

**Deployment Best Practices  
for  
Municipal Wireless Networks**

March 2007

## Table of Contents

Introduction.....	3
Municipal Wireless Mesh Design Principles .....	4
The SkyGateway .....	6
First-Hop SkyExtenders.....	11
The Second-Hop SkyExtender .....	14
The Third-Hop and beyond SkyExtender .....	17
The SkyConnectors .....	19
Analyze the Target Area of Coverage .....	20
Identify a Central Location for the SkyGateways.....	23
Next-Hop Considerations .....	26
How Many SkyGateways?.....	28
Deployment Strategies - SkyGateway Placement.....	29
Perform a Site Survey to Ascertain Usable Frequencies .....	30
Analyze Equipment on a Single Frequency.....	32
Review the Topology for Self-Interference Candidates .....	47
Fill In Remaining Coverage Areas with Additional Nodes .....	52
Validate Link Topology for Areas of Concern .....	55
Monitor for Stability.....	57
Summary .....	59

## Introduction

The installation of a SkyPilot Mesh Network, like any network, requires careful planning of the installation, integration and operation of the network. The addition of wireless into the mix adds some additional planning considerations including RF path, site preparation, node mounting, cabling, shielding, grounding, etc.

This group of best practices has been published to help the operator mitigate potential issues that could impact the end-user experience. The issues that typically impact wireless installations are interference, multipath, line of sight (both visual and Fresnel), self-interference and adjacent channel interference. All of these topics will be covered in some detail as we progress through the material. An explanation of each of these issues along with examples of how to compensate, or correct, for them will also be included.

Beyond the considerations of the RF signal path are those of bandwidth management, traffic shaping, network management, etc. This paper will focus specifically on the proper planning, installation, deployment and management of the SkyPilot Municipal Mesh Network. Other papers will include best practice approaches to lightning protection and grounding methodologies, network traffic and capacity planning and fault management practices and procedures.

The topic of municipal deployment strategies will commence with an overview of the basic practices, and principles surrounding an effective deployment of SkyPilot hardware - a dive into the theory of municipal mesh network design. It will then progress toward applying that theory to a planning phase wherein the proposed network is placed on a map of the coverage area. Once the theoretical structure has been completed it's time to engage in site surveys which include location of appropriate mount points, verification of signal path and RF spectral analysis. Then, get ready to install the hardware and deploy it. The final sections will cover strategies for initial and ongoing topology analysis to ensure that the network performs up to design standards consistently.

There is a lot of material to cover and some of this may seem somewhat tedious. However, spending sufficient time to learn and grasp this material and information will make the deployment of SkyPilot Municipal Broadband Mesh Network a nearly painless process.

Let's review the topical areas that we'll be covering within this best practices paper:

- Design Principles – An Introduction
- Mapping the Proposed Territory
- Site Survey Principles and Procedures
- Installation of Equipment
- Deployment Practices
- Topology Analysis and Review
- Topology Monitoring and Maintenance

## Municipal Wireless Mesh Design Principles

Suffice it to say, any “Best Practices” paper ends up discussing the theoretical concepts behind these practices so that the operator can better understand what he/she is impacting when they do something in the wireless mesh. This white paper is no different, as certain design principles need to be understood in order to effectively deploy a properly functioning municipal wireless mesh network.

The SkyPilot product has many features built-in, that make the deployment of a mesh network much less painful than other products on the market – including directional antennas which significantly reduce the impact of a transmitting node on its neighbors, an internal bandwidth management protocol which fairly distributes communication opportunities among all active child nodes of a parent, prioritization of traffic and bandwidth rate control.

When you’ve completed reading this section you will understand the following topics:

- The SkyGateway
  - Best location
  - Spatial diversity
  - Introduction to self-interference
  - Managing hop counts
- The 1<sup>st</sup> Hop SkyExtender
  - Best location
  - Locating child nodes
  - Reducing self-interference
- The 2<sup>nd</sup> Hop SkyExtender
  - Best location
  - Locating child nodes
  - Reducing self-interference
- The third-hop and beyond SkyExtender
  - Best location
  - Locating child nodes
  - Reducing self-interference
- The SkyConnectors
  - Best location
  - Reducing self-interference
- Spacing requirements for nodes



## The SkyGateway

The first item to be introduced in the theory of mesh network design is the SkyGateway. The SkyGateway is, for all practical purposes and intents, the most important node in the mesh network.

The SkyGateway is the source of all traffic destined for the SkyPilot mesh and it's the termination point for all traffic coming from the SkyPilot mesh. As such, it's location within the mesh is crucial for the mesh to function holistically. The first-hop links, to the SkyGateway are the most critical links within the mesh network. Any elements of design compromised around the SkyGateway will compromise the entire mesh network, so it's important to spend some time discussing this very important aspect.

When you've completed reading this section you will have come away with an understanding of the following topical areas:

- Best location
- Spatial diversity
- Introduction to self-interference
- Managing hop counts

## Locating the SkyGateway

This section will cover the location of the SkyGateway with respect to all other nodes within the mesh. At the risk of sounding like a broken record we're going to repeat that the node that requires the most careful attention to placement within the SkyPilot network is the SkyGateway – we can't say this enough. Compromising the location of the SkyGateway within the mesh can lead to increased self-interference, multipath, lack of system capacity due to low link modulation rates or other issues.

The goal, with the installation of any wireless node is to strategically locate it in such a manner as to reduce the probability of it interfering with any other node on the network. To this end, the **SkyGateway should be located at the center of the target service area.** Why the center? When the SkyGateway is located at the center of the network, it can utilize a concept called **spatial diversity**. We're going to cover this topic in greater detail in the following section.

It is important to note that, based on projected and actual traffic rates the design "center" of the network may not always be the physical center of the network. This is an important concept to grasp and is covered, more, in the section on spatial diversity, which follows. The design end-goal is to **surround the SkyGateway with first-hop nodes that are active in all directions.**

One final consideration, with regards to putting the SkyGateway at the center of the network has to do with **minimizing hop-count across the network.** When the SkyGateway is located at the edge of the network it requires a larger number of hops to reach the far-end of the network. With the SkyGateway located in the center of the network, hop-count will be minimized.

## What is Spatial Diversity and why is it Important?

Spatial diversity is a topic that is commonly seen in discussions centered on the elimination of self-interference within a wireless network. Self-interference is characterized by the fact that when a node transmits, within a wireless network, it has potential to interfere with the transmission of other wireless nodes within that network. If the interference, with other links, is strong enough it will cause data loss on those links at a minimum, lowered modulation rates (which translates to lower throughput) and possibly a loss of link, altogether. Spatial diversity is only available on wireless nodes which have directional antennas. Without directionality to RF transmission, there can be no spatial diversity.

### Concepts Required to Understand Spatial Diversity

In order to understand spatial diversity better, we need to understand some concepts about the SkyPilot protocol and node behavior within that protocol. More detail, on these topics, may be found in the Support section at [www.skypilot.com](http://www.skypilot.com). The information, presented here, will be a cursory overview for the purposes of explaining how spatial diversity can be accomplished within the SkyPilot mesh network.


There are two types of nodes within the SkyPilot mesh – a parent node and a child node. The SkyGateway can only function as a parent node and the SkyConnector can only function as a child node. The SkyExtender, however, may function either as a parent or a child node – depending on its current operational needs within the mesh.

A parent node communicates with child nodes and either transmits data to the child, or allows the child to transmit data to the parent. The parent node uses a scheduling agent to determine which child nodes receive bandwidth, how much bandwidth, when the bandwidth will be allocated and in which direction that bandwidth will travel (either to, or from, the child node).



A child node must first listen for its parent. If the parent requires the attention of the child, it will communicate that intention to the child during the listening period and the child will respond accordingly – either receiving data, or transmitting it. In the case of the SkyExtender, if the parent node, to the SkyExtender, does not require the time/attention of the SkyExtender, the SkyExtender can turn itself into a parent node and communicate directly with its child nodes. Each SkyExtender, within the SkyPilot mesh, is a self-managing entity; managing bandwidth allocation between itself and all child nodes with active links to it. It performs this management without the aid/support of the parent node. This is an important concept to remember when dealing with self-interference issues – every SkyExtender has the potential to be a parent and transmit independently of the SkyGateway, or its parent node.

### Spatial Diversity Defined

All of that said, we need to define what the concept of spatial diversity means. In this case, a picture is worth a thousand words, and it is presented in the diagram to the right. In short, spatial diversity refers to the location of **active links**, with respect to the parent node. Theoretically perfect spatial diversity has **active first-hop nodes** located at 90 degree angles to each other, surrounding the parent node, as shown to the right. This provides maximum separation of antenna coverage areas. The SkyPilot hardware has eight embedded directional antennas and this arrangement utilizes



those antennas in the best manner possible. If each node, located around a parent node experiences the same utilization rate, the parent will spend equal time in all directions, thus reducing the potential for self-interference in areas of the network to which it's not communicating.



The illustrations, to the left, and to the right, show what spatial diversity is not. These are examples of poor spatial diversity.

As can be seen in the diagram on the left, all of the SkyExtenders are located on one side of the SkyGateway, forcing the SkyGateway to transmit continuously in a single direction. This can lead to several issues including increased incidence of self-interference (a function of how active the SkyGateway is), increased incidence of multipath, especially standing waves and other issues within the mesh.

A similar scenario can be seen in the diagram to the right. This diagram indicates a seemingly “better” spatial diversity in that the first-hop nodes are located in a roughly 90 degree pattern around the SkyGateway. The lack of spatial diversity, in this instance, is not due to node location its due to node activity level. The nodes on the right are inactive as compared to the nodes on the left. As in the previous example, this leaves the SkyGateway actively transmitting on a single side leading to the same issues noted above, but with less frequency, as a function of the amount of activity engaged with the nodes located to the SkyGateways right side.

These diagrams lead us directly to the topic of self-interference, which will be covered in an overview format, in this paper. You may find more detailed information on self-interference in the support section of [www.skypilot.com](http://www.skypilot.com).

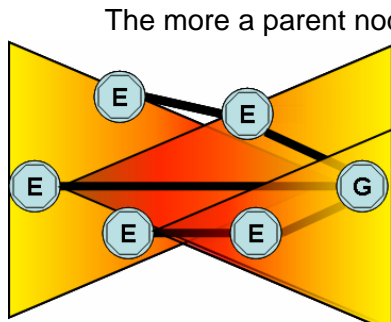
## Spatial Diversity and Multipath and Self-Interference

Why should child nodes be located around the parent node, as opposed to being located on a single side? As discussed above there are two reasons: reduction of self-interference and reduction of multipath. The more times a parent node transmits from a given antenna the greater the probability that those transmissions will interfere with other nodes in that region of the network. The more child nodes with active links to a single antenna the more time the parent node will spend transmitting to that sector, thus increasing the potential for self-interference, as well as the potential for significantly more multipath in that sector.

By locating **active** first-hop nodes in a spatially diverse manner around the parent node we significantly reduce the probability of the parent node transmitting in a given area for too long. The parent node will be actively scheduling transmissions to all child nodes located around it. If transmissions, from interfering nodes, are received on antennas that are located opposite the currently active antenna, rear-side rejection comes into play, lowering the impact of the actual interference. This is also the case for interference entering a node from the sides. This statement



will become clearer as we examine the placement of first and second-hop nodes with respect to each other and the SkyGateway.



less constant in a given area.

The more a parent node transmits around itself, in a spatially diverse manner, the less chance that multipath will “build up” on a given side. Remember that multipath translates into a delay spread which interferes with the received signal. The more transmissions that occur, in a given region, the greater chance there is for multipath, and multipath standing waves, to develop as a function of the amount of time spent in that sector. Multipath reflections are more likely to interfere with reception when the parent node stays actively engaged on a single side. This is due to an issue called “standing waves”. Standing waves develop when transmission is more or


Moving from the topic of multipath to self-interference, the illustration to the left depicts what occurs, in the situation where the SkyGateway is located on one side of a mesh network. This not only applies to the location of the SkyGateway with respect to its active child nodes, but also directly applies to the location of the SkyExtender with respect to its active child nodes.

The areas with the darker red color are locations where the signals are the hottest, in terms of transmission power. The signal power degrades toward the yellow end of the spectrum, as a function of distance between transmitting node and receiving node.

As you can see, when the SkyGateway transmits to the SkyExtender in the center of the network, that transmission is going to blanket, or interfere with, the transmissions between the other first-hop SkyExtenders and their child nodes (remembering that each SkyExtender can function as a parent node when its parent node does not have anything for it to do). This is self-interference.

The same conditions exist when second-hop SkyExtenders, in this network, transmit to their first-hop parents. They will cause self-interference with transmissions between the other first-hop SkyExtenders and the SkyGateway. Because they’re using the same antenna, or one adjacent, the rejection of the interfering RF signal will be low.

In the event that a service area requires more capacity a second, third, or more, SkyGateway could be added to fulfill the additional capacity requirements. This solution can also be utilized to obtain greater coverage density in municipal deployments. The following diagram is an example of how this would be accomplished. The nodes in blue are from a network with one primary frequency. The nodes, in red, form another network, on a different frequency. This type of methodology will ensure that adequate coverage is maintained, while maintaining both spatial and frequency diversity.



## Managing Hop Counts

The final topic of discussion is the management of hop counts within the mesh. By locating the SkyGateway at the very center of the mesh, the hop count is automatically minimized by design. The center of the network is equidistant from all points surrounding it. Locating the SkyGateway on one side of the mesh will create an automatically higher hop count to reach the far end, in addition to contributing to the issues, mentioned above.

It is recommended that the hop count be maintained at, or below, four hop counts, from the SkyGateway. Each hop adds additional latency to the network and also increases the load for the first and second-hop aggregation points. High latency impacts the end-user experience and ultimately the trust in, and acceptance of, your deployment. Managing hop-count and adding additional SkyGateways, and overlay networks will ensure the best possible end-user experience.

## First-Hop SkyExtenders

Now that you have a good grasp on the location and function of the SkyGateway in the mesh, we can begin to focus on the principles surrounding the placement of the first-hop SkyExtenders. The first-hop SkyExtenders, as a group, determines what we call system capacity. Again, as with the SkyGateway, the placement and location of the first-hop SkyExtenders, with respect to the SkyGateway AND their child nodes, is crucial to the stable and efficient operation of the mesh network.


When we're completed with this section you will have a good understanding of the following elements, with direct application to the first-hop SkyExtenders in the mesh network:

- System Capacity
- Best location
- Locating child nodes
- Reducing self-interference

### A Note about System Capacity

It was stated in the introductory paragraph, to this section, that the first-hop SkyExtenders determine the system capacity, or the capacity of the mesh. What is meant by this statement? We're going to make a couple of quick assumptions here to illustrate this point, so that the math is easier. The concepts are the same in the real system with actual numbers.

For sake of this argument, we're going to assume that each of the SkyExtenders, connected to the SkyGateway, is bearing equal load in terms of network traffic. We're also going to assume that the link speeds, listed, are actual data rates, rather than raw modulation rates. Assuming a 4-second interval, where each first-hop SkyExtender utilizes one second, we calculate an average data rate of:  $(54\text{Mbps} + 36\text{Mbps} + 12\text{Mbps} + 6\text{Mbps}) / 4 \text{ seconds} = 27\text{Mbps}$ . We have to average the total rate of transfer, for each link, across the 4-second span of time during which transmission was occurring.




Contrast this result to the rate achieved, in the same time period, if all links were 36Mbps, rather than the diverse rates shown above and we achieve an average rate of 36Mbps, nearly 10Mbps more than the above average rate of transfer. The average rate of transfer is higher, even though some of the link rates are lower, or the same in some cases, as the previous example.

The first-hop links, as can be seen from the above examples, determine the capacity of the mesh network. The higher the modulation rates that can be achieved, on these links, the greater the capacity of the mesh connected to a SkyGateway.

## Location of First-hop SkyExtenders


As noted above, the location of the first-hop SkyExtenders, with respect to their child nodes and to their parent SkyGateway is as crucial to network performance as the location of the SkyGateway itself. This is because the first-hop SkyExtenders are aggregation points for the entire mesh network, transmitting to, and receiving from, all nodes in their related quadrant (see spatial diversity above).



The **active** first-hop SkyExtenders should be located in a spatially diverse manner about the SkyGateway. In theory, there should be four **active** first-hop SkyExtenders per SkyGateway deployed and these four SkyExtenders should be located at 90 degree angles from each other, as shown in the diagram. As noted above, in order to avoid issues with multipath and self-interference, the first-hop SkyExtenders should not be located on one side of the SkyGateway. This arrangement provides the most strategic deployment method for the first-hop SkyExtenders.

Another deployment rule-of-thumb, with regards to the first-hop SkyExtenders is that they should be located relatively equidistant from the SkyGateway. This provides link modulation rates and RSSI values that are similar between all of the first-hop links. Significant variations, between the SkyGateway and the first-hop SkyExtender links, in terms of RSSI and/or modulation rates, will end up compromising both system capacity and the ability to differentiate between signal and interference. Always shoot for first-hop links that are roughly equidistant and have comparable RSSI and modulation rate.

## A Quick Note about Nodes Extremely Close to the SkyGateway




Because of the density of a mesh deployment, there will be a “first- tier” ring around the SkyGateways that is closest to the SkyGateway mount location. It is HIGHLY recommended that these first-tier SkyExtenders be parented to a single SkyGateway and that these SkyExtenders have NO child nodes linked to them.

This will avoid any unnecessary self-interference occurring in the mesh. Once we’re beyond this initial “first tier” ring, standard deployment practices will allow the addition of child nodes to the first-hop nodes. This only applies to SkyExtenders located

extremely close to the SkyGateways. The expansion of the concept of the close-in tier and the first-hop nodes attached to other SkyGateways is illustrated in the diagram, to the right.

Notice, from this diagram, the spatial diversity employed with the first-hop SkyExtenders in the blue and the yellow networks. The first-hop SkyExtenders are located at 90 degree angles to each other and their spacing, in the coverage area, is equidistant from the SkyGateway. This is the type of installation/deployment we’re shooting for when determining the placement of the first-hop nodes.



If further density is necessary, additional SkyGateways can be added at the center of the mesh and additional first-hop nodes can be linked to the new SkyGateways. The new first-hop SkyExtenders will be interleaved between the already existing units, on this diagram.

## **Self-Interference and First-hop SkyExtenders**

The first-hop SkyExtenders can contribute to, and be impacted by, self-interference on the mesh network. In order to sufficiently mitigate the effects of self-interference on these links **good spatial diversity must be maintained around the SkyGateway**. Lack of proper spatial diversity will have the SkyGateway transmitting more in one sector than the others, thus increasing the chance that self-interference, or significant multipath interference will result.

The **SkyExtenders, in the first-hop, should be located in a relatively equidistant spacing about the SkyGateway**. This will ensure that the link modulation rates and the link RSSI values are equal and reduce the susceptibility to interference that may be generated by the mesh network. Target the RSSI and the modulation rates as high as possible for these first-hop links as they determine system capacity.

## The Second-Hop SkyExtender

The next step in the design theory of municipal networks, to cover, is the aspect of second-hop SkyExtenders. The second-hop SkyExtenders are two physical hop counts out from the SkyGateway. Their parent nodes will be the first-hop SkyExtenders with active links to the SkyGateway. As these nodes still serve as aggregation points, attention to their placement, with respect to the SkyGateway and other SkyExtenders on the same frequency but in different topology clusters, will need to have strict attention paid to it.


After reading this section you will understand the following concepts as they relate to the Second-hop SkyExtender:

- Best location
- Locating child nodes
- Reducing self-interference


### Locating the Second-Hop SkyExtender

As mentioned in the opening paragraph, the location of the second-hop SkyExtenders is also key to the efficient operation and deployment of the mesh network. This section is going to explore some rules of thumb surrounding the placement of the second-hop SkyExtenders with relation to their parent node, each other and their child nodes.

More than likely there will be multiple SkyExtenders functioning as second-hop SkyExtenders, connected to the first-hop SkyExtender parents. Again, as with the SkyGateway, we're looking to provide as much spatial diversity as we can between the second-hop SkyExtender and the child nodes that will be connecting to it.



Perfect spatial diversity, at this tier, would look like the diagram, to the left. However, a more real-world spatial diversity, in a network would look more like the diagram to the right.



This combination of parent node, and attached child nodes, is what is referred to, in this document, as a topology cluster.


### Links to Parent Nodes and Child Nodes

The discussion of spatial diversity, in the second-hop nodes is an excellent segue into the discussion of some rules surrounding the placement of all second-hop nodes. These rules, when followed, will reduce the amount of self-interference, multipath build-up and other potential issues within the mesh network. While it's realized that, in most real-world networks, you're not going to be able to achieve perfect node placement, working to get as closely as possible to these deployment guidelines will yield a highly successful municipal mesh deployment.

The first, and most critical rule to follow is: **avoid forming child links on the antenna used to link to a parent node**. Forming a child link, on the antenna used to connect to a parent node, will create conditions for self-interference that cannot be avoided, within the mesh.




The example to the right is an illustration of this point. Second-hop SkyExtender  $E_2$  is located between the first-hop SkyExtender  $E_1$  and the SkyGateway. Whenever  $E_1$  transmits to  $E_2$ , it will interfere with any transmissions being received by the SkyGateway on the antenna sector(s) facing  $E_1$  and  $E_2$ . The closer the proximity of  $E_1$  and  $E_2$  to the SkyGateway the more destructive this interference will be. If there are child nodes located between  $E_2$  and  $E_1$  they will create additional self-interference.




There are ways to rectify this situation. Let's take a look at a few of them. Assuming that RSSI levels and modulation rates between the second-hop SkyExtender and the SkyGateway were similar to the RSSI level and modulation rates of the first-hop SkyExtender to the SkyGateway the scenario could be resolved simply by parenting the current second-hop SkyExtender to a different SkyGateway, as is shown to the left.

It is important to maintain that "balance" between the RSSI and the modulation rate, on every link. Thus if there are some significant differences the best solution would be to add another SkyExtender to the mix, as shown on the right. The addition of the new SkyExtender, to this situation will allow for better modulation rates and RSSI values between the first-hop SkyExtenders and the SkyGateway thus offering the highest resilience to interference. This also will provide slightly greater coverage density in the area, another added benefit.




When determining positions of child nodes in the second-hop, and beyond, the rule of thumb for placement of the child nodes, is that **child nodes should be angled back, and away, from the parent node**. This allows the best spatial diversity to be achieved with the second-hop nodes.



What does this look like, in practice? As you'll see from the diagram, on the left, the SkyExtenders in the second-hop (labeled  $E_2$ ) are angled back, and away, from the parent node ( $E_1$ ). This provides spatial diversity and significantly lowers the chances of self-interference occurring. In practice, the municipal network design would attempt to keep the topology cluster ( $E_1$  – the parent node, and  $E_2$  the child nodes, limited to three nodes to reduce the amount


of self-interference that could be generated within the cluster, or with other clusters on the same frequency, but located elsewhere in the mesh.

An example of how a "complete" second tier would look is illustrated in the diagram, to the right. This arrangement provides 360 degree coverage of the area to be serviced. It also meets the requirements for spatial diversity in that each of the child nodes is angled back, and away, from the parent node link.



If there is a need to increase the coverage density, it can be accomplished by adding another SkyGateway and SkyExtenders which will link to the new

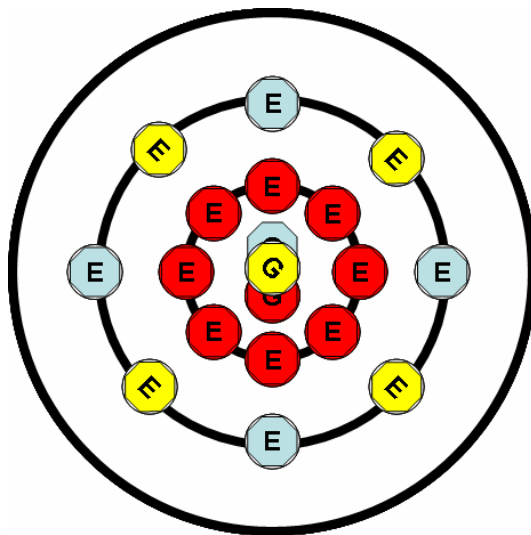
SkyGateway. This interleaves the coverage areas of the SkyGateways and provides significantly denser coverage than would otherwise be possible with a mesh consisting of a single SkyGateway.



An example of interleaving another mesh network is shown in the picture, to the left. In this example there are three interleaved mesh networks used to increase coverage density, reduce self-interference, multipath, standing waves, etc. You can see that the rules of spatial diversity, parent/child node link formation, etc. have all been followed in this example. In some municipal mesh networks, this interleaving scenario is carried out with 4-5 SkyGateways at a single location.

Should coverage be required in the immediate vicinity of the SkyGateway a fourth SkyGateway should be added and links formed to first-hop ONLY nodes in order to serialize transmission to those nodes, as per the previous discussion.

It's worthwhile to note that this strict attention to detail is really most applicable to the first and second-hop nodes within the mesh. These are the most active nodes within the mesh and they generate the greatest potential for self-interference, multipath (standing waves), etc, if care isn't paid. Beyond the second-hop, the changes from exact placement with other nodes in the mesh, to more of a ensuring that placement doesn't cause interference with first and second-hop nodes.



and attention focus respect to focus on self-nodes.

generate the greatest potential for self-interference, multipath (standing waves), etc, if care isn't paid. Beyond the second-hop, the changes from exact placement with other nodes in the mesh, to more of a ensuring that placement doesn't cause interference with first and second-hop nodes.



## The Third-Hop and beyond SkyExtender


We're finally to the "final" tier in terms of mesh network design. We call this the final tier mainly because the rules that follow, from here on out, apply to all nodes beyond the third-hop. At this hop count we're not so much concerned about the perfect relationship between nodes as what we are the amount of self-interference generated within the first and second-hop nodes. The activity level, in the third-hop and beyond becomes significantly lower, as these nodes are aggregating less, and less, user traffic. Thus, if self-interference occurs among the third and greater hop nodes, it's impacting a small amount of user traffic, if at all. If self-interference occurs with a first-hop node, it's impacting the aggregated traffic from the entire mesh behind it.

This brings us to a quick point about the SkyPilot nodes. A node will receives data, traffic from multiple child nodes. The data is concatenated and sent to the parent node. This concatenation of data from multiple child nodes increases link transmission efficiency. Nodes that are multiple hop-counts from the SkyGateway are tasked primarily with handling small bursts of traffic from individual users. Nodes that are close to the SkyGateway, in terms of hop-count, are managing aggregated traffic from their child nodes.. Losing data, at these nodes, impacts many nodes in the system. Losing a radio frame at these hops will impact a large number of users, in the mesh. This is why we spent so much time emphasizing careful node placement, mesh interleaving and other topics, at the start of this section.

When you've completed the section on the third-hop and beyond SkyExtender you will have garnered enough information to understand the following topical areas:

- Best location
- Locating child nodes
- Reducing self-interference

### Locating the Third-hop SkyExtenders



what this means.

As has been noted, above, the critical factor to pay attention to, with respect to locating SkyExtenders in the third tier, and beyond, is the elimination of potential self-interference with nodes operating in the first and second tiers of the mesh. There are a few rules of thumb that need to be considered when placing the third-hop SkyExtenders.

The first rule has already been covered in the second-hop SkyExtenders. As a child node, the third-hop SkyExtenders should be located back, and away, from the antenna that their parent's parent is linked to. While confusing in writing, the illustration, to the left, shows

### Reducing Self-Interference

As can be seen, the SkyExtenders in the third-hop, labeled (E<sub>3</sub>) are located "back" and "away" from the link between E<sub>2</sub> and E<sub>1</sub> providing not only spatial diversity, but also the avoidance

of any self-interference with the closer hops in the mesh. The amount of spatial diversity, to be required here, is not as great as the links closer to the SkyGateway.

The most important rule of thumb is “avoid forming child links on the antenna used to link to the parent node. This applies to the location of the third-hop, and beyond, nodes with respect to their parents. In this situation, putting a third-hop node anywhere between  $E_1$  and  $E_2$  would be detrimental to the performance of the first-hop links on the network.

## The SkyConnectors

In some implementations the use of a SkyConnector may be required to reach customers that may want a much higher Quality of Service, or may not otherwise be able to obtain wireless service, due to path obstructions, or issues, that cannot be overcome by the lower powered 2.4GHz side of the SkyExtender. This section will cover the basic rules surrounding the placement of the SkyConnector within the mesh.

When you have finished reading this section you will have an understanding of:

- Best location
- Reducing self-interference

### Best Place to Locate SkyConnectors

SkyConnectors, like their SkyExtender counterparts, have some rules that surround their placement to ensure the best performance within the mesh. The vertical beam width and horizontal beam width on the SkyConnector has slightly less reach than a SkyExtender so the impact to the mesh isn't quite as bad if the placement isn't perfect.

It is important to remember that when a SkyExtender transmits to the SkyConnector, it transmits at the same power level that it would utilize when transmitting to another SkyExtender. Thus it is important to locate SkyConnectors in such a manner as to avoid interference with other nodes, when the SkyExtender is transmitting to the SkyConnector.

### Reduction of Self-Interference Generated by a SkyConnector

The rules about locating a SkyConnector are the same as those for locating a SkyExtender; the SkyConnector should be located back, and away from the link used to connect to the parent node. It should be located in an area such that when the parent SkyExtender transmits to it, the transmission will not interfere with other SkyExtenders, on the same frequency. While this may be hard to avoid, in some deployments, this is the intended end-goal.


## Analyze the Target Area of Coverage

The first step, in network planning, is to analyze the target area of coverage. Personally, we like to start with a Google Earth map of the territory we're going to cover, a very large monitor with incredibly high resolution, several cups of coffee and a few minutes of time to spare. The goals, in analysis of the target areas break into the following categories;

- Map the boundary of your mesh network.
- Locate the points of interest for the mesh
- Plan the average coverage area for a single node

## Mapping the Target Coverage Area

Mapping the coverage area translates into figuring out where the boundary of the mesh is going to be. Place that boundary, clearly, on the map of your target area. An example of this phase is shown in the photo, below. Note this display is taken directly out of Google Earth. For more accurate views of targeted coverage areas, you may want to, at this phase of the network planning/design process, and want to use software such as EDX, or the freeware version that provides similar functionality, called Radio Mobile.



In this, and the examples following, we will be utilizing the coverage area, shown to the left. This area is approximately 2.5 miles square. This target area, for our municipal network, is a largely suburban area, with light young trees and a combination of homes and apartment/condos.

## Locating Points of Interest

In this initial phase of planning we're working toward identifying key features, on the map, that will assist in the planning phase of network design. Once the strategic coverage area has been mapped, the next step is to go through that map and locate points of interest, including major roads, intersections and towers and buildings with substantial height above the surroundings.

Several of these key points of interest have been circled on the map to the right.





Major streets are highlighted in pink, major intersections, buildings and light poles have been highlighted in yellow.

Why these key areas? Wide streets allow the wireless operator to take advantage of longer distances with minimal obstructions. This has the direct benefit of allowing higher RSSI values (RSSI is a function of obstructions in the Fresnel Zone Line of Sight) and as a direct result higher modulation rates. Major roads, and large intersections house potential mount points for the SkyExtenders used to service the mesh network, that are well above the surroundings, and provide clear 360 degree coverage. Thus they are good potential targets to mount first and second-hop SkyExtenders for the mesh. Towers, taller light poles, and other geographical landmarks with substantial height above the surroundings also serve as good potential mount points for the SkyGateways, as well as first-hop SkyExtenders. These, too, provide minimal obstructions in the transmission path and will allow maximum RSSI and modulation rates for given links.

## Mapping SkyExtender Coverage Area

After the boundaries of the SkyPilot network are determined, the next step in the planning process is to determine the approximate coverage radius of a single node on the network. As previously mentioned, it is recommended that SkyExtender coverage be confirmed by performing tests in the field using an area of the network considered to represent typical, or average, topology. Coverage should not be driven by mesh node spacing, but rather by coverage requirement for 2.4 GHz connectivity. Node spacing is then a factor of achieving desired reception for the weakest links in the coverage area, which is usually a laptop user with wireless network card.

It's also important to note that customer requirements and expectations will drive node spacing on the network. In more densely populated urban areas where customer expectations about coverage and signal quality are high the nodes will have to be spaced closer together to eliminate/reduce any potential dead spots in the area. In less densely populated areas where customer expectations are lower, the nodes can be spaced at greater distances.

In less densely populated areas, you may be able to increase the center-to-center spacing considerably, and in more densely developed areas, with high absorption, diffraction and reflection characteristics you may have to reduce that value to reflect the installation environment.

Once this radius of coverage has been established, potential node locations with respect to the coverage radius can be plotted onto the map using squares to identify the range. By plotting the coverage area for each



node onto the map you can then survey the surroundings for the best potential mount points (i.e. utility poles, building rooftops, and other high mount points that will provide good RF line of sight and adequate coverage. This mapping process prepares you for the next step – the site survey.

## Identify a Central Location for the SkyGateways

Now that you have a general idea of where the mesh nodes are going to be placed – we’re not entirely certain, yet, as we haven’t left the office to survey the mount points, it’s time to review the network and determine the best location for the SkyGateways that will be servicing the mesh.

When you have completed reading this section you will understand the following concepts:

- How to identify a location for SkyGateway placement
- What site considerations need review for the SkyGateway
- How Many SkyGateways Can Be Placed

## Identifying the Best Location for a SkyGateway

As discussed in the theory outline the best location for a SkyGateway is going to be at the very center of the mesh. This works well, in theory, but rarely works out to where there’s a building, tower, or other object with just the right location and mount opportunities to house the SkyGateways that will be servicing your mesh. In practice you’re always going to be a little left, right, above or below the mesh center with the SkyGateway placement. There are several considerations that go into the mount point, for the SkyGateway, and those considerations will be covered, in this section.

### Location, Location, Location

Of course, theoretical placement is best, but when you can’t do that the location that you choose needs to be as central as you can possibly make it to provide good spatial diversity. The central location is chosen to minimize the amount of self-interference on the network, reduce hop counts to the edge of the network, etc. The diagram, below, shows the choice made between actual center and realistic center of the mesh.

Now that you have the target cover area identified, look at the initial placement choices. Are there areas where coverage may not be as dense due to lack of people, buildings or other reasons for not providing coverage in an area? Do you project a “side” of the mesh being significantly more active than another, even through you may have the same number of SkyExtenders servicing both regions?

Traffic load considerations can provide you with some answers related to where the SkyGateway should be located within the mesh. In reality, you want the SkyGateways to be located at the “active



center” of the mesh, rather than to one side of the “active center” (see page 5). We realize that traffic rates / conditions are variable from day-to-day and even hour-to-hour in these networks, what we’re looking for, here, is the “active center” on a day-to-day basis.

The reason we’re focusing on the “active center” of the network just a bit, is because this may provide some extra leeway in terms of locating the SkyGateway at, or near, the center of the mesh, as it’s been designed/planned. The center of the mesh is the starting point, let’s continue to look at other considerations that come into play with the central location that has to support multiple SkyGateways.

### **Mounting Considerations**

The next point of consideration is the support of multiple SkyGateways at a given mesh center. As mentioned above, to obtain the coverage density required for many large municipal wireless mesh installations, there will need to be four or five SkyGateways at a given mesh center. The site chosen, to mount the SkyGateways must have the wherewithal to support the installation of five SkyGateways.

What considerations are we talking about? We’re discussing the elements of a successful wireless path from the SkyGateway to its first-hop SkyExtenders. We’re also discussing the elements of proper spacing to reduce inter-node interference and other issues that might arise from mounting nodes too closely. Let’s bullet point individual considerations for the SkyGateway mounting point:

- **Center** of the mesh
- Minimal obstructions for **360 degree coverage** area
- Minimum of **12 feet vertical spacing** between SkyGateways
- **Higher elevation** than surroundings
- **Support for at least five SkyGateways** based on above criteria
- **Support for data backhaul** to the location that will support traffic rates for SkyGateways
- **Chosen height must low enough to avoid shadowing** large portions of the mesh.

Of all the requirements, for the SkyGateway mounting considerations, there are two that require the most attention – 360 degree coverage area and the ability to get data connections to the SkyGateway.

The 360 degree requirement is driven from the need for spatial diversity on the network. We need to have SkyExtenders located around the SkyGateway, preferably with 90 degree separation to maintain spatial diversity. This is a critical design requirement and works to eliminate self-interference, multipath, standing waves and various other RF problems which will reduce network performance, and efficiency.

The mounting height of the SkyGateway also needs to be addressed. The SkyGateway, of all components in the mesh, needs to be mounted the highest in order to provide the greatest RF reach and coverage. The SkyGateway and SkyExtenders have a six (6) degree tilt that can be used to your advantage to increase the height, but be careful not to raise the SkyGateway so high that the child nodes are going to be shadowed by the SkyGateway.



It is also a good idea to identify how you're going to provide data backhaul, from this site, to the Internet. A wireless, or other backhaul, must be provided to allow end-user data access. The amount of data bandwidth provided will be a function of the total number of SkyGateways at the location, and the calculated average rate of throughput from each of the SkyGateways.

Finally, at least one of the SkyGateways should be dedicated to servicing only first-hop SkyExtenders located in close proximity to the mounting location for the SkyGateways. As mentioned above, servicing only the first-hop SkyExtenders, at this close range, should eliminate the problems with self-interference that would normally be experienced if the SkyExtenders that were in very close range had child links to second-hop SkyExtenders.


## Next-Hop Considerations

When the SkyGateway mounting location is being considered it's also important to consider the mounting locations for the first-hop SkyExtenders that will link with the SkyGateway. As the initial site survey is being performed, for the SkyGateway, spend some additional time at the site to identify potential mount points for the first-hop SkyExtenders.

These mounting points should conform to the following requirements:

- **Clear** Fresnel **Zone Line of Sight** from SkyGateway to SkyExtender locations
- **Equidistant link spacing** between SkyGateway and first-hop SkyExtenders
- **One location per direction** (e.g. N / S / E / W) to ensure spatial diversity
- Clear Fresnel Zone Line of Sight from first to second-hop SkyExtenders

What types of mounting points should we be looking for, for the first-hop SkyExtenders? Apartments, towers, tall light poles or traffic signals at intersections, parking structures, office buildings, municipal infrastructure buildings, etc. are all types of structures for potential mount points of the first-hop SkyExtenders. Most of these locations will provide adequate Fresnel Zone Line of Sight to minimize multipath, sufficient height above the terrain so that the second-hop SkyExtenders can be easily reached, etc.



Another consideration, for the location of first-hop SkyExtenders is along major, wide, streets that provide minimal obstruction and clear Fresnel Zone Line of Sight. Locating the first-hop SkyExtenders along these major thoroughfares will provide, on average, much better RSSI levels and modulation rates than if they were located along narrower streets with more obstructions.

If you're looking to provide additional coverage density, you will need to incorporate additional SkyGateways, and expand the number of first-hop SkyExtenders and mount points. Any additional SkyGateways follow the same rules of spatial diversity, Fresnel Zone Line of Sight, equidistant spacing, as the initial SkyGateway installation.

An example of this configuration is shown, to the left. There are two complete mesh networks (red and blue) and a third overlay network, started (green).

Note that the red network runs true North / South / East / West and the blue overlay, runs at 45 degrees to that, covering Northeast, Southeast, Southwest and Northwest. An overlay mesh is used to fill-in dead spots not covered by the initial mesh deployment. The number of SkyGateways and SkyExtenders required in the mesh will be determined by mesh density requirements.

## Considerations for Second-Hop SkyExtenders


Second-hop SkyExtenders require similar attention as the first-hop SkyExtenders. When you've completed the initial site survey for the SkyGateways, the next step is to perform a site survey for the first-hop SkyExtenders. During this site survey, spend time looking for mount points for the second-hop SkyExtenders, keeping in mind that second-hop SkyExtenders need to be pointed back, and away from their parent SkyExtender **AND** the SkyGateway, and that you should avoid forming links, to child nodes, on the same antenna used to communicate with the parent node. The second-hop SkyExtenders should be located at 45, or 180 degree angles to each other to maximize RF isolation.

Possible mount points for the second-hop SkyExtenders should follow the same criteria as first-hops. There should be clear Fresnel Zone Line of Sight, relatively equidistant link spacing from the parent node, one location per direction to ensure good spatial diversity, etc.

## How Many SkyGateways?

The last consideration is how many SkyGateways should, or can, be located at the SkyGateway mount point? We recommend that four (4) to five (5) SkyGateways be mounted at the center of the mesh before moving on to another mesh source. This number of SkyGateways should more than adequately fulfill the density requirements to service a given area with a functional mesh that is the best compromise between complete coverage and significant numbers of dead-spots, in a dense urban deployment.

In less dense urban areas the number of SkyGateways required to service a given area may drop considerably. The difference in requirements might be easily seen by viewing the differences in two downtowns. The more rural version, below, would require significantly fewer nodes to service than the more densely populated area in the photo above.



Secondarily, density of the mesh will be determined by the total area that is required to be covered, and the number of “dead spots” within the mesh that don’t require coverage.

## Deployment Strategies - SkyGateway Placement

Now that you're through the planning phase of the mesh it's time to install equipment, perform the final site survey's to determine usable frequencies and bring the network online, tune it up and bring it to full operation. This final stage before full operation is the deployment stage.

The deployment process involves several steps, and we're going to be covering them, in some detail. Most of these will be related to RF planning, link evaluation, link tuning, etc. For detailed information on physical equipment installation and grounding, please refer to the support section of our website at [www.skypilot.com](http://www.skypilot.com).

After you've completed this section you will have a solid foundation for understanding the following topics:

- RF Site Survey Suggestions
- Analyzing the initial deployment and building topology clusters
- Reviewing links for self-interference
- Filling any remaining gaps in coverage
- Validating topology for additional areas of concern
- Monitoring for stability
- WE'RE OPERATIONAL!!!!

## Perform a Site Survey to Ascertain Usable Frequencies

An RF site survey should be performed, in addition to the site survey performed above to ascertain an acceptable site for mounting, to determine operating frequencies for the mesh. The period during which the mesh nodes are being installed is a good time to focus on an RF site survey to scope out usable frequencies in the 5.8GHz range for the backhaul, and the 2.4GHz range for the wireless access points. This process is used only to determine the frequencies that will be used to start the next section “Analyze on a Single Frequency”. You will use the frequencies, determined in this process, to bring up the SkyPilot Mesh.

The RF site survey consists of using a spectrum analysis tool, such as Cognio, to evaluate the RF energy present at a given location. Many of these tools not only allow real-time analysis but also capture for replay and later viewing. Site survey's should include measurements at various times during the day – e.g. a measurement in the early morning, morning, afternoon, evening and late evening, etc. to determine the frequency, time and intensity of interference at a given location.

When you've completed reading this section you should know, and understand:

- The most important RF site survey's to complete
- When the RF site survey should be completed
- How long should the RF site survey should be performed

## The Important Areas for RF Site Surveys

We'd like to think that every node on the network should have an RF site survey performed. RF site surveys are most important at the SkyGateway location, and the first and second-hop SkyExtender locations. As noted above, the links to the first-hop SkyExtenders determine the capacity of the mesh. The SkyGateway and first-hop SkyExtenders need frequency spectrum that's free and clear to maintain appropriate link quality. If minimal interference exists it may be overcome by having high RSSI values (signal strength) on the links. Higher RSSI values equate to better interference rejection. Target link RSSI values will be covered shortly.

## When Should the Data be Captured and for How Long?

SkyPilot encourages you to spend time, at each of the major locations, especially those of the SkyGateway and the first-hop SkyExtenders, performing spectrum analysis. The spectrum should be sampled at all hours of the day to ensure link quality/integrity round the clock.

If this cannot be feasibly accomplished you should at least get a clear picture of the interference that exists at these key sites during the projected peak periods of network usage – i.e. network busy time typically from 8AM to 5PM, Monday through Friday.

Captures should be run for a minimum of one-half hour at each site. This allows you to determine if measured events are merely anomalies, randomly occurring events, or persistent

spectral events in that location. Anything that appears to be persistent across the measurement period should be marked as unusable spectrum.

### **Which Spectrum Analyzer?**

SkyPilot Networks, Inc. uses the Cognio spectrum analysis tool, internally. Many of our customers also use this tool. It performs all the functions necessary to a successful site survey and even allows you to capture the data for later playback, viewing and manipulation. There are many other wireless tools out there. We certainly suggest that you spend some time investigating the available options and choose the tool(s) that is right for your particular application and needs.



## Analyze Equipment on a Single Frequency

Once the spectral analysis is complete and operating frequencies have been chosen it's time start up the equipment and begin gathering information about link quality between nodes, on the selected frequencies. The recommended method, for doing this, involves starting up a single SkyGateway and ALL of the SkyExtender nodes in the mesh, analyzing performance and bringing down the SkyGateway. Bring up the next SkyGateway, analyze the performance, and repeat until all SkyGateways have been cycled.

During this phase of deployment, you're working to determine the best links and frequencies to use from SkyGateway to first-hop, first-hop to second-hop, second-hop to third-hop, etc. Once we've analyzed the best links in formation, we can start to assign preferred parents, to ensure those links are formed each time a node is brought online.

When you've completed reading this section you will understand:

- The process by which individual SkyGateways are used to test.
- How to identify and record link quality
- How to identify ideal first and second-hop candidates
- How to partition the network by SkyGateway
- How to identify areas that need attention

## What Is a Preferred Parent?

We mentioned the concept of a "preferred parent" in the introduction to this section. What, exactly, is a "preferred parent?" When a node begins its link formation process it will attempt to form a link with the node, in the mesh, that seems best at the time. This preferred link, from the child node standpoint, is chosen based on number of hops to the SkyGateway (route cost), RSSI and modulation rate. The best combination of all the factors combined becomes the preferred link for the child node.

In some instances the child node can choose the "wrong" parent. This "wrong" parent may cause self-interference with other nodes, when it transmits to the child, or the opposite may occur when the child transmits to the parent node. This situation may not directly impact communication between the child node and the parent node but it will impact the performance of other links, within the mesh.

To prevent this from occurring and to allow a high degree of fine tuning, within the mesh, SkyPilot has added the concept of the "Preferred Parent". When a child node is provisioned, either manually or through SkyEMS, an option to enter the "Preferred Parent" is offered. When the child begins its authorization process and downloads its configuration file, the Preferred Parent will be part of the configuration. If the child has formed this primary link with a node other than the preferred parent, the child will drop the current link and optimize the new link with its preferred parent – even if that parent is on a frequency different from what the child is currently on.



## **The Process for Determining Link Quality to the First-hop**

As mentioned, above, it is best to bring a single SkyGateway online at a time, during the initial link evaluation phase. This allows each SkyGateway to form links with potential first-hop clients, on the various frequencies chosen for operation. There are two methods by which this can be done.

The first step is to, power down all but a single SkyGateway at the central location. If there are five of them, four of them should be powered down during the initial testing phase of the process. This ensures that all potential first-hop links will connect to the correct SkyGateway.

If powering individual SkyGateways down requires too much time and effort, there is another method to ensure that devices form links to only a single SkyGateway at a time. This can be accomplished by configuring the domain and operating frequency, for all SkyExtenders and the SkyGateway currently in-use, to a value not used by the other SkyGateways.

### **How much of the Mesh needs to be Installed?**

At this particular stage, in the process, we're looking to evaluate only the potential first-hop and the potential second-hop SkyExtender links. The priority is on the first-hop SkyExtender links, as they determine the capacity of the entire mesh. It is highly recommended that all, or most, of the first-hop SkyExtenders are installed and ready to form links with the SkyGateways. A large percentage, but not all, of the SkyExtenders that are slated for the second-hop links should also be installed and ready for operation.

SkyExtenders and SkyGateways can be physically installed in any order, in the mesh. Once a SkyExtender is installed, and powered up, it will listen forever for HELLO messages from a potential parent link. Thus, while the work at the SkyGateway site is ongoing, installation of SkyExtenders can be worked in parallel, targeting the work on the first and second-hop SkyExtenders as a priority.

### **Once Installed Provision the Nodes**

Once the nodes are installed at their various locations they must be provisioned. Nodes cannot establish permanent links with each other until they are provisioned to do so. Each node requires the configuration of frequency, domain and other information. This includes the SkyGateway.

The provisioning of nodes, including the SkyGateway, can be accomplished in one of two ways. The first method by which this can be accomplished is via the Command Line Interface (CLI) of the individual nodes. This is called manual provisioning and is a labor intensive method of provisioning nodes. To use this method, the operator must access each, individual, node via a serial interface cable. Parameters for frequency, IP address, domain ID (used for the RF domain) and various other system functions will need to be individually configured. As changes need to be made, on the network, all affected nodes will require individual attention.

We suggest the provisioning of the individual nodes be performed through SkyEMS. When, and if, changes need to be made they can be accomplished quickly (i.e. when changing SkyGateways for the next step in the node analysis phase). As with the installation of the SkyExtenders at their various locations, the configuration of nodes, in SkyEMS, can be accomplished during any part of the initial phases of planning and physical installation.

Within SkyEMS a single node profile can be created, for this initial phase of testing, to accommodate the SkyExtenders. Another node profile will need to be created to accommodate the SkyGateways. Each SkyExtender and SkyGateway, however, requires a unique “Node Maintenance” entry. The “Node Maintenance” entry uniquely defines the individual node attributes (i.e. MAC Address), etc. and assigns the appropriate “Node Profile”.

Once each of the nodes is correctly provisioned, the next step, in the link analysis process is to ensure that all nodes in the mesh can establish, and maintain, appropriate links at the backhaul frequency (5.8GHz range). For this step, it is recommended that only the node profile, for the 5.8GHz be provisioned, on the SkyExtenders.

## Where to Get Information on Link Quality

Once the potential first-hop links are online, the first step in the deployment process is to analyze the links that form within the mesh. This analysis phase forms the basis for many of the decisions to be made in the further steps of this process including the need for more SkyExtenders, changing node placement/location, etc. This step, along with a few others, is crucial to the success of the deployment process, and many other steps along the way hinge on the success of this step.

### nodeInfo – A Wealth of Information

SkyPilot provides, on its support site, a script called nodeInfo. The nodeInfo script performs many functions including; polling each node for GPS coordinates and generating a Google Earth compatible KML file with each node located on the file; polling each node for link data, including local and remote antenna, local and remote RSSI values and local and remote modulation rates; polling each node for uptime and reboot reasons; allowing the user to perform a systematic reload of node profiles for all, or selected, nodes on the network; and various other functions. The nodeInfo script is not included as part of the SkyPilot EMS package but can be downloaded for free, and installed, at any time.

## Link Evaluation

The first step in the process is to **bring a single SkyGateway online** and at a minimum all potential first-hop SkyExtenders with it. Allow time for the links to form, and the mesh to settle out. If there are SkyConnectors in the field, do not start these up at this time. Once the mesh has settled, all SkyExtenders should have formed links with another node somewhere in the mesh. This initial stage may, or may not, be what you expected it to be, but we have a starting point from which to work.

An example of a network with all potential first-hop nodes installed is shown in the Google Earth diagram,



below. At the center of this network are three SkyGateways. In this example only a single SkyGateway, at a time, will be brought online to form links with all potential first-hop nodes in the mesh. All links will be evaluated at the first-chosen frequency. A new frequency will be selected, for the mesh, and all potential first-hop links will be evaluated at the new frequency. This process will be repeated until all chosen frequencies, in this case three of them, have been evaluated. The first SkyGateway will be decommissioned and the next SkyGateway brought online and the links evaluated in the same manner, as the first.

The link evaluation process is, we'll admit, a bit tedious. As mentioned above, we provide the `nodeInfo` script to help you along, but it still requires some data gathering and analysis. During this step the goal is to gather and record the data about link quality between all the nodes in the mesh. The `nodeInfo` script produces a file with the extension `*links.csv*` as part of the title. This file generates output that looks like the example below.

Current MAC	Current IP	Link MAC	Node Type	Status	LRSSI	RRSSI	LTxMod	RTxMod	LAnt	RAnt	Range	Connect Time	Downed Time	First Opt Time
00:0a:db:01:0a:3c	172.16.1.8	00:0a:db:01:0b:f0	gw	act path	32	33	18	18	0	6	1350	JUL 24 2006 09:22:13	JUL 24 2006 09:19:51	JUL 24 2006 09:20:49
00:0a:db:01:0a:3c	172.16.1.8	00:0a:db:01:19:89	ext	standby-o	19	23	18	24	6	5	6450	AUG 01 2006 16:11:05	AUG 01 2006 21:02:22	JUL 26 2006 13:26:42
00:0a:db:01:0a:3c	172.16.1.8	00:0a:db:01:19:ad	ext	standby-h	28	30	18	24	0	4	0	JUL 24 2006 09:20:11	JUL 24 2006 09:20:11	JUL 24 2006 09:20:11
00:0a:db:01:0a:3c	172.16.1.8	00:0a:db:01:19:c2	ext	standby-h	21	20	12	12	7	3	0	JUL 24 2006 10:55:04	JUL 24 2006 10:55:04	JUL 24 2006 10:55:04
00:0a:db:01:0a:3c	172.16.1.8	00:0a:db:01:0b:5b	cpe-o	standby-o	17	16	9	12	5	0	8100	JUL 27 2006 09:50:04	JUL 27 2006 09:50:31	JUL 26 2006 12:34:33
00:0a:db:01:0a:3c	172.16.1.8	00:0a:db:01:21:98	cpe-o	act mgmt	39	42	36	36	1	0	450	JUL 24 2006 09:24:26	JUL 24 2006 09:20:30	JUL 24 2006 09:23:10
00:0a:db:01:0a:3c	172.16.1.8	00:0a:db:01:22:da	cpe-o	standby-o	30	0	24	24	6	0	600	JUL 24 2006 09:20:51	JUL 24 2006 09:20:51	JUL 24 2006 09:22:20
00:0a:db:01:0a:3c	172.16.1.8	00:0a:db:01:30:b9	cpe-o	standby-o	22	0	6	24	6	0	900	JUL 24 2006 09:21:44	JUL 24 2006 09:21:44	JUL 24 2006 09:22:50
00:0a:db:01:0a:3c	172.16.1.8	00:0a:db:03:00:73	cpe-o	act mgmt	26	28	18	18	1	0	300	JUL 29 2006 09:51:54	JUL 29 2006 09:50:17	JUL 29 2006 09:50:21
00:0a:db:01:0a:3c	172.16.1.8	00:0a:db:01:20:e1	cpe-i	act mgmt	19	18	18	12	5	0	150	AUG 02 2006 07:13:56	AUG 02 2006 07:09:19	AUG 02 2006 07:10:08
00:0a:db:01:0a:3c	172.16.1.8	00:0a:db:01:20:ed	cpe-i	act mgmt	30	35	36	24	2	0	450	JUL 24 2006 09:23:46	JUL 24 2006 09:22:17	JUL 24 2006 09:22:18
00:0a:db:01:0a:3c	172.16.1.8	00:0a:db:01:21:01	cpe-i	act mgmt	35	33	36	24	7	0	150	JUL 24 2006 09:23:42	JUL 24 2006 09:22:16	JUL 24 2006 09:22:16
00:0a:db:01:0a:3c	172.16.1.8	00:0a:db:01:21:06	cpe-i	act mgmt	24	26	36	18	3	0	450	JUL 24 2006 09:24:40	JUL 24 2006 09:21:45	JUL 24 2006 09:23:13
00:0a:db:01:0a:3c	172.16.1.8	00:0a:db:01:21:43	cpe-i	act mgmt	36	42	36	36	4	0	300	JUL 24 2006 09:23:42	JUL 24 2006 09:20:39	JUL 24 2006 09:22:16
00:0a:db:01:0a:3c	172.16.1.8	00:0a:db:01:21:44	cpe-i	act mgmt	32	30	36	36	2	0	300	JUL 24 2006 09:23:36	JUL 24 2006 09:22:16	JUL 24 2006 09:22:16

As can be seen, this file includes the node identifier (both MAC and IP), the MAC address of the node it currently sees a link on, the type of node the link has been formed with (typically gw = SkyGateway, ext = SkyExtender, cpe-i – SkyConnector Indoor or cpe-o – SkyConnector Outdoor), the status of the link (act path = current active link to parent node, standby-o or standby-h signifying a standby-link that has either gone through optimization (o) or not (h), act mgmt - link formed with child node, act path – link formed with parent node), RSSI values (LRSSI – local or receive RSSI, and RRSSI – far-end RSSI value), modulation rates (LTxMod – local modulation used to transmit to far-end, and RTxMod – remote transmit modulation used to transmit to this node), antenna (LAnt – local antenna used for this link, RAnt – remote antenna used for this link), Range (distance between the two nodes in meters) and some miscellaneous fields about the link up/down events, in terms of time of day.

Out of all that stuff, what's important in this file? These are the fields that you need to pay attention to:

- Current MAC (identifies which node you're logging information for)
- Link MAC (identifies the far-end node on the link, also important)
- LRSSI (shows the receive signal strength for this node)
- RRSSI (shows the receive signal strength for the far-end node)
- LTxMod (shows the transmission modulation rate to the far-end node)
- RTxMod (shows the receive modulation rate for this node)
- LAnt (shows which antenna is being used on the local node)

- Range (shows the distance between two nodes in meters)
- RAnt (shows which antenna is being used on the far-end node)

Let's take an example, of potential first-hop links to a SkyGateway, and we'll apply some of the rules we've already learned, and a few that we're going to learn. Here is a partial \*links.csv\* file for our first run at the SkyGateway. This is theoretical data and may not reflect what you see in your network. It's only being used as an example. We have also removed some of the time fields, off the end of the file, as they don't contain any data necessary to perform this analysis.

Current MAC	Current IP	Link MAC	Node Type	Status	LRSSI	RRSSI	LTxMod	RTxMod	LAnt	RAnt	Range
00:0a:db:01:22:e7	172.16.1.4	00:0a:db:01:0b:f0	gw	act path	8	9	6	6	0	4	350
00:0a:db:01:11:ad	172.16.1.6	00:0a:db:01:0b:f0	gw	act path	34	36	6	12	4	0	450
00:0a:db:01:19:c2	172.16.1.7	00:0a:db:01:0b:f0	gw	act path	33	35	36	36	4	4	300
00:0a:db:01:0a:3c	172.16.1.8	00:0a:db:01:0b:f0	gw	act path	40	37	36	36	0	6	350
00:0a:db:01:11:a5	172.16.1.9	00:0a:db:01:0b:f0	gw	act path	30	32	36	48	3	0	350
00:0a:db:01:0b:5b	172.16.1.10	00:0a:db:01:0b:f0	gw	act path	26	28	9	12	0	6	450
00:0a:db:01:21:98	172.16.1.11	00:0a:db:01:0b:f0	gw	act path	23	25	12	18	0	6	400
00:0a:db:01:21:78	172.16.1.16	00:0a:db:01:0b:f0	gw	act path	20	22	6	24	0	0	450

As can be seen, we have several SkyExtenders that have formed links to the SkyGateway. All of them have formed primary links. One assumption that we're going to make is that all of these nodes are potential first-hop candidates for this SkyGateway.

What can be noticed, here? Link distances are approximately equidistant and range from 300 meters to approximately 450 meters. Our rule for first-hop equidistance has been satisfied, in this example. All nodes are showing the SkyGateway as the chosen active path, but modulation rates and RSSI values vary significantly node-to-node. Finally, most of the SkyExtenders, in this example have chosen one of two antennas on the SkyGateway, as their primary antenna. This will allow us to split the network between the various SkyGateways, when that time comes.

We're going to reproduce the diagram again so you don't have to page back and forth to continue with the example.

Current MAC	Current IP	Link MAC	Node Type	Status	LRSSI	RRSSI	LTxMod	RTxMod	LAnt	RAnt	Range
00:0a:db:01:22:e7	172.16.1.4	00:0a:db:01:0b:f0	gw	act path	8	9	6	6	0	4	350
00:0a:db:01:11:ad	172.16.1.6	00:0a:db:01:0b:f0	gw	act path	34	36	6	12	4	0	450
00:0a:db:01:19:c2	172.16.1.7	00:0a:db:01:0b:f0	gw	act path	33	35	36	36	4	4	300
00:0a:db:01:0a:3c	172.16.1.8	00:0a:db:01:0b:f0	gw	act path	40	37	36	36	0	6	350
00:0a:db:01:11:a5	172.16.1.9	00:0a:db:01:0b:f0	gw	act path	30	32	36	48	3	0	350
00:0a:db:01:0b:5b	172.16.1.10	00:0a:db:01:0b:f0	gw	act path	26	28	9	12	0	6	450
00:0a:db:01:21:98	172.16.1.11	00:0a:db:01:0b:f0	gw	act path	23	25	12	18	0	6	400
00:0a:db:01:21:78	172.16.1.16	00:0a:db:01:0b:f0	gw	act path	20	22	6	24	0	0	450

Now that we've locked our potential first-hop candidates in, what can we tell about the links to the SkyGateway? The top link, in this example, has both low modulation rates and low RSSI values. This node should not be considered as a candidate for a first-hop link, to this SkyGateway. If similar link stats show up on all SkyGateways, for this node, it's time to consider repositioning the node higher or lower on its mount point, or changing the mounting location altogether. A node in this close range to a potential parent link should host both high RSSI and modulation rate. Failure

to achieve both indicates path obstructions and/or the presence of interference, most likely multipath.

Three links come close to the link criteria established earlier, to be discussed in further detail in the information following, that we'd like to meet for first-hop links. They are highlighted in yellow. The best of the three, is the link 00:0a:db:01:0a:3c. This link hosts high modulation rates (36Mbps / 36Mbps) and high RSSI values (40 and 37). This link would be the best first-hop candidate for antenna 0. The first highlighted link displayed, 00:0a:db:01:19:c2, presents the best candidate for antenna 4. As antenna 4 is located 180 degrees from antenna 0, this also represents perfect spatial diversity, which makes both of these links ideal choice as a first-hop candidate.

Note that the SkyExtenders, in this example are either using antenna 0, 4 or 6. When we finally go to "break up" all the nodes into their final topology clusters, this becomes important. Let's say that we have four active SkyGateways serving the mesh. We have at least 3 first-hop candidates, in the same region, with links to this SkyGateway on different antennas, allowing for spatial diversity and the ability to interleave nodes for greater depth of coverage. This provides us with three distinct frequency first-hop nodes, allowing for easy division of the network into three different meshes, governed by three different SkyGateways.

Repeat this exercise with each frequency chosen. The reason for repeating this exercise for each chosen frequency, on each SkyGateway, is to determine the best operating frequency for that SkyGateway. As a secondary goal you're attempting to determine the best operating frequency for each link to the SkyGateways. You will find that some links will do better at different frequencies, than others. Store the data for each link in a format that can be easily read and interpreted, which will be in the next step – choosing links.

We suggest highlighting, or marking, the best link(s) to this SkyGateway and crossing them off the list of potential first-hop candidates for the other SkyGateways in the mesh. That way the only candidates requiring evaluation for the supporting SkyGateways are those that remain. This will make the process of link evaluation and selection much simpler.

Once the process has been completed for the current SkyGateway, shut it down and bring up the next SkyGateway. Go through each frequency, noting the link conditions for the first and second-hop candidates, and save them for later use. Cycle through all frequencies and repeat for each SkyGateway in the mesh.

### **First and Second-Hop Candidates**

Now that you've gathered and organized your data for the various links, it's time to define the process of selecting the first and second-hop candidates for the mesh. These first and second-hop links will have threshold RSSI values, and modulation rates as a first priority and spatial diversity as a close second priority. Finally, these links should be relatively equidistant from the SkyGateway but that should not take precedence over link quality (RSSI and modulation rate) or spatial diversity. It is more important to select high quality links than it is to select perfect spacing, when faced with that choice. Our design goal is resilience to interference from any source and mitigation of self-interference.




## What Qualifies as a First-Hop Link?

First-hop links are the most important links on your network. Can we stress that some more? These links determine the system capacity and, as such, need to have the highest RSSI and modulation rates of all the links in the mesh.


The **target RSSI for first-hop SkyExtenders should be a minimum of 35** to achieve a **target modulation rate of 36Mbps, at a minimum.**

In the example, shown below, there are only a few candidates for first-hop links that meet these criteria. One of the chosen links has greater range/distance from the SkyGateway, but is well within the bounds of meeting the criteria. Other links closer to SkyGateway do not meet the first-hop link criteria.

Current MAC	Current IP	Link MAC	Node Type	Status	LRSSI	RRSSI	LTxMod	RTxMod	LAnt	RAnt	Range
00:0a:db:01:22:e7	172.16.1.4	00:0a:db:01:0b:f0	gw	act path	35	34	36	36	0	4	350
00:0a:db:01:11:ad	172.16.1.6	00:0a:db:01:0b:f0	gw	act path	27	28	36	36	4	0	475
00:0a:db:01:19:c2	172.16.1.7	00:0a:db:01:0b:f0	gw	act path	37	34	36	36	4	4	500
00:0a:db:01:0a:3c	172.16.1.8	00:0a:db:01:0b:f0	gw	act path	32	32	24	24	0	6	450
00:0a:db:01:11:a5	172.16.1.9	00:0a:db:01:0b:f0	gw	act path	23	25	36	36	3	0	900
00:0a:db:01:0b:5b	172.16.1.10	00:0a:db:01:0b:f0	gw	act path	26	28	9	12	0	6	600



Notice that both of the potential first-hop links, in this example, are connected to the SkyGateway on the same antenna. When this situation occurs review these links on other SkyGateways. Choose the best alternative SkyGateway as a parent for one of these links so that it will serve as a first-hop link for different overlay mesh in the network.



It is also important, at this stage of the evaluation process, to note the direction of each of the chosen first-hop link candidates, in terms of North, South, East and West of the SkyGateway. **The end-**

**goal, of this process, is to not only choose first-hop links for each SkyGateway that are the best in terms of modulation rate and RSSI, but also to choose first-hop links that provide spatial diversity for each SkyGateway.** When the process is complete, the SkyGateways should be transmitting in each of the eight potential directions.

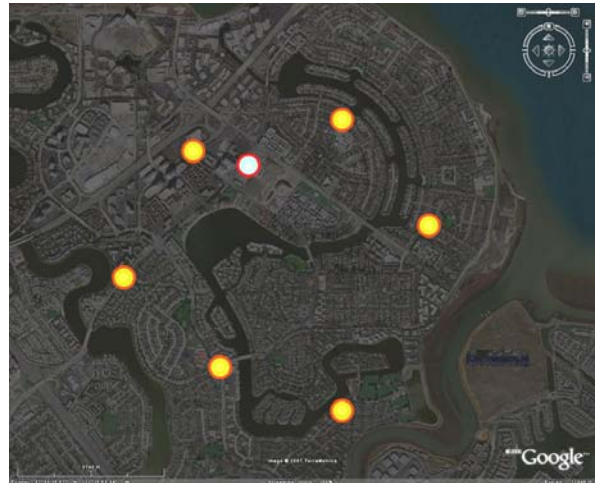
Depending on coverage density and customer speed requirements, a given direction may be covered by several different SkyGateways to ensure adequate throughput for all users accessing the network. As an example, SkyGateway 1 might cover nodes close in to the North, South, East and West. SkyGateway 2 might cover nodes that are further out, to the North, South, East and West. Both SkyGateways will be configured to different frequencies to ensure no interference.

The example, below, displays the end result of the link analysis exercise, in a real network. The SkyExtenders, shown on this page in yellow, are relatively equidistant from the SkyGateway (shown in blue), but not perfectly. You'll note that the two SkyExtenders above the SkyGateway are much closer to the SkyGateway than those below it are.

You can also see, from this diagram, that good spatial diversity was maintained as well as the physical geography and the mesh topology allowed it. The SkyExtenders are located in 45, or 90, degree increments from each other to reduce, as much as possible, the amount of self-interference in the mesh.

Once all of the first-hop candidates have been chosen and are ready to link to their respective SkyGateway, edit the individual node profiles, for the SkyExtenders, and set the primary frequency and preferred parent options to the preferred SkyGateway.

Configure the individual node profiles for the SkyGateways for their primary operational frequency and other parameters, too. Once you have completed the configuration of all SkyGateways and first-hop SkyExtenders it's time to let the nodes come up. Reboot all the SkyGateways on the network and wait for the first-hop nodes to form links to their preferred parent.



Once this process has completed, re-run the nodeInfo script, making sure that you at least get the KML (Google Earth) output. What we're looking for, in this first pass, is to ensure that we have achieved adequate spatial diversity with all SkyGateways in the mesh, that the links still meet the minimum requirements for first-hop links and that we have adequate coverage in the first-hop, for the network. An example of a network needing more coverage is shown, above, and a network that provides adequate coverage is shown, left.



## A Final Note About An Important Aspect of Link Quality

There is one final note about link quality that was not brought up, in the material above. The SkyPilot nodes, as you may be aware, utilize GPS pulses to synchronize data transmissions. If a SkyGateway, or SkyExtender, cannot receive the GPS signal adequately it will not be able to accurately time its transmission cycles. This causes data to be lost, on the network, causing retransmissions, latency and packet loss.

To determine if a given node is able to consistently receive the GPS sync pulses, log in to the node in question. Perform the command:

```
show int rf
```

You will be prompted with the following output:

```
> show int rf
```

```
radio:
```

```
MAC Address: 00:0a:db:01:0e:ad
Default Gateway: 00:0a:db:01:0e:ad
9023 PPS received
4808660 frames received
5207559 frames sent
654465 CRC errors
313451 decrypt CRC errors
0 Michael errors
0 RX - wrong antenna
0 RX - too small
0 RX - wrong destination
0 RX - no link
0 RX - drops
0 RX - mblk events
0 TX - errors
0 TX - timeouts
0 memory events
0 descriptor events
0 netJob events
2 tau sync events
36 PPS misses
0 PPS timeouts
36400 discards
0 late responses
0 no frame events
0 busy events
```

Compare the PPS misses and PPS timeouts to the PPS receive value. The percentage of missed plus receive should not be significant when compared to the number of received PPS. Occasional misses are acceptable and will not impact link performance. However, a significant number of misses, on a consistent basis will impact link performance.

## Choosing Second-Hop Links

Second-hop SkyExtender-to-SkyExtender links similar rules as the first-hop SkyExtender to SkyGateway links with the exception of placement with respect to their parent nodes. These rules were covered in a previous section in some detail, but will be reviewed here to illustrate the



process of picking good second-hop candidates for the mesh. The end-goal, in this pass, is to select the best frequency, link and parent node for the second-hop SkyExtenders.

Current MAC	Current IP	Link MAC	Node Type	Status	LRSSI	RRSSI	LTxMod	RTxMod	LAnt	RAnt	Range
00:0a:db:03:02:7f	172.16.3.1	00:0a:db:01:68:9f	ext-d	act mgmt	36	34	36	36	1	1	350
00:0a:db:03:0b:3d	172.16.3.10	00:0a:db:01:68:9f	ext-d	act path	35	36	36	36	0	0	450
00:0a:db:03:0b:62	172.16.3.12	00:0a:db:01:68:9f	ext-d	act path	34	41	36	48	0	3	300
00:0a:db:03:06:7f	172.16.3.44	00:0a:db:01:68:9f	ext-d	act path	41	46	36	48	0	6	300
00:0a:db:03:07:cb	172.16.3.55	00:0a:db:01:68:9f	ext-d	act path	33	40	36	54	0	0	300

In the example, above, there are three potential link candidates for second-hops on the node with the MAC address of 00:0a:db:01:68:9f. We'll start by drawing attention to the red box on the right of the example. It will be noticed that all three of these candidates use the same antenna, on this link. Thus we should choose only one of these as a second-hop SkyExtender for this node. The other two should be preferred parented to other first-hop nodes. I .

As with the first-hop SkyExtender links, you're looking for a **target RSSI of 35** and a **target modulation rate of 36Mbps**. Values below this criterion should not be considered as good choices for second-hop links, unless you have no other choice.

Comparing the link modulation rates, and RSSI values, for this example the bottom link (MAC 00:0a:db:03:06:7f) is the best choice. The RSSI values are well within range, in fact 10dB higher than the target value and the remote modulation rate is 48Mbps, better than the 36Mbps target value. Second and third choices extend above this link and are marked in progressively lighter colors.

Some other rules specific to second-hop links that need to be reviewed, in brief:

- Do not link a child node on the same antenna used to link to the parent node.
- Second-hop links should not be formed back toward the SkyGateway.
- Second-hop links should be located at 45 or 180 degree angles

Once the appropriate second-hop links have been chosen for the various mesh overlays, edit the node-profiles for each of the second-hop links and configure the preferred parent and primary frequency fields to ensure the appropriate links are formed.

### Looking for Areas Needing Attention

As a final pass at the network, carefully review everything that's been done, to this point. By now, you should have all of the SkyGateways assigned to different frequencies and all first-hop nodes should have preferred SkyGateway assignments, and all second-hop candidates should have preferred parents configured based on link quality.

Your next step is to review the configuration and highlight any directions with poor first-hop links for fill-in to avoid over-burdening SkyExtenders. As first-hop links were assigned, and chosen, there may have been a few nodes with either low RSSI or low modulation rate, or both. These nodes, as a result of their susceptibility to interference (low RSSI value) or inability to move data quickly (low modulation rate) will be burdened more than their first-hop peers. If repositioning the node did not resolve these issues, then it's time to consider placing an intermediary node

between the current first-hop and the SkyGateway that will improve both the RSSI and the modulation rate. The current first-hop node can then serve as a second-hop, improving both modulation rate and RSSI on the link.

### **Why the High First and Second-Hop RSSI and Modulation Rates?**

The question is posed many times, why do we suggest such high modulation rates and RSSI values on the first and second-hop links in the mesh. We have covered, in some detail, the reason for the high modulation rate. The modulation, on the first-hop links, determines the capacity of the mesh. Low values here translate into low throughput for every user on the mesh.

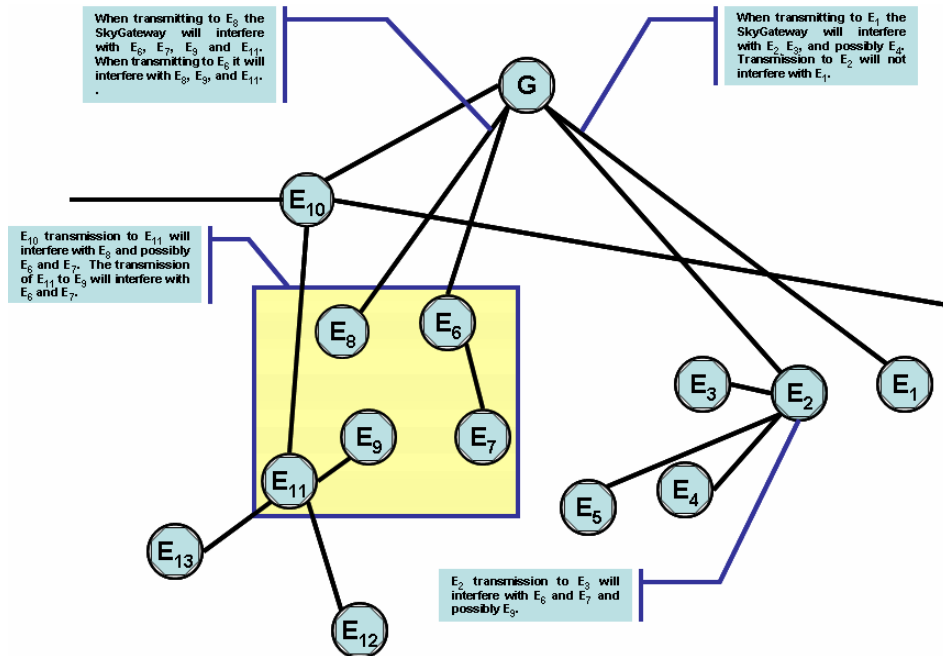
Let's now take a look at the reason for the high RSSI values on these critical links in the mesh. RSSI stands for Received Signal Strength Indicator. It represents the strength of the signal being received by a node. The higher this value, the "louder" the signal appears to the receiving node. If interference does show up, in the channel, the signal is so much louder than the interference, with a high RSSI, that the interference will not cause issues. On the reverse of that, if the RSSI value is very low, it takes very little interference to wreak havoc with the radio reception and cause all received data to be lost.

As your first and second-hop links are the critical links within the mesh, the focus on maintaining high RSSI values and high modulation rates is a priority. This improves mesh resilience and end-user perception and performance.

## What If I Cannot Achieve Spatial Diversity In Some Areas of My Network?


In some areas of the network, it's going to be nearly impossible to provide good spatial diversity. This is, in most instances, due to a lack of resources upon which a SkyGateway, or SkyExtender, can be mounted. Not all areas, within a given municipality, provide for perfect installations.

This next diagram displays a section of an already deployed network where the location of the SkyGateway did not allow spatial diversity to be maintained. As can be seen from this diagram, all of the SkyExtenders are located on a single side of the SkyGateway. With the close proximity of the SkyGateway to the SkyExtenders, in this mesh, self-interference will be strong in all areas of the network.



Notice that many of the second and third-hop SkyExtenders are located in the vicinity of other first-hop nodes, further increasing the chances of self-interference. Also note that child nodes, in this example are not pointing back, and away from the parent node, as presented above. The potential areas of self-interference, and this isn't all of them, are highlighted, on the diagram, above.

To properly deal with this type of situation it is better to “serialize” all SkyExtenders to the SkyGateway;



in other words, assign all the SkyExtenders to the SkyGateway and eliminate the second-hop nodes. If this presents too much of a load to the SkyGateway (i.e. you cannot serve up enough bandwidth to individual subscribers in this region) then the only choice, to keep self-interference at an absolute minimum, would be to add another SkyGateway (which is what this operator did) to the location and divide the SkyExtenders between the two SkyGateways.

As can be seen from the example, to the left, two frequencies are used to segment the SkyPilot network. These frequencies are interleaved to produce maximum separation and reduce interference even further. In the instance where there are two child nodes ( $E_2$  and  $E_4$ ) transmission to either node will not cause self-interference with any other node.

There still is, in this diagram, some potential for self interference when the SkyGateway transmits to  $E_6$ . It is possible that this transmission may interfere with data transmission between  $E_4$ , or  $E_2$ , and  $E_3$ . The node performance would have to be watched, over time. If it's significantly degraded, then  $E_6$  could be assigned to the other SkyGateway to prevent further self-interference.

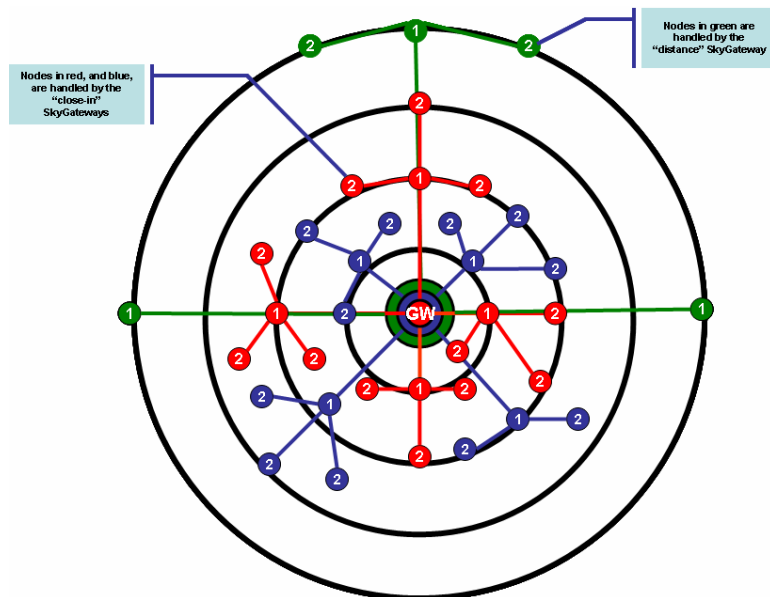
## How Many Hops?

Another question that often surfaces is: "How many hops from the SkyGateway should I consider?" While we have had some customers test our products and successfully transfer files with 12-13 hops away from the SkyGateway this is not a recommended network configuration.

Each hop, added to the SkyPilot mesh, adds additional latency, approximately 10mS per hop. As latency increases, if nothing else at the client side, or the server side, changes throughput decreases. Data throughput is, primarily, a function of TCP Window Size and Latency. Of course this is assuming that there is no data loss in the transmission path, no network bottlenecks between client and server, etc. Thus adding nodes will effectively reduce TCP throughput unless someone changes registry, or kernel, parameters for TCP window size, etc.

In a typical mesh deployment, SkyPilot suggests that mesh networks go no more than four hops from the SkyGateway. If additional coverage, beyond this distance, is required, then a near/far SkyGateway scenario can be enacted, as per the following example.

**NOTE:** Additional SkyGateways and SkyExtenders have been removed from this diagram to highlight the concept of a SkyGateway dedicated to the service of SkyExtenders that are further out and another to SkyExtenders that are much closer to the network center.



---

**NOW THAT YOU'RE DONE... REBOOT THE SKYGATEWAYS! LET THE MESH  
CONFIGURE ITSELF!**

## Distance versus Modulation Rate

The following table, derived from actual field data, presents modulation rate versus average RSSI versus distance. These measurements were taken in suburban areas with light to medium foliage. The actual results achieved will vary from region to region, but these values should be representative of what can be achieved in many municipal wireless deployments.

If there are no trees, in the area, you will typically have the ability to maintain much higher modulation rates. A large number of trees, or high levels of RF interference will tend to yield lower modulation rates than presented.

		Local RSSI	Remote RSSI	Link Distance
Modulation Rate	54	44	37	386
	48	40	43	369
	36	34	37	429
	24	25	25	475
	18	23	25	548
	12	23	21	566
	9	20	20	725
	6	17	20	813

Modulation Rate versus RSSI versus Distance - Comparison Table

## Device Links Drive Spacing

As can be seen from the above chart, modulation rates are a function of distance and, in the end, greater distances between devices have a large impact on the speed that any given subscriber will be able to achieve when utilizing the network. In that light, end-user throughput requirements will drive modulation rate requirements which, in turn, will drive distance requirements, between nodes in the mesh.

This same rule of thumb applies to the 2.4 side of the network. The backhaul distances and modulation rates will determine how much bandwidth will be available to any 2.4MHz user connected to the SkyExtender Access Point. Conversely, the distance between the user and the SkyExtender Access Point will further determine the end-user experience, beyond the capabilities of the backhaul.

If the 2.4MHz access points are spaced in such a way as to accommodate high powered Wi-Fi CPE's, customers using laptop computers to access the network will suffer from speed and distance limitations. It is important to pay close attention to end-user to be serviced, within a given region, and drive node spacing requirements from the 2.4GHz coverage density required and secondarily from the modulation rates required to provide end-user throughput at acceptable levels for that region. The weakest link in the chain (whether it's 2.4GHz access or 5.8GHz backhaul) will drive node spacing requirements. Typically, that link is going to be laptop wireless network card power, and not the backhaul spacing, although in some instances, the backhaul modulation rate requirements may drive the spacing.



## Review the Topology for Self-Interference Candidates

We have a working mesh, at this point. The best first and second-hop candidates have been identified and preferred parented to the appropriate nodes. Frequencies have been chosen for best link operation, and have been used to partition the network out by using mesh overlays. Now it's time to review the network for potential self-interference candidates.

Self-interference is generated whenever an actively transmitting node overlaps with an actively receiving node, neither of which are communicating with each other. When this condition occurs, the end-result is lowered modulation rates despite seemingly high RSSI values.

The nodeInfo script, discussed above, can also be used to generate self-interference output. The output is generated in two different formats – a CSV file for manipulation and viewing in Excel, or other application that supports CSV and; a KML file that can be loaded into Google Earth. The KML file not only displays the self-interference candidates, but it also displays each node of the mesh in its appropriate location on the Google Earth map for your location, making the identification of self-interference a more visually pleasing experience.

When you've completed this section you will understand:

- What self-interference is
- What tools can be used to identify self-interference
- How to determine how bad self-interference is
- What to do to avoid self-interference

## A Review of Self-Interference

The previous section on self-interference covered the how's and the why's of self-interference in some detail. We'll review the pertinent points, here, and then move quickly into the identification and elimination of self-interference in the mesh.

Self-interference occurs whenever a transmitting node has a coverage area which overlaps with one, or more, receiving nodes, that are not direct children, or a parent of the transmitting node. If these nodes are actively receiving data during the same period in which the transmitting, or interfering, node is active, self-interference is the result. Depending on the RSSI value of the self-interference, the received data could be dropped, link modulation rates lowered, etc.

Self-interference is a probabilistic function, meaning that self-interference is usually not an issue for inactive nodes. The more active the interferer node is the greater the probability is that it will interfere with another node.

By this definition it's easy to see that the SkyGateway, being the busiest node in the mesh, can be the greatest candidate for generating self-interference. The next-in-line candidates for the generation of self-interference are the first-hop SkyExtenders, then second-hop SkyExtenders, etc. These are also the most "susceptible" devices in terms of total impact to mesh performance; interfering with SkyGateway reception interferes with everyone on the network. Interfering with a sixth-hop SkyExtender interferes with only a small body of users on the mesh, as compared to all potential users within a mesh.



## Identifying Self-Interference

We know that it's important to eliminate self-interference, to the greatest extent possible; how can it be identified? The nodeInfo script, provided from the support section of [www.skypilot.com](http://www.skypilot.com) provides us with some tools to assist in the identification of self-interference on the network. When the self-interference option is run, two files can be generated that provide self-interference information: a KML file and a CSV file.

As mentioned above, the KML file is Google Earth compatible and can be loaded directly into Google Earth. This shows the position of all mesh nodes on the Google Earth map. Clicking on any of the nodes will bring up a page that displays node data (i.e. MAC Address, IP Address, etc) and link data (all active and standby links) and self-interference data (which links are experiencing self-interference). An example of what this looks like can be viewed on the following page.

The CSV file presents similar data, but in a text-based format that can be loaded into Microsoft Excel, or similar program, to be manipulated and viewed based on various sorting criteria. Displayed, below, is an examples of the CSV file. We will discuss, in some detail, each of the displayed data and their meaning. The next page will display both output formats and define the data presented in each.

## Self-Interference Output - Descriptions

Listening node	Mac	IP	Node Type	Transmitting Node	Tx Mac	Tx IP	Tx Node Type	Link State	LAnt	LRSSI	RTxMod	Interference	RSSI	SNR	Interferer	Interferer MAC	Interferer IP	Interferer Node Type	Interferer Link	Interferer Link MAC	Interferer Link IP	Interferer Link Node Type	Interferer Link State	Interferer LAnt
skypilot_mgw	00:0a:db:01:0e:ad	10.0.0.20	gw	CUSTOMER1_CPE	00:0a:db:01:32:0f	10.0.0.96	cpe-o	act mgmt	6	28	36	20	8	skypilot_ext3	00:0a:db:01:19:e4	10.0.0.23	ext	CUSTOMER1_CPE	00:0a:db:01:0a:82	10.0.0.55	cpe-o	act mgmt	3	
skypilot_mgw	00:0a:db:01:0e:ad	10.0.0.20	gw	CUSTOMER1_CPE	00:0a:db:01:1c:36	10.0.0.97	cpe-o	act mgmt	7	21	24	19	2	CUSTOMER1_CPE	00:0a:db:03:16:69	10.0.0.195	cpe-o	skypilot_ext5	00:0a:db:01:19:a3	10.0.0.25	ext	act path	0	
skypilot_mgw	00:0a:db:01:0e:ad	10.0.0.20	gw	CUSTOMER1_CPE	00:0a:db:01:1c:36	10.0.0.97	cpe-o	act mgmt	7	21	24	12	9	skypilot_ext3	00:0a:db:01:19:e4	10.0.0.23	ext	CUSTOMER1_CPE	00:0a:db:01:0a:82	10.0.0.55	cpe-o	act mgmt	3	
skypilot_mgw	00:0a:db:01:0e:ad	10.0.0.20	gw	skypilot_ext1	00:0a:db:01:0a:36	10.0.0.21	ext	act mgmt	5	22	36	36	13	skypilot_ext3	00:0a:db:01:19:e4	10.0.0.23	ext	CUSTOMER1_CPE	00:0a:db:01:0a:82	10.0.0.55	cpe-o	act mgmt	3	
skypilot_mgw	00:0a:db:01:0e:ad	10.0.0.20	gw	skypilot_ext2	00:0a:db:01:0a:4a	10.0.0.22	ext	act mgmt	5	25	6	25	0	skypilot_ext1	00:0a:db:01:0a:36	10.0.0.21	ext	skypilot_ext10	00:0a:db:03:03:cc	10.0.0.30	ext	act mgmt	0	
skypilot_mgw	00:0a:db:01:0e:ad	10.0.0.20	gw	skypilot_ext2	00:0a:db:01:0a:4a	10.0.0.22	ext	act mgmt	5	25	6	25	0	skypilot_ext1	00:0a:db:01:0a:36	10.0.0.21	ext	CUSTOMER1_CPE	00:0a:db:01:17:72	10.0.0.70	cpe-o	act mgmt	0	
skypilot_mgw	00:0a:db:01:0e:ad	10.0.0.20	gw	skypilot_ext2	00:0a:db:01:0a:4a	10.0.0.22	ext	act mgmt	5	25	6	25	0	skypilot_ext1	00:0a:db:01:0a:36	10.0.0.21	ext	skypilot_ext3	00:0a:db:01:19:e4	10.0.0.23	ext	act mgmt	0	

### Listening Node

The column labeled as Listening node in the CSV file is found, in the Google Earth output as the link with the red lettering indented and immediately following it. The Listening node is the node that is being interfered with. The one suffering from interference. Data such as the MAC Address, IP Address and node type are presented, to assist in the identification of the node.

### Transmitting Node

The transmitting node communicates with the listening node when the potential for self-interference manifests itself, and only works in this direction. MAC Address, IP Address and node type are presented.

### Link Stats – LAnt, LRSSI, RTxMod, Interference RSSI

The data presented in these four columns displays the impacted antenna and link quality indicators (RSSI, modulation rate). The Interference RSSI value displays the interference power, for this link. If this value is greater than the actual link RSSI, the interference is stronger/louder than the actual signal.

### SNR

The value, in this column is derived by subtracting link RSSI from Interference RSSI. If the value is a negative number, or 0 it indicates that the interference is stronger, in power, than the actual signal. If the value is a positive number, the signal is stronger than the interference. The value, in this column, should be greater than 13 in order for a 36Mbps link to remain stable in the face of interference. A value lower than 13 indicates that the self-interference be strong is strong enough to compromise the link quality.

### Interferer

The interferer is the node that has the potential to cause the self-interference when the transmitting and listening are engaged in communication. This node will be the one causing the actual self-interference, if and when it occurs. MAC Address, IP Address and Node Type are presented for further identification.

### Interferer Link

This data identifies the node that the interferer is transmitting to when the self-interference is generated with the Listening Node. MAC Address, IP Address and Node Type are presented for further identification.

```

skypilot_ext5
Telnet - Web Interface
IP : 10.0.0.25
MAC Address : 00:0a:db:01:13:a3
Active Image Name : SkyExt.1.3.bin
Active Image State : Accepted
Backup Image Name : SkyExt.1.2p3.bin
Backup Image State : Accepted
Active Image Bank : B
Node type : SkyExtender
Switchboard Version: D
BootROM : 001.06.000.3
Uptime : 4 days 20:10:39
Reboot Reason : Registration timeout

MAC Address      Hostname      IP Address      MType State      LRSSI RSSI  LTMd  RTMod  LAnt  RAnt  Range
-----
00:0a:db:02:00:4e      cpe-o act mgmt      26  26  36  36  2  0  3900
-----
Tx Interferer  Hostname      Rx Interferer  Hostname      SNR
-----
00:0a:db:01:1f:99      00:0a:db:01:22:54 CUSTOMER1_CPE  1
00:0a:db:01:1f:99      00:0a:db:01:15:0b CUSTOMER1_CPE  1

00:0a:db:01:17:7a CUSTOMER1_CPE  10.0.0.88      cpe-o act mgmt      13  14  18  6  7  0  9800
-----
Tx Interferer  Hostname      Rx Interferer  Hostname      SNR
-----
00:0a:db:01:06:f4 skypilot_ext7      00:0a:db:03:07:21 CUSTOMER1_CPE  0
00:0a:db:01:0a:4a skypilot_ext2      00:0a:db:01:15:07 CUSTOMER1_CPE  -6
00:0a:db:01:0a:4a skypilot_ext2      00:0a:db:01:15:4e CUSTOMER1_CPE  -6
00:0a:db:01:0a:4a skypilot_ext2      00:0a:db:03:09:cf CUSTOMER1_CPE  0
00:0a:db:01:0a:4a skypilot_ext2      00:0a:db:03:07:27 CUSTOMER1_CPE  -6
00:0a:db:01:0a:4a skypilot_ext2      00:0a:db:03:16:b2 CUSTOMER1_CPE  -6
00:0a:db:01:0a:4a skypilot_ext2      00:0a:db:03:09:e0 CUSTOMER1_CPE  0
00:0a:db:01:0a:4a skypilot_ext2      00:0a:db:03:0e:97 CUSTOMER1_CPE  0
00:0a:db:01:0a:4a skypilot_ext2      00:0a:db:03:15:28 CUSTOMER1_CPE  -6
00:0a:db:01:0a:4a skypilot_ext2      00:0a:db:01:15:04 CUSTOMER1_CPE  -6
00:0a:db:01:0a:4a skypilot_ext2      00:0a:db:03:17:09 CUSTOMER1_CPE  -6
00:0a:db:01:0a:4a skypilot_ext2      00:0a:db:03:04:1b CUSTOMER1_CPE  -6
00:0a:db:01:0a:4a skypilot_ext2      00:0a:db:01:15:4e CUSTOMER1_CPE  -6
00:0a:db:01:0a:4a skypilot_ext2      00:0a:db:03:16:f9 CUSTOMER1_CPE  0
00:0a:db:01:0a:4a skypilot_ext2      00:0a:db:01:0c:44 CUSTOMER1_CPE  7
00:0a:db:01:0a:4a skypilot_ext2      00:0a:db:01:09:f9 CUSTOMER1_CPE  -6
00:0a:db:01:0a:4a skypilot_ext2      00:0a:db:03:04:23 CUSTOMER1_CPE  -6
00:0a:db:01:0a:36 skypilot_ext1      00:0a:db:01:1d:6b skypilot_ext7  7
00:0a:db:01:0a:36 skypilot_ext1      00:0a:db:03:0e:cf CUSTOMER1_CPE  -4
00:0a:db:01:0a:36 skypilot_ext1      00:0a:db:01:15:58 CUSTOMER1_CPE  -4
00:0a:db:01:0a:36 skypilot_ext1      00:0a:db:01:0a:4a skypilot_ext2  -4
00:0a:db:01:0a:36 skypilot_ext1      00:0a:db:03:0d:1e CUSTOMER1_CPE  -4
00:0a:db:01:0a:36 skypilot_ext1      00:0a:db:03:08:e3 CUSTOMER1_CPE  7
00:0a:db:01:0a:36 skypilot_ext1      00:0a:db:03:0a:36 CUSTOMER1_CPE  -4
00:0a:db:01:1f:99 skypilot_ext1      00:0a:db:01:32:16 CUSTOMER1_CPE  -4
00:0a:db:01:22:13 CUSTOMER1_CPE      00:0a:db:01:0a:4a skypilot_ext2  5

00:0a:db:03:0c:c0 CUSTOMER1_CPE  10.0.0.146     cpe-o auth fail      7  23  23  6  6  2  0  10500
-----
Tx Interferer  Hostname      Rx Interferer  Hostname      SNR
-----
00:0a:db:01:08:b8 skypilot_ext7      00:0a:db:03:0d:97 CUSTOMER1_CPE  -4
00:0a:db:01:08:b8 skypilot_ext7      00:0a:db:01:15:63 CUSTOMER1_CPE  -10
00:0a:db:01:08:b8 skypilot_ext7      00:0a:db:03:09:46 CUSTOMER1_CPE  -4
00:0a:db:01:08:b8 skypilot_ext7      00:0a:db:03:17:ab CUSTOMER1_CPE  -4
00:0a:db:01:08:b8 skypilot_ext7      00:0a:db:03:15:66  1
00:0a:db:01:1f:99 skypilot_ext1      00:0a:db:01:22:54 CUSTOMER1_CPE  -6
00:0a:db:01:1f:99 skypilot_ext1      00:0a:db:01:15:0b CUSTOMER1_CPE  -6

```

## Sorting Through the Self-Interference File

Start by sorting the data presented by “Listening node” and then by the “SNR” column, in ascending order. This will present an ordered list of nodes experiencing interference, ranked by worst case interferers for each node. Compare the XLS file output with the Google Earth KML file output and determine the directionality of the interferer with respect to the impacted link(s).

## Avoiding Self-Interference

When confronted with the possibility of self-interference you should be able to go back to the design rules (i.e. child nodes should be pointed back and away from the parent node, and avoid forming child links on the antenna used to link to a parent node) and determine if the interferer meets this criteria, or not. If it does not, re-evaluate the link of the interferer, or the node impacted by the interference to determine if there is a better link to be chosen or formed that will reduce the impact of self-interference. Repeat the process for each link experiencing significant self-interference. Re-route for better antenna diversity or change frequencies for appropriate nodes.

Once links have been reworked and re-established, re-run the self-interference output and determine if the changes implemented have reduced the amount of, and/or number of instances of self-interference in the network. Repeat the process until the network has as much spatial diversity and as little self-interference as can be accomplished. The more that networks can be interleaved, i.e. having more SkyGateways at the central location servicing nodes, the greater the spacing between individual SkyExtenders is, on a given frequency, and the less chance there will be for self-interference, or the self-interference will have much less of an impact due to distance between similar frequency SkyExtender nodes.

## Fill In Remaining Coverage Areas with Additional Nodes

Now that the core of the mesh is as functional as you can make it, it's time to fill in the gaps, which will typically show up on the edges of the mesh. There are a few ways to do this including running diagnostics using nodeInfo and driving to various locations and testing for coverage. The latter, of course, being the best test of coverage.

When you have completed reading this section you will have an understanding of:

- Identifying areas requiring additional coverage
- Adding new nodes for additional coverage
- Analysis of the additions

## Identifying Areas Requiring Additional Coverage

In order to start identifying the areas that might be lacking coverage, we suggest you start by producing the KML file output to review the coverage area and help identify any areas that need



additional nodes. It is always best to verify coverage, in a questionable area, by performing a site survey, and actively testing that area. This ensures that adequate, or inadequate, coverage is available in that area of the mesh.

As can be seen in the photo to the left, the SkyGateway being analyzed is “lopsided” in terms of the amount of SkyExtenders being served by the unit. One sector, on this unit is covering 11 SkyExtenders and the other sector is covering two.

Throughput on the 11 SkyExtenders will be degraded as the SkyPilot bandwidth scheduler is a round-robin type of scheduler. Under high load conditions these 11 SkyExtenders, in this example, will get the same number of transmit/receive opportunities as the leg with only two SkyExtenders. Thus the link with two SkyExtenders will achieve better performance, and throughput during periods of active use.

During periods of light use, the SkyGateway will, more than likely, be spending much of its time transmitting in the direction of the 11 SkyExtenders. It may interfere with transmissions between the second-hop and third-hop SkyExtenders,



during some of its transmission intervals. The amount of self-interference it generates will be proportional to the amount of traffic in these various sections of the network.

If this were the actual case that we're trying to make the point about, the following actions could be taken to rectify the situation:

- Create more spatial diversity around the SkyGateway, if possible, by assigning SkyExtenders in the 4 quadrants. This will allow the SkyGateway to balance the load more equally across all child nodes in the mesh. It will also reduce the potential for self-interference as the SkyGateway will be spending more time transmitting in directions other than the
- More evenly space the load on the SkyGateway so that the load is distributed in multiple directions, not just one.

The diagram, to the right, shows the same SkyGateway with additional SkyExtenders added to balance the load and provide better spatial diversity. This map is taken from a live network so, as you can see, the “balance” in all directions is not perfect, but it is good and it achieves the end result that we’re looking for.



## Adding New Nodes

After the areas requiring additional coverage have been identified, add the additional nodes to the mesh to fill the gaps in coverage. During the initial phase configure the new nodes with the target primary frequency, but permit all other frequencies on the mesh as “allowed frequencies”. This will allow the SkyPilot protocol to make the initial topology choices for us. There may be interference, multipath, or other issues at the new location that come into play that you may not be immediately aware of. Allowing the protocol to make the decision, as to what the preferred parent link will be, enables the network topology to overcome most of the issues at the new sites.

## Testing the New Additions

Once the new nodes are operational, test for end-user coverage and add additional nodes as required. Again, as mentioned above, the average spacing to gain adequate 2.4GHz coverage in most areas is approximately 500 feet center-to-center for SkyExtenders. In some areas of a given downtown coverage may require spacing closer than that to sufficiently cover the areas that need to be covered. In less densely populated areas, you may be able to space nodes at a much greater distance, center-to-center.

---

After the links have stabilized and coverage requirements have been met, re-run self-interference analysis for adjusting frequency or enforcing specific links through a preferred parent. At this point, the network should be operational and stable, covering the targeted area sufficiently.

## Validate Link Topology for Areas of Concern

Now that links have been formed and the network is operational, it's time to focus energy, over the next short time period evaluating the topology and determining how it's operating. The focus starts on the first-hop links and extends to second-hop links and, finally, out from the second-hop if there's a specific need to do so.

When you have completed reading this section you will understand the following topics:

- How to identify and resolve weak first and second-hop links
- How to reduce hop counts on links with many hops
- How to identify and resolve external interference with topology clusters

### Identifying and Correcting Weak First and Second-Hop Links

Initially, the focus here is to identify the first and second-hop nodes with modulations below 24-36Mbps or RSSI lower than 25. These weak links can be readily identified using the nodeInfo script and generating the "links" output for the run. This allows you to easily, and quickly filter and sort the information using MS Excel. Links can be viewed, on an individual basis, through the KML file in Google Earth.

#### RSSI Is First

The first point of interest is the RSSI value. If the RSSI is high (25 or greater) then the focus should be on modulation rate. RSSI determines the ability of the link to hear, or differentiate, the correct signal over an interfering signal. A low RSSI value means that the link is susceptible to external interference – whether that interference is generated by other nodes on the mesh, or some outside source. Either way, a low RSSI value indicates susceptibility to these interference sources.

If the node in question is located very near its parent node, the RSSI should be high. If it's not, the path between the node in question and the parent node should be checked for obstructions. Clear Fresnel Zone Line of Sight needs to be maintained.

If sufficient Fresnel Zone Line of Sight exists for the nodes, the next step is to survey the area and look for potential objects that could generate multipath (i.e. reflections) and standing waves. This can be large reflective plate glass windows, metal roofs, metal sided buildings, etc.

If a significant contributor to multipath is found, there are only a few solutions. The first is to attempt to raise/lower the node to better the multipath/standing wave issues. In many cases, either raising or lowering the node will get the antennas out of the path of the reflections, or it will change the angle of incidence causing the reflections to no longer converge on the node. If raising or lowering the node does not resolve the issue, the only other choice is to relocate the node to a more strategic mount point.

#### Modulation Rate

Typically, when the modulation rate is low and the RSSI value is high it indicates the presence of interference. That interference can be interference from other nodes in the vicinity

which fail to follow the design guidelines presented above, or it can be from other sources. If the modulation rates are fluctuating throughout the day and dropping when mesh traffic levels are high and increasing when the mesh traffic levels drop off) then self-interference is likely the culprit.

To determine if self-interference is the issue, we suggest running the nodeInfo with self-interference output enabled. Pay special attention to the nodes (Listening node) that are changing modulation rates and determine, from the output which nodes are interfering with them. Changing a preferred parent so that the links conform more closely to the rules presented earlier in this paper should resolve these issues.

If the issue is interference from another source (i.e. another point-to-point microwave link, etc. then the issue can only be resolved by either changing the frequency on the interfering link, or changing the preferred parent/frequency of the affected node(s).

## **Identifying and Correcting Links with Large Hop Counts**

The next area of focus is nodes with poor link options resulting in hop counts beyond five links. Why these nodes are in this condition; self-interference when linked to their preferred parent; poor parent choices? Review by running the nodeInfo script and generating the “links” output. Evaluate the links formed and determine if there are better alternatives. Make changes to the node profile and reevaluate the changes once completed.

## **Identifying and Correcting Topology Clusters Experiencing External Interference**

The last focus of attention is those topology clusters suffering from interference from external forces that may function better as a cluster on an alternate frequency. How is this identified? Running the self-interference script should display little to no self-interference potential. Modulation rates may be low, RSSI values may be low, etc. Changing node elevations (i.e. changing the multipath characteristics), etc. have no impact on the link quality. In these instances, try assigning all nodes, in that cluster, to an alternate frequency to see if link performance doesn't improve. In most cases, if the interference is coming from outside the network, these changes should significantly improve both link RSSI and modulation rates.

## Monitor for Stability

Once the entire network is up, and operational, it's time to start monitoring the network for stability. Instability presents itself in two forms: device failures, or loss of link due to interference. Both conditions compromise data throughput potential for end-users. This section will take you through some practices to implement in terms of watching the network for stability.

When this section has been read in its entirety you will understand:

- How to identify devices with low availability
- How test average throughput and latency
- Determining which links are changing
- Active path monitoring through SkyControl

## Nodes with Low Availability

The first item on the list of things to watch is for nodes with low availability. These are nodes that may be losing links on a frequent basis, rebooting, etc. SkyEMS displays, at the bottom of the screen, a list of alarm events. If a given link is showing up, in this area, frequently this would be an indicator of low availability. Other monitoring agents, whether they be ping-based, or SNMP-based can also provide information about device availability.

Once these nodes are identified work through normal troubleshooting procedures, such as determining if the issue is related to power interruption (total outage), self-interference (link outage), external interference (link outage), etc. to determine the cause of the issue with the node. Once the issue has been identified, work to resolve it adding a UPS, changing frequencies, preferred parents, or replacing the unit, whichever is deemed necessary.

## Data Throughput Stability and Latency

Another important aspect of stability is how much data throughput a customer can get at any given time on the network. SkyPilot highly suggests testing the throughput to key nodes at regular intervals to ensure that throughput can be maintained. As the network load, from the customer base increases, the throughput to key nodes will decrease as a function of that network load. At some point, the end-user will obtain an unacceptably low throughput/performance rate. It is at that point, that the network should be expanded to provide more units, in a given area, for coverage. This expansion may require the addition of another SkyGateway and more SkyExtenders to provide additional coverage density.

Throughput also correlates to latency, in most cases. Using a tool like SmokePing, or similar application, to measure/monitor latency to strategic nodes along the network path is suggested. As the latency begins to increase to these strategic points run throughput tests to these locations. If the targeted nodes are showing less than acceptable results, determine the cause of the reduction in throughput and work to correct the issue.

In many instances, the cause of the increased latency will be a significant increase in network traffic along that path. As the latency increases beyond a threshold, the end-user throughput will begin to degrade. It will be at that point that additional equipment will be necessary to share the load to that region of the network.

## Link Stability

Periodically, you should run the `nodeInfo` script and dump out the `*.kml` file and the `*links.csv` file to keep an eye on the topology and how the active links are forming between nodes. You're particularly interested, here, in the links that may be fluctuating, or going from node to node over time. In other words, you're looking for link stability.

SkyControl, part of the SkyPilot EMS, has what is called "Active Path Polling". This utility polls the various SkyPilot nodes for information on the active links including RSSI and modulation rates. Enabling this feature will allow you to locate and identify links that vary widely in terms of RSSI and modulation rates. When modulation rates, alone, are varying this typically indicates that there is some sort of interference causing the modulation rates to change. If the RSSI values are changing significantly, this typically indicates that there may be multipath issues impacting the links. Either way, this utility allows you to target specific areas of the SkyPilot network that may be experiencing issues and allows you to work directly with those areas.



## Summary

A lot of material has been covered in the pages previous to this one. Following the outlined procedures should yield a highly successful municipal mesh deployment with few, if any, issues to attend to. The goal, of all the information, is to provide an approach which reduces self-interference and increases network resilience and stability. Doing it right the first time should help to avoid issues as the network grows and expands beyond the first deployment.

The basic strategy of municipal wireless design boils down the following summary:

- Map the coverage area and determine the boundaries.
- Locate strategic mount points within the coverage area.
- Locate wide streets, etc. and try to use these for first-hop links.
- Determine average spacing between nodes based on coverage requirements.
- Map node spacing onto map area using squares to represent coverage area.
- Locate the SkyGateway(s) at the center of the mesh.
- Locate the SkyGateway(s) on the highest mount point available with 360 coverage.
- Use a single SkyGateway to service extremely close first-hop nodes.
- Locate first-hop SkyExtenders with clear line of sight to parent and child nodes
- Locate second-hop SkyExtenders with clear line of sight to parent and child nodes.
- Make every effort to locate child nodes back, and away, from the parent nodes.
- Avoid linking a child node to the same antenna as the parent node.
- Perform RF site surveys to determine operating frequencies while installation of nodes is taking place.
- Power up a single SkyGateway and all potential first-hop nodes.
- Using the \*links.csv file determine the best first-hop nodes for the SkyGateway.
- Focus on RSSI values being higher than 25.
- Focus on modulation rates being 36Mbps at a minimum.
- Focus on spatial diversity.
- Change operating frequencies and note link quality and repeat.
- Choose the best operating frequency for the SkyGateway.
- Repeat for other SkyGateways.
- Once completed assign preferred parents.
- Bring up potential second-hop nodes and perform similar work with them.
- Run nodeInfo and look at the \*.KML file to determine areas lacking coverage.
- Add new nodes and re-run nodeInfo to verify coverage.
- Look for links that have hop counts greater than 4 hops and parent the far children elsewhere.
- Look at the \*links.csv file and determine any links that need attention (low RSSI or mod rate)
- Resolve issues with those links.
- Monitor the network for stability and changes.
- Work to resolve changes as they occur.

© 2007 SkyPilot Networks, Inc. All rights reserved. SkyGateway, SkyExtender, SkyConnector, SkyAccess, SkyControl, SyncMesh, SkyPilot, SkyPilot Networks, the SkyPilot logo, and other designated trademarks, trade names, logos, and brands are the property of SkyPilot Networks, Inc. or their respective owners. Product specifications are subject to change without notice. This material is provided for informational purposes only; SkyPilot assumes no liability related to its use and expressly disclaims any implied warranties of merchantability or fitness for any particular purpose.



*Leading the Mesh Revolution*

1100 Island Drive  
Redwood City, CA 94065  
408.764.8000  
US Toll Free 866 SKYPILOT  
sales@skypilot.com