System Architecture and Protocol Guide

Table of Contents

Introduction **3**

System Components 3

Network Topologies **4**

Point-to-Point/Multipoint 4
Mesh 4

RF Management **7**

Eight-way Antenna Array 7
Spectral Management 8
Link Management 9
Per-link Data Rates for Different Modulation Types 9
SkyPilot Synchronous Mesh Protocol 10

Quality of Service **12**

Security Provisions **14**

Traffic Filtering 14
VLANs 15
Encryption and Authentication 15

Network Management **16**

Introduction

The SkyPilot Carrier-Class Broadband Wireless System supports an innovative architecture offering point-to-point, point-to-multipoint and mesh topologies. Automatic network discovery and connectivity allows an operator to quickly build a network offering high levels of coverage and redundancy.

The system implements a sophisticated and flexible Layer 2 network. Rate control and packet prioritization enables multiple service offerings, including high-quality voice over IP (VoIP) support. Packet filtering combined with wireless link encryption provides very high levels of security.

Full authentication and management of subscriber nodes is supported, with each network element providing an SNMP Agent and a Command Line Interface (CLI). Software images are upgraded over the air.

This document provides a description of the SkyPilot system architecture and protocols. The hardware components are described along with details of network topologies and routing, spectral and link management, network-wide synchronization, quality of service, security provisions, and network management.

System Components

The SkyPilot Carrier-Class Broadband Wireless System has four hardware components that together support a highly flexible wireless network architecture:

- ***SkyGateway***
  The SkyGateway provides an interface between the wireless network and the Internet or other Local/Wide Area Network (LAN/WAN). The SkyGateway links directly to SkyExtenders and/or SkyConnectors as customer premises equipment (CPE) in point-to-point/multipoint configurations, or indirectly via multiple SkyExtender hops in a full mesh network topology. The SkyGateway has a compact, lightweight, fully integrated mechanical design with a 10/100Base-T interface and an integral GPS receiver to provide a highly accurate system timing source.

- ***SkyExtender***
  The SkyExtender enables a network operator to expand coverage and increase system capacity. The SkyExtender can link directly to a SkyGateway and/or one or more SkyExtenders and/or SkyConnectors. The SkyExtender hardware is identical to that of the SkyGateway, with the software image differentiating functionality. The 10/100Base-T Ethernet interface may be used to provide a subscriber interface, allowing the SkyExtender to perform double-duty as CPE that extends the network's reach from the SkyGateway(s).

- ***SkyExtender DualBand***
  The SkyExtender DualBand has the same backhaul features as a SkyExtender, and adds an integral Wi-Fi access point to afford service providers with greater flexibility when deploying wireless access solutions in HotZones or HotRegions. As the name implies, the SkyExtender DualBand has separate radios for backhaul (at 5 GHz) and access (industry-standard IEEE 802.11b/g at 2.4 GHz). The separation of access from backhaul enables the use of multiple non-overlapping channels when deploying a "honeycomb" of Wi-Fi cells, thus enhancing throughput, subscriber density, network capacity and scalability.

- ***SkyConnector***
  The SkyConnector comes in both outdoor and indoor variants, and provides an interface between a subscriber and the wireless SkyPilot network. The link may be indirect via a SkyExtender or direct to a SkyGateway. The subscriber interface is standard 10/100Base-T Ethernet. The SkyConnector is fully integrated and houses a single antenna with a 28° horizontal coverage.

Please refer to the individual SkyPilot equipment datasheets for a comprehensive description of the features and functionality of each unit.

Network Topologies

The SkyPilot Carrier-Class Broadband Wireless System supports point-to-point, point-to-multipoint and mesh topologies. An operator has the flexibility to start with a simple point-to-point/multipoint configuration and add SkyExtenders as subscriber demand requires to create a more robust mesh topology over time.

Point-to-Point/Multipoint

In a typical point-to-multipoint configuration, a SkyGateway is installed on a cell-tower, tall building or other height-advantaged location. All SkyConnectors deployed within range then communicate with the SkyGateway's centralized location. Note that the compact nature of the SkyGateway means that smaller networks can easily be supported; simply install the SkyGateway at any appropriate location where wired or wireless backhaul can be provided.

This simple, omnidirectional configuration can be extended outward in any direction using SkyExtenders. The links or "hops" from the central SkyGateway to the SkyExtenders can be point-to-point or point-to-multipoint. At some point (no pun intended), the network topology becomes a mesh with the potential for multiple, redundant links among SkyExtenders and one or more SkyGateways.

Mesh

A SkyPilot mesh network utilizing SkyExtenders offers a number of significant advantages over point-to-point and point-to-multipoint topologies:

- ***Coverage***
  In a point-to-multipoint wireless network, each subscriber must link to a single, central base station location. In most deployment scenarios, large buildings, hills and other line-of-sight obstructions result in a lack of service availability to a large percentage of prospective subscribers. The SkyPilot mesh topology offers service to any subscriber within range of the SkyGateway(s) *or* any SkyExtenders. This topology enables subscribers to "connect around" obstructions that would defeat a point-to-multipoint system, and serves to extend the overall range or reach of the network.

- ***Capacity***
  The SkyPilot system supports variable modulation rates, with a continuous link optimization process that selects the highest data rate possible for a given link. As SkyExtenders are added to the network and the mesh grows, the number of possible links increases and the average distance between them can decrease, which together result in increased modulation rates and greater overall network capacity.

- ***Redundancy***
  The SkyExtenders automatically establish multiple links with other SkyExtenders and/or SkyGateways. Though only a single link per node is used for routing traffic at any moment in time, the remaining links are retained and available for immediate use in the event of a primary link failure or degradation. This level of redundancy extends throughout the mesh where multiple SkyGateways can be deployed in a load-balancing arrangement with automatic failover.

- ***Spectral Efficiency***
  All communications throughout a SkyPilot mesh are directed and coordinated (see sections on ***RF Management*** and the ***SkyPilot Synchronous Protocol***), allowing for greater spectral re-use to achieve better performance and enhanced scalability.

Mesh Routing

All subscriber traffic is routed between a SkyGateway and the subscriber's node (a SkyConnector, SkyExtender DualBand, or SkyExtender providing subscriber connectivity). While traffic can be exchanged among nodes within the mesh, it must ultimately be forwarded by a router or switch to an external network, which also adds an additional level of security. The SkyPilot system software allows for the optional switching of subscriber traffic by a SkyGateway.

The SkyPilot mesh routing protocol, which is independent of the networking interfaces at both the SkyGateway and the SkyExtender/SkyConnector nodes, utilizes a cost-based algorithm that factors link capacity (modulation levels) and hop count. Each SkyGateway periodically generates a routing cost message. Every directly-linked SkyExtender and SkyConnector receiving

this message records the link on which the message was received, the SkyGateway identifier and the associated cost. These SkyExtender and SkyConnector nodes utilize this cost data to select the optimal, lowest cost link on which to forward data.

Every SkyExtender forwards a message specifying its lowest cost path to a given SkyGateway. In the case of a network with multiple SkyGateways, each SkyExtender forwards a cost message for each SkyGateway. The SkyExtenders forward the cost message(s) to any and all neighboring nodes.

Routing Cost

The routing cost is calculated as follows:

1) A SkyGateway generates a routing cost message based on the cost table shown below. The values in this table represent a weighted inverse relationship to modulation level with double the weight given to downstream capacity. This weighting towards higher modulation has the effect of focusing system traffic on the SkyGateway links with the greatest capacity, thus maximizing the overall capacity of the entire SkyPilot network.

SkyGateway Link Cost Table

Raw Mod Rate	6Mbps	9Mbps	12Mbps	18Mbps	24Mbps	36Mbps	48Mbps	54Mbps
Downstream	85	56	41	27	20	13	9	8
Upstream	43	28	21	14	10	7	5	4

For example, consider the cost calculated for a single SkyGateway link with the same overall aggregate capacity. The total cost of a link with 36 Mbps downstream modulation (cost of 13) and 18 Mbps upstream modulation (cost of 14) link is **27**; the total cost of a link with 24 Mbps downstream modulation (cost of 20) and 24 Mbps upstream modulation (cost of 10) is **30**; and the total cost of a link with 12 Mbps downstream modulation (cost of 41) and 36 Mbps upstream modulation (cost of 7) is **48**. In this example, the first link with a cost of 27 will be selected first.

2) A SkyExtender selects the lowest cost to a given SkyGateway. For each SkyExtender link, the cost is calculated using a formula and values from the table below. The total cost of the link is the sum of the lowest incoming link cost increased by 10% plus the outgoing link cost determined by its downstream and upstream rates. The algorithm is as follows:

Total_link_cost = (lowest_incoming_cost * 1.1) + outgoing_link_cost **(from the table below)**

SkyExtender Link Cost Equation and Table

Raw Mod Rate	6Mbps	9Mbps	12Mbps	18Mbps	24Mbps	36Mbps	48Mbps	54Mbps
Downstream	72	48	36	24	18	12	9	8
Upstream	36	24	18	12	9	6	5	4

The figure below illustrates an example of routing costs in a relatively small network with seven SkyExtenders and a single SkyGateway. As each SkyExtender is added to the network, it determines its route selection based on either the direct link cost (first table) or the cost calculation above.

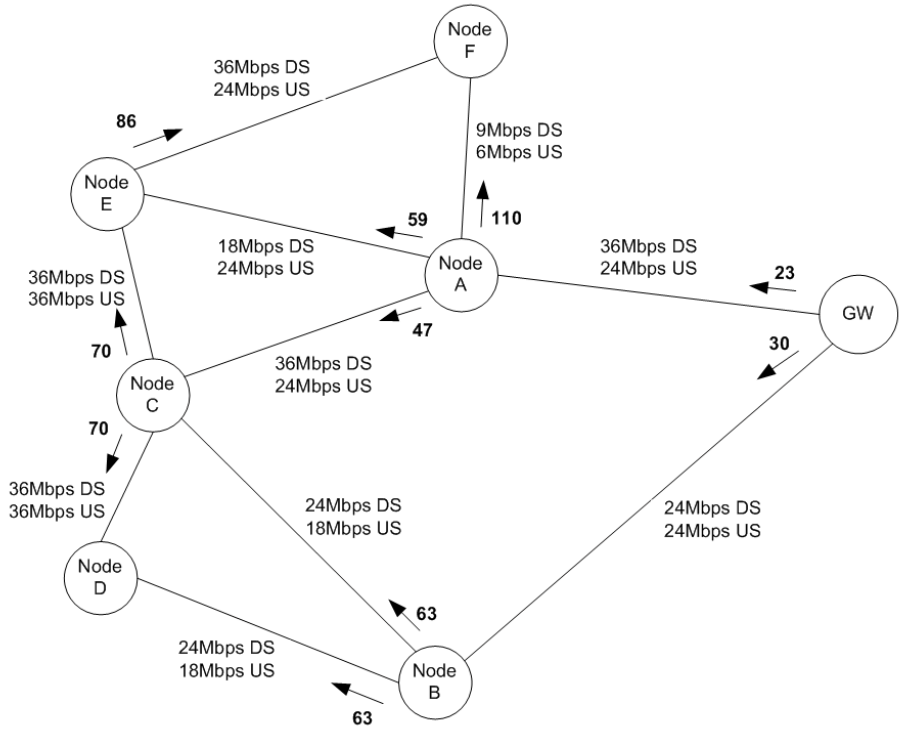

Routing Example

Example of mesh routing based on lowest link cost.

- SkyExtender A – This node, the first one deployed, receives a cost of **23** directly from the SkyGateway. The node selects to route over this 36/24 Mbps link.
- SkyExtender B – This node receives a cost of **30** directly from the SkyGateway, and it selects to route over this 24/24Mbps link.
- SkyExtender C – This node receives two cost messages, one from SkyExtender A advertising a cost of (23*1.1) + 21 = **47** and one from SkyExtender B advertising a cost of (30*1.1) +30 = **63**. This node selects to route via SkyExtender A and stores the route via SkyExtender B as a backup route. Note that when SkyExtender C is added to the network, it forwards a cost related to the SkyGateway/SkyExtender B route to SkyExtender A, as well as a cost related to the SkyGateway/SkyExtender A route to SkyExtender B, thereby providing an alternate route to these nodes.
- SkyExtender D – This node receives a cost of **63** from SkyExtender B and a cost of **70** from SkyExtender C. Therefore, it selects to route using two hops via SkyExtender B.
- SkyExtender E – This node receives a cost of **59** first from SkyExtender A then a cost of **70** from SkyExtender C. Therefore, it selects to route using two hops via SkyExtender A.
- SkyExtender F – This node receives a cost of **110** from SkyExtender A and a cost of **86** from SkyExtender E. In this case, SkyExtender F selects to route using a three-hop route via SkyExtender E and SkyExtender A.

Rerouting

A rerouting decision is made either due to an active route failure or because a significantly lower cost route is consistently available. There are four scenarios that cause a reroute operation:

1. A failure in the link or a SkyGateway/SkyExtender system along the selected route.
2. A failure in the link or a SkyGateway/SkyExtender further along the route to the SkyGateway.
3. One or more links along the route degrade in performance, resulting in reduced capacity and, therefore, an increased cost.
4. A lower cost route becomes available. Typically this is through improved link quality or the addition of SkyExtender(s) and/or SkyGateway(s) to the network.

A link (or node) failure is detected by three consecutive failures to receive a “keep-alive” message. These link keep-alive messages are transmitted every 1.5 seconds. Therefore, a link is designated as having failed after a maximum of 4.5 seconds. If desired, the system can be configured to further reduce the maximum time it takes to detect a failure.

A SkyExtender that detects a link failure attempts switching to the next lowest cost route available and activate the associated link. This activation includes the exchange of scheduling information and takes approximately 1 second. Once that activation is successfully completed, the SkyExtender then transmits the cost associated with this new route.

If the SkyExtender detecting a link failure has no alternate routes or the activation of the associated link(s) fails, then it transmits an “infinite” cost message and starts a frequency hunt for another route. Upon *receiving* an infinite cost message related to an active route, a SkyExtender will immediately attempt to reroute.

A SkyExtender will change to a lower cost route only when the cost is at least 15 units lower than the current active route.

RF Management

The SkyPilot Carrier-Class Broadband Wireless System affords an unprecedented level of control over radio frequency (RF) utilization. At the foundation of these capabilities is an innovative 8-way antenna array that supports point-to-point, point-to-multipoint and omnidirectional communications. The array enables more effective utilization of available spectrum on a network-wide basis, and even permits granular management of individual links among all nodes.

Eight-way Antenna Array

The SkyGateway, SkyExtender and SkyExtender DualBand systems all have an integral 8-way antenna array providing a full 360° of coverage in 45° segments. This design enables a granular control over transmission direction, power levels and more, which together allow the network topology to be optimized for maximum performance with minimal interference.

Within this array, each of the eight antennas provides a focused 45° horizontal beam; all eight antennas together provide omnidirectional 360° coverage wherever desirable. Internal to each system is a high-speed switching fabric that alternates among the antennas to provide focused, directional transmissions. When utilized individually, a single antenna establishes a point-to-point link that can transmit at higher power levels than permitted for point-to-multipoint or omnidirectional communications. This allows the system to deliver as much as 44.5 dBm / 28.2 W EIRP, resulting in much longer range (up to 10 miles or 16 kilometers between nodes) and/or greater penetration through obstacles that attenuate the signal.

The directional capabilities of the antenna array also permit more effective utilization of available spectrum by allowing a single node to communicate at different frequencies with different neighboring nodes. With point-to-multipoint and omnidirectional systems, by contrast, all nodes must utilize a common frequency. By extension, a mesh network utilizing omnidirectional antennas must employ the same frequency end-to-end. The SkyPilot directional design overcomes this limitation, which can severely cripple the performance and scalability of first-generation mesh solutions.

By directing its RF energy in an optimal fashion, the SkyPilot design greatly reduces susceptibility to both internal and external interference. Mitigating self-interference allows for spectral re-use within the network, which increases aggregate capacity and permits densities not possible with other wireless broadband solutions. In addition, the system schedules transmissions only on active sectors (see section on ***SkyPilot Synchronous Protocol***), thereby further enhancing performance and quality of service.

Spectral Management

SkyExtenders and SkyConnectors require no pre-configuration and automatically search or “hunt” among all supported frequencies and, in the case of the SkyExtender, all antennas. When a SkyExtender or SkyConnector is powered up, it checks non-volatile memory for a *Preferred Frequency*. This is the frequency that should be favored in any search, and is set after a successful connection to the network or by operator configuration. In addition to the Preferred Frequency, there is also a *Frequency Hunt List*. This list, defined by the operator, specifies the frequencies that should be searched if the Preferred Frequency is unavailable. The Frequency Hunt List is *not* stored in non-volatile memory and, therefore, is *not* preserved through a power reset.

If there is a Preferred Frequency defined, the SkyExtender or SkyConnector will first attempt to listen for Hello messages on that frequency. A Hello message is transmitted on each of the eight antenna approximately once per second. These messages provide system timing information, as well as timing for the associated Hello Response message. A SkyConnector will listen for 2 seconds on its single antenna. A SkyExtender will listen for 2 seconds on each of its eight antennas for a total of 16 seconds at that frequency. If no Hello messages are heard, the node will then check the Frequency Hunt List. If a list is defined, it will start at the lowest frequency in that list, testing each in turn while alternating with the Preferred Frequency. For instance, if an operator defines a Preferred Frequency of 5765 MHz and a Frequency Hunt List of 5775, 5800 and 5825 MHz, then the sequence becomes 5765, 5775, 5765, 5800, 5765, 5825 MHz. This search of the Frequency Hunt List is attempted a total of two times. The SkyConnector listens for 2 seconds at each frequency, while the SkyExtender listens for 16 seconds.

After checking the Preferred Frequency and any frequencies defined in the Frequency Hunt List, the node searches the complete frequency band indicated by the product model number. In the US, this is the UNII band frequencies of 5745-5835 MHz. The search starts with the non-overlapping channels (5745, 5765, 5785, 5805, 5825 MHz). If there is a defined Preferred Frequency, the node continues to alternate between listening on each of these channels and that stored frequency. This search of the non-overlapping channels is attempted a total of two times.

If the non-overlapping channel frequency hunt fails, then the node will scan all frequencies with an increment of 5 MHz. In the US this means 5745, 5750, 5755, etc. up to 5835 MHz. During this search the node continues to alternate with the Preferred Frequency. If, after two attempts this hunt also fails, then the Preferred Frequency is deleted and the node will continue hunt the full frequency range with a 5 MHz increment.

If a node hears a Hello message at any frequency, the search will be stopped for a time equal to four complete listening periods. This translates into 64 seconds for SkyExtenders and 8 seconds for SkyConnectors. This allows the node time to gather enough Hello messages to start the link optimization process. If a connection is made (see ***Link Management***), this unit will change to the discovery phase and its frequency hunt will be halted.

The discovery phase lasts 45 seconds and allows a node time to complete link establishment, authorization and configuration, and possibly to receive Hello messages from other nodes (that may offer lower cost routes). The following may also occur during the discovery phase:

- If, at the completion of the discovery phase timer no link has been establish and no link establishment is in progress, then the node will resume frequency hunting.
- If, at the completion of the discovery phase timer a link establishment is in progress, then the node will restart the discovery timer with a 90 second duration.
- If, at the completion of the discovery phase timer a link has been successfully established and the node authenticated, then a final 45 second discovery timer is started and the node is marked as being connection-ready.
- If, at the completion of the discovery phase timer a node is marked as being connection-ready, then the best route is selected (see ***Mesh Routing***) and the node is placed online.

Link Management

The SkyPilot system supports variable modulation types and encoding rates. Differing rates can be utilized for each individual link within the network, including each of the multiple individual links possible for any given node. The various modulation types and rates are shown in the table below. Note that the payload traffic capacity will be less than the stated data rates due to packet header and protocol overhead.

Per-link Data Rates for Different Modulation Types

Raw Data Rate (Mbps)	Modulation Type	Coding Rate	Coded bits per Subcarrier	Coded bits per OFDM Symbol	Data bits per OFDM Symbol
6	BPSK	½	1	48	24
9	BPSK	¾	1	48	36
12	QPSK	½	2	96	48
18	QPSK	¾	2	96	72
24	16QAM	½	4	192	96
36	16QAM	¾	4	192	144
48	64QAM	½	6	288	192
54	64QAM	¾	6	288	216

Each link within a SkyPilot network is optimized for maximum performance. This optimization occurs both at the point of link creation and as an ongoing process in the case of active links utilized for subscriber traffic.

RSSI vs. Modulation Level

The modulation level of a link will generally relate directly to the Received Signal Strength Indicator (RSSI). RSSI values are available via the SkyPilot MIB or CLI. Interference within a channel can cause a reduced modulation level.

The chart below shows test data mapping packet throughput against the Receive Signal Strength Indicator (RSSI). 10,000 packets were transmitted for each combination of modulation and RSSI level. Note that this test was performed in the lab, free from interference and multipath reflections, for both SkyGateway/SkyExtender and SkyExtender/SkyExtender links. The thick horizontal line indicates a 90% packet success rate. The current link optimization algorithm requires a 95% packet success rate in order to select the modulation level.

Based on these test results, the table below can be used to map RSSI to modulation levels:

RSSI	Modulation Level (Mbps)
7	6
8	9
9-11	12
12-13	18
14-18	24
19-30	36
31+	48

Initial Link Optimization

The SkyPilot Synchronous Mesh Protocol (see separate section) implements an Automatic Repeat reQuest (ARQ) mechanism that allows for the rapid retransmission of errored packets. This feature operates transparently to the all protocols (such as TCP/IP) transmitted over a SkyPilot system. When a link is first established, typically when a new SkyExtender or SkyConnector is added to a network or after a network reset, a series of test packets are exchanged in order to select the optimal modulation type and rate. The ARQ protocol is disabled and the two nodes constituting the link transmit 90 packets at each of the modulation levels in turn. The number of packets received at each modulation rate is then returned to the transmitting node. *Using these results, the nodes select the highest modulation and coding level that supports a > 95% successful reception rate*.

Continuous Link Maintenance

Every active link used to transfer subscriber traffic (see the ***Mesh Routing*** section for details about link selection) is monitored for basic connectivity and to ensure the optimal modulation levels.

The general link connectivity is continuously tested through the use of “keep-alive” messages. These messages are exchanged every 1.5 seconds. This message transfer tests not only the RF connectivity, but also management control level connectivity. After three consecutive message failures, the link is considered to be down, and a reroute or Frequency Hunt is initiated (see the ***Spectral Management*** section for details).

The modulation level of each link is continuously optimized. This is achieved by monitoring the packet retransmission rates over the link. A sliding window of 576 packet transactions is monitored. This means that at any moment in time the previous 576 packet transmissions/receptions are analyzed. The following actions are taken based on the failure rate experienced:

- If the failure rate exceeds 50 packets (8.68%), then the modulation is reduced by one level. Note that a packet failure is defined as a single packet not being received successfully.
- If the packet failure rate is less than 10 (1.74%) then the modulation is increased one level.

SkyPilot Synchronous Mesh Protocol

The SkyPilot Carrier-Class Broadband Wireless System implements a synchronized variant of the asynchronous IEEE 802.11 protocol. This SkyPilot protocol enhancement utilizes the baseband capabilities of 802.11, including all modulation levels and OFDM. In addition, the SkyPilot protocol utilizes the 802.11 frame format.

The SkyPilot enhanced synchronous protocol, called Time Division Duplex (TTD), provides the following advantages when compared to the asynchronous 802.11 protocol:

- Predictable, low latency
- Scalability
- Accurate, granular QoS control
- Spectral efficiency

IEEE 802.11

The 802.11 standard was created to provide a low cost, easily implemented Wireless Local Area Network (WLAN). The standard defines a Carrier Sense Multiple Access (CSMA) protocol that is used to control access to the channel. The CSMA mechanism requires that all (or at least the vast majority of) nodes within the network are able to hear the transmission of all other nodes. When standard 802.11 devices are unable to hear their peers, as is very often the case in large-scale outdoor deployments, multiple nodes attempt to transmit at the same time, causing unwanted collisions and a loss of capacity. Testing has shown that such a network, with as little as four active 802.11 nodes, can experience heavy degradation of system throughput. This issue is often referred to as the “hidden node problem.”

To address the “hidden node problem,” a small number of vendors have implemented polling schemes based on the standard 802.11 Request to Send/Clear to Send (RTS/CTS) mechanisms to allow a central base station to control CPE traffic. This method successfully avoids the high level of data traffic collisions, but provides only a crude level of bandwidth and latency control. SkyPilot has developed this approach further, providing a fully scheduled MAC layer protocol with a very fine granularity of bandwidth control.

SkyPilot Mini-Slots

All SkyPilot nodes are synchronized to the same 1 second timing frame. The SkyGateway and SkyExtender nodes house a GPS receiver that provides a standard, highly accurate 1 Pulse Per Second (PPS) clock. SkyConnectors derive the 1 second frame from timing messages exchanged with the connected SkyGateway or SkyExtender.

In order to allow for the most efficient utilization of the wireless channel capacity, the timing of the upstream transmission is adjusted such that packets are received at the SkyGateway or SkyExtender regardless of the distance from the transmitting node.

Each 1 second frame is divided into 10 mini-slots, each with a 100 µs in duration. These smaller time increments allow bandwidth to be allocated for low latency and low jitter at a very granular level.

Bandwidth Allocation

Bandwidth management within a mesh network presents a particular challenge owing to the high number of interconnections. The SkyPilot system implements a very efficient distributed dynamic bandwidth allocation scheme, with each SkyGateway or SkyExtender responsible for scheduling bandwidth for all nodes linked directly. These distributed schedules are coordinated through the enforcement of a parent/child relationship. For each node, the node one hop closer to the SkyGateway is considered a scheduling parent. This parent and child relationship is repeated in a hierarchical fashion, with the SkyGateway as the final bandwidth-scheduling parent in the hierarchy.

Each parent distributes a bandwidth schedule to its child nodes. This schedule details the mini-slots in which communication from the parent to the child node *may* occur. A child node must *always* listen on the appropriate antenna at these scheduled times for transmission from a parent node.

Contention Slots

In addition to the downstream transmission slots, the bandwidth schedule also details contention slot opportunities. These contention slots are used by the child node to indicate upstream bandwidth needs.

In order to obtain upstream bandwidth, a child node transmits a Bandwidth Request message in the scheduled contention slot. This Bandwidth Request message details the amount of upstream data in each upstream queue.

Multiple child nodes can transmit in the same contention slot, resulting in a collision. To address this issue, each child node implements a randomized backoff. After each Bandwidth Request transmission, a child node randomizes the number of contention slots opportunities before attempting to issue the next Bandwidth Request. Note that a child node will *not* transmit a Bandwidth Request message if it has recently communicated with its parent.

Traffic Control with TDD

A parent node controls the transfer of traffic to and from its child nodes. A parent node makes the following real-time scheduling decisions:

Transmitting Data to Child

Ethernet packets are retrieved from the output queues for transmission to child nodes. Various QoS mechanisms (see ***Quality of Service***) dictate the ordering of packet transmissions to the various child nodes. Typically, the system transmits an Ethernet packet in its entirety; however, if the parent must interrupt the transmission due to traffic shaping mechanisms or to listen to its own parent node, it has the ability to fragment the packet.

Having made the decision to transmit to a given child node, a parent starts the transmission at the next scheduled mini-slot with that child. The transmission continues until complete. While this transmission is in progress, the parent node does not communicate in mini-slots scheduled with other child nodes. Transmitting a large Ethernet frame spans many mini-slots.

Receiving Data from a Child

If a child node has data queued for transmission upstream and is not in active communication with its parent, it will signal the need by transmitting a Bandwidth Request message in the next contention slot. The details of the Bandwidth Request, including the priority of the data queued, is stored by the parent node and accessed by the scheduler.

In the current SkyPilot system, child status is polled by the parent. Having made the decision to poll status from a given child node, a parent transmits a Poll Request message in the next scheduled mini-slot with that child. Upon being polled, a child node immediately transmits a packet up to the maximum size specified in the Poll Request message.

The polling of a child continues as long as upstream bandwidth is queued plus an additional integration period of 0.5 seconds after the queue is emptied. This continued periodic polling allows an active node to make a low-latency upstream transfer by avoiding the need for the child node to contend. The parent also utilizes any unused bandwidth to poll recently active nodes beyond the 0.5 second integration period.

In addition to the dedicated Poll Request message, a parent node can also include a Poll Request with a data frame transmitted to a child. A future revision of the SkyPilot protocol may replace the Poll Request messages with upstream bandwidth grants that directly allocate one or more transmit opportunities to a child node. Both of these mechanisms increase bandwidth efficiency by removing the Poll Request message overhead and propagation delay.

Quality of Service

The SkyPilot Carrier-Class Broadband Wireless System supports the following Quality of Service (QoS) features:

- ***Subscriber Rate Control***
  An operator can define maximum upstream and downstream data rates for a subscriber. This rate control is implemented at the ingress points; i.e. at the SkyGateway in the downstream direction, and at the SkyConnector/SkyExtender in the upstream direction.

- ***Packet Prioritization***
  An operator may prioritize individual subscriber packets based on a range of parameters. This classification is completed at the ingress points and the packets marked as having a high or low priority. This prioritization is then maintained by any intermediary SkyExtenders.

Downstream QoS

The downstream QoS functionality is implemented by the SkyGateway and any intermediary SkyExtenders. The SkyGateway limits and shapes the user data received over the local 10/100Base-T Ethernet interface using a token bucket-based algorithm. The SkyGateway maintains a token bucket for every SkyConnector, every SkyExtender DualBand, and every SkyExtender providing subscriber termination). The depth of each token bucket is defined by the maximum allowable burst size, and is filled with tokens at a constant rate derived from the configured maximum data rate.

The SkyGateway classifies all Ethernet packets received on the 10/100Base-T Ethernet interface in accordance with the configuration specified by the operator. For each individual subscriber (SkyConnector or SkyExtender) the operator may designate packets as having a high priority using one or more of the following means:

- ***Source/Destination IP Address***
  A single IP address, a number of IP addresses, an IP address range or a number of IP address ranges. This classifier can be used to prioritize packets exchanged between a subscriber and a voice gateway.

- ***IP Type of Service (ToS)***
  An IP header field that is used by certain protocols/equipment to indicate packet priority and type.

- ***IP Protocol Type***
  An IP header field used to indicate the packet protocol type.

- ***IEEE 802. 1p***
  A Layer 2 protocol that is used to mark the priority of individual Ethernet packets.

The operator can use a range of classifiers to control prioritize packets. For instance, an operator may assign a high priority to packets matching a range of IP addresses *and* having a defined ToS field setting. Any packets *not* matching this classification will be assigned a low priority. The resulting classification is tagged (in the SkyPilot protocol header) so that any intermediary SkyExtenders can maintain the prioritization.

Upon receiving an Ethernet packet, the SkyGateway first performs the traffic classification. Having designated a packet as having a high or low priority, the subscriber's token bucket is checked to see if it has enough tokens for packet transmission. If there are no tokens in the bucket, tokens may be borrowed. For low priority packets, tokens can be borrowed up to double the configured downstream rate. For high priority packets, tokens can be borrowed up to three times the configured downstream rate. If a token is available, the packet is queued to the downstream link—high priority packets to the high priority queue and low priority packets for the low priority queue. If there are no tokens available, then the packet will be dropped.

The figure below illustrates the SkyGateway QoS mechanisms:

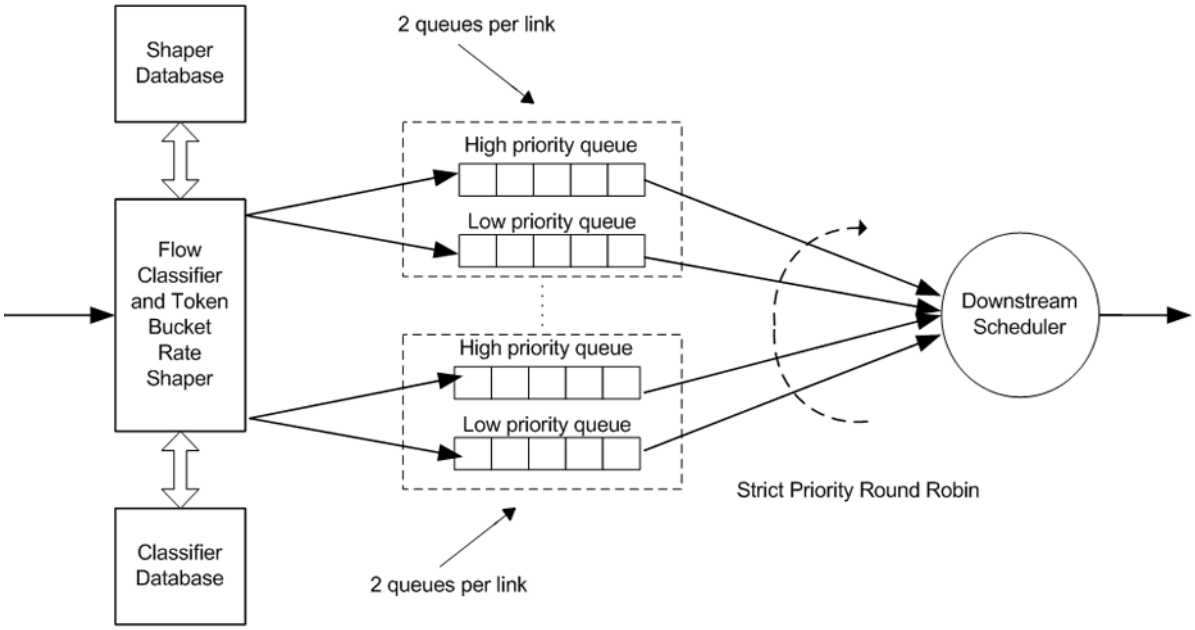

Flow Classification and Rate Shaping

Note that broadcast and multicast packets are also rate-controlled. The default configured downstream rate is 128 kbps.

Upstream QoS

In the upstream direction, the SkyConnector or SkyExtender limits and shapes the user data received over the local 10/100Base-T Ethernet interfaces. The SkyConnector and SkyExtenders perform the same packet classification and token bucket shaping as the SkyGateway.

Security Provisions

A SkyPilot Carrier-Class Broadband Wireless System operates as an intelligent, virtual Ethernet switch, with a full learning bridge implemented at both the subscriber interface (SkyConnector or SkyExtender) and the networking interface on the SkyGateway. This Layer 2 network architecture allows subscribers to have full mobility between SkyGateways without the need to update IP addresses. The architecture also allows for enforcement of same security provisions that are possible with a physical Ethernet switch. The diagram below shows how Virtual LANs might be supported on a SkyPilot network.

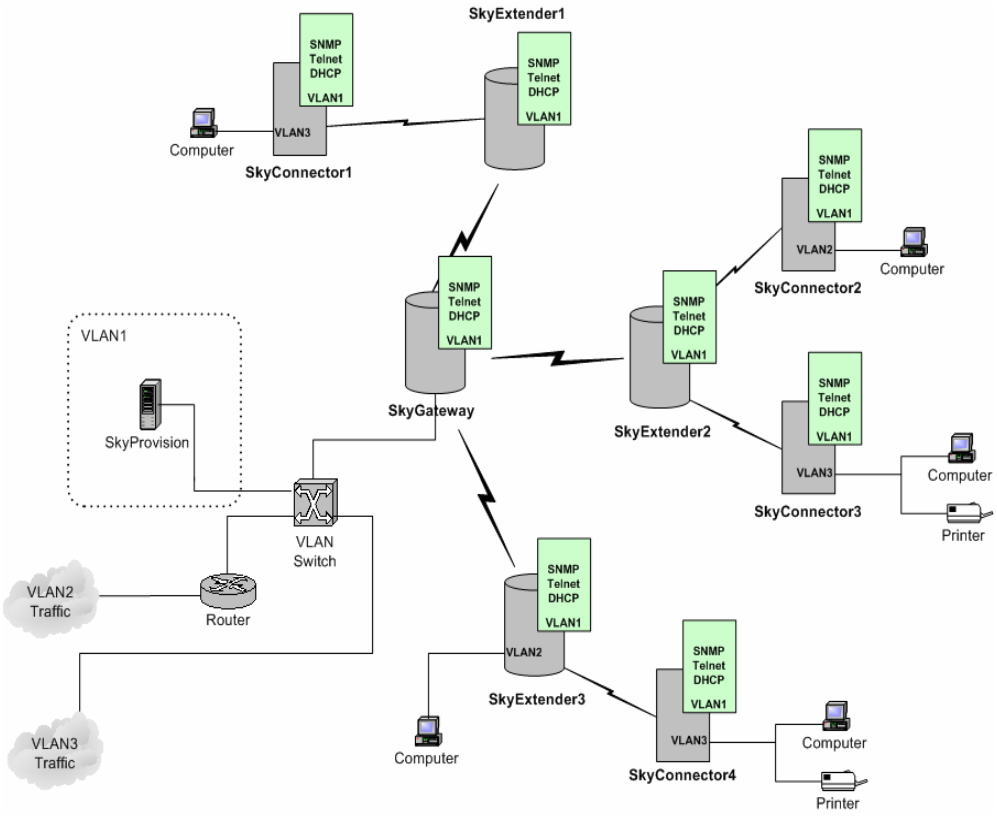

Here is an example of a SkyPilot network supporting standard VLANs as a virtual Ethernet switch.

Packet-level traffic filtering allows an operator to control the subscriber access, improve system security and manage IP address allocation. VLANs provide additional network security and the means to partition traffic for differing subscriber groups. The peer-to-peer switching control provides additional security and interconnectivity options. Finally, encryption and authentication provide confidentiality and integrity for all user traffic. Each of these provisions is explained below.

Traffic Filtering

Each packet received on the SkyConnector or SkyExtender subscriber interface is examined, and any configured filters are applied. Multiple filters can be applied simultaneously with flexible logic. Packets can be discarded only if *all* filters are true, or packets can be filtered if *any* of the filters is true. The following filter types are supported:

- EtherType
- Source/Destination IP Address Range
- IP Protocol Type
- IP Port ID

Typical uses of the filtering capability include protocol access restrictions, for instance filtering all non-IP protocols and source IP address ranges. The latter filter allows an operator to control the IP addresses utilized by a given customer.

VLANs

The SkyPilot system allows an operator to partition subscriber traffic through the use of Virtual LANs. VLANs limit the scope of broadcast and multicast traffic, and facilitate the segmentation of traffic in the backbone network.

An individual subscriber can be configured to a single VLAN. In this mode all Ethernet packets received on the subscriber interface of the SkyConnector or SkyExtender are tagged with the configured VLAN ID. These packets are then forwarded to the SkyGateway, which transparently switches the frames without modifying the VLAN tag. In this mode, the SkyConnector or SkyExtender will strip the VLAN tag (i.e. convert to standard Ethernet format) before transmitting the packet on the subscriber interface.

In VLAN mode, any 802.1q format packets received by the SkyConnector or SkyExtender on the subscriber interface will be checked for a VLAN ID. If the ID matches the configured VLAN ID, then the packets will be forwarded to the SkyGateway. If the ID does *not* match, the packet will be discarded. If the VLAN mode is *not* configured, then a SkyConnector or SkyExtender will transparently switch 802.1q VLAN tagged Ethernet frames.

Peer to Peer Switching

The SkyPilot system permits control over Layer 2 switching by the SkyGateway. An operator can enable or disable the Layer 2 switching between SkyConnectors and SkyExtenders on a per VLAN basis. With switching enabled, an operator can offer low latency Layer 2 VLAN connectivity. With switching disabled, an operator can utilize an external switch or router for high levels of network security and flexibility.

Encryption and Authentication

The SkyPilot network provides for subscriber confidentiality, message integrity, and endpoint and network authentication. Confidentiality prevents unwanted listeners of network traffic. Message integrity ensures that network traffic has been delivered unaltered to the intended recipient. Endpoint authentication provides trust between two nodes on a network. Without this endpoint authentication, confidentiality and message integrity are difficult to achieve. Endpoint authentication also ensures that only authorized nodes join the network. Network authentication ensures that nodes join only trusted networks.

Network Key

The SkyPilot system provides support for a shared network key. This static key is installed (via the CLI) by a network operator prior to deployment. Only nodes that share the network key are authorized to join the network. Note that, due to security concerns, the system key cannot be changed remotely.

Identification Certificate

All SkyPilot nodes have a unique identification certificate installed during manufacture. These are part of a certificate chain that is signed all the way to the SkyPilot root certificate authority, and uniquely identify each SkyPilot node.

Authentication

When a SkyPilot node connects to the network via a SkyExtender or directly via a SkyGateway, the following peer-to-peer authentication protocol is completed:

1. The authentication negotiation is initiated by the node with the numerically smaller MAC address (Node A). This node sends an SP_AUTH_HELLO message that contains its Identification Certificate to the node with the numerically larger MAC address (Node B).
2. Upon receiving the SP_AUTH_HELLO message, Node B verifies that the ID certificate in the message has been signed by the correct certificate authority (CA) using its root public key. This step provides assurance to Node B that Node A possesses a valid certificate. If Node B successfully verifies the Identification Certificate, it will compute a random number. This random number is one of the ingredients used in generating the secret session key.
3. Node B will encrypt the random number it generated with the shared network key. It will then encrypt the ciphertext with Node A's public key. The purpose of this dual encryption is to verify that Node A has the same shared static key as Node B, and to ensure that only Node A can decrypt the message. The final ciphertext is combined with Node B's Identification Certificate and placed in an SP_AUTH_CHALLENGE message. This message is sent to Node A.

4. When Node A receives the SP_AUTH_CHALLENGE message, it verifies Node B's Identification Certificate using its root public key. Node A decrypts Node B's random number using its private key and shared static key. Node A then computes its own random number. Using the two random numbers it computes the secret session key.
5. Node A encrypts its random number using the shared static key. It then encrypts this ciphertext along with Node B's decrypted random number using Node B's public key. The resulting ciphertext is placed in an SP_AUTH_RESPONSECHALLENGE message and sent to Node B.
6. When Node B receives the SP_AUTH_RESPONSECHALLENGE message, it decrypts the random numbers. Node B verifies that this decrypted random number matches the one it generated in Step 2. If the random numbers match, then Node B will generate the secret session key using the two random numbers. Node B will encrypt Node A's decrypted random number using Node A's public key. Node B will send Node A an SP_AUTH_RESPONSE message containing the ciphertext.
7. Upon receiving the SP_AUTH_RESPONSE message Node A decrypts the random number using its private key. It compares the decrypted random number with the one it generated in Step 4.

If at any step this authentication protocol fails, the link will be placed in the authorization failed state and no user data will be exchanged.

Encryption

Having successfully completed the authentication protocol, the nodes load the session keys into hardware. All packets passed between these two nodes are then encrypted utilizing these 128-bit Advanced Encryption Standard (AES) session keys.

Network Management

The management of all nodes in a SkyPilot Carrier-Class Broadband Wireless System can be handled through either the local Command Line Interface (CLI) and/or through the SkyProvision element management system (EMS). Note that the SkyProvision EMS supports all configuration parameters, while the CLI supports a subset. Use of the SkyProvision EMS is recommended, therefore, to utilize the full potential of the SkyPilot system.

All SkyPilot nodes have an SNMP Agent that can be used to monitor link status and traffic statistics, and to query the configuration. The SNMP protocol, along with Telnet used for remote CLI access, run on top of an IP stack. Therefore each SkyPilot node requires the allocation of an IP address. This address can be configured manually via the CLI, or more typically, allocated via DHCP.

A DHCP server is bundled with the SkyProvision EMS application. It is recommended that this server be used for the configuration of the SkyPilot management IP addresses. This DHCP server may also optionally be used for subscriber IP address allocation. When using the SkyProvision EMS, the following configuration sequence occurs once a node has successfully completed its link layer authentication (see ***Encryption and Authentication***):

1. The node issues a DHCP_REQUEST message to obtain a management IP address.
2. The DHCP_RESPONSE contains both the allocated IP address and the IP address of the SkyProvision server application.
3. The node utilizes HTTP to retrieve an XML format configuration file from the server.
4. The XML file is parsed and all configuration parameters applied. The parameters include Ethernet filters, QoS settings, Ethernet port state (which is used to control the subscriber's provisioned state) and any VLAN settings.
5. The XML configuration file also defines the primary and back-up software versions. If the versions specified do *not* match those in memory, then the node initiates an FTP Get to retrieve the new software image.

CLI

The Command Line Interface (CLI) gives an operator the ability to manage and monitor a SkyPilot node locally or remotely. The CLI supports a subset of the management parameters available through the configuration file. This subset is defined in the document titled "***Managing Your SkyPilot Network***".

The CLI may be used to configure a small network during testing or early deployment, and *must* be used to configure any SkyGateway Management VLAN ID. If the SkyProvision EMS is utilized, any configuration completed via the CLI *must* also be made in the SkyProvision database, otherwise the intended configuration will be lost on the next reboot of the node.

SNMP

Each SkyPilot node has an SNMP Agent. This agent allows an operator to query system configuration, and to monitor status and statistics. SNMP provides a predominately read-only interface with only the **reload** and **reboot** write attributes. Reload causes the configuration file to be downloaded and any changes in provisioning parameters to be completed. Reboot will cause the node to be reset and reconnect with the network.

The SkyPilot nodes support standard MIB-II, EtherLike and Bridge MIBs, along with a proprietary SkyPilot MIB. To obtain a copy of the SkyPilot MIB please contact SkyPilot Customer Support.

Software Download

The SkyPilot nodes store two software images in non-volatile flash memory. One of the images is designed as *active*; the other as *standby*. The node always runs the active image. A software download is invoked by changing the revision of the standby image. Once that image is successfully downloaded it can be activated by switching the active and standby revions.

SkyPilot™ NETWORKS

2055 Laurelwood Road
Santa Clara, CA 95054-2747
408.764.8000
US Toll Free 866 SKYPILOT
sales@skypilot.com