



CHAPTER 6

RFID Tag Considerations

This chapter has the following main sections:

- [RFID Tag Technology, page 6-1](#)
- [Using Wi-Fi RFID Tags with the Cisco UWN, page 6-15](#)
- [Tag Telemetry and Notification Considerations, page 6-27](#)
- [Chokepoint Considerations, page 6-31](#)

RFID Tag Technology

The majority of RFID tags produced today are *passive* RFID tags, comprised basically of a micro-circuit and an antenna. They are referred to as passive tags because the only time at which they are actively communicating is when they are within relatively close proximity of a passive RFID tag reader or *interrogator*.

Another type of common RFID tag in the marketplace today is known as the *active* RFID tag, which usually contains a battery that directly powers RF communication. This onboard power source allows an active RFID tag to transmit information about itself at great range, either by constantly *beaconing* this information to a RFID tag reader or by transmitting only when it is prompted to do so. Active tags are usually larger in size and can contain substantially more information (because of higher amounts of memory) than do pure passive tag designs. The tables shown in [Figure 6-1](#) provide a quick reference of common comparisons between active and passive RFID tags. Within these basic categories of RFID tags can be found subcategories such as *semi-passive* RFID tags.



Note

The terms *beacon* and *beaconing* have been used in the RFID industry for some time, predating the establishment of the formal 802.11 standards. When an active RFID tag periodically beacons, it is simply transmitting a tag message (much like any other messages the tag might send) at a set interval. Despite the use of similar terminology, this should *not* be confused with an 802.11 Beacon. An 802.11 Beacon is a management frame that the 802.11 access point (or the beacon sender in an IBSS) transmits to provide time synchronization and PHY-specific parameters in order to facilitate mobile stations locating and identifying a BSS or IBSS.

Figure 6-1 Active and Passive RFID Comparison

	Active RFID	Passive RFID
Tag Power Source	Internal to tag	Energy transferred from the reader via RF
Tag Battery	Yes	No
Availability of Tag Power	Continuous	Only within field of reader
Required Signal Strength from Reader to Tag	Very Low	Very High (must power the tag)
Available Signal Strength from Tag to Reader	High	Very Low

	Active RFID	Passive RFID
Communication Range	Long range (100m or more)	Short or very short range (3m or less)
Sensor Capability	Ability to continuously monitor and record sensor input; data/time stamp for sensor events	Ability to read and transfer sensor values only when tag is powered by reader; no date/time stamp
Data Storage	Large read/write data storage (128KB) with sophisticated data search and access capabilities available	Small read/write data storage (e.g. 128 bytes)

190588

Recent market developments have brought yet another category of RFID tag into the spotlight. Known as hybrid or *multimode* tags, these combine several different tag technologies into a versatile package that can be tracked by one or more location technologies. Multimode RFID tags are typically low power, small form factor devices that allow a single physical tag to assume multiple personalities and perform tasks that previously would have required several individual physical tags to be attached to the asset. A multimode tag, for example, may combine multiple active tag subcategories along with a passive tag into a single homogenous product.

Passive RFID Tags

Passive RFID tags typically do not possess an onboard source of power. Instead, the passive RFID tag receives its power from the energizing electromagnetic field of an RFID reader (or interrogator). The energy coupled from the electromagnetic field undergoes rectification and voltage multiplication in order to allow it to be used to power the passive tag's microelectronics. In the typical passive RFID tag design, the tag cannot communicate with host applications unless it is within the range of an RFID reader.

Interrogators come in many forms, with two common examples being handheld reader-interrogators (shown on the left in Figure 6-2) and large stationary models capable of reading many tags simultaneously as they pass (shown in the center of Figure 6-2). Embedded sub-miniature passive RFID readers and tags (shown on the right in Figure 6-2) can be used in applications requiring immediate action verification. Examples of this might include immediate verification of proper supply-line hose connections. In these types of applications, passive RFID tags and microreaders embedded into hose plugs and receptacles ensure that the proper supply hoses are connected to the proper material sources at all times. Should an incorrect connection be made, the mismatch is detected and the system refuses to open an electromagnetic flow control.


Figure 6-2 *Passive RFID Interrogators*



190589

Passive RFID tags (shown in Figure 6-3) consist of a coil and a microcircuit that includes basic modulation circuitry, an antenna, and non-volatile memory.

Figure 6-3 *Passive RFID Tags*



190590


Passive RFID tags vary in how they communicate data to RFID readers and how they receive power from the RFID reader's inductive or electromagnetic field. This is commonly performed via two basic methods:

- Load modulation and inductive coupling in the near field—In this approach (see Figure 6-4), the RFID reader provides a short-range alternating current magnetic field that the passive RFID tag uses for both power and as a communication medium. Via a technique known as *inductive (or near-field) coupling*¹, this magnetic field induces a voltage in the antenna coil of the RFID tag, which in turn powers the tag. The tag transmits its information to the RFID reader by taking advantage of the fact that each time the tag draws energy from the RFID reader's magnetic field, the RFID reader itself can detect a corresponding voltage drop across its antenna leads. Capitalizing on this phenomenon, the tag can communicate binary information to the reader by switching ON and OFF a load resistor to perform *load modulation*. When the tag performs load modulation, the RFID reader detects this action as amplitude modulation of the signal voltage at the reader's antenna. Load modulation and inductive coupling can be found among passive RFID tags using frequencies from 125 to 135 kHz and 13.56 MHz. Limitations that exist with regard to the use of such low frequencies include the

1. A technique based on Faraday's principle of magnetic induction.

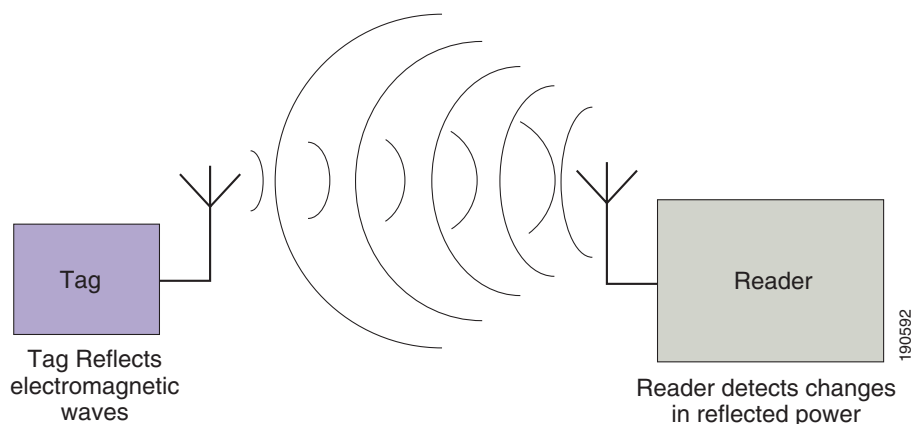
necessity to use larger antennas, low data rate and bandwidth and a rather dramatic decay in the strength of the electromagnetic field ($1/r^6$), where r represents the distance between a low frequency interrogator and a passive RFID tag.

Figure 6-4 *Passive Tag Load Modulation*



- Backscatter modulation and electromagnetic coupling in the far field—In this approach (shown in [Figure 6-5](#)), the RFID reader provides a medium-range electromagnetic field that the passive RFID tag uses for both power and a communication medium. Via a technique known as *electromagnetic (or far-field) coupling*, the passive RFID tag draws energy from the electromagnetic field of the RFID reader. However, the energy contained in the incoming electromagnetic field is partially reflected back to the RFID reader by the passive tag antenna. The precise characteristics of this reflection depend on the load (resistance) connected to the antenna. The tag varies the size of the load that is placed in parallel with the antenna in order to apply amplitude modulation to the reflected electromagnetic waves, thereby enabling it to communicate information payloads back to the RFID reader via *backscatter modulation*. Tags using backscatter modulation and electromagnetic coupling typically provide longer range than inductively coupled tags, and can be found most commonly among passive RFID tags operating at 868 MHz and higher frequencies. Far field coupled tags typically provide significantly longer range than inductively coupled tags, principally due to the much slower rate of attenuation ($1/r^2$) associated with the electromagnetic far-field. Antennas used for tag employing far field coupling are typically smaller than their inductively coupled counterparts.

Figure 6-5 *Passive Tag Backscatter Modulation*



Note that neither of these two techniques allows passive RFID tags to communicate *directly* with 802.11 infrastructure access points. All communication from the passive RFID tag occurs via the RFID reader.

Passive RFID tags are less costly to manufacture than active RFID tags and require almost zero maintenance. These traits of long-life and low-cost make passive RFID tags attractive to retailers and manufacturers for unit, case, and pallet-level tagging in *open-loop* supply chains. Open-loop supply chains typically allow little to no regulation of whether RFID tags leave the control of the tag owner or originator. Because of their dependence on external reader energy fields and their low reflected power output, passive RFID tags have a much shorter read range (from a few inches for tags using load modulation up to a few meters for those using backscatter modulation) as well as lower read reliability when compared to active RFID tags.

The passive RFID tag is available commercially packaged in a wide variety of designs, from mounting on a simple substrate to creating a classic “hard” tag sandwiched between adhesive and paper (commonly referred to as an RFID “smart” label). The form factor used depends primarily on the application intended for the passive RFID tag and can represent the bulk of the passive RFID tag cost.


Semi-Passive RFID Tags

Semi-passive RFID tags overcome two key disadvantages of pure passive RFID tag designs:

- The lack of a continuous source of power for onboard telemetry and sensor asset monitoring circuits.
- Short range.

Semi-passive tags differ from passive tags in that they use an onboard battery to provide power to communication and ancillary support circuits, such as temperature and shock monitoring. It is interesting to note that although they employ an onboard power source, semi-passive RFID tags do not use it to directly generate RF electromagnetic energy. Rather, these tags typically make use of backscatter modulation and reflect electromagnetic energy from the RFID reader to generate a tag response similar to that of standard passive tags (see [Figure 6-6](#)). The onboard battery is used only to provide power for telemetry and backscatter enabling circuits on the tag, not to generate RF energy directly.

Figure 6-6 Backscatter Modulation in Semi-Passive RFID Tags



Semi-passive RFID tags operating in the ISM band (shown in [Figure 6-7](#)) can have a range of up to 30 meters with onboard lithium cell batteries lasting several years. Range is vastly improved over conventional passive RFID tags primarily because of the use of a backscatter-optimized antenna in the semi-passive design. Unlike a conventional backscatter-modulated passive RFID tag, the antenna contained in a semi-passive tag is dedicated to backscatter modulation and there is no dependence on the semi-passive RFID tag antenna to be a reliable conduit of power for the tag. Therefore, the semi-passive tag antenna can be optimized to make most efficient use of the backscatter technique and provide far better performance than purely passive RFID tag antenna designs.

Figure 6-7 Semi-Passive RFID Tags



Several varieties of semi-passive RFID tags exist, with and without onboard NVRAM, real time clocks, and various types of environmental sensors. Semi-passive RFID tags also support interfaces to tamper indicators, shock sensors, and so on. Common applications of semi-passive RFID tags include but are not limited to vehicle asset tracking, security access systems, supply chain automation, cold storage management, and hierarchical asset tracking systems.

Active RFID Tags

Active tags are typically used in real-time tracking of high-value assets in *closed-loop* systems (that is, systems in which the tags are not intended to physically leave the control premises of the tag owner or originator). Higher value assets can usually justify the higher cost of the active tag, and presents strong motivation for tag reuse. Medical equipment, electronic test gear, computer equipment, reusable shipping containers, and assembly line material-in-process are all excellent examples of applications for active tag technology. Active RFID tags (see [Figure 6-8](#)) can provide tracking in terms of *presence* (positive or negative indication of whether an asset is present in a particular area) or real-time location. Active RFID tags are usually physically larger than passive RFID tags. Most RTLS systems are based on the use of active RFID tag technology.

Figure 6-8 Active RFID Tags



Active tags can contain 512 KB or more of RAM, which enables the active tag to store information from attached assets for transmission at the next beacon interval or when polled. This large memory capacity also makes active RFID preferable to passive RFID in situations when the RFID tag cannot simply be used as a “license plate” or reference, to enable an immediate lookup in a host database. A good example of this might be a remote military installation where a host database may or may not be available at all times. By storing critical asset data directly on the tag itself, this information can be retrieved directly from the tag and used regardless of the availability of the host system.

Active RFID tags can be found operating at frequencies including 303, 315, 418, 433, 868, 915, and 2400 MHz with read ranges of 60 to 300 feet. Active RFID tag technology typically display very high read rates and read reliability because of their higher transmitter output, optimized antenna, and reliable source of onboard power. Active RFID tag cost can vary significantly depending on the amount of memory, the battery life required, and whether the tag includes added value features such as onboard temperature sensors, motion detection, or telemetry interfaces. The durability of the tag housing also affects price, with the more durable or specialized housings required for specific tag applications coming at increased cost. As with most electronic components of this nature, prices for active tags can be expected to decline as technological advances, production efficiencies, and product commoditization all exert a downward influence on market pricing.

Beaconing Active RFID Tags

Beaconing active RFID tags are used in many RTLS systems and are primarily useful when the location of an asset needs to be tracked anywhere and anytime via the use of location receivers. With a beaconing active RFID tag, a short message payload containing the unique identifier of the RFID tag is emitted at pre-programmed intervals. This interval is programmed into the tag by the tag owner or user, and it can be set appropriately depending on how often tag RSSI updates are required. A shorter tag transmission interval typically results in shorter tag battery life but may improve tag location accuracy in some cases, since tag RSSI is reported more often. Longer tag transmission intervals increase tag battery life but as tag RSSI is reported less often, the frequency of location update will be less.

802.11 Active RFID Tags

802.11 (Wi-Fi) active RFID tags (shown in Figure 6-9) are designed to operate in the unlicensed ISM bands of 2.4 to 2.4835 GHz or 5.8 to 5.825 GHz. Currently manufactured 802.11 Wi-Fi active RFID tags available at publication are limited to 2.4 GHz.

These tags exhibit the characteristics of active RFID tags, but also comply with applicable IEEE 802.11 standards and protocols. Wi-Fi RFID tags can readily communicate directly with standard Wi-Fi infrastructure without any special hardware or firmware modifications and can co-exist alongside Wi-Fi clients such as laptops, VoWLAN phones, and so on. When powered on, assets equipped with 802.11

Wi-Fi client radios can be tracked natively without the need to have an asset tag attached. Other assets lacking an internal 802.11 Wi-Fi client radio can be tracked via a physically attached 802.11 active RFID tag. A physically attached 802.11 active RFID tag also makes it possible to use the location-aware Cisco UWN to track assets with integrated Wi-Fi client radios when those radios are powered off.

Figure 6-9 802.11 Wi-Fi Active RFID Tags



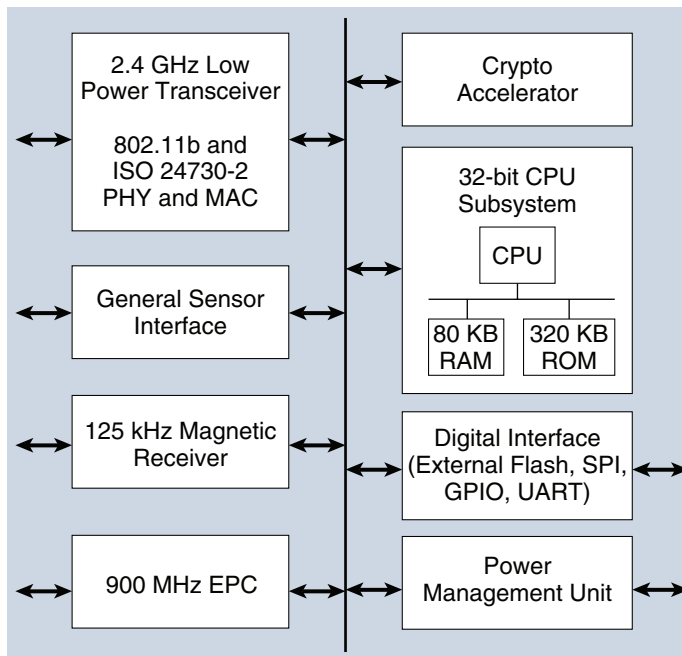
Multimode RFID Tags

As mentioned previously, transponder active RFID tags offer the combination of a primary tag operational mode with a secondary method of communication that can be used for a plethora of added value functions, such as activation, deactivation, behavior modification and so on. This type of tag has been used for quite some time in highway toll plaza applications, for example, where tags are triggered to transmit when in proximity of high speed activators, thereby triggering a debit to the user's account for the toll charge.

A relatively new development has been the introduction of *multimode* RFID tags that leverage multiple location technologies. Multimode tags offer the functional equivalent of having assets equipped with several individual tags in one physical package. This can be very useful when assets must travel outside of a single enterprise closed loop system into other systems, where the same type of location tracking technology may not be in use. For example, consider the case where reusable shipping containers must be tracked at a manufacturer, a distributor and a retailer using a combination of ISO24730-2 TDoA, 802.11 Wi-Fi Active RFID and passive RFID. A multimode tag could offer all three of these technologies in a single small form factor, low power draw package. Such a device may also include the capability to use tag magnetic signaling proximity communication devices as well. This can offer distinct advantages in terms of management, maintenance and overall ease of deployment, especially when compared to equipping assets with three or more physically separate RFID tags.

Multimode tags of this nature have been made much more feasible by the availability of highly integrated tag OEM silicon that combines two or more distinct RFID tag technologies into a single chip or chipset. This is exemplified by the G2C501 from G2 Microsystems (shown in Figure 6-10), which is a complete Wi-Fi system-on-chip (SoC) that includes 802.11b Wi-Fi active RFID, 900 MHz EPC Global Gen 1 Class 0 passive RFID, 2.4 GHz ISO24730-2 TDoA, a 32-bit CPU, crypto accelerator, real-time clock and sensor interfaces.

Figure 6-10 G2C501 RFID System-On-A-Chip (SoC)



The use of highly integrated tag silicon offers many advantages to the tag vendor, including:

- Small form factor
- Low power consumption
- Well documented software and hardware interfaces
- Flexible support for multiple location technologies

A good example of a multimode tag that capitalizes on such capabilities is the WhereNet IV asset tag from WhereNet Corporation (<http://www.wherenet.com>), shown in the lower left hand quadrant of Figure 6-9. The WhereNet IV combines a Cisco Compatible Extensions compliant 802.11 Wi-Fi active tag implementation along with 125kHz magnetic signaling and ISO 24730-2 capabilities in a small, highly integrated design.

Chokepoint Triggers

Chokepoint triggers are proximity communication devices that trigger asset tags to alter their configuration or behavior when the asset tag enters the chokepoint trigger's area of operation. This alteration could be as simple as causing the asset tag to transmit its unique identifier, or more complex, including causing the tag to change its internal configuration or status. One of the prime functions of a chokepoint trigger is to stimulate the asset tag such that it provides indication to the RTLS that the tag


has entered (or exited) the confines of an area known as a *chokepoint*. Chokepoints are tightly defined physical areas (such as entrances, exits or other types of constrictions) that provide passage between connected regions. Figure 6-11 illustrates some common examples of chokepoints.



Note

While *chokepoint triggers* are typically deployed within *chokepoints*, it is often commonplace to hear the term *chokepoint* used to refer to a chokepoint trigger.

Figure 6-11 Common Chokepoint Areas



Outdoor chokepoint locations may include a fenced gate, bridge, toll plaza, or similar passageway. Indoor chokepoint locations includes connecting entrances or exits between:

- A building's interior rooms or floors such as doorways, ramps, gates, stairwells, elevator entrances, and so on).
- Adjacent structures (such as passageways or tunnels) or the interior and exterior of structures (main and auxiliary entrances, loading docks, fire exits, and so on).

Chokepoint triggers can initiate behavioral changes in tags that can immediately alert the location system that the tagged asset has entered or exited the chokepoint area. Due to the comparatively modest range of chokepoint triggers in relation to the overall area covered by an RTLS, the RTLS is able to deterministically localize the asset to the confines of the chokepoint area relatively quickly and with excellent reliability. In addition to displaying the chokepoint area on floor maps, the RTLS can use the

detection of assets within chokepoints to trigger events in external systems. These can include database updates, notification alerts, or alarms. When properly augmented by appropriate application software, chokepoint applications may include:

- Tracking of high value assets—Chokepoint location tracking can help ensure that valuable assets intended for a particular area stay within such areas. If these assets are detected as being removed via entrances or exits, for example, the RTLS is alerted.
- Manufacturing process control—Equipment, parts, and finished products can be precisely tracked as they move between the various production stations. This helps ensure not only that all required process stations are visited, but that they are visited in the proper sequence.
- Inventory control—By strategically equipping all distribution center entrances and exits with chokepoint location tracking capabilities, inventory databases can be automatically updated as product enters or leaves the distribution center.
- Security—The movement of tagged assets can be tracked and monitored to protect against unauthorized removal from the premises or unauthorized movement within the facility itself structure.

Low power, short range chokepoint triggers make it possible to expand usage beyond traditional entry and exit passages. Low output power enables customization of the chokepoint trigger's effective range to better correspond to very small, tightly defined areas such as shelves, racks, storage bins, workstations and patient beds. The movement of assets into or away from such limited areas can be then be precisely monitored (such as the placement or removal of equipment in a rack, for example) in a similar fashion to that of the higher power chokepoint triggers described earlier.

The specific changes in tag behavior that can be enacted by a chokepoint are vendor dependent. Tag behavior modification may include, but are not limited to:

- Immediate tag multicast message transmission
- Tag reactivation
- Tag deactivation
- Tag transmission interval change
- Indicator lamp activation
- Storage of floor or cell identifiers
- Appending of additional messages to tag multicast messages, such as:
 - Chokepoint identification
 - Pre-configured message data
 - Telemetry data

Not every active tag vendor supports the use of chokepoint triggers with their tags. Of those that do, the use of chokepoint triggers tends to be tag vendor specific. Each vendor offering asset tags that are compliant with the Cisco Compatible Extensions for Wi-Fi Tags specification usually supplies chokepoint triggers that are designed specifically for compatibility with those tags. At the current time, chokepoint triggers are not interoperable between asset tags from different manufacturers.

Range may vary between models and manufacturers, with those chokepoint triggers used with asset tags compliant with the Cisco Compatible Extensions for Wi-Fi tags specification typically possessing effective ranges between 10 inches and approximately 25 feet. These products operate using low frequency magnetic signaling. Range tends to be predictable, with excellent penetration of typical building materials and their contents.

Figure 6-12 depicts low frequency, magnetic signaling-based chokepoint trigger devices from AeroScout and WhereNet. AeroScout refers to their chokepoint triggers as *Exciters* and WhereNet refers to their products as *WherePorts*. The AeroScout EX-2000 Exciter and the WhereNet WherePort products are larger footprint models, capable of providing the maximum possible range for large chokepoint areas or room-based presence detection applications. These products are intended for vehicular doorways, gates and other large chokepoint areas, with adjustable ranges that can exceed 20 feet. The compact AeroScout EX-3100 and EX-3200 Exciters are intended for short range use in smaller chokepoints such as doorways, shelves and racks. The range of these products spans from 8 inches to a maximum of 6.5 and 9.75 feet, respectively.

Figure 6-12 AeroScout Exciters and WhereNet WherePorts



Additional information on these products can be found at the following vendor web sites:

<http://www.aeroscout.com/content.asp?page=exciter>

http://www.wherenet.com/products_whereport.shtml



Note

The Cisco WCS is used to define chokepoint triggers to the location-aware Cisco UWN, but cannot be used to configure the chokepoint triggers themselves at this time. This must be accomplished using software provided by the vendor of the chokepoint trigger (the AeroScout Network Exciter Manager (ANEM) and the WhereNet SystemBuilder / WhereWand are two examples). Chokepoint triggers that have been added to WCS without proper configuration by the vendor's chokepoint management software may not function properly.

Once configured, chokepoint triggers can operate in one of two modes:

- An *online mode*, where their status is monitored by software supplied by the chokepoint trigger vendor via an Ethernet or serial data connection.
- An *offline mode*, where the configured chokepoint trigger operates with only a power connection required.

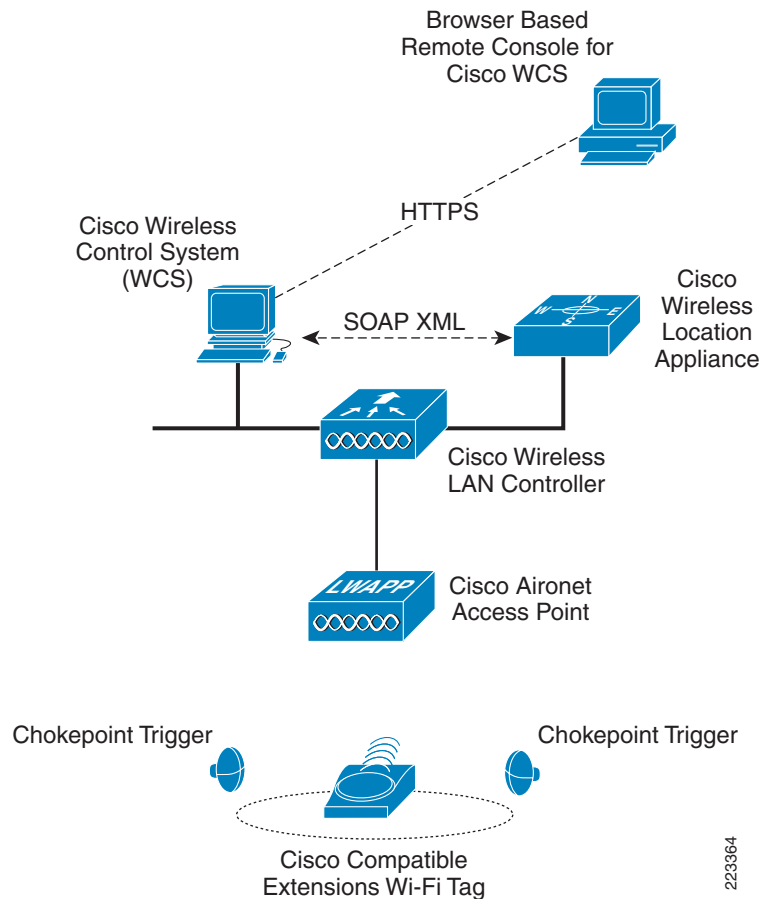
Chokepoint triggers are identified by unique addresses that enables tags receiving their transmission to clearly identify the chokepoint trigger responsible for stimulating them. This identifier is typically the MAC address of the chokepoint trigger for Ethernet-based models, but could be any locally administered and assigned identifier (such as a “Transmit ID” of a WhereNet WherePort). In Release 4.1 of the location-aware Cisco UWN (shown in Figure 6-13), when an asset tag compatible with the Cisco

Compatible Extensions for Wi-Fi Tags specification enters the effective range of a chokepoint trigger, the tag is stimulated by the chokepoint trigger and identifies the source of such stimulation to the location-aware Cisco UWN using a tag multicast frame that is sent via using 802.11. All access points detecting this tag multicast frame forwards it to their registered controller, which in turn results in the generation of LOCP Measurement Notification frames destined for the location appliance.

**Note**

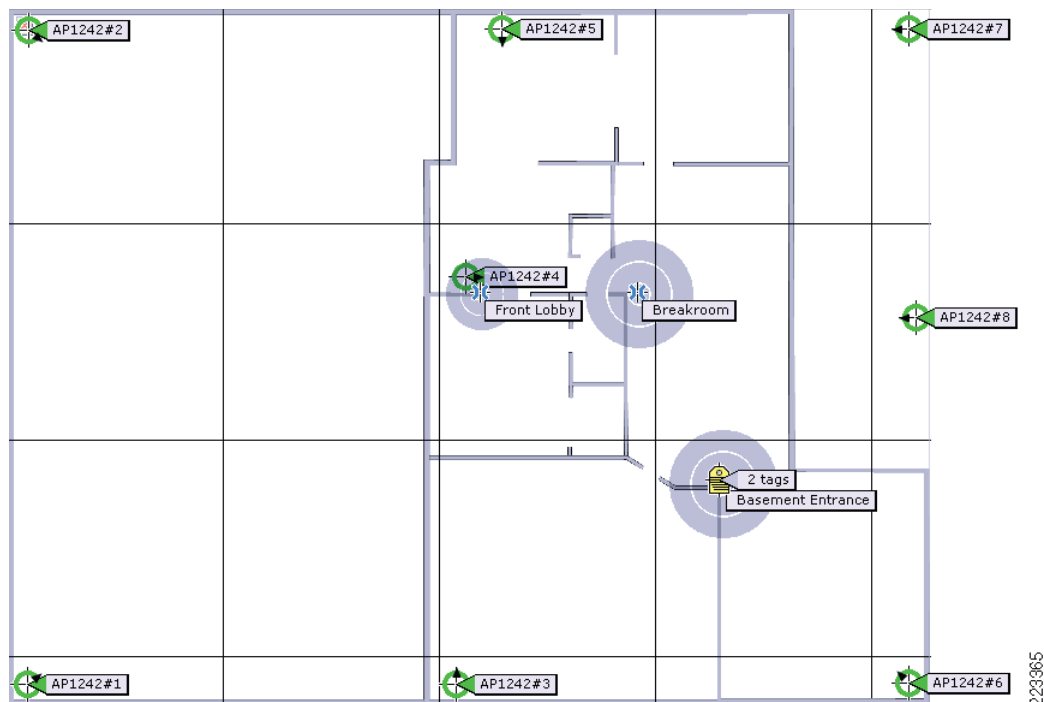
Communication between chokepoint triggers and asset tags is unidirectional, from the chokepoint trigger to the asset tag. In software Release 4.1, there is no direct communication between chokepoint triggers and the location-aware Cisco UWN.

Figure 6-13 Location-Aware Cisco UWN with Chokepoint Triggers



The location appliance uses the information provided to it by the LOCP Measurement Notification to indicate that the tag's current location is within the configured range of the specified chokepoint. This information is placed in the appropriate location appliance databases and made available to location clients via the location appliance API. Location clients may display chokepoint location information on floor maps. An example is the WCS floor map shown in [Figure 6-14](#), where we can see two RFID tags located at the chokepoint labeled *Basement Entrance*). The location appliance can also trigger alerts and other asynchronous northbound notifications to WCS and external applications using email, syslog, SOAP, or SNMP traps.

Figure 6-14 WCS Floor Map With Chokepoints



In Release 4.1 of the Cisco UWN software, after a tag has left the range of a chokepoint trigger, the location appliance continues to indicate the tag's location as being within the configured range of the chokepoint trigger until one of the following events occur:

- The tag indicates that it is now out of range of that chokepoint trigger.
- The value configured for the Chokepoint Out of Range Timeout expires (shown in [Figure 6-15](#), default 60 seconds).

After one of these events occur, the location appliance uses RF Fingerprinting to calculate the location of the device until such point that it enters into another chokepoint area and into the stimulation zone of another chokepoint trigger. If the device is then stimulated by a subsequent chokepoint trigger and successfully reports this stimulation to the Cisco UWN, the location appliance then places the tracked device at the location of the new chokepoint.

Figure 6-15 Chokepoint Out of Range Timeout

Location Server > Location Parameters > 'AeS_Loc2'

Location Parameters

Enable calculation time ?	<input type="checkbox"/> Enable
Enable OW Location ?	<input type="checkbox"/> Enable
Relative discard RSSI time ?	<input type="text" value="3"/> minutes.
Absolute discard RSSI time ?	<input type="text" value="60"/> minutes.
RSSI Cutoff ?	<input type="text" value="-75"/> dBm
Smooth Location Positions ?	<input type="text" value="More smoothing (new value weighted less)"/>
Chokepoint Usage ?	<input checked="" type="checkbox"/> Enable
Chokepoint Out of Range Timeout ?	<input type="text" value="60"/> seconds.

223366

Using Wi-Fi RFID Tags with the Cisco UWN

Compatible RFID Tags

An often asked question revolves around whether the Cisco Location Appliance can be leveraged to track RFID tags that already are being deployed by product and durable goods manufacturers as part of a larger business initiative. Often applied *en masse* to manufactured or distributed goods, these tags are most commonly passive RFID designs, but in the case of some durable high-cost goods, active RFID may also be used. In many cases, products and goods are being tagged at the time of production or initial distribution in compliance with mandates set forth by large commercial or governmental entities.

The answer depends on the type of RFID tag being used. As of Cisco UWN software Release 4.1, only 802.11 Wi-Fi active RFID tags (or multimode asset tags containing 802.11 Wi-Fi active RFID capabilities) can communicate directly with Wi-Fi access points (including Cisco Wi-Fi access points). At this time, most commonly available “pure” passive RFID tags or non-Wi-Fi active RFID tags are not capable of communicating with the location-aware Cisco UWN and the Cisco Wireless Location Appliance. Of the available 802.11 Wi-Fi active tag designs currently on the market, not all are compliant with the Cisco Compatible Extensions for Wi-Fi Tags specification. Non-compliant asset tags from PanGo / InnerWireless and AeroScout Ltd. can be recognized by the location-aware Cisco UWN. However, these tags will not be able to make use of the advanced features in the Cisco Compatible Extensions for Wi-Fi Tags specification and introduced in Release 4.1. Non-compliant asset tags from vendors other than PanGo Networks and AeroScout are not supported for use with the Cisco Wireless Location Appliance.

To determine whether a Wi-Fi active RFID tag is compatible with the Cisco Compatible Extensions for Wi-Fi Tags specification and capable of taking advantage of the advanced features of the location-aware Cisco UWN, the Cisco Compatible Extensions website (http://www.cisco.com/web/partners/pr46/pr147/ccx_wifi_tags.html) should be consulted. A current listing of all tags and tag vendors compatible with the Cisco Compatible Extensions for Wi-Fi Tags specification may be found there.

The listing of the tag and tag vendor on the Cisco Compatible Extensions website indicate that the asset tag has passed stringent validation testing as part of the Cisco Compatible Extensions Program for Wi-Fi tags. The Cisco Compatible Extensions program for Wi-Fi tags allows customers with a location-aware Cisco Unified Wireless Network to benefit from the latest innovation and technology advancements offered by Cisco's technology partners. Registered channel partners may view the guidelines for the Cisco Compatible Extensions Program for Wi-Fi Tags at the following URL:

<http://www.cisco.com/web/partners/downloads/partner/WWChannels/download/wifiguide.pdf>.

In some cases, passive or non-802.11 active RFID reader interrogators may be deployed in an environment that is also serviced by a Cisco LWAPP-enabled wireless network, independently of the location tracking capabilities of the Cisco UWN and the location appliance. These reader/interrogators may be using traditional wired Ethernet as their uplink to the network, or they may have an integrated Wi-Fi client radio (such as the case of portable RFID interrogators like those shown in Figure 6-16). Although it is not possible at this time to track the individual passive RFID tags associated with these portable RFID tag readers using the Cisco location appliance, tracking the portable readers themselves is typically feasible because of their use of industry standard 802.11 client radios. As long as these readers act as standard WLAN clients and authenticate/associate to WLAN SSIDs serviced by controllers defined to the location appliance, they are treated just as other WLAN clients and are indicated on floor maps by a blue rectangular icon.

Figure 6-16 Portable RFID Interrogators with Integrated Wi-Fi Uplink



Using 802.11b Tags in an 802.11g Environment

Another common question that often arises is about the potential performance impact of using an 802.11b asset tag in a network that otherwise consists entirely of 802.11g clients and access points. The crux of such discussions is typically centered around whether or not protection mechanisms (such as RTS-CTS or CTS-to-self) are initiated by the 802.11g network to assure compatibility between the 802.11b asset tags and the 802.11g network.



Note

For an explanation of 802.11g performance, capacity, and protection mechanisms, see the whitepaper entitled *Capacity, Coverage and Deployment Considerations for IEEE 802.11g* at the following URL: http://www.cisco.com/en/US/products/hw/wireless/ps430/products_white_paper09186a00801d61a3.shtml.

A popular point of discussion often revolves around whether these protection mechanisms are initiated upon the introduction of one or more of the following to the all-802.11g wireless infrastructure:

- An 802.11b asset tag that is transmitting tag layer two multicast messages.
- An 802.11b asset tag (acting as a WLAN client) that is issuing probe requests.
- An 802.11b asset tag (acting as a WLAN client) that actively associates.

First and foremost, it should be clearly understood that 802.11b asset tags that transmit tag messages using Layer 2 multicasts (and do not attempt to associate to any WLANs) will *not* cause the initiation of any 802.11g protection modes under any circumstances. This includes asset tags operating in strict compliance with version 1 of the Cisco Compatible Extensions for Wi-Fi tags specification.

Laboratory research and analysis have shown that protection mechanisms are not initiated throughout an entire network of access points if an 802.11b asset tag or WLAN client is simply powered on. In fact, the following are observed:

- A probe request from an 802.11b asset tag that is *not associated* to any access point on a particular channel does not in and of itself cause the initiation of protection mode by an 802.11g access point that detects it.
- Protection mode is not initiated until the 802.11b asset tag successfully associates to either the cell in question or an adjacent cell on the same channel. At that point, the target cell as well as any other cells on the same channel and RF-adjacent to the target cell initiate protection mode.
- Access points that are not on the same channel as the 802.11b asset tag or not RF-adjacent to it does not initiate protection mode.

Some 802.11b asset tags may, as an optional feature, periodically probe and attempt to briefly associate to the wireless infrastructure in order to conduct over-the-air firmware or configuration updates. The observations stated above would apply to these tags, but only during the brief periods during which these extended modes of communication are in use.

Enabling Asset Tag Tracking



Note

Beginning with the Cisco UWN Release 4.1, it is no longer necessary to enable asset tag tracking in WLAN controllers using the **config status rfid enable** CLI command. RFID tag data collection in controllers containing Release 4.1 is now enabled by default.

Enable Asset Tag RF Data Timeout

The *RFID Data Timeout* parameter sets a static time value (in seconds) that must elapse without any access points on the controller detecting an asset tag, before that asset tag is removed from the internal tables of the controller. For general usage, it is recommended that this parameter be set to a minimum of three times (and a maximum of eight times) the longest tag transmission interval found in the general tag population. This should be inclusive of stationary as well as any “in-motion” transmission intervals. The valid range of values for this parameter is 60-7200 seconds and the default value is 1200 seconds.

For example, for a tag with a constant transmission interval of 60 seconds, you may choose to set the RFID data timeout to 480:

```
(Cisco Controller) >config rfid timeout 480
(cisco Controller) >

(Cisco Controller) >show rfid config
RFID Tag data Collection..... Enabled
```

RFID data timeout..... 480 seconds

To ensure proper collection of updated asset tag RSSI from WLAN controllers, it is recommended that the RFID data timeout always be greater than the asset tag polling interval on the location appliance, which is discussed in the next section.

Enable Asset Tag Polling

To use the location appliance for asset tag tracking, SNMP asset tag polling must be explicitly enabled via the Locate > Location Server > Polling Parameters GUI panel. To enable it, use the checkbox indicated by the red rectangle in Figure 6-17.

Figure 6-17 Enabling RFID Tag Polling


The screenshot shows the Cisco Wireless Control System interface. The breadcrumb navigation is 'Location Server > Polling Parameters >'. The 'Polling Parameters' section includes the following settings:

Enable	Polling	Interval (secs)
<input checked="" type="checkbox"/>	Client Stations	120
<input checked="" type="checkbox"/>	Rogues	600
<input type="checkbox"/>	Asset Tags	120
<input type="checkbox"/>	Statistics	900

Other settings shown include: Retry Count: 3, Timeout (secs): 5. There are 'Save' and 'Cancel' buttons at the bottom.

The default polling interval value represents the time period between the start of subsequent polling cycles in which the location appliance polls the controller using SNMP. For example, if a polling cycle requires 30 seconds to complete and the polling interval is 300 seconds, polling cycles start every 330 seconds, as shown in Figure 6-18.

Figure 6-18 Polling Interval



Depending on the degree of asset movement, updated tag RSSI information obtained via shorter polling intervals may be translated into more frequent location updates in some cases. However, depending on the time lag between the asset tag polling interval configured on the location appliance and the average transmission interval configured amongst the general tag population, a risk of reduced asset tag polling efficiency may occur. In extreme cases of deployments with a large number of WLAN controllers, a too

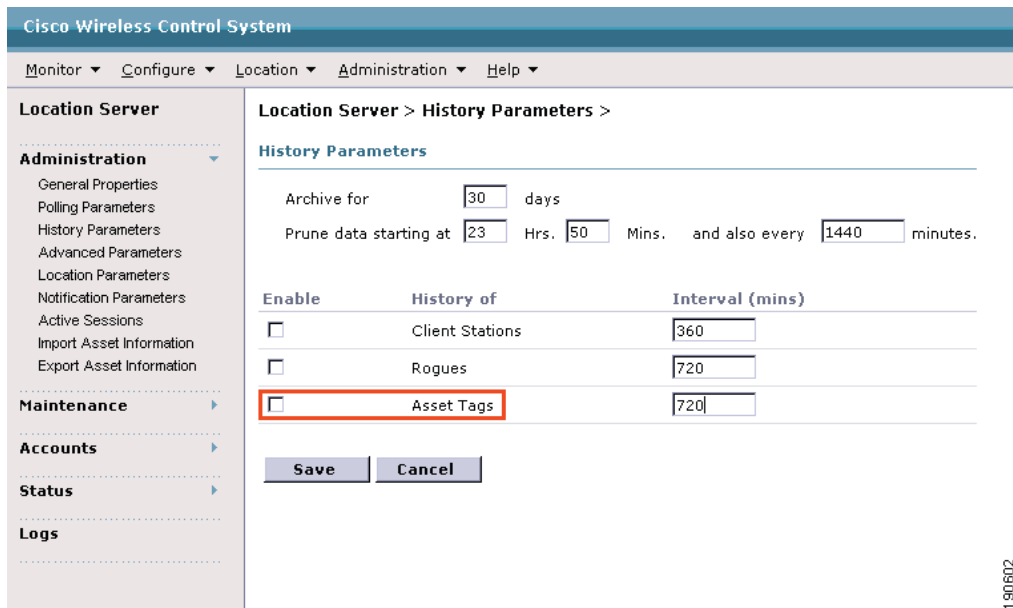
short asset tag polling interval could burden both the location appliance as well as the WLAN controllers with almost constant (and often times unproductive) polling. This wastes resources that could have been put to use more productively, and could negatively impact performance.

In general, for a given population of asset tags with the same transmission interval, the most productive and efficient polling is found to occur when the location appliance's asset tag polling interval is configured to be greater than or equal to the asset tag's transmission interval. For example, in a population of 100 asset tags each with a transmission interval of 60 seconds, if the location appliance's asset tag polling interval is left at the default of 120 seconds (twice the tag transmission interval) it is likely that controllers will receive updated RSSI from all 100 tags at least once (and most likely twice) within the 120 second time interval. Setting the asset tag polling interval to 30 seconds in an attempt to increase the frequency of tag location updates might indeed accomplish this goal for some tags, however, overall polling efficiency is likely to decline.

In a population of asset tags that are configured with mixed transmission intervals, a tradeoff typically is required between the desire to acquire frequently updated RSSI information from tags possessing the shortest transmission intervals versus overall polling efficiency for the general tag population. Shorter asset tag polling intervals can be configured to favor tags that transmit multicast frames more frequently, but depending on the number of WLAN controllers deployed, asset tag polling intervals should not be set so short that the location appliance is spending the bulk of its time constantly polling controllers, which could impact performance in an environment with many controllers present. Remember that the speed at which location updates are displayed on location client screens depends not only on the frequency of updates between controllers and the location appliance, but also upon the frequency with which the location client polls the location appliance for updates.

Recording of asset tag location history is disabled by default. If location trending and the analysis of past asset tag location history is desired, location history recording should be enabled via the Location > History Parameters screen, as shown in Figure 6-19. Enable the **Asset Tags** line item and specify the history archival interval between writes of historical data to the database (default is 720 seconds). Note that the recording of location history is not mandatory to perform asset tag tracking, but is often desirable, as it allows the location appliance to “playback” the history of locations the asset tag has visited.

Figure 6-19 Enabling RFID Tag History



190602

Enable Asset Tag Display

For WCS to display the location of asset tags, asset tag display must be explicitly enabled via Monitor > Maps > Campus > Building > Floor, as shown in Figure 6-20. To enable the display of asset tags, make sure that **802.11 Tags** is selected from the dropdown Layers menu. Refresh or reload the WCS floor map page and yellow tag icons is used on the floor map to denote the current location of any detected asset tags.

Figure 6-20 Enabling Display of Asset Tags on WCS



Configuring Asset Tags

In order to communicate with the location-aware Cisco UWN, asset tags must be properly configured for parameters such as channels, transmission interval, and data formats. In this section, we examine the basic parameter settings necessary for AeroScout tags to be recognized by the UWN and properly localized.



Note

AeroScout asset tags are highlighted in this section only as an example of how to configure asset tags that are compliant with the Cisco Compatible Extensions for Wi-Fi Tags specification. Keep in mind that each vendor's asset tags require configuration using vendor-specific tools. Users of AeroScout, InnerWireless (PanGo), WhereNet, G2 or other asset tag vendors offering similar products should always consult their vendor's product documentation for appropriate configuration guidelines.

In comparison to the earlier 2.x versions of AeroScout Tag Manager, version 3.x introduces several new features designed to support AeroScout asset tags that are compliant with the Cisco Compatible Extensions for Wi-Fi Tags specification, including the recently introduced AeroScout T3 asset tags.

This section outlines the steps necessary to configure AeroScout asset tags for basic communication with the location-aware Cisco UWN. It does not attempt to serve as a substitute for the much more comprehensive vendor documentation offered by AeroScout in this regard. The following AeroScout documents should serve as the primary reference materials with regard to configuration of AeroScout asset tags using Tag Manager:

- AeroScout Tag Manager Quick Start
- AeroScout Tag Manager 3.0 User Guide

In order to take advantage of the new capabilities introduced by the Cisco Compatible Extensions for Wi-Fi Tags specification, AeroScout asset tags should contain the following tag firmware levels (see [Figure 6-21](#)):

- AeroScout T2—Firmware Release 4.3x or greater
- AeroScout T3—Firmware Release 6.0x or greater

AeroScout asset tags with firmware releases prior to those listed will still interoperate with software Release 4.1 of the location-aware Cisco UWN. However, tags not meeting these specifications will not take advantage of the capabilities introduced by the Cisco Compatible Extensions for Wi-Fi Tags specification that are present in software Release 4.1.

Figure 6-21 AeroScout T2 and T3 Asset Tags



AeroScout asset tags contain both a 2.4 GHz IEEE 802.11b transceiver as well as a low-frequency, short-range 125 kHz magnetic signaling receiver. 2.4 GHz output power is configurable up to a maximum of +19dBm (81mW). During tag configuration, AeroScout asset tags use their 802.11b interface to reply to commands and data received from a programming device known as a *Tag Activator*, which is an Ethernet addressable, low-frequency 125 kHz magnetic signaling transmitter housed in combination with a 802.11b receiver. Tag Activators are designed to be used in conjunction with Windows-based tag configuration software known as *Tag Manager*.

It is important to note that AeroScout asset tags are only capable of receiving information from Tag Activators via their magnetic signaling 125 kHz receiver. AeroScout asset tags are not equipped with a magnetic signaling transmitters, and Tag Activators are not equipped with magnetic signaling receivers. AeroScout asset tags receive commands and data from Tag Activators via magnetic signaling, and respond back to the Tag Manager application confirming those transmissions using their 802.11b capabilities and the 802.11b receiver in the Tag Activator.

The AeroScout Tag Activator (shown in [Figure 6-22](#)) can be powered via 802.3af Ethernet or an external 5VDC power source. The Tag Activator works in conjunction with AeroScout Tag Manager software to configure, program, activate, or deactivate up to 50 AeroScout asset tags simultaneously at a range of up to approximately three feet. The use of a Tag Activator is completely non-intrusive in relation to the AeroScout tag hardware. There are no cables that interconnect the two, and the use of the Tag Activator eliminates disturbing the environmental seal of the tag casing for configuration modifications. Minimal disruption of tag seals is an advantage if the asset tag is intended for use in harsh or wet environments where tight environmental sealing is required.

Figure 6-22 AeroScout Tag Activator



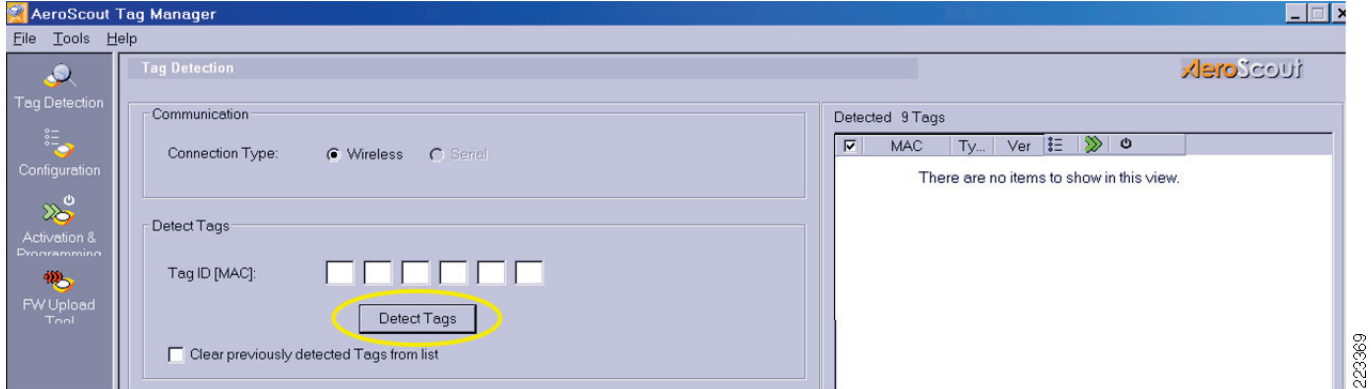
The following AeroScout document should serve as the primary reference with regard to the AeroScout Tag Activator:

- *AeroScout Tag Activator User's Guide*

In order to configure AeroScout T2 or T3 asset tags for basic communication with software Release 4.1, the following steps should be followed:

1. Deploy the AeroScout Tag Activator in accordance with the vendor's recommendations as outlined in the *AeroScout Tag Activator User's Guide*. The AeroScout tag activator may be powered directly from a 802.3af compliant switch or from a non-802.3af switch using the provided AC power supply included with the product. **Spanning tree portfast** should be configured on any Cisco switch port to which the AeroScout Tag Activator is attached to avoid potential instability.
2. Configure the AeroScout Tag Manager to communicate with the Tag Activator as per the vendor's recommendations as outlined in the *AeroScout Tag Activator User's Guide* and the *AeroScout Tag Manager version 3.0, Quick Start Guide*. Ensure that the Tag Activator is properly recognized by the Tag Manager.
3. Place up to 50 AeroScout tags within about three feet of the Tag Activator and detect the tags using the "Detect Tags" feature as shown in [Figure 6-23](#).

Figure 6-23 Detecting Tags using Tag Manager v3.04



- Once the tags have been detected (Figure 6-24), select all tags by clicking on their checkboxes, as shown in the right hand column of the screen depicted in Figure 6-25.

Figure 6-24 Successful Tag Detection using Tag Manager v3.04

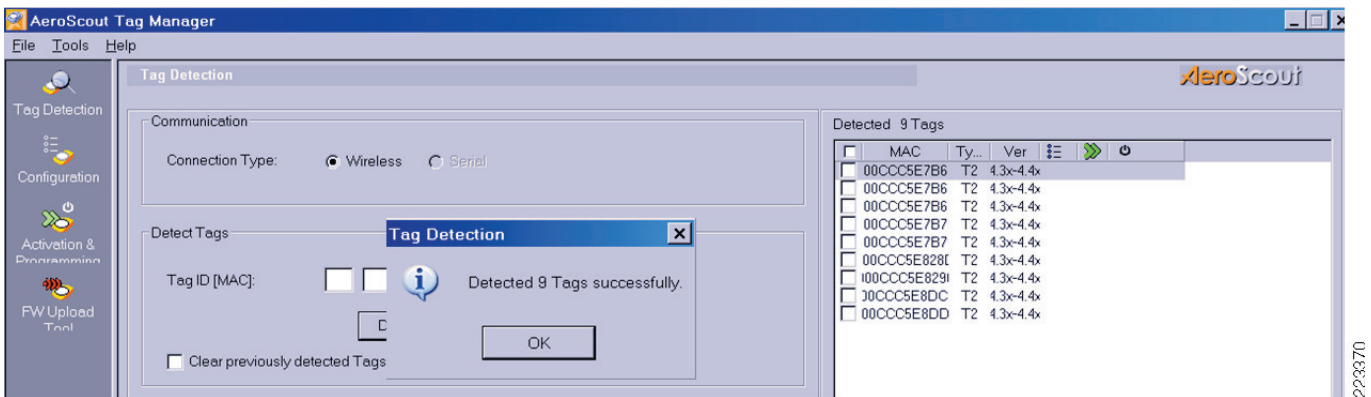
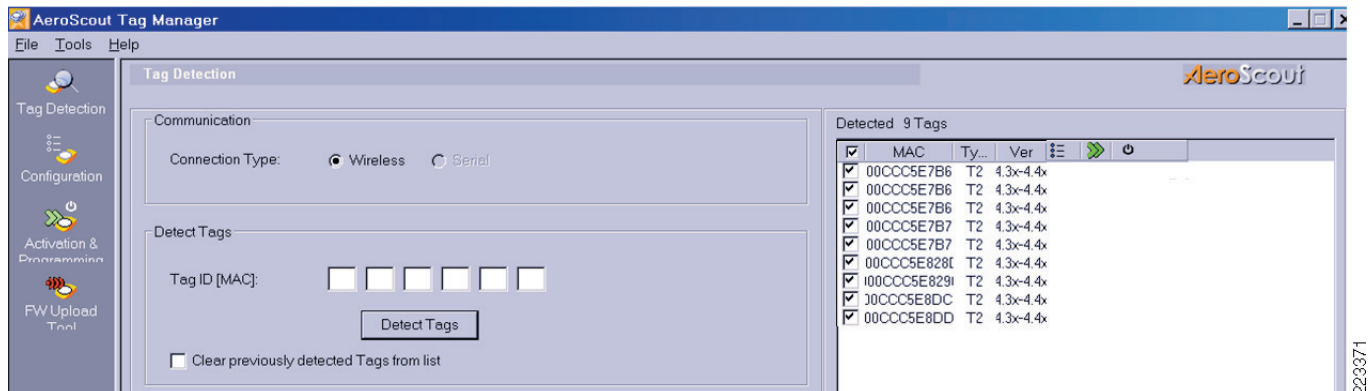


Figure 6-25 Selecting Tags to Configure

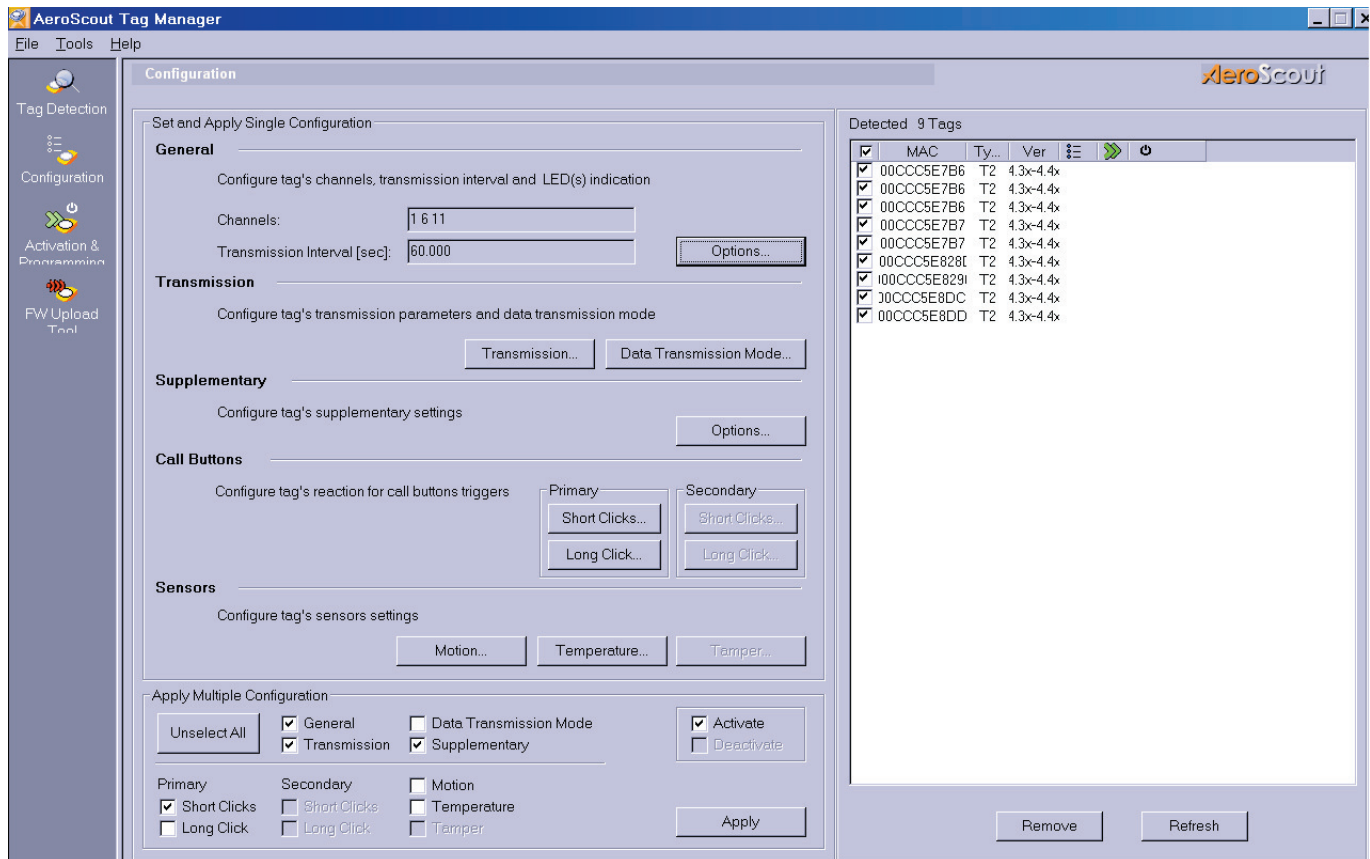


- Select the configuration option from the left hand column of the Tag Detection menu, which yields the Tag Configuration menu (shown in Figure 6-26).

**Note**

When making minor modifications to preconfigured tags, it is recommended that the current configuration of the tag be imported into Tag Manager and used as a configuration template, with any modifications then applied to that configuration. The result can then be applied to one or more tags. To do this, after selecting the **Configuration** menu option, place the mouse cursor over the tag that you would like to use as a template. Right click, and select **Get Tag Configuration**, respond **Yes** when asked to proceed.

Figure 6-26 Tag Manager 3.04 Configuration Panel



6. Configure each parameter subcategory for basic operation of T2 or T3 tags with the Cisco UWN software Release 4.1. If you have selected both T2 and T3 tags, note that only the configuration options that apply to *both* tag models are available. Once all parameters in a configuration group have been configured, they may be applied to the selected tags by clicking on the **Apply** button that appears within each group. Alternatively, you may delay applying changes until all groups have been configured (use the **Apply Multiple Configuration** option shown at the bottom of [Figure 6-26](#)). All parameters selected are applied to all selected asset tags and will override any other values that may be present.

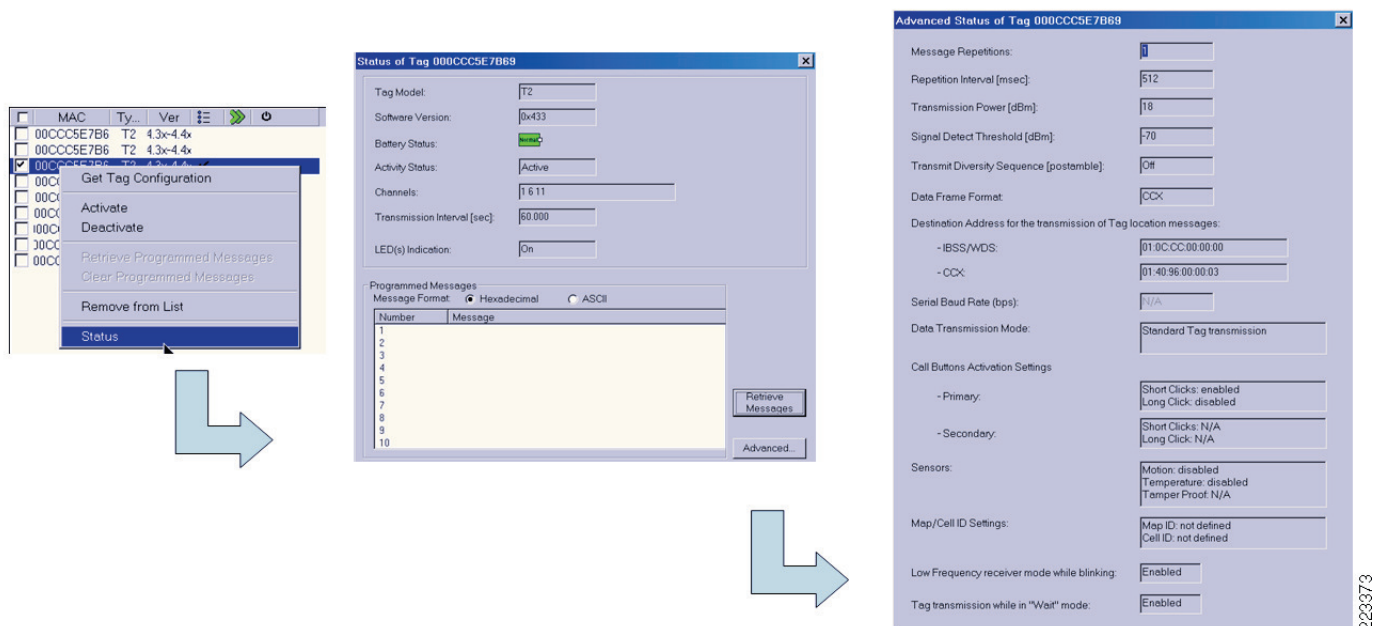
- a. General Parameters:

- Channel Selection—It is recommended that tags be configured for the standard set of 802.11b non-overlapping channels, typically channels 1, 6 and 11 (or otherwise depending on your regulatory domain).

- LED Indication—In most cases, it is useful to have visual indication of when the tag is using its communication interfaces. In cases where there are reasons why such indication is undesirable, such as in a light sensitive, security or other “stealth” application, the LED can be disabled.
 - Transmission Interval When Not In Motion—Select an appropriate tag transmission interval for your asset tagging application, in seconds or milliseconds. Typically tags are configured to transmit less frequently when stationary using this parameter setting as compared to when they are in motion. In-motion transmission intervals are set using the Motion Sensor category settings.
- b. Transmission Parameters:
- Message repetitions—Standard operation for the AeroScout tag is to transmit a single multicast transmission on all defined channels. This parameter controls the number of times each transmitted message is repeated, per channel. It is generally recommended that this parameter be raised from the default value of one to a value of three. Doing this helps protect against lost tag transmissions, which results in lost RSSI readings. Lost RSSI readings is a confirmed cause of degraded location accuracy, especially in environments where there is a significant likelihood of tag transmissions being interfered with or dropped due to congestion or interference. Avoid configuring an excessive number of message repetitions, as there are few conditions where a message repetition factor greater than 3 would be truly required. The setting of three message repetitions works very well for the majority of environments. Setting this parameter above a value of 5 is typically not considered necessary.
 - Message Repetitions Interval—The delay between subsequent message repetitions on the same channel, specified as either 128, 256 or 512 milliseconds. The default value is 512 milliseconds.
 - Transmission Power (dBm)—The default value for transmission power is typically +18dBm on T2 model AeroScout asset tags. The location-aware Cisco UWN is capable of discerning the transmission power used by tags compliant with the Cisco Compatible Extensions for Wi-Fi Tags specification.
 - Data Rate—Data rates of 2 Mbps can only be specified for T3 tags. Although the message payloads and frame sizes associated with asset tags are very small, the use of a faster transmission speed can allow T3 tags to transmit their payloads faster and free the channel for use by other stations sooner. This can also reduce battery consumption since each frame’s transmission time is shorter.
 - Data Frame Format—This parameter should be changed from the default value of **IBSS** to **CCX**.
 - Destination Address—This value must be specified as 01:40:96:00:00:03 for use with software Release 4.1 and later releases.
- c. Data Transmission Mode Parameters:
- Normal Tag Transmission (without additional message)—Select this parameter unless you have valid reasons to configure it otherwise. For example, the location client you are using in conjunction with your asset tags may be able to process additional stored messages on your tag, sent as part of tag payloads, or you may be using an AeroScout T2 telemetry tag that allows for telemetry to be read directly from sensors onboard custom-integrated host peripheral devices.
- d. Supplementary Settings:
- CCX Options—Transmit Out of Range Chokepoint Group should be enabled.
- e. Call Buttons Primary —Configure these options if you wish to use call button signaling (Panic Button alerting) with software Release 4.1.
- Short Clicks (button depression that last less than 2 seconds):
 - Enable Short Clicks should be checked

- Number of Short Clicks: 1
 - Tag Reaction Parameters: Send Standard Tag Transmission
 - Message Repetition: 1
 - Long Clicks (button depression that lasts at least 2 seconds):
 - Enable Long Clicks should be checked
 - Number of Long Clicks: 1
 - Tag Reaction Parameters: Send Standard Tag Transmission
- f. Call Buttons - Secondary—These are identical options to those listed for “Call Buttons – Primary” but are only available if you are using T3 asset tags.
- g. Sensors:
- Motion—These options can be used to enable the on-board motion sensor if desired.
 - Temperature—These options can be used to enable on-board temperature sensors if desired. Note that the on-board temperature sensor is not supported in T2 tags with v4.3x firmware.
 - Tamper—This option can be enabled for T3 tags only. Enabling this option allows tag tamper indication to be sent to the Cisco UWN.
7. In some cases, the existing configuration of an AeroScout asset tag may be in question and need verification. Using Tag Manager v3.04, this is a straightforward process. Simply right-click on any detected tag and click on **Status** from the pop-up menu. This brings up a listing of basic tag configuration parameters, with further detail available by selecting Advanced Configuration as shown in Figure 6-27.

Figure 6-27 Retrieving The Configuration of a Single Tag



The preceding quick, seven-step configuration guide is just a short synopsis of the required steps to configure and activate AeroScout tags for use with the Cisco UWN software Release 4.1. Refer to the *AeroScout Tag Manager v3.0 User's Guide* for more detailed information as well as information on several other useful configuration options in the Tag Manager.

Tag Telemetry and Notification Considerations

Beginning with software Release 4.1, the location aware Cisco UWN will recognize tag telemetry and high priority notifications transmitted by Wi-Fi Tags specification may transmit tag telemetry and high-priority notifications to the location-aware Cisco UWN. This information is passed from WLAN controllers to the Cisco Wireless Location Appliance using the Location Control Protocol (LOCP), which is described in [Cisco Location Control Protocol \(LOCP\), page 3-36](#).

This section provides initial best practice recommendations and other information and should be kept in mind when designing solutions that are dependent on telemetry and high-priority notification functions found in Cisco UWN software Release 4.1.

Deploying Tag Telemetry

Active RFID tags supplied by tag vendors in compliance with the Cisco Compatible Extensions for Wi-Fi Tags specification may include the ability to accept telemetry data from onboard sensors or from sensors integrated into the asset to which the tag is attached. If configured to do so, these active RFID tags can pass this telemetry data as part of the tag transmissions that are sent to the Cisco UWN at periodic transmission intervals, or when entering into the stimulation zone of chokepoint triggers.

For example, an asset tag connected to the fuel level sensor of a forklift may be able to pass fuel level telemetry via the Cisco UWN to the location appliance and its location clients (which could include WCS and third party location clients). The ability of the asset tag to perform these telemetry functions is dependent upon the asset tag manufacturer, and typically requires the appropriate level of integration and physical connectivity between the tag and sensors found aboard the attached asset. Note that some asset tags are available with their own onboard sensors, which can measure certain ambient environmental characteristics (such as temperature and humidity) external to tagged assets without any dependence on embedded sensors.

Onboard tag sensors, for example, might be appropriate where the primary concern surrounds general environmental conditions effecting both the asset tag as well as the asset to which it is attached. Thus, an asset tag equipped with onboard temperature sensors would be appropriate in detecting whether an attached asset was incorrectly stored in temperatures outside recommended ranges. Embedded sensors within the asset itself would be more appropriate when the goal is to alert the system administrator to an internal condition resulting from improper use that could result in costly damage to the asset if not addressed promptly. A good example of this might be an engine providing indication of an insufficient internal lubrication, which could result in costly repairs.

As described in the section entitled [Asset Tag Telemetry Using LOCP, page 3-38](#), beginning with the Cisco UWN software Release 4.1 all tag telemetry sent by tags compliant with the Cisco Compatible Extensions for Wi-Fi Tags specification is aggregated by WLAN controllers and passed to the Cisco Wireless Location Appliance. In software Release 4.1, LOCP uses a polled mechanism to collect tag telemetry after the fact, the timing of which is tied to the traditional SNMP polling mechanism used to gather asset tag RSSI information. The location appliance updates the telemetry information for each asset tag in its databases with that received from the most recently responding WLAN controller that has

included telemetry information for that specific tag's MAC address. If archiving of tag historical information has been enabled on the location appliance, tag telemetry information is included along with other tag information (shown in Figure 6-28).

Figure 6-28 Archive Playback of Tag Telemetry and "Emergency" Data

Aeroscout Tag 00:0c:cc:5c:05:13 -- Select a command -- **GO**

Asset Name _____ Asset Group _____
 Asset Category _____ MAC Address 00:0c:cc:5c:05:13

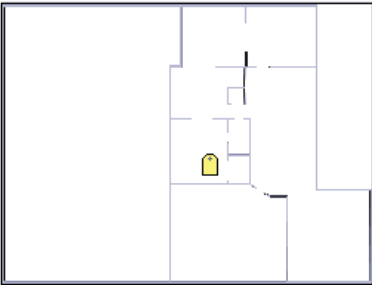
From : Tue Apr 24 23:01:53 EDT 2007
 To : Tue May 15 19:13:23 EDT 2007

Time Stamp	Floor	Battery Status
17 Sat Apr 28 10:03:04 EDT 2007	Lab Annex #2 Alpharetta Campus>AP1242 Building>Test Lab Annex #2	80 %
18 Fri Apr 27 21:09:29 EDT 2007	Alpharetta Campus>AP1242 Building>Test Lab Annex #2	80 %
19 Fri Apr 27 19:09:29 EDT 2007	Alpharetta Campus>AP1242 Building>Test Lab Annex #2	80 %

Change selection every 2 secs **Play** **Stop**

Location

Location Calculated Fri Apr 27 19:08:25 EDT 2007
 Floor Alpharetta Campus>AP1242 Building>Test Lab Annex #2



[Enlarge](#)

Tag Statistics

Data Collected Fri Apr 27 19:07:37 EDT 2007
 Bytes received 343228
 Packets received 6601

Telemetry Data

MOTIONPROB : Movement

Emergency Data

Reason: Tampering
 Tamper State: Active

Tag Properties

Data Collected Fri Apr 27 19:07:25 EDT 2007
 Controller 10.1.96.18
 Battery Status 80 %

The default configuration of some active RFID tags may provide for transmitting only one tag transmission per channel per transmission interval. While this setting can help optimize the battery life of the tag in some cases, this single transmission per channel may not always be successfully detected by the expected number of access points, especially in RF-noisy or congested environments. This can result in missing RSSI readings, which can cause location inaccuracy.

Therefore, in such environments it is recommended that tags be configured to transmit multiple transmission repetitions per channel at each transmission interval, which should aid in improving tag detection and location accuracy as well as increasing the reliability of tag telemetry as well. It is recommended that the tag vendor's configuration software should be used to set the number of tag transmissions to three (but not more than five) per channel per transmission interval.

Although it is unlikely that LOCP telemetry collection will burden modern wired and wireless networks, nevertheless it is good practice for the network designer to understand the nature of the traffic that can be expected in their designs. The following traffic and frame size information has been observed during LOCP telemetry testing in support of this document:

- *Echo Request*—Sent periodically by the location appliance to each defined WLAN controller based on the configuration of the Echo Interval parameter (Location Servers > Advanced > LOCP Parameters). LOCP Echo Request Ethernet frames are 100 bytes in length and are transmitted to TCP destination port 16113.
- *Echo Response*—Sent periodically by each WLAN controller in response to an Echo Request (see above). Like Echo Requests, LOCP Echo Response Ethernet frames are 100 bytes in length.
- *Information Request*—Sent periodically by the Location appliance to each WLAN controller to request information. LOCP Information Request Ethernet frames are 106 bytes in length are transmitted to TCP destination port 16113. LOCP Information Requests are the primary mechanism used in software Release 4.1 to conduct LOCP polling.
- *Information Response*—Sent periodically by each WLAN controller in response to the receipt of a LOCP Information Request frame (LOCP Polling). The basic size of a LOCP Information Request Ethernet frame for a controller that has not detected any tags is 113 bytes. If one tag is detected, this frame size will increase to 144 bytes and for two tags it will increase to 175 bytes (these frame sizes do not include any telemetry data). Frame sizes will increase based on the number of tags currently active in the controller's database as well as the amount of telemetry that has been collected. Support for fragmentation and reassembly of combined tag payloads is inherently to LOCP.

To ensure proper LOCP operation between the location appliance and any WLAN controllers defined to it, ensure that port 16113 is not blocked by any firewalls or other security devices.

When designing solutions that will rely on the reporting and collection of tag telemetry with Release 4.1, there are a few considerations that should be kept in mind:

1. *Telemetry Timing*—Since in Release 4.1 telemetry is aggregated on a per-tag basis by WLAN controllers and passed to the location appliance only during a periodic LOCP polling cycle, users of software Release 4.1 should not rely on the receipt of tag telemetry to be real-time in nature. It is reasonable to expect that there will be a delay between the time the tag sends the telemetry information and the time it is updated in the location appliance database and made available to location clients.
2. *Northbound Asynchronous Notifications*—In Release 4.1 of the location-aware Cisco UWN, the location appliance does not issue asynchronous northbound notifications (in the form of email, SNMP, SOAP or UDP-Syslog messages) for telemetry received from tags. Therefore, any external applications (such as paging systems, text messaging, enterprise management consoles and so on) relying on northbound notifications in these formats must receive them from an alternate source having visibility to tag telemetry, such as a third-party location client.

Battery telemetry, however, is an exception. In this case, the location appliance will trigger northbound asynchronous notifications based on remaining battery life for tags compliant with the Cisco Compatible Extensions for Wi-Fi Tags specification. These notifications are generated as per the following trigger condition definitions:

- Battery Level is Low—Reported battery life remaining is 30%
- Battery Level is Medium—75% battery remaining > 30%
- Battery Level is Normal—Battery remaining is > 75%

Deploying Tag High-Priority Notifications

Beginning with software Release 4.1 of the Cisco UWN, asset tags compliant with the Cisco Compatible Extensions for Wi-Fi Tags specification may transmit high-priority and vendor-specific notifications to the location-aware Cisco UWN. This information is transmitted as part of a tag transmission that is sent on-demand, and is passed from WLAN controllers to the Cisco Wireless Location Appliance using LOCP. Keep in mind that the format of the tag message sent by the tag when a high-priority type event occurs is very similar to the standard tag multicast transmission sent during each tag transmission interval, except that it contains additional information that conveys the nature of the high-priority event.

It is important to note that information contained in the tag notifications received over RF by the WLAN controller is passed (with minimal delay) to the location appliance in the form of LOCP Information Notifications. Thus, for example, when a call button is depressed on an asset tag that is compatible with the Cisco Compatible Extensions for Wi-Fi Tags specification, a LOCP Information Notification is transmitted by the WLAN controller to the location appliance very shortly after the tag notification has been received by the controller's registered access points. Once received by the location appliance, the updated call button status is reflected in the location appliance database (for example, "panic button depressed") and made available to location clients. If archiving of tag historical information has been enabled on the location appliance, tag "emergency" information is archived along with other tag information (shown in [Figure 6-28](#)).

The basic size of a LOCP Information Notification Ethernet frame is approximately 130 bytes. Frame sizes can be larger based on additional information included in the frame, such as tampering information or vendor-specific data. In Release 4.1, LOCP Information Notifications are not aggregated by WLAN controllers. WLAN controllers will transmit a LOCP Information Notification frame to the location appliance for *each* tag high-priority notification received via *each* of its registered access points (including any high-priority notification repetitions).

Expressed mathematically, it can be stated that for each notification event coming from a tag, the total number of LOCP Information Notifications that can be expected to be transmitted from a WLAN controller to the location appliance can be calculated as:

$$LOCP\ Information\ Notifications_{TOTAL} = Detecting\ APs_{TOTAL} * High-Priority\ Notification\ Repetitions_{PER\ CHANNEL}$$

where *High-Priority Notification Repetitions_{PER CHANNEL}* represents the total number of high-priority notifications that are sent by the tag on a single RF channel. Note that the number of high-priority notification repetitions per channel should not be confused with the standard setting for tag message notifications per channel, which applies to tag transmissions that are sent periodically based on the expiration of a tag transmission interval. It should also be noted that this calculation yields the maximum possible value for *LOCP Information Notifications_{TOTAL}* as it assumes that all notification repetitions coming from the tag are successfully detected by all access points included within *Detecting APs_{TOTAL}* and none are dropped due to interference, contention or other RF anomalies.

Using our formula, we can calculate the expected number of LOCP Information Notifications that will be generated if the call button is depressed once on an asset tag compliant with the Cisco Compatible Extensions for Wi-Fi Tags specification within the following Release 4.1 environment:

- Two WLAN controllers
- Three access points registered to each controller, for a total of six detecting access points.
- Tags send one notification for each call button depression on each of channels 1, 6 and 11

Substituting this information into the aforementioned equation, we see that $6 * 1$ or 6 total LOCP Information Notifications will be transmitted from the WLAN controllers to the location appliance in this example. Note that although both WLAN controllers will be sources of LOCP Notifications in this example, the number of WLAN controllers present in the environment has no bearing on the number of LOCP Notifications that will be sent to the location appliance. We could have substituted three WLAN controllers with two access points registered to each in this example, and the calculated value for *LOCP*

*Information Notifications*_{TOTAL} would have been the same. It is the number of access points that detect the tag multicast transmissions bearing the high-priority notification information sent that is pertinent to the number of LOCP Notifications that will be generated from controllers to the location appliance.

To ensure proper LOCP operation between the location appliance and any WLAN controllers defined to it, always ensure that port 16113 is not blocked by any firewalls or other security devices.

Configuring Tags for Telemetry and Notifications

While the support of tag telemetry and notifications are basic components of the Cisco Compatible Extensions for Wi-Fi Tags specification, each tag vendor uses their GUI or CLI-based tag software to enable, disable or otherwise customize precisely how these features are supported in their products. While a limited amount of AeroScout tag configuration information has been already provided in prior sections of this document, more comprehensive information specifically relating to the configuration of external telemetry sensors and asset tags is available from asset tag vendors, but is beyond the scope of this document.

Readers seeking such information are directed to the following sources of information:

- *AeroScout T2 Tag User Guide*
- *AeroScout Tag Manager User Guide version 3.0*
- <http://www.aeroscout.com> or your AeroScout account and technical support team

For asset tags from other vendors that are compatible with the Cisco Compatible Extensions for Wi-Fi Tags specification, it is recommended to contact those vendors directly. These would include:

- InnerWireless (formerly PanGo Networks) <http://www.innerwireless.com>
- WhereNet <http://www.wherenet.com>

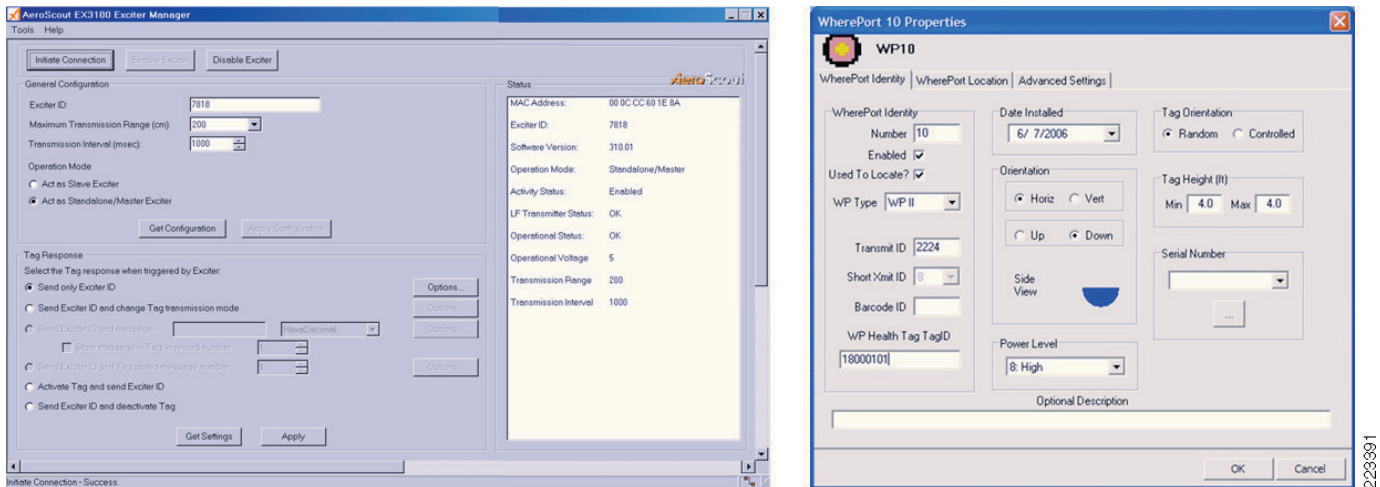
Chokepoint Considerations

Configuring Chokepoint Triggers

In order to use chokepoint triggers with the Cisco UWN, they must be properly configured using the appropriate vendor-supplied software utility, defined to WCS, placed on floor maps and synchronized as part of an updated network design to the location appliance. After all of this is complete, the location appliance will be able to recognize that asset tags compliant with the Cisco Compatible Extensions for Wi-Fi Tags specification have been stimulated by a particular chokepoint trigger MAC address and proceed to localize the asset tag. Location clients may then display the asset tag's location at the chokepoint icon associated with the chokepoint trigger's MAC address.

Various chokepoint trigger specific parameters such as transmission range, IP address, transmission interval, transmission repetitions and so on are set using vendor-specific utilities. For non IP-addressable AeroScout EX-3100 series Exciters, the AeroScout Exciter Manager standalone software utility must be used (shown on the left in [Figure 6-29](#)). For WhereNet WherePort chokepoint triggers, the WhereNet System Builder (shown on the right in [Figure 6-29](#)) and the WhereNet WhereWand are used.

Figure 6-29 Vendor-Specific Configuration Utilities



Note that each vendor maintains their set of software tools necessary for configuration of their chokepoint triggers. These software configuration tools are not interoperable between vendors (for example, AeroScout software configuration tools cannot be used to configure WhereNet chokepoint triggers or vice-versa).

In general, the individual configuration of each vendor's chokepoint trigger device is beyond the scope of this white paper. This document does, however, attempt to shed light on specific chokepoint trigger configuration parameters that are of particular significance in solving design challenges. As necessary, the topical sections of this document make reference to such parameters as necessary. However, complete and detailed configuration information relating to the specific configuration of each vendor's chokepoint trigger can be found in the appropriate vendor's documentation:

Available from AeroScout Corporation:

- AeroScout EX-3100 Exciters:
 - *AeroScout Exciter EX-3100 User Manual*
 - *AeroScout EX-3100 Exciter Manager User's Manual*
- AeroScout EX-3200 Exciters:
 - *AeroScout EX-3200 User Guide*
- AeroScout EX-2000 Exciters:
 - *AeroScout Exciter EX-2000 User Guide*

The following reference manuals are recommended for configuration of AeroScout EX-2000 and EX-3200 Exciters, using either the AeroScout System Manager or the AeroScout Network Exciter Manager (ANEM). The AeroScout Network Exciter Manager is a standalone Exciter software configuration utility specifically designed for users of AeroScout Exciters and the Cisco UWN.

- *AeroScout Engine Version 3.2 User's Guide*
- *AeroScout Network Exciter Manager (ANEM) User's Guide*

Technical documentation for WhereNet WherePort chokepoint triggers and the necessary software and hardware for configuration of WherePorts is available from WhereNet Corporation (<http://www.wherenet.com>) via your WhereNet account representative.

Defining Chokepoint Triggers to the Cisco UWN

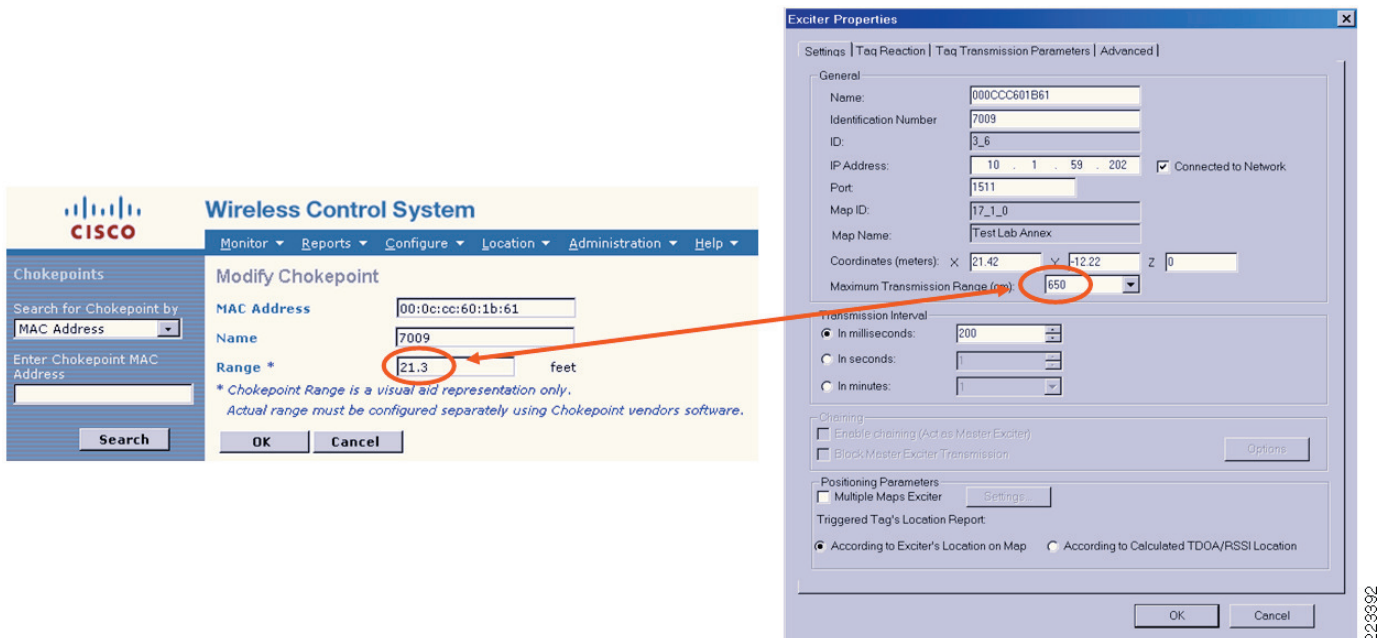
As mentioned earlier, after chokepoint triggers have been individually configured using the configuration tools supplied by the vendor, they must be defined to WCS, placed on appropriate floor maps and synchronized with the location appliance as part of an updated network design. Only then can they be used to track asset tags compliant with the Cisco Compatible Extensions for Wi-Fi Tags specification.

Complete step-by-step guidance regarding how to define compatible chokepoint triggers to WCS and the location appliance can be found at the following location:

http://www.cisco.com/en/US/products/ps6386/products_configuration_guide_chapter09186a008082d7d2.html#wp1058654.

When defining chokepoint triggers, it should be noted that the range is specified in both the vendor's configuration program as well as in WCS (shown in Figure 6-30). However, it is the range configuration parameter specified in the vendor's configuration program that actually sets the transmission range of the chokepoint trigger, not the range setting in WCS. The value that is specified for the range of the chokepoint trigger in WCS simply sets the size of the gray concentric rings that appear surrounding each chokepoint icon on WCS floor maps. These concentric rings are visual aids placed simply to serve as a convenient reminder of the range associated with the chokepoint trigger.

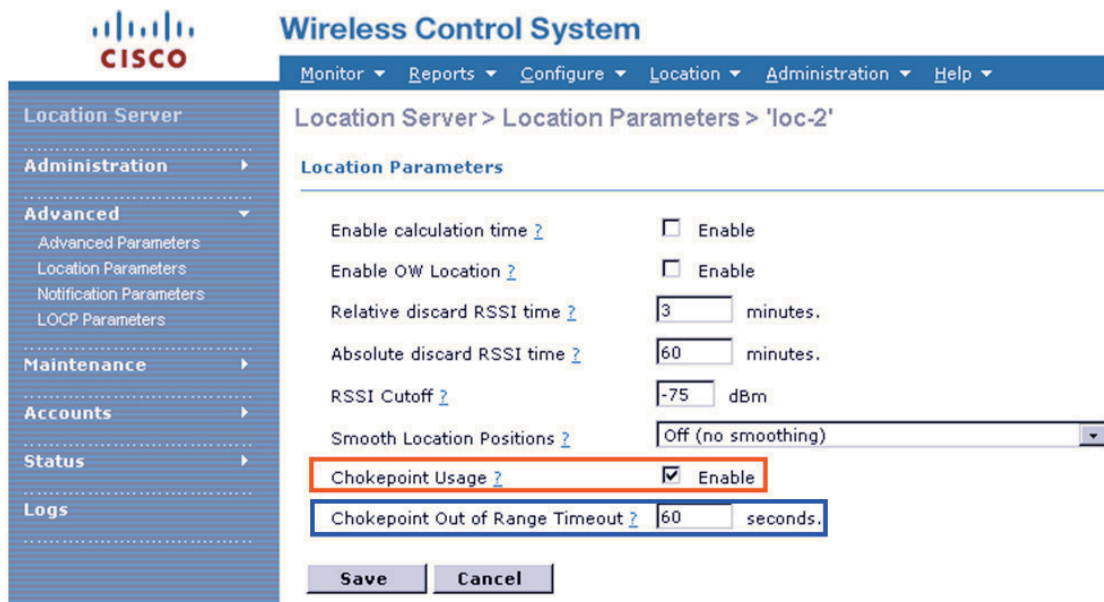
Figure 6-30 WCS and Vendor Range Parameters Compared



Note that these concentric rings do not represent any type of “special” area. For example, when RF Fingerprinting is being used to as the means of localizing tags instead of chokepoint location, tags may be placed by the system anywhere on floor maps (including within these gray concentric rings) if that is the location deemed to be correct by the location appliance.

There are also two additional parameters regarding the use of chokepoints that are found on the Location > Location Servers > Advanced > Location Parameters menu screen, as shown in Figure 6-31:

Figure 6-31 Chokepoint Advanced Location Parameters



223393

- Chokepoint Usage**—This checkbox (shown within the red rectangle in Figure 6-31) must be enabled in order for the location appliance to use chokepoint location techniques to localize tags. This occurs when it receives incoming LOCP Measurement Notifications indicating that a tagged asset has been stimulated by a chokepoint trigger. With regard to the chokepoint capabilities contained with the Cisco location appliance and the location-aware Cisco UWN, these techniques are only used with asset tags that are compliant with the Cisco Compatible Extensions for Wi-Fi Tag specification. If this parameter is disabled, the appliance will use the standard mechanism of RSSI based RF Fingerprinting to calculate tag location at all times.
- Chokepoint Out of Range Timeout**—This parameter (shown within the blue rectangle in Figure 6-31) specifies the timer used to age the last “in-range” report received from for an asset tag that is being localized using chokepoint location techniques. It assures that any tags no longer transmitting frames indicating they are within range of a chokepoint trigger are removed from that chokepoint in the active location database, once the Chokepoint Out of Range Timeout has expired. These tags are assumed to have left the chokepoint and are reverted back to being localized using standard RF Fingerprinting techniques.

Chokepoint Trigger Traffic Considerations

Beginning with Cisco UWN software Release 4.1, tags compliant with the Cisco Compatible Extensions for Wi-Fi Tags specification can use a consistent method to inform the UWN that they are within (or have left) the proximity of a chokepoint trigger. Once received by access points and forwarded to registered controllers, this information is passed to the location appliance using LOCP Measurement Notifications, which have already been described in [Cisco Location Control Protocol \(LOCP\)](#), page 3-36.

The length of each 802.11 tag multicast frame transmitted in response to stimulation received from a chokepoint trigger is approximately 63 bytes, which includes only a single chokepoint MAC address and does not include any historical chokepoint information. The length of the frame could increase due to the inclusion of a historical list of chokepoints traversed, or it may be larger than 63 bytes due to vendor-specific information that may be included in the frame. For example, during lab testing with

AeroScout T2 tags, it was observed that the typical size of the tag multicast frame emitted when in proximity of a chokepoint trigger is approximately 71 bytes, slightly larger than the multicast frame transmitted by these same tags during routine periodic transmissions (56 bytes). This 71-byte length is greater than the expected 63 bytes, and upon further examination it is discovered that eight additional bytes of vendor-specific information is included.

The Cisco Compatible Extensions for Wi-Fi Tags specification also allows asset tags to communicate historical information about the chokepoints they traverse to the Cisco UWN. This could increase the size of the frame by approximately 10 bytes per chokepoint trigger encountered depending on the number of historical entries maintained. The basic size of a LOCP Measurement Notification Ethernet frame is approximately 160 bytes. Frame sizes may be larger based on additional information included in the frame, such as historical chokepoint information.

In software Release 4.1, LOCP Measurement Notifications are not aggregated by WLAN controllers. WLAN controllers will transmit a LOCP Measurement Notification frame to the location appliance for each incoming tag multicast transmission, received by each of its registered access points, that indicates that the tag has been successfully stimulated by a chokepoint trigger. Therefore, the number of LOCP Measurement Notifications generated by one or more WLAN controllers for a single tag transmitting multicast frames indicating that the tag has been stimulated by a chokepoint trigger, is dependent upon:

- the number of registered access points that are within range of the tag and that have detected the tag's chokepoint-related transmissions.
- the number of times the tag will transmit a multicast frame on each configured 802.11 channel in response to chokepoint trigger stimulation.

This can be expressed mathematically as:

$$LOCP\ Measurement\ Notifications_{CHOKEPOINT} = Detecting\ APs_{TOTAL} * 802.11\ Repetitions_{PER\ CHANNEL}$$

Note the following considerations:

1. This calculation yields the number of *LOCP Measurement Notifications* that result from a single tag reacting to a single chokepoint stimulation event.
2. Chokepoint triggers by default transmit multiple stimulation packets over their magnetic signaling medium. This could result in multiple stimulation events, which is highly dependent on the amount of time spent within the chokepoint stimulation zone and other factors.
3. This calculation yields a maximized value for *LOCP Measurement Notifications* as it assumes that all frames transmitted by the tag are successfully detected by the number of access points specified in *Detecting APs_{TOTAL}* (none are dropped due to interference, contention or other RF anomalies).

In the majority of cases:

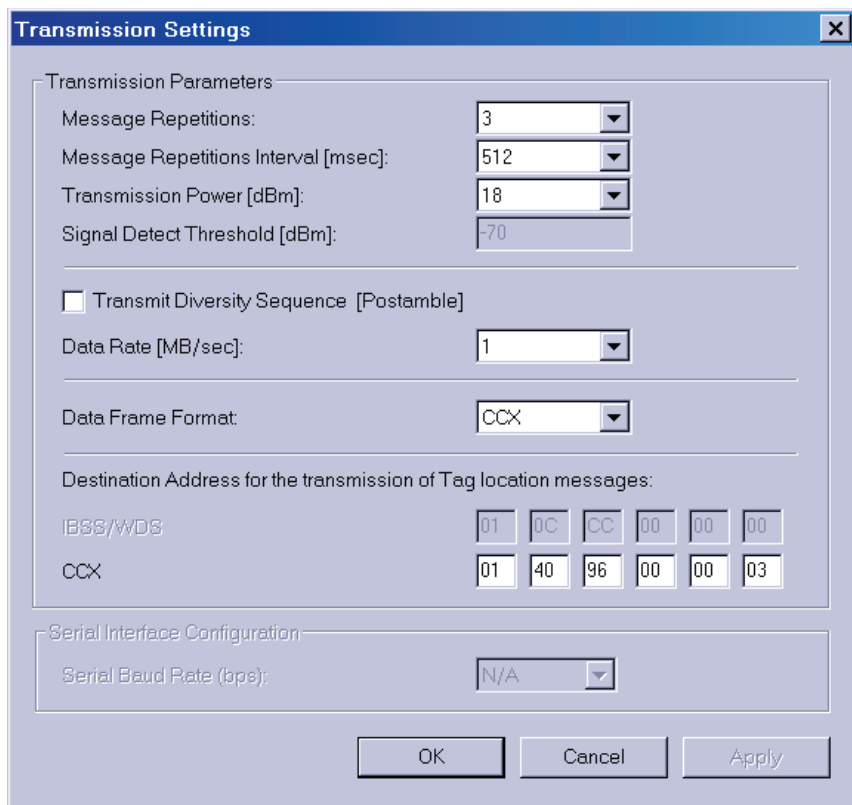
- Chokepoints are deployed in areas where the surrounding access point spacing meets the requirements discussed in [Access Point Placement, page 5-5](#).
- Access points and tags are configured to operate on the non-overlapping 2.4 GHz channels (channels 1 (2412 GHz), 6 (2437 GHz) and 11 (2462 GHz) in the Americas, for example).

Most chokepoint triggers assume a default value of one for the number of times they repeat, per channel, tag multicast transmissions indicating that the tag has been successfully stimulated by a chokepoint trigger. Only a single tag multicast transmission frame containing the stimulating chokepoint trigger's MAC address need be received in order to result in the generation of a *LOCP Measurement Notification*. Because of this, the default value for the number of times these chokepoint-related tag transmissions are repeated per channel is usually sufficient, especially since this tag transmission will typically be repeated across three 2.4 GHz channels, resulting in more than one access point receiving the tag transmission, even without increasing the repetition count. However, in some cases where interference or congestion may be extremely high, it may make sense to increase the repetition count slightly. In other cases involving tagged assets traversing through chokepoint areas at high speed or at fringe distances from

chokepoint triggers, this parameter can be used to increase the likelihood of reliable stimulation. In all cases, however, such increases should be done judiciously given the ability of this parameter to affect the amount of traffic added per stimulated tag in large tag environments.

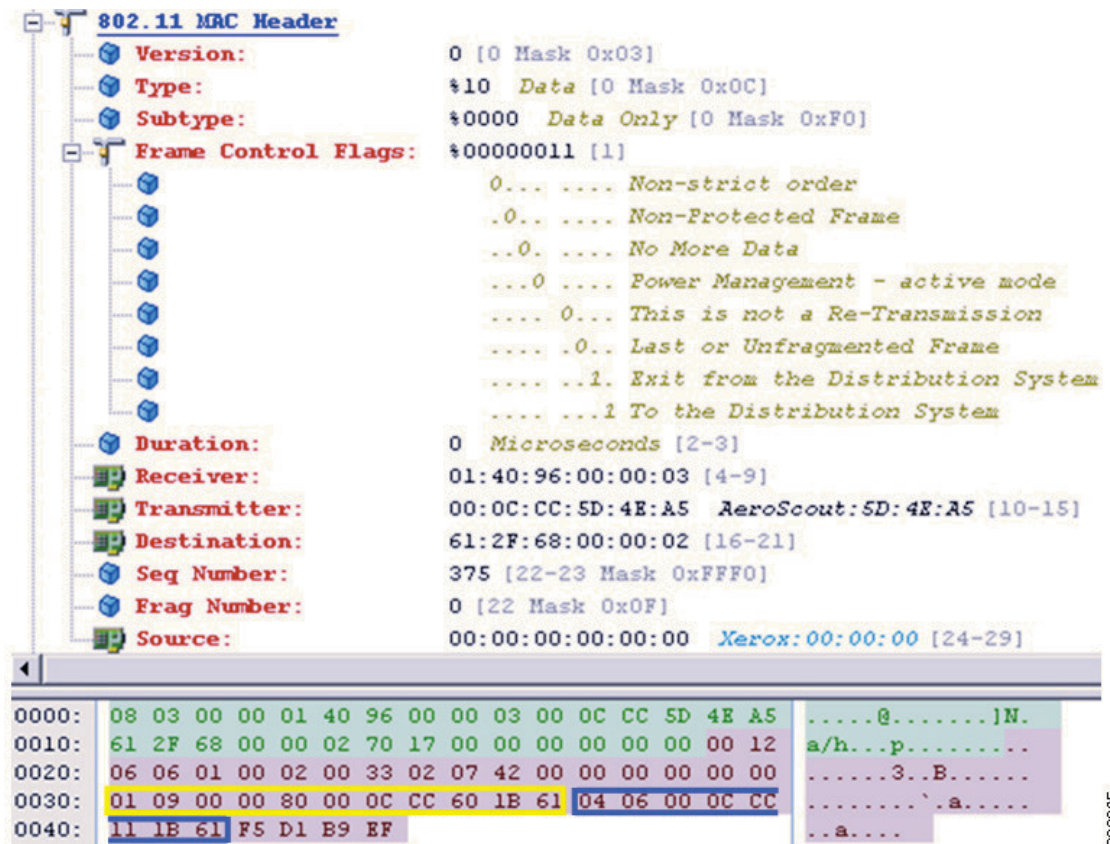
It should be noted that the repetition count that applies to the tag multicast frames sent in response to chokepoint stimulation is usually managed independently of the repetition count for other tag events such as telemetry, high-priority notifications or periodic tag transmissions sent as a result of the tag's configured transmission interval. When configuring tags and chokepoint triggers, it is important to maintain this distinction. For example, with AeroScout tags the repetition count that applies to the 802.11 frames sent by a tag in response to a chokepoint stimulation event is known as the "Tag Repetition of an Exciter event" parameter. It is configured on a per-Exciter basis using the AeroScout System Manager, Exciter Manager or ANEM utility. In contrast, the tag repetition parameter used for non-Exciter related events is known as the Message Repetitions transmission parameter. It is set on a per-tag basis using the Transmission Settings panel of the AeroScout Tag Manager, as shown in Figure 6-32.

Figure 6-32 Transmission Settings Panel in AeroScout Tag Manager (not used for Exciter Events)



As mentioned earlier, the length of each 802.11 multicast tag frame transmitted by a tag in response to chokepoint stimulation is approximately 63 bytes, which includes only a single chokepoint MAC address and does not include any historical chokepoint information. We mentioned earlier that it was observed during testing that the 802.11 multicast frame transmitted by an AeroScout T2 tag also contains eight bytes of vendor-specific information. Figure 6-33 illustrates this, with the mandatory chokepoint information contained within the yellow rectangle and the additional vendor-specific information contained within the blue rectangle.

Figure 6-33 Vendor-Specific Information Included in Tag Checkpoint Transmission



Although both the standard checkpoint information as well as the optional vendor-specific information travels from the tag to the access point contained within the same tag multicast frame, in software Release 4.1 the WLAN controller parses this into two separate LOCP notifications:

- An LOCP Measurement Notification containing the checkpoint group information that is 160 bytes in length.
- An LOCP Information Notification containing the vendor-specific information that is 138 bytes in length.

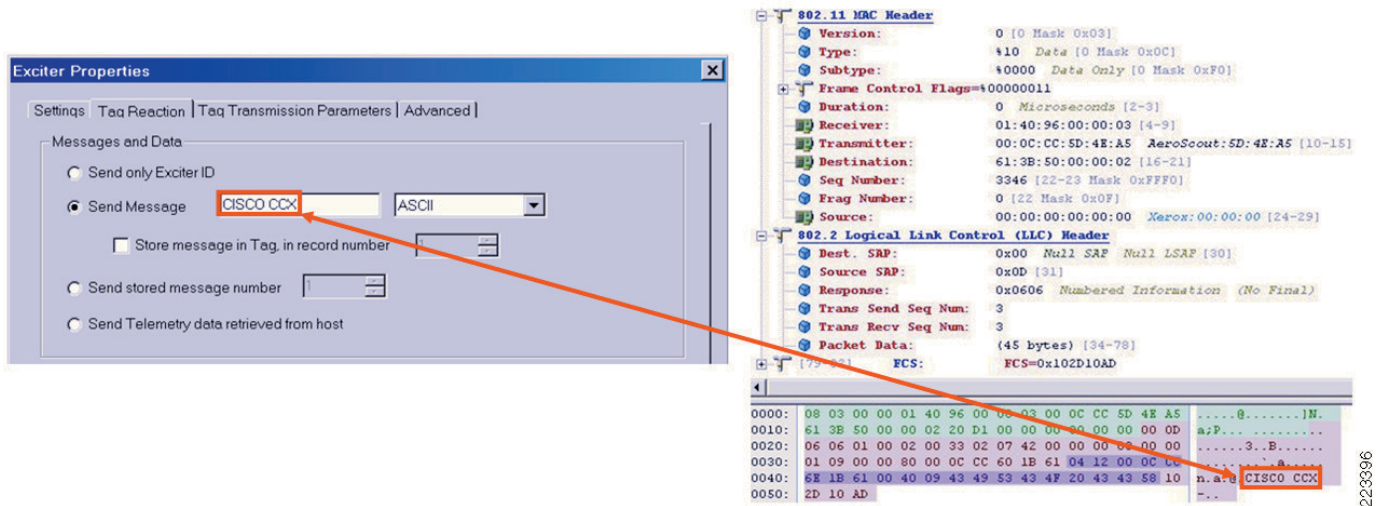
The precise composition of the vendor-specific fields varies depending on the checkpoint and tag vendor. For example, AeroScout allows for additional message information to be appended to the Exciter ID via the Tag Reaction tab of the Exciter Properties menu in the AeroScout System Manager and AeroScout Network Exciter Manager. These capabilities are also available via the Exciter Manager utility for users of the AeroScout EX-3100 Exciter.

Vendor specific information can be:

- Directly entered and saved on a per-Exciter basis.
- Saved to tag memory for later reuse.
- Consist of one of ten preconfigured messages programmed into tags.
- Emanate from a host attached to the tag.

Figure 6-34 illustrates the use of this capability for an AeroScout EX-2000 Exciter and AeroScout tags. In this figure the Exciter instructs the tag to append vendor-specific information in addition to the vendor-specific Exciter ID to each tag transmission frame sent as a result of stimulation received from the Exciter.

Figure 6-34 AeroScout Vendor-Specific Information Options



In Figure 6-34 we see the message “CISCO CCX” being defined to the Exciter as well as the complete 83 byte message transmitted by the tag when stimulated by the Exciter. This 83-byte message includes the standard information regarding the MAC address of the stimulating Exciter as well as the vendor specific information. Note that the text defined to the Exciter in the AeroScout System Manager is seen transmitted by the tag at offset x0046 in the trace (you can see the ASCII text “CISCO CCX” shown at the right in Figure 6-34). Every access point receiving this information will forward it to their registered controller where a 160-byte LOCP Measurement Notification as well as a 148-byte LOCP Information Notification will be sent to the location appliance. Although this information was hard-coded at the Exciter, the Exciter could have just as easily instructed the tag to instead include telemetry data that it retrieved from the asset (host) that it is attached to, such as embedded sensor data.

Keep in mind that results of our test observations obviously are, in this case, AeroScout specific, as other vendors may or may not opt to allow the inclusion of vendor-specific information to the same degree.