion process resets the Home ID and Node ID to zero. The only exception is a controller, if you exclude a controller the Home ID reverts to the ID programmed during manufacturing.

Notes:

- Any Primary Controller can exclude a node.
- The node does **not** have to exist on the controller's network to send it an exclusion command.

Network Wide Inclusion (NWI)

NWI is the process in which you can include devices into the primary controller that are outside the controllers range. Before NWI, all devices had to be included within close proximity to the primary controller and then the routing table reconfigured after installing. NWI allows the installer to start at the controller and include the devices that are within its range. Once those have been included in the network, then the installer can include the next level of devices that are not within direct range of the primary controller. The image below represents an ideal installation using NWI. The key is nodes 1-4 **must** support NWI in order to pass the inclusion message on to the primary controller during the inclusion process.

As depicted by the following figure, using NWI the installer will include in Slave Nodes 1-4 (red circle) first, then proceed to include nodes 5-12 (green circle). The nodes within range of the primary controller marked with a red circle and those that not within range are marked with the green circle. The red circle has a diameter of 30' from the primary controller and the green circle has a diameter of > 75' from the primary controller.

Example would be Node 7 is out of the range of the primary controller, so using the NWI feature of node two we will "bounce" the enrollment process off Node 2 to get a successful inclusion into the network.

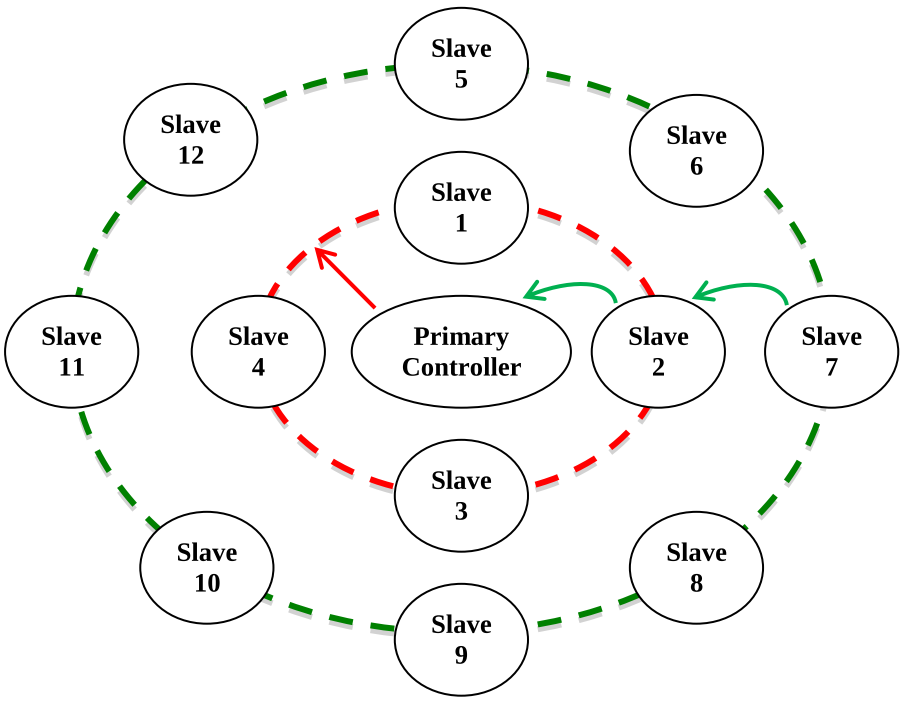

General Installation Guidelines

The following section contains important information essential to all applications concerning a Z-Wave installation. Honeywell recommends these guidelines are followed and understood before installing or servicing a Z-Wave installation.

General Information

- **Range**: 75 feet in a perfect non-intrusive environment, with an optimum range of 30 feet.
- **Maximum house size:** 7500 feet^2 using one controller

Environmental Considerations

RF interference on a Z-Wave network can severely impact the functionality of the devices. The following items must be taken into consideration when designing or troubleshooting a network.

- 900 MHz Cordless Telephones
- Wireless Speaker Expanders or Extenders
- IR to RF Remote Control Extenders
- Older Baby Monitors

Different building materials:

- Metal foil backed insulation
- Metal foil backed wall paper
- Concrete with rebar
- Plaster wall construction

Construction

Z-Wave is intended for residential and light commercial applications. Most residential construction has three basic shapes the standard Box shape, U shape, and L Shape. The ideal installation is a box or rectangular shaped house with the primary controller centrally located.

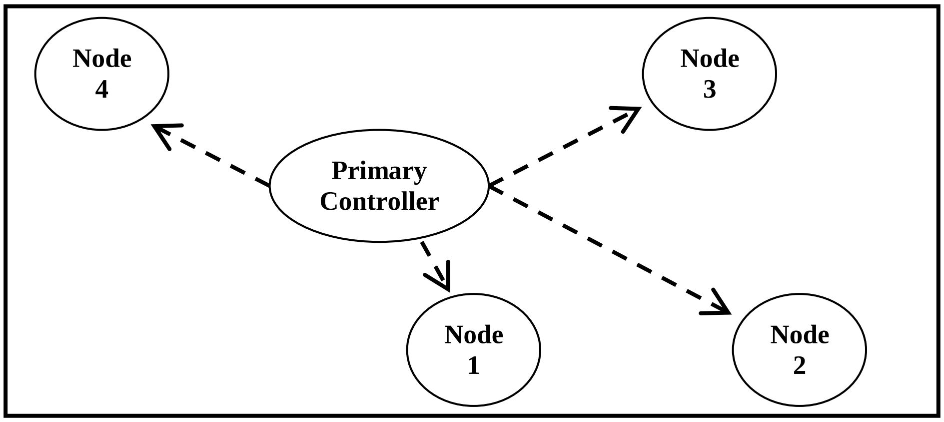

However, because Z-Wave devices are low powered there are certain construction elements that need to be taken into consideration when designing the mesh network.

Calculating Loss

Construction Materials:

As with any installation with rf devices, the installer or system designer must consider a wide variety of installation material that can interrupt the flow of data from one node to the next. Recall that the **optimum rang is 30'** for node-to-node transmission. The **maximum is 75'** in a perfect open environment as shown in the image above. Finally, keep in mind that node transmission is **Omni-directional**. See the following Chart:

Material	Thickness (Inches)	Signal Loss
Glass	.25	10%
Drywall	< 4	30%
Wood	3	30%
Stone	10.5	70%
Concrete	4	70%
Concrete	8	90%
Reinforced Concrete	4.5	95%
Concrete	12	98%

Example 1:

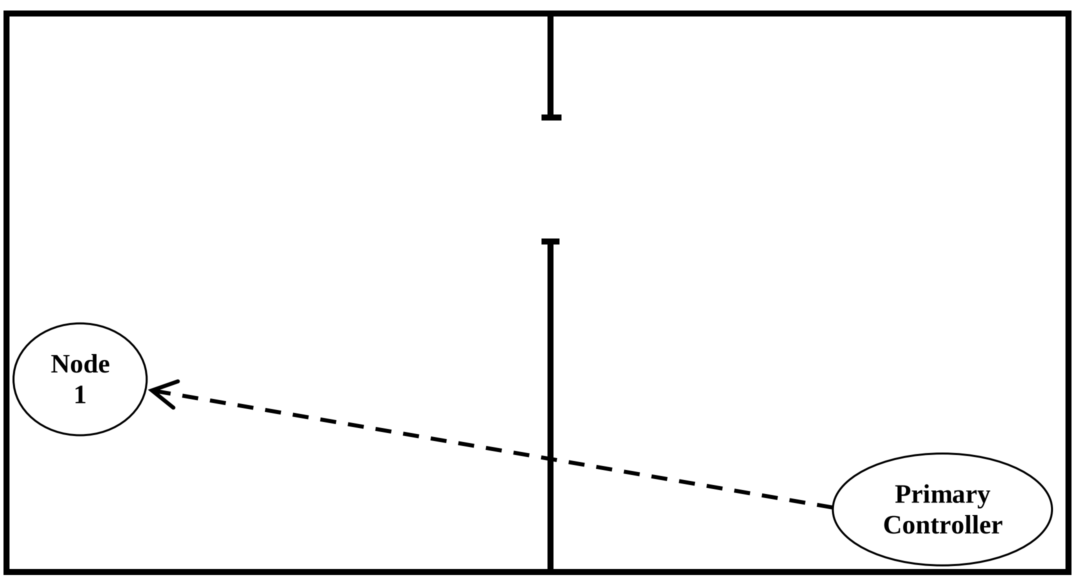

In this example, we have a simple system with one Primary Controller and one Node. For a drywall installation, we calculate the loss as follows:

Dry Wall Installation: 30ft. x 25% loss = 7.5ft (There will be a loss of 7.5ft from the primary controller to the node.)

Concrete, Plaster, Stone, Cinder Block: 30 ft. x 50% loss = 15ft (There will be a loss of 15ft from the primary controller to the node.)

What we determine from these calculations is the maximum optimal distance the message will travel is 22.5ft if using drywall (30ft – 7.5ft = 22.5 ft) and 15' if using concrete, plaster, stone, or cinder block (30 ft – 15 ft = 15 ft). If the distance to Node 1 exceeds these limitations, you will need to add another node on the interior wall to repeat the message.

Example 2:

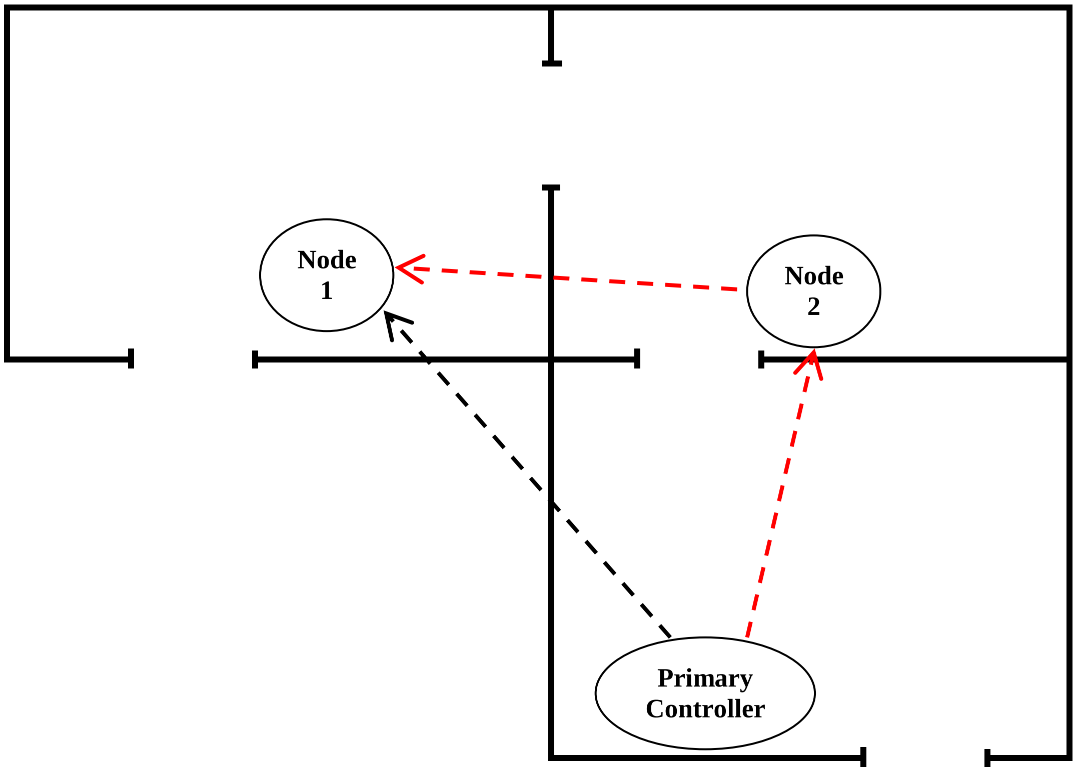

In this example, the message has to pass through two exterior walls in order to reach node 1. There for the calculation will be as follows:

Dry Wall Installation:	(30ft. x 25% = 7.5ft) x 2 walls = 15ft (There will be a loss of 15ft from the primary controller to the node.)
Concrete, Plaster, Stone, Cinder Block:	(30ft. x 50% = 15ft) x 2 walls = 30ft (There will be a loss of 30ft from the primary controller to the node.)

If this home has been constructed with brick, or if there is foil insulation, then the message will be trying to transmit through two walls that are reducing the signal by 50% each time. This installation will suggest another node to be placed on an interior wall to repeat the signal to node 1. Taking into consideration the optimum range of 30ft. it would NOT be recommended to communicate through a concrete wall.

Designing Device Routing

All installations must be well designed to create variable routing paths for different nodes. A solid design element is to avoid any one node being a key point of transmission to other nodes. This can create bottlenecks and latency issues. Use the following floor plans as considerations for any installation:

Good System Design Plan:

A good floor plan has a well-distributed mesh network, has high node density to compensate for failed devices and creates multiple routing options.

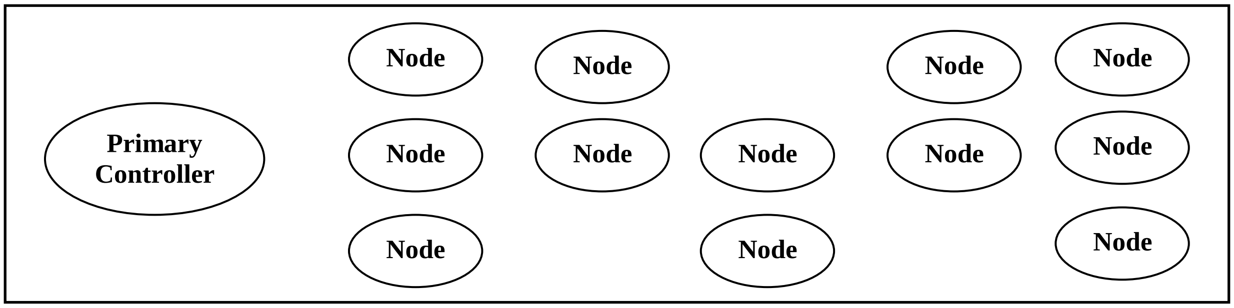

Poor System Design Plan:

A poor system design contains holes in the routing options, sparse node density, routing dependency through a central device, and large distances between devices.

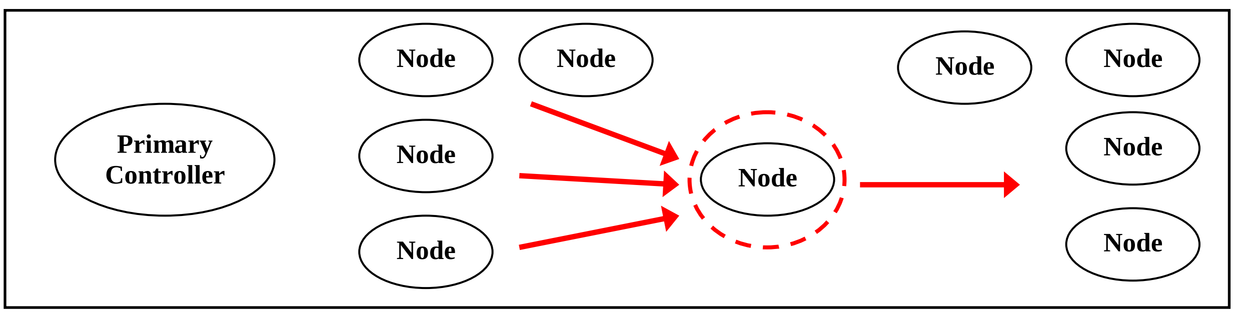

Why is this important to recognize?

This is important to recognize because the nodes to the left of the node surrounded by the red circle are dependent on that node. In large applications, this can cause latency issues and bottlenecks in data transfer that could result in loss of information. In addition, it contains the routing path to the nodes on the right and because this is a key element to the design if the product were to fail, it would require a complete system rebuild. All devices on the right would have to be excluded then re-included.

What is Z-Wave Plus

Z-Wave Plus is essentially just an extension of Z-Wave. Z-Wave Plus is a new certification, also known as 500 Series, 5th Generation, Z-Wave for Gen5 or just plain Gen5.

Z-Wave Plus extended features include:

Z-Wave Plus in Practice

If you have a complete Z-Wave Plus system, where every device is of the new Gen5 variety, you will get all the benefits that the new features offer such as extended battery life, much longer range and increased bandwidth.

Below are some of the effects when using Z-Wave and Z-Wave Plus devices, sensors and controllers based on 300-series and 400 series Z-Wave chips.

Battery Life

The battery life for Z-Wave Plus devices is significantly improved over previous generations. It doesn't rely on any other devices in the system, therefore, you will see the extended battery life offered by Z-Wave Plus devices- up to 50% longer than existing devices.

Range

The range of the devices relies heavily on other devices in the network. If you are using a mix of Z-Wave Plus and existing devices you will not see the increase in range offered by Z-Wave Plus.

Controller

The main consideration is the Z-Wave controller- if the controller is not Z-Wave Plus enabled then all devices added to the that controller's network will default to acting as Z-Wave. This is because Z-Wave Plus is back-wards compatible with Z-Wave devices, when it is installed with Z-Wave devices it behaves just like a Z-Wave device as those existing devices have no way to communicate with it using Z-Wave Plus commands.

If you do have a Z-Wave Plus controller, then you will start to see the benefits of Z-Wave Plus, but make sure that the other devices in the 'route' to the Z-Wave plus device are also compatible. Otherwise, it will again default to operating as a plain Z-Wave device.

Second Generation of Z-Wave Plus

Second generation of S2 will include Z-Wave DSK and SmartStart. S2 security allows an including controller to verify that a joining node(device) is indeed the physical device that it claims to be. Depending on the user interface, an including controller may allow the installer to enter a Device-Specific Key (DSK) string of digits or scanned as a QR code.

The DSK of the device is used only during inclusion, where the device is granted one or more of the network keys by the gateway(controller). These keys are used to encrypt the communication, and only shared with authenticated devices.

Below is an example of a controller prompting for the DSK of a Z-Wave Plug in module. The DSK can be found on the QR code of the product.

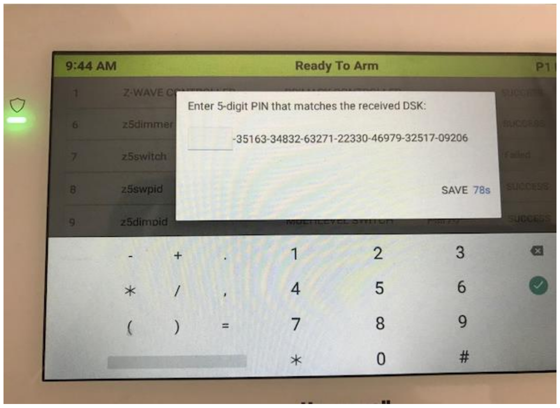

A QR code reader may also be used to determine the DSK of the Z-Wave device. The QR code reader will display the following string of digits, the 5 Digit DSK will be Digits 13-17 as seen below.

900117186003617753516334832632712233
046979325170920600100409701792022000
099203041234101332

Z-Wave SmartStart removes the need for initiating the end device to start inclusion. Inclusion is initiated automatically on power-ON, and repeated at dynamic intervals for as long as the device is not included into a Z-Wave network. As the new device announces itself on power-ON, the protocol will provide notifications, and the gateway can initiate the inclusion process in the background, without the need for user interaction or any interruption of normal operation. This improvement also removes the possibility of other devices being included, as the SmartStart inclusion process only includes authenticated devices.

Provisioning List

Z-Wave SmartStart is designed to facilitate the inclusion of multiple devices simultaneously, enabled by the Provisioning List, and by storing unique end device DSKs in the gateway. By simply scanning the end device QR code before shipment to the end user, the Provisioning List is created in the gateway. This ensures that the gateway is ready to include the end device out of the box as soon as it is powered on, eliminating the need for any further user interaction on the end device. The order in which the devices are powered is irrelevant; the gateway will include them as long as they are on the Provisioning List.

Removing SmartStart devices from a Z-Wave network

Because SmartStart always includes end devices securely, it is possible to do remote reset of an end device when it is removed from the provisioning list. The gateway can automatically perform these steps. If a device is reset without being removed from the gateway's provisioning list, it will automatically be included again the next time it is powered on.