# Protection Ripple in ERP 802.11 WLANs

White Paper

June 2004

Planet3 Wireless, Inc.
Devin Akin, CTO
Devin@cwnp.com

 

## Understanding Use of 802.11g Protection Mechanisms

The 802.11g amendment to 802.11-1999 (R2003) clearly says that access points (APs) should signal to all associated stations in the basic service set (BSS) to use protection mechanisms (RTS/CTS or CTS-to-Self) when a NonERP (802.11b) station (STA) associates to the AP.

IEEE 802.11g, Section 7.3.2.13 – "*If one or more NonERP STAs are associated in the BSS, the **Use_Protection** bit shall be set to 1 in transmitted ERP Information Elements.*"

Use of these protection mechanisms can easily cause more than a 50% loss in overall WLAN throughput in the BSS. Latency is also increased significantly, more so with RTS/CTS than with CTS-to-Self.

The same section of the 802.11g amendment also states, "*The **NonERP_Present** bit shall be set to 1 when a NonERP STA is associated with the BSS. Examples of when the **NonERP_Present** bit may additionally be set to 1 include, but are not limited to, when*

> *a) A NonERP infrastructure or independent BSS is overlapping (a NonERP BSS may be detected by the reception of a Beacon where the supported rates contain only Clause 15 or Clause 18 rates).*
> *b) In an IBSS, if a Beacon frame is received from one of the IBSS participants where the supported rate set contains only Clause 15 or Clause 18 rates.*
> *c) A management frame (excluding a Probe Request) is received where the supported rate set includes only Clause 15 or Clause 18 rates.*"

This paragraph means that if a STA or AP hears a Beacon that has a supported rate set of 11, 5.5, 2, and 1 Mbps (802.11b) or only 2 and 1 Mbps (802.11) sent by a nearby AP or a STA that is part of an IBSS, it may enable the **NonERP_Present** bit in its own Beacons. Do not forget that it is the **NonERP_Present** bit that triggers protection within a BSS by instructing the AP that hears it in the ERP Information Element of a Beacon to enable the **Use_Protection** bit.

To summarize:

1) If a NonERP STA associates to an ERP AP, the ERP AP will enable the **NonERP_Present** bit in its own Beacons, enabling protection mechanisms in its BSS.
2) If an ERP AP hears a Beacon with an 802.11b or 802.11 supported rate set from another AP or an IBSS STA, it will enable the **NonERP_Present** bit in its own Beacons, enabling protection mechanisms in its BSS.

## The Perfect Site Survey

Let's first consider that most 802.11b and 802.11g site surveys use channels 1, 6, & 11 in an appropriate pattern as shown below.

 

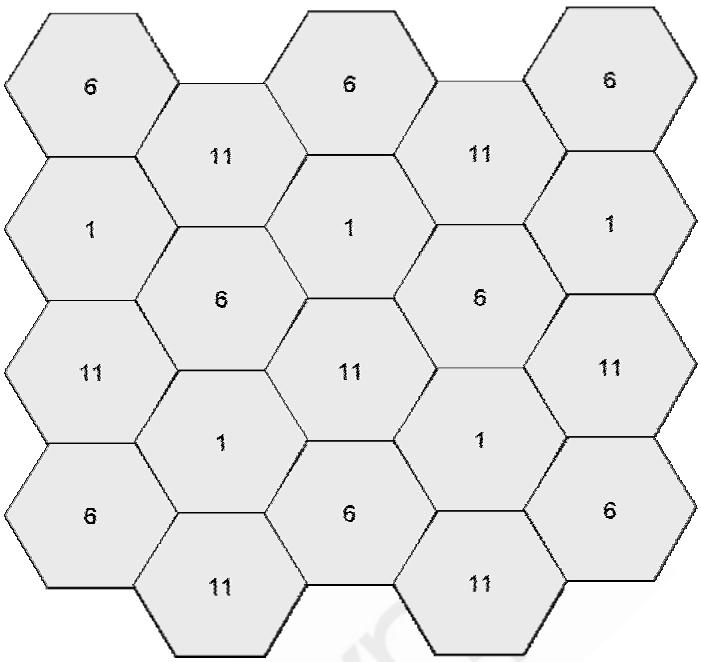


Second, let's consider, for the sake of argument, that this is a purely ERP, sometimes called "Pure G", environment, meaning that all APs and STAs are 802.11g capable and are using the ERP-OFDM PHY. Further, no STAs or APs have been manually configured to use RTS/CTS with a threshold. This yields a scenario where no RTS/CTS or CTS-to-Self frames should be in use on the wireless medium.

Now let's say we have an installation that has been configured properly to abide by a high quality site survey aimed at giving WLAN users data rates no lower than 24 Mbps. These cells would be relatively small as you can imagine. APs would be fairly close together, and the output power would be reasonably adjusted according to the physical environment. Proper antennas would be used for appropriate coverage and overlap to allow for seamless roaming between cells. Keep in mind that the RF signals from APs do not stop at the surveyed "boundary", but rather keep going – decreasing in power as they go. The survey is performed such that a high quality 24 Mbps data link is attainable at the boundary set by the surveyor.

Enter the NonERP STA. This STA could join any BSS, and the results will be the same. The 802.11 standard says that a Beacon must be sent at one of the basic (required) data rates, and typically the lowest basic rate is used by the AP. The basic rate set is configurable on the AP. Since most 802.11g APs are configured to support 802.11b stations (with rates of 1, 2, 5.5, & 11 Mbps), Beacons will generally be sent at 1 Mbps (if 1 Mbps is configured to be the lowest basic rate of course). As soon as a NonERP STA associates to an ERP AP, the AP will signal for protection in the ERP Information Element of its Beacon using the **Use_Protection** bit, and it will immediately enable the **NonERP_Present** bit which instructs other nearby APs to enable **Use_Protection**. Frames sent at 1 Mbps can be read reliably at a considerable distance by other wireless STAs and APs on the same channel, which means that enabling protection in one BSS may cause the enabling of protection in other BSSs across a large physical area. See below for an example.

 

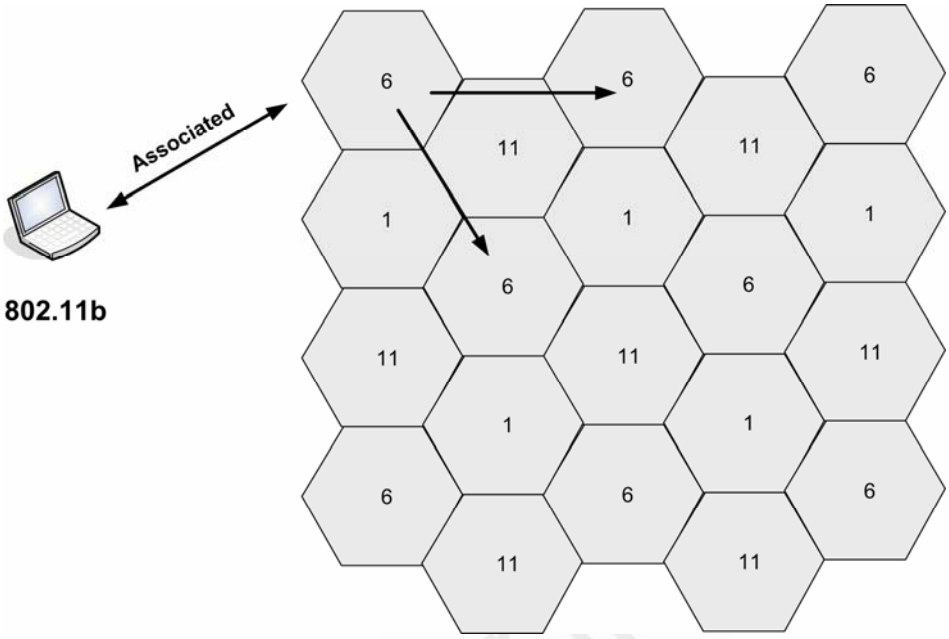


Another example of a situation where protection might be enabled is when a NonERP (802.11b) AP is placed within earshot of an ERP (802.11g) AP. A client station doesn't have to associate at all when this happens. The mere fact that the 802.11b AP is present in the area, sending Beacons with NonERP rate sets (11, 5.5, 2, & 1 Mbps), is enough to cause the ERP AP to enable the **NonERP_Present** bit, which in turn enables the **Use_Protection** bit. In a scenario like this, the limits of the throughput damage caused by the 802.11b AP would be based on how many 802.11g APs can hear its Beacons. If the output power of the 802.11b AP was high, the damage to the enterprise WLAN could be severe. Also consider for a moment a scenario such as an apartment complex, where everyone is switching to 802.11g APs except three people, which use channels 1, 6, and 11. Everyone's 802.11g networks in the building would essentially be reduced to 802.11b speeds.

This is certainly an eye-opener, but the worst is yet to come. What could possibly be worse, you ask? Let's consider that two access points are in the same room, transmitting at 100 mW, and are configured for channel 1 and channel 6 respectively. We have often seen the detrimental effects of adjacent channel interference in our lab environment. Using 802.11b WLANs with high output power (30 – 100 mW) at close proximity on adjacent channels yields a significant throughput loss on each channel due to interference. This scenario is shown below:

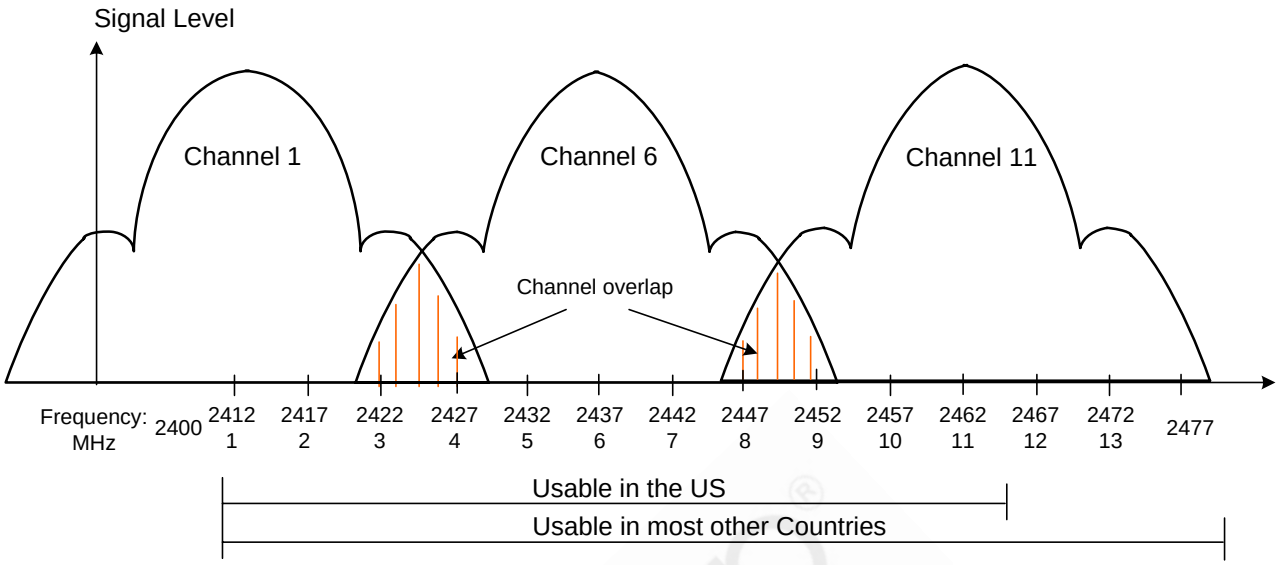


This graphic illustrates how an AP on channel 1 can hear transmissions (including Beacons) from an AP on channel 6. An AP with protection disabled only needs to receive one valid (uncorrupted) Beacon from an AP with the **NonERP_Present** bit enabled to be pushed into using protection itself. Having tested this theory in the lab at length, I can tell you that two APs, one configured for channel 1 and the other configured for channel 11 (yes, channel 11), can hear each other's Beacons without a problem so long as they are reasonably close (in the same room perhaps) and/or has the power turned up (30 – 100 mW). This tells us that channels 1, 6 & 11, which are used throughout most enterprises, are highly susceptible to this problem. See the example below for an illustration of how Beacons from an AP on one channel may affect Beacons from APs on other channels.

 

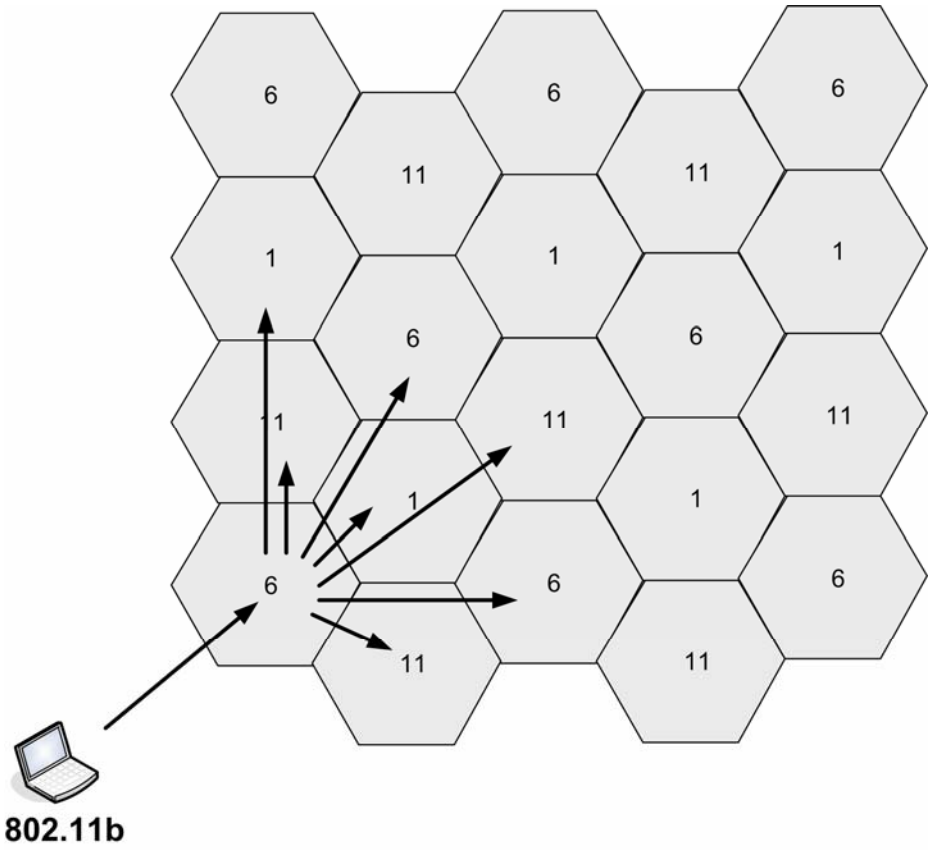


This problem would become much worse when the WLAN was improperly designed, such as when APs have too much output power, are too close together, and are not reusing channels appropriately. Keep in mind that SSID configuration or even security solutions like 802.1X/EAP on the APs will not affect this situation in any way. Whether this is a scenario of multiple independent BSSs or a single ESS really has no bearing.

## Disabling Protection and Roaming

The 802.11g amendment says that APs and STAs may ***disable*** protection if they receive an MMDPU (in this case a Beacon) with the **Use_Protection** bit set to 0 from within their own BSS. Depending on how the NonERP station leaves a BSS (roaming, powering down, going out of range, improperly removing a radio card, etc) will determine whether or not that AP disables use of protection and sets the **NonERP_Present** bit to 0 (disabled) in a timely manner. APs normally disable protection and disables the **NonERP_Present** bit immediately after the last NonERP STA has left the BSS in a manner that lets the AP know it is leaving. Simply removing a PC Card from a laptop will typically result in the association staying valid on the AP for a specified timeout period (30 – 120 seconds is usually the default setting for this).

Section 7.3.2.13 of the 802.11g amendment says that enabling protection may be triggered by a nearby AP's Beacons if the **NonERP_Present** bit is enabled, and disabling protection works in the same way, but in reverse. NonERP APs can cause ERP APs to disable protection, except for when there is a compelling reason not to, such as when a NonERP STA is associated to an ERP AP that is hearing Beacons from another ERP AP which has the **NonERP_Present**

disabled. The fact that a NonERP STA is associated to the ERP AP would override what another ERP AP's Beacons are saying.

This situation of constantly enabling and disabling protection in each BSS due either to NonERP STAs joining and leaving, due to instructions heard from other ERP APs in Beacons' ERP IE, or hearing Beacons from NonERP APs is what we refer to as "Protection Ripple." This ripple follows each NonERP STA as it roams throughout an enterprise WLAN's coverage area.

Suppose that a NonERP STA successfully roams from one ERP AP to another in an ESS. This will cause the new AP to enable protection (if it wasn't already enabled from having heard the old AP's Beacons), and the new AP will notify the old AP that it has now reassociated this particular client. This will cause the old AP to drop the association with the NonERP STA immediately. When the new AP enables the NonERP_Present and Use_Protection bits, its Beacons may then cause the old AP to immediately re-enable protection (depending on whether the old and new APs can hear each other's Beacons) even though the NonERP STA has left its BSS. This means that everywhere a NonERP STA goes in an ERP enterprise WLAN, protection mechanisms are not only enabled on the associated AP, but are also triggered elsewhere depending on which ERP APs can hear which other ERP APs. And please, do not mix 802.11b and 802.11g APs within an enterprise deployment – this is a recipe for protection mechanism disaster.

The chance that any single AP can hear one or more APs is very good in most enterprises. When a NonERP STA station roams, it causes a wave of **Use_Protection = 0** from the old ERP AP (to the nearby APs that can hear its Beacons) immediately followed by a wave of **Use_Protection = 1** from the new ERP AP across the enterprise WLAN as shown below. Imagine how many times this is happening when there are many NonERP clients roaming about the enterprise!

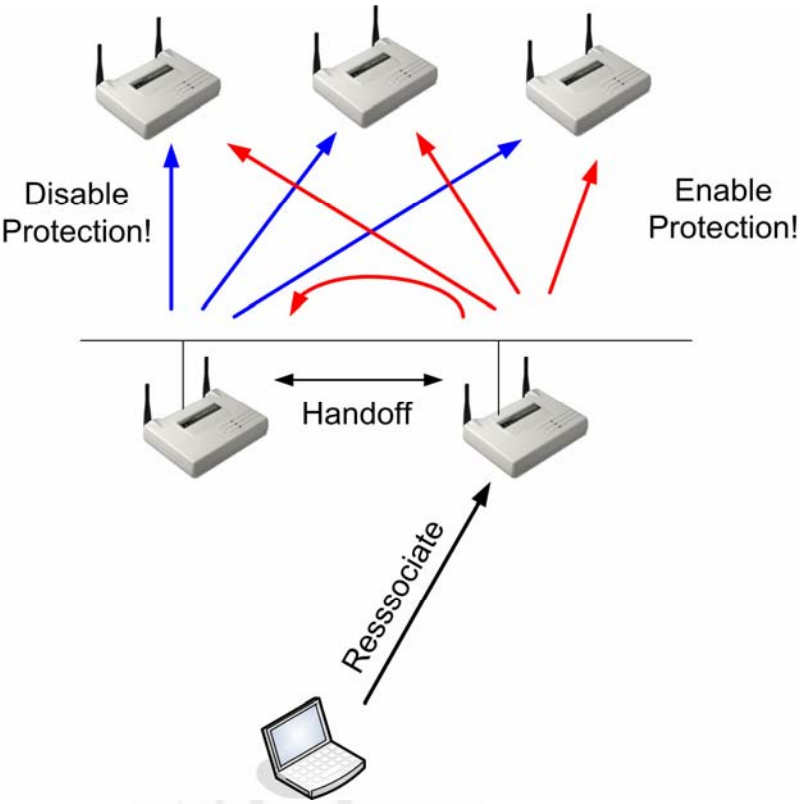


## Protection Ripple in Secure Environments

Let's now consider that many AP vendors and WLAN switch vendors use their APs as Wireless Intrusion Detection System (WIDS) sensors as well. For example, an AP may temporarily switch to scanning, then back to normal AP mode. Vendors will often explain that placing APs so that they can hear each other (even if only barely) makes sure you have good IDS coverage. In addition, many of the enterprise AP and WLAN Switch vendors are launching automated site survey tools. These tools often require that the APs can hear other APs so that an AP failure can be automatically compensated for by other APs. These types of features are great for some things, but they prove the point that it's common for APs to hear other APs and this ripple problem could easily be more severe in most ERP WLANs than ever realized.

In a WLAN secured by 802.1X/EAP, protection ripple could become very problematic because Open System Authentication is used to connect the STA to the AP prior to EAP authentication. See below for a screenshot that demonstrates the difference between connecting to the AP and connecting to the network. As you can see in this station's configuration screen, the station must use Open System Authentication in order to successfully 802.11 authenticate and associate to the AP prior to EAP authenticating and associating to the network using a user name and password. When configuring 802.1X/EAP authentication on access points or WLAN switches, you must enable Open System Authentication and EAP.

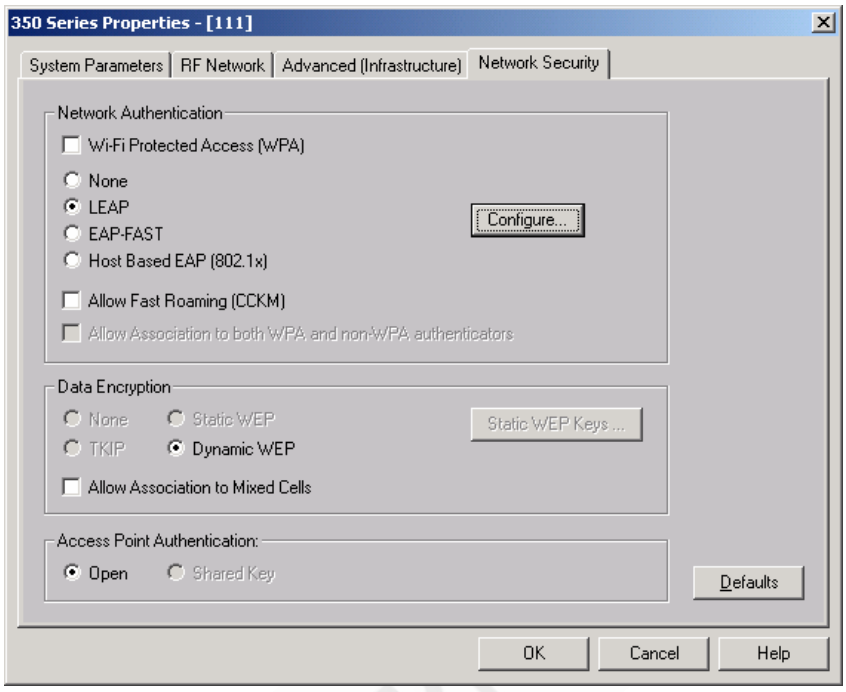


Below is a capture of a NonERP STA 802.1X/LEAP authentication to one of two APs configured exactly alike (channel, SSID, and LEAP user) in the same lab environment. The NonERP STA only EAP authenticates to a single AP of course, and here are the results.

| Packet | Source Physical | Dest. Physical | BSSID | Channel | Data Rate | Protocol |
|---|---|---|---|---|---|---|
| 2 | 00:0D:ED:A5:51:70 | FF:FF:FF:FF:FF:FF | 00:0D:ED:A5:51:70 | 11 | 1.0 | 802.11 Beacon |
| 3 | 00:0D:ED:A5:4F:70 | FF:FF:FF:FF:FF:FF | 00:0D:ED:A5:4F:70 | 11 | 1.0 | 802.11 Beacon |
| 4 | 00:0D:ED:A5:51:70 | FF:FF:FF:FF:FF:FF | 00:0D:ED:A5:51:70 | 11 | 1.0 | 802.11 Beacon |
| 5 | 00:0D:ED:A5:4F:70 | FF:FF:FF:FF:FF:FF | 00:0D:ED:A5:4F:70 | 11 | 1.0 | 802.11 Beacon |
| 6 | 00:0A:8A:47:BC:50 | 00:0D:ED:A5:4F:70 | 00:0D:ED:A5:4F:70 | 11 | 1.0 | 802.11 Auth |
| 7 | 00:0D:ED:A5:4F:70 | 00:0A:8A:47:BC:50 | | 11 | 1.0 | 802.11 Ack |
| 8 | 00:0D:ED:A5:4F:70 | 00:0A:8A:47:BC:50 | 00:0D:ED:A5:4F:70 | 11 | 11.0 | 802.11 Auth |
| 9 | 00:0A:8A:47:BC:50 | 00:0D:ED:A5:4F:70 | | 11 | 11.0 | 802.11 Ack |
| 10 | 00:0A:8A:47:BC:50 | 00:0D:ED:A5:4F:70 | 00:0D:ED:A5:4F:70 | 11 | 1.0 | 802.11 Assoc Req |
| 11 | 00:0D:ED:A5:4F:70 | 00:0A:8A:47:BC:50 | | 11 | 1.0 | 802.11 Ack |
| 12 | 00:0D:ED:A5:4F:70 | 00:0A:8A:47:BC:50 | 00:0D:ED:A5:4F:70 | 11 | 11.0 | 802.11 Assoc Rsp |
| 13 | 00:0A:8A:47:BC:50 | 00:0D:ED:A5:4F:70 | | 11 | 11.0 | 802.11 Ack |
| 14 | 00:0A:8A:47:BC:50 | 00:0D:ED:A5:4F:70 | 00:0D:ED:A5:4F:70 | 11 | 1.0 | EAPOL-Start |
| 15 | 00:0D:ED:A5:4F:70 | 00:0A:8A:47:BC:50 | | 11 | 1.0 | 802.11 Ack |
| 16 | 00:0D:ED:A5:4F:70 | 00:0A:8A:47:BC:50 | 00:0D:ED:A5:4F:70 | 11 | 11.0 | EAP Request |
| 17 | 00:0A:8A:47:BC:50 | 00:0D:ED:A5:4F:70 | | 11 | 11.0 | 802.11 Ack |
| 18 | 00:0D:ED:A5:4F:70 | 00:0A:8A:47:BC:50 | 00:0D:ED:A5:4F:70 | 11 | 11.0 | EAP Request |
| 19 | 00:0A:8A:47:BC:50 | 00:0D:ED:A5:4F:70 | | 11 | 11.0 | 802.11 Ack |
| 20 | 00:0A:8A:47:BC:50 | 00:0D:ED:A5:4F:70 | 00:0D:ED:A5:4F:70 | 11 | 1.0 | EAP Response |
| 21 | 00:0D:ED:A5:4F:70 | 00:0A:8A:47:BC:50 | | 11 | 1.0 | 802.11 Ack |
| 22 | 00:0A:8A:47:BC:50 | 00:0D:ED:A5:4F:70 | 00:0D:ED:A5:4F:70 | 11 | 1.0 | EAP Response |
| 23 | 00:0D:ED:A5:4F:70 | 00:0A:8A:47:BC:50 | | 11 | 1.0 | 802.11 Ack |
| 24 | 00:0D:ED:A5:4F:70 | 00:0A:8A:47:BC:50 | 00:0D:ED:A5:4F:70 | 11 | 11.0 | EAP Request |
| 25 | 00:0A:8A:47:BC:50 | 00:0D:ED:A5:4F:70 | | 11 | 11.0 | 802.11 Ack |
| 26 | 00:0D:ED:A5:51:70 | FF:FF:FF:FF:FF:FF | 00:0D:ED:A5:51:70 | 11 | 1.0 | 802.11 Beacon |
| 27 | 00:0A:8A:47:BC:50 | 00:0D:ED:A5:4F:70 | 00:0D:ED:A5:4F:70 | 11 | 1.0 | EAP Response |
| 28 | 00:0D:ED:A5:4F:70 | 00:0A:8A:47:BC:50 | | 11 | 1.0 | 802.11 Ack |
| 29 | 00:0D:ED:A5:4F:70 | 00:0A:8A:47:BC:50 | 00:0D:ED:A5:4F:70 | 11 | 11.0 | EAP Success |
| 30 | 00:0A:8A:47:BC:50 | 00:0D:ED:A5:4F:70 | | 11 | 11.0 | 802.11 Ack |
| 31 | 00:0A:8A:47:BC:50 | 00:0D:ED:A5:4F:70 | 00:0D:ED:A5:4F:70 | 11 | 1.0 | EAP Request |
| 32 | 00:0D:ED:A5:4F:70 | 00:0A:8A:47:BC:50 | | 11 | 1.0 | 802.11 Ack |
| 33 | 00:0D:ED:A5:4F:70 | FF:FF:FF:FF:FF:FF | 00:0D:ED:A5:4F:70 | 11 | 1.0 | 802.11 Beacon |
| 34 | 00:0D:ED:A5:4F:70 | 00:0A:8A:47:BC:50 | 00:0D:ED:A5:4F:70 | 11 | 11.0 | EAP Response |
| 35 | 00:0A:8A:47:BC:50 | 00:0D:ED:A5:4F:70 | | 11 | 11.0 | 802.11 Ack |
| 36 | 00:0D:ED:A5:4F:70 | 00:0A:8A:47:BC:50 | 00:0D:ED:A5:4F:70 | 11 | 11.0 | EAPOL-Key |
| 37 | 00:0A:8A:47:BC:50 | 00:0D:ED:A5:4F:70 | | 11 | 11.0 | 802.11 Ack |
| 38 | 00:0D:ED:A5:4F:70 | 00:0A:8A:47:BC:50 | 00:0D:ED:A5:4F:70 | 11 | 11.0 | EAPOL-Key |
| 39 | 00:0A:8A:47:BC:50 | 00:0D:ED:A5:4F:70 | | 11 | 11.0 | 802.11 Ack |
| 40 | 00:0D:ED:A5:51:70 | FF:FF:FF:FF:FF:FF | 00:0D:ED:A5:51:70 | 11 | 1.0 | 802.11 Beacon |
| 41 | 00:0D:ED:A5:4F:70 | FF:FF:FF:FF:FF:FF | 00:0D:ED:A5:4F:70 | 11 | 1.0 | 802.11 Beacon |
| 42 | 00:0D:ED:A5:51:70 | FF:FF:FF:FF:FF:FF | 00:0D:ED:A5:51:70 | 11 | 1.0 | 802.11 Beacon |

In the graphic above, you will notice Beacons originating from two APs. The NonERP STA has chosen to EAP authenticate against AP 4F:70. Frames 2-5 each show protection disabled as shown below.

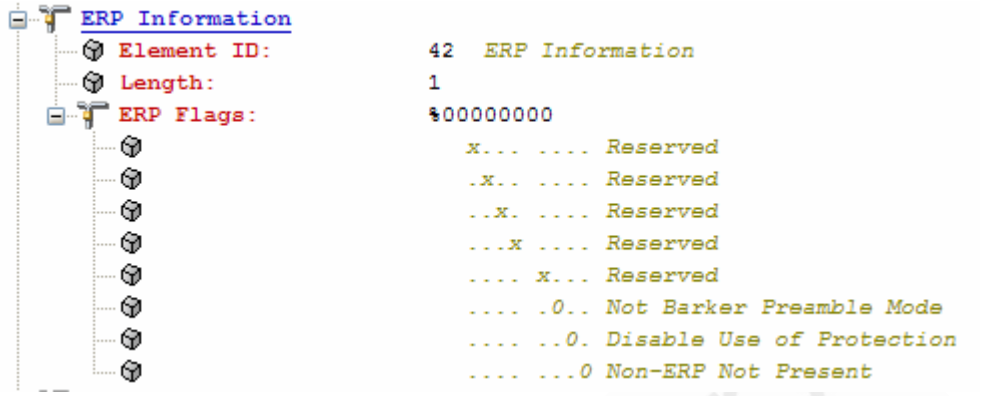


Frame 26 also continues to show protection disabled – because the NonERP STA is not associating to it. In frame 33, sent by 4F:70, protection is enabled (as shown below) even though the 802.1X/EAP authentication has not completed. At this point, the NonERP STA has only performed Open System Authentication and Association (as shown in frames 6-13).

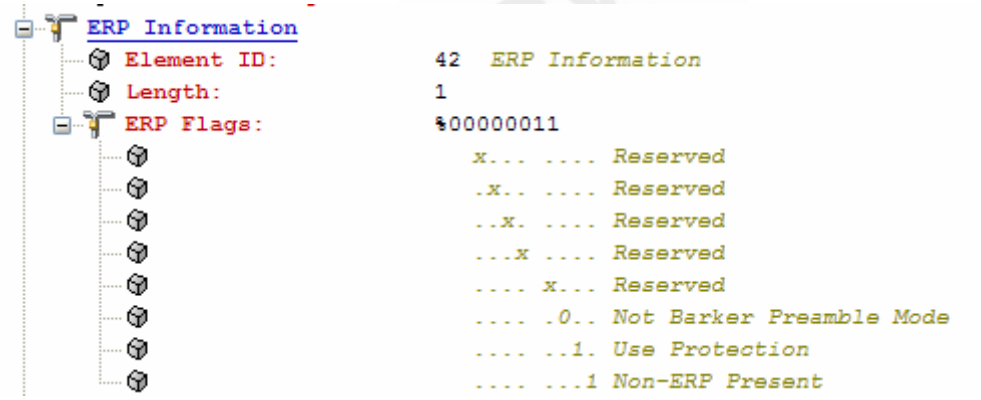


Frame 40 is the first Beacon transmitted by 51:70 after 4F:70 enabled the **Use_Protection** and **NonERP_Present** bits in the ERP Information Element in its Beacons. Notice below that 51:70 has now enabled the **Use_Protection** bit in its own Beacons.

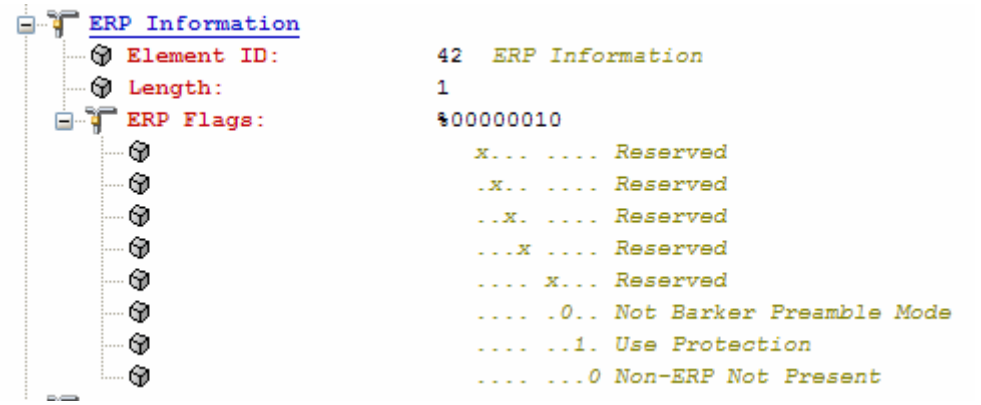

This demonstration shows us that even those NonERP STAs with an invalid 802.1X/EAP login would be successful at enabling protection on an ERP AP secured by 802.1X/EAP. To finalize our testing, we configured one ERP AP to a different channel and set password to a wrong value. We retested this scenario and saw that the NonERP STA was unable to EAP authenticate, but it still triggered protection mode on the AP because it had 802.11 Open System authenticated successfully, which would always be the case. This in turn enabled protection on the second ERP AP (now operating on a different channel than the first AP). Even ERP APs which deny access to NonERP STAs unintentionally become a weapon of mass destruction on the enterprise WLAN, proliferating the use of protection mechanisms across the network instantly to any ERP APs within earshot.

Let's go a few steps deeper, shall we? Not only is an ERP AP required to set the **NonERP_Present** bit to 1 (which triggers the enabling of the **Use_Protection** bit) when a NonERP STA associates, but even before that, the Association Request frame from a NonERP STA meets the specification of section 7.3.2.13 by having a supported rate set of only 11, 5.5, 2 and 1 Mbps. The 802.11g amendment says that the **NonERP_Present** bit may be set to 1 when, "*A management frame (excluding a Probe Request) is received where the supported rate set includes only clause 15 or clause 18 rates.*" Clause 15 is 802.11 DSSS, and Clause 18 is 802.11b DSSS. To nail down the exact event that is causing protection to be enabled on 802.1X/EAP APs, realize that the first frame that the ERP AP will receive that meets this requirement is the Association Request frame. This means that a NonERP station coming in contact with an 802.1X/EAP AP will cause the ERP AP to enable protection even before it is 802.11 associated. Remember that Open System is required in 802.1X/EAP networks. The NonERP STA only needs to send the Association Request frame to the ERP AP to get the party started.

Let's now look at the automatic cleanup – proliferation of "**Use_Protection = 0**." Below is a frame capture of the same two ERP APs used above. AP 4F:70 currently has a NonERP STA associated.

| Packet | Source Physical | Dest. Physical | BSSID | Channel | Data Rate | Protocol |
|---|---|---|---|---|---|---|
| 1 | 00:0D:ED:A5:4F:70 | FF:FF:FF:FF:FF:FF | 00:0D:ED:A5:4F:70 | 11 | 1.0 | 802.11 Beacon |
| 2 | 00:0D:ED:A5:51:70 | FF:FF:FF:FF:FF:FF | 00:0D:ED:A5:51:70 | 11 | 1.0 | 802.11 Beacon |
| 3 | 00:0D:ED:A5:4F:70 | FF:FF:FF:FF:FF:FF | 00:0D:ED:A5:4F:70 | 11 | 1.0 | 802.11 Beacon |
| 4 | 00:0D:ED:A5:51:70 | FF:FF:FF:FF:FF:FF | 00:0D:ED:A5:51:70 | 11 | 1.0 | 802.11 Beacon |
| 5 | 00:0D:ED:A5:4F:70 | FF:FF:FF:FF:FF:FF | 00:0D:ED:A5:4F:70 | 11 | 1.0 | 802.11 Beacon |
| 6 | 00:0D:ED:A5:51:70 | FF:FF:FF:FF:FF:FF | 00:0D:ED:A5:51:70 | 11 | 1.0 | 802.11 Beacon |
| 7 | 00:0D:ED:A5:4F:70 | FF:FF:FF:FF:FF:FF | 00:0D:ED:A5:4F:70 | 11 | 1.0 | 802.11 Beacon |
| 8 | 00:0D:ED:A5:4F:70 | 00:40:96:A1:9A:F9 | 00:0D:ED:A5:4F:70 | 11 | 11.0 | 802.11 Disassoc |
| 9 | 00:0D:ED:A5:4F:70 | 00:40:96:A1:9A:F9 | 00:0D:ED:A5:4F:70 | 11 | 11.0 | 802.11 Disassoc |
| 10 | 00:0D:ED:A5:4F:70 | 00:40:96:A1:9A:F9 | 00:0D:ED:A5:4F:70 | 11 | 11.0 | 802.11 Disassoc |
| 11 | 00:0D:ED:A5:4F:70 | 00:40:96:A1:9A:F9 | 00:0D:ED:A5:4F:70 | 11 | 11.0 | 802.11 Disassoc |
| 12 | 00:0D:ED:A5:4F:70 | 00:40:96:A1:9A:F9 | 00:0D:ED:A5:4F:70 | 11 | 5.5 | 802.11 Disassoc |
| 13 | 00:0D:ED:A5:4F:70 | 00:40:96:A1:9A:F9 | 00:0D:ED:A5:4F:70 | 11 | 5.5 | 802.11 Disassoc |
| 14 | 00:0D:ED:A5:4F:70 | 00:40:96:A1:9A:F9 | 00:0D:ED:A5:4F:70 | 11 | 2.0 | 802.11 Disassoc |
| 15 | 00:0D:ED:A5:4F:70 | 00:40:96:A1:9A:F9 | 00:0D:ED:A5:4F:70 | 11 | 2.0 | 802.11 Disassoc |
| 16 | 00:0D:ED:A5:51:70 | FF:FF:FF:FF:FF:FF | 00:0D:ED:A5:51:70 | 11 | 1.0 | 802.11 Beacon |
| 17 | 00:0D:ED:A5:4F:70 | FF:FF:FF:FF:FF:FF | 00:0D:ED:A5:4F:70 | 11 | 1.0 | 802.11 Beacon |
| 18 | 00:0D:ED:A5:51:70 | FF:FF:FF:FF:FF:FF | 00:0D:ED:A5:51:70 | 11 | 1.0 | 802.11 Beacon |
| 19 | 00:0D:ED:A5:4F:70 | FF:FF:FF:FF:FF:FF | 00:0D:ED:A5:4F:70 | 11 | 1.0 | 802.11 Beacon |
| 20 | 00:0D:ED:A5:51:70 | FF:FF:FF:FF:FF:FF | 00:0D:ED:A5:51:70 | 11 | 1.0 | 802.11 Beacon |

Taking a look at frames 6 & 7 respectively, we notice that 51:70 is announcing **Use_Protection = 1**, and 4F:70 is announcing **Use_Protection = 1** and a **NonERP_Present = 1**.

 

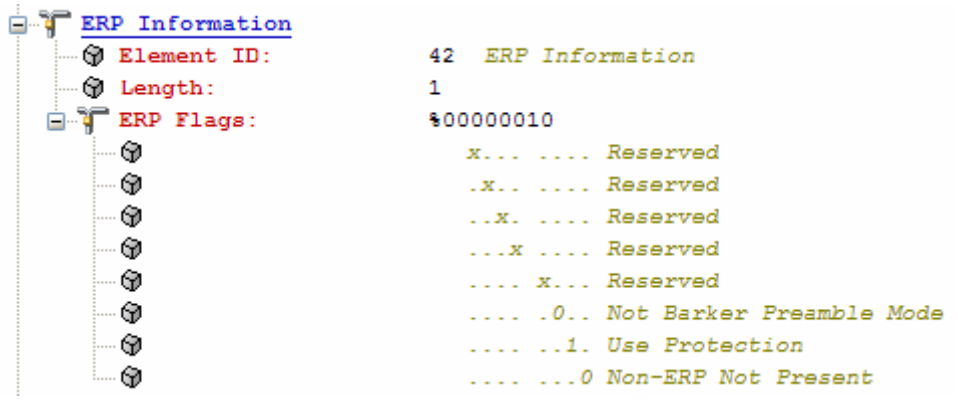


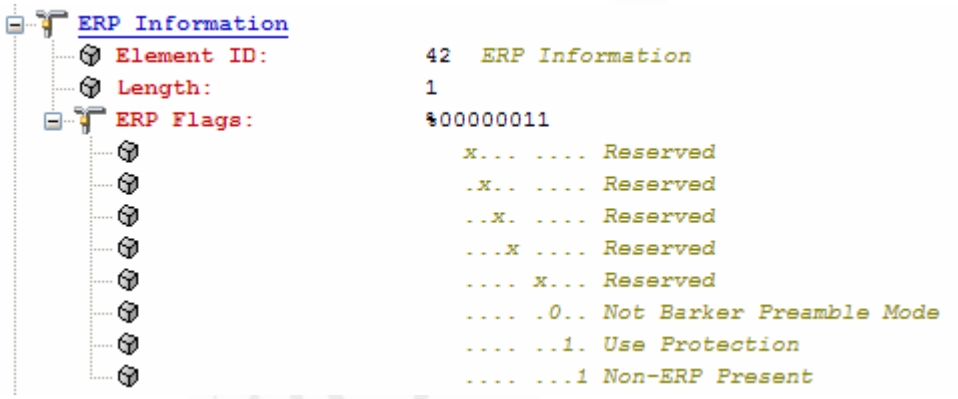


It is 4F:70 that has instigated protection on the WLAN. We used the Disassociate feature in the 4F:70 AP to kick the NonERP STA off the WLAN (frames 8-15). Now looking at frame 16, we notice that even though 4F:70 has given the NonERP STA the boot, nobody has yet told 51:70 about it.

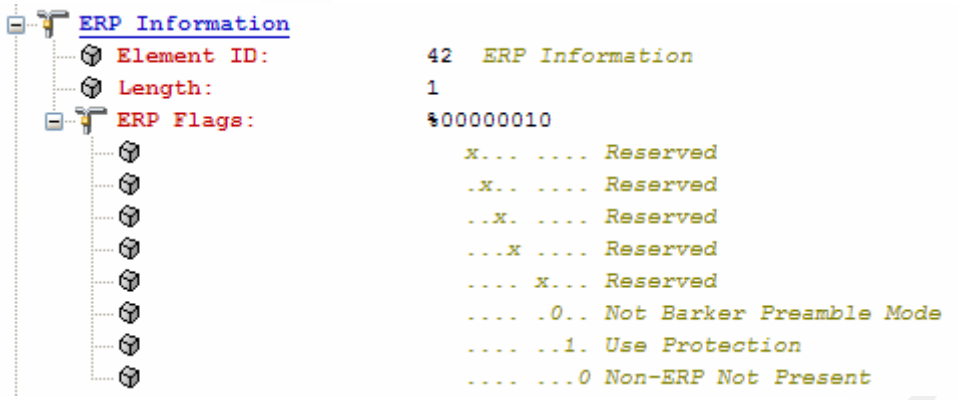


We can see in frame 17 below that 4F:70 announces **Use_Protection = 0**. The very next Beacon from 51:70 shows that it has followed suit.

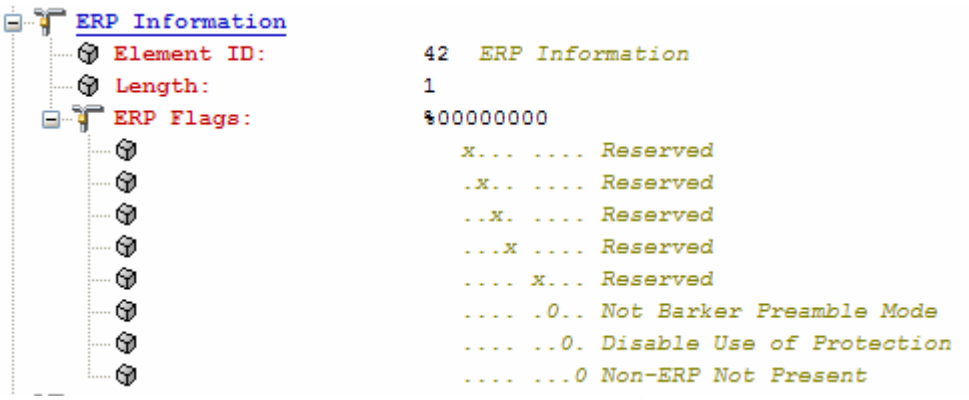


What can we do about this problem? One way to keep this protection ripple situation from getting out of control (its modus operandi) is to disable broadcasting of the SSID in the Beacon and disable responses to Probe Request frames with blank (null) SSID fields. By doing so, STAs that do not specifically have the correct SSID configured will not successfully authenticate using Open System Authentication. If they cannot authenticate, they will not send Association Request frames. Keep in mind that finding a network's SSID is not very difficult with a wireless protocol analyzer, so this precaution is only to prevent accidental triggering of protection. You could rectify this problem by simply disabling 802.11b rate support on all access points. The headache then is that PDAs and other small computing devices which do not support 802.11g would be eliminated from the WLAN, which may be unacceptable in many organizations. There are other adverse effects that disabling of clause 15 and clause 18 data rate sets might have on an ESS, but that is beyond the scope of this whitepaper.

In closing, I'm sure I should sum up what all of this means in a nutshell. It means, "Do not use 802.11b anywhere near 802.11g networks if you hope to realize the performance gain you expected when you upgraded your network from 802.11b to 802.11g."