


The 5GHz “Problem” For Wi-Fi Networks: DFS

[Skip to main content](#)



[WifiNigel](#)

The 5GHz “Problem” For Wi-Fi Networks: DFS

- [May 27, 2018](#)

Wi-Fi networking provides us with 2 bands for the operation of wireless LAN networks: the 2.4Ghz band and the 5GHz band. The 2.4GHz band has a reputation of being something of a “sewer” of a band, due to its limited number of usable channels, the number of Wi-Fi devices already using the band, and the high levels of non-Wi-Fi interference that it experiences. Many wireless LAN professionals will generally advise that you put your “important stuff” on the 5GHz band

whenever possible. 5GHz has far more channels available, a corresponding lower number of devices per channel, and generally suffers much lower non-Wi-Fi interference. However, beneath the headline of "2.4Ghz = bad, 5Ghz = good", there lurks a shadowy figure that can be troublesome if you're not aware of its potential impact: DFS.

Background

Wi-Fi networks operate in areas of RF spectrum that require no licence to operate. This is in contrast to many other areas of the radio spectrum that generally require some form of (paid-for) licence to operate radio equipment.

All wireless services are generally subject to a range of enforceable technical restrictions to ensure they operate in a manner that will minimize interference to other wireless services. This may include restrictions on parameters such as RF transmit power levels and limiting the spectral characteristics of transmitted signals (e.g. channel widths used, spectral masks etc.).

Even though they may be licence-exempt, Wi-Fi networks are still subject to restrictions to minimize their impact on other wireless services and equipment in the same areas of spectrum used by WLANs.

One particular service that shares spectrum with wireless LANs is radar. Some types of radar installation operate in the 5GHz band that is used by Wi-Fi network. This means that they may use some of the same

radar stations that have been deployed; there are many radar installations do not use 5GHz.

However, due to the coexistence of both radar and Wi-Fi networks in the same area of spectrum, the Wi-Fi standard (IEEE 802.11) was designed to incorporate a spectrum sharing mechanism on 5GHz to ensure that Wi-Fi networks do not operate on frequencies (hence causing interference) that are used by nearby radar stations. This mechanism is known as Dynamic Frequency Selection (DFS) and is designed to mitigate interference to 5GHz radar by WLANs.

How Does DFS Work?

DFS operation is as follows:

Channel Availability

Before an AP will use a channel that may be impacted by radar, it will perform a "Channel Availability Check" to check for radar signals on that channel. The AP will listen for 60 seconds for the presence of radar signals. If no radar is detected, then the channel is designated as being an "Available Channel".

When powering up an AP that uses a DFS channel, you will see that the 2.4GHz radio becomes available as soon as the AP has completed its boot sequence, but the 5GHz radio may not be available for another minute. This is due to the AP performing its channel availability check, if the AP is trying to use a DFS-impacted 5GHz channel.

In some regions, where channels 120 – 128 are allowed for use by Wi-Fi networks, there may be an increased channel availability check of 10 minutes. This means that the 5GHz radio is not available until 10 minutes after the access point has booted up. This extended checking period is due to weather radar restrictions on those channels.

In-Service Monitoring

Once an AP is operating on a DFS channel, it has to monitor for the presence of radar signals appearing on that channel. This is known as “In-Service Monitoring”.

The AP must continuously monitor its channel for the presence of radar signals.

Channel Shutdown

If a radar signal is detected, then the AP must cease transmissions on the channel within the “Channel Move Time”, which is 10 seconds in the EU/UK.

At the end of this period, the AP will have ceased transmissions and moved to a new channel.

Prior to moving channels, some WLAN solutions may provide a “Channel Switch Announcement” 802.11 frame to connected clients to advise them which channel the AP will be moving to. Support for this on both WLAN infrastructure and client equipment seems to be optional from my own observations and should not be relied upon as a reliable method for clients to find the AP on its new channel.

Experience shows that there are variations between WLAN solutions around which channels an AP will choose to move to when radar is detected. In some solutions, APs that detect radar will move to channel 36 exclusively. In other solutions, APs will choose to move to any of the available non-DFS channels. Some will jump to any available 5GHz channel they find (DFS or non-DFS). Behaviour in this area seems to be inconsistent and is not defined within the 802.11 standard.

Non-Occupancy Period

Once radar has been detected on a channel, then the "Non-Occupancy Period" begins. This is a 30-minute period in which no further transmissions will be made by the AP on the affected channel.

At the end of the 30-minute period, most APs will attempt to return to their original channel, subject to a channel availability check. (Again behaviour in this area varies between vendors)

Radar Signal Characteristics


Radar signals themselves are very short duration pulses of Radio Frequency energy. In contrast to WLAN signals, they have no specific framing format, which makes their identification quite challenging.

Looking at the testing methodology in ETSI EN 301 893 V2.1.1 (Annex D), test pulses sent to WLAN gear under test may vary between 0.5 and 30 micro-seconds and be subject to a variety of test patterns. The table below is an extract from the document:

Radar test signal # (see note 1 to note 3)	Pulse width W (μ s)		Pulse repetition frequency PRF (PPS)		Number of different PRFs	Pulses per burst for each PRF (PPB) (see note 5)
	Min	Max	Min	Max		
1	0,5	5	200	1 000	1	10 (see note 6)
2	0,5	15	200	1 600	1	15 (see note 6)
3	0,5	15	2 300	4 000	1	25
4	20	30	2 000	4 000	1	20
5	0,5	2	300	400	2/3	10 (see note 6)
6	0,5	2	400	1 200	2/3	15 (see note 6)

(Credit: Extract from ETSI standard: EN 301 893 V2.1.1)

The diagram below shows a single burst pattern that may be used to test WLAN devices:



(Credit: Extract from ETSI standard: EN 301 893 V2.1.1)

There is little doubt, that compared to the detection of well-structured, longer duration 802.11 frames, WLAN equipment has been set quite a challenge in reliably detecting radar signals (which can lead to annoying side-effects, discussed later).

Are All 5GHz Channels Subject to DFS?

No, not all channels in the 5GHz band are subject to DFS. The channels that are exempt vary from country to country, as dictated by local regulations. In the UK/EU, channels 36, 40, 44 and 48 are not subject to DFS. However, all remaining channels are subject to DFS. In the USA, channels 36 - 48, together with 149 - 165 are exempt from DFS operation, with all remaining channels requiring DFS operation. (Check your local spectrum regulation authority for the latest information for your country).

Channels that are not subject to DFS operate without having to perform any radar checks. Therefore, they are not subject to any disruptions from local radar equipment (or any other sources RF interference that may cause false-positive detection)

What Happens To Clients During a DFS Event?

Devices that are subject to DFS checks are divided in to two roles: master and slave. It is the role of the master device to advise slave devices when radar has been detected and that a channel shutdown is required. In WLANs, the access point is usually the master device, with the associated clients designated as slaves.

Once radar is detected, it is the duty of the master device to advise the slaves that a channel change is imminent via a channel switch announcement message. This message should advise slaves (clients) which channel the AP intends to move to.

What Is The Impact on Client Applications During a DFS Event?

Once a radar signal has been detected, the impact on clients due to the required channel change is “variable”.

WLAN systems may or may not send a channel switch announcement (CSA). If no announcement is received by a client (or is lost in transit), then the client will be forced to go through its probing process to find a suitable BSSID with which to associate. Depending on the network configuration and client capabilities (e.g. 802.11k/v/r), the time to re-associate with the network will vary. Note that even if a CSA is received, a client may still choose to go through its own AP discovery process based on probing or 80.211k information it has received.

Once the move to a new channel has been completed, there will then be the usual delays in the resumption of application data flow due to processes such network access authentication and DHCP exchanges – these will again vary with network configuration.

Whatever the configuration of the WLAN and client capabilities, the move to a new channel will not be without some connectivity impact. This impact may be unnoticeable for users who are using non-real-time applications (e.g. mail, web browsing), but will certainly have an impact on latency sensitive, real-time applications (e.g. voice, video).

Although DFS is, in theory, a great idea to protect systems that share the 5GHz spectrum, it has a major pitfall: false positives.

Detecting a radar signal signature is quite a tricky business. Due to the variety of radar signatures that may be detected, together with the short-duration nature of radar signals, false positive events may be quite frequent in some WLAN systems.

A false positive means that an AP is fooled into thinking that a radar signal is present by a non-radar RF signal. This causes a channel change, when one is not needed. This obviously leads to un-necessary WLAN disruption, that has varying impact on clients, depending on the applications in-use.

Theories around the exact cause of false positive events seem to be numerous, depending on who you speak with. I’ve heard the following possible causes cited:

- Transient conditions due to high densities of clients
- Bad client drivers causing short term RF spikes
- Co-channel interference from distant APs on the same channel
- Local non-Wi-Fi equipment interference

Whatever the cause, the false positives observed generally tend to be observed during times of increased user presence (i.e. they seem more likely during “office hours” as user numbers & activity increase).

How Do I Detect DFS Events?

To find out if your network is being impacted by DFS events, you need to check the trap logs or syslog messages from your wireless system.

All systems should report when a radar hit has been detected. This will generally be recorded in the logs of the AP, wireless controller or management system. Often this will be forwarded as an SNMP trap to a management system or perhaps as a syslog message to your logging server.

If you have log analysis and trending capabilities, it is well worth monitoring radar events to look for patterns of behaviour (e.g. particular sites, event times and channels)

What Do “Real” DFS Events Look Like?

You might be scratching your head at this point wondering how you can tell the difference between “real” DFS events and false positives.

In my experience, DFS events caused by genuine radar systems tend to be limited to a specific subset of channels on the 5GHz band. For instance, you may check your system logs and find that in a particular building, only channels 116 and 120 (for example) are reporting DFS events (i.e. radar hits). Also, these tend to be at a consistent rate throughout the day.

In contrast, false positives tend to be spread across a very wide portion of the 5GHz band and will vary in frequency throughout the day. They also generally fall to very low levels outside of office hours and at

weekend (depending on the working patterns of your particular establishment).

How Can I Mitigate DFS Events on a Wi-Fi Network?

There are a few options available to try to mitigate the impact of DFS events on a WLAN:

1. For genuine radar events that are impacting a subset of the 5GHz band, simply exclude the impacted channels from any channel planning. If using static channel planning, then avoid using affected channels. If using an auto-RF mechanism (i.e. automated channel planning), then exclude the affected channels from those available to the configuration of the auto-RF process
2. Trying to mitigate false positives is a little more tricky. Options include:
 1. If you have sufficient non-DFS channels, do not use DFS channels at all in channel planning. This option is very much dependent on WLAN capacity requirements and the local regulatory domain in which the network operates
 2. Work with the WLAN vendor to find out if they have a more recent version of operating code that is less susceptible to DFS false positives. I have seen this approach used many times, with varying degrees of success
 3. Work with a vendor or VAR to try to identify any local sources of interference that may create false positives. Occasionally, it may be possible to identify a particular client type or item of non-Wi-Fi equipment that is causing false positives, If you have a support contract, assistance from a suitably qualified

WLAN expert armed with a spectrum analyser and able to perform log analysis may be invaluable in tracking down offenders.

4. If your vendor is unable to fix the issue, it may be worth trying an alternative vendor. Though this may seem extreme, I have seen huge variations between vendors and their susceptibility to DFS false positives. A limited-scope proof of concept costs little to deploy and can provide amazing leverage with your existing vendor...

What will be the impact of DFS on my Wi-Fi Network?

The impact of DFS events on your network (both "real" & false positive) will always be the same: a DFS event is detected and an AP will change channels. This will also cause the associated wireless clients to change channels.

The actual impact on the end-user will vary depending on what they are doing on their client.

Many applications that aren't latency sensitive will simply continue with little obvious impact on service. If a user is browsing the web, sending email or even streaming a video file (assuming some buffering), they will generally not notice their client jump between channels as their associated AP changes channels. This assumes your

WLAN is correctly designed so that clients have viable alternative APs available.

If clients are using real-time, latency sensitive applications, then they are much more likely to observe some sort of negative impact. The transition to a new channel is likely to be quite long (in WLAN terms). It will vary depending on the required operations (e.g. channel probing, 802.1X exchanges, DHCP exchanges etc.), but will generally be long enough to have an impact of real-time applications such as voice and real-time video. The use of enhanced WLAN features such as 802.11r/k may help to ensure that clients can significantly speed up the AP selection and roaming process.

These considerations provide a useful indication as to whether DFS events are going to provide problems on your WLAN and whether you should consider the impact of DFS on your wireless network.

Conclusion

In many Enterprise wireless WLANs, there will generally be a requirement to use as many unique 5GHz channels as possible. This provides opportunities to mitigate co-channel interference and increase capacity through the use of channel bonding (if required).

However, understanding and verifying the impact (if any) of radar detection is important to ensure the requirements of our WLAN design are not compromised.


References

Here are some other great sources of information you might like to look at for more information about DFS:

- Designing with DFS Channels (Devin Akin Webinar):
<https://www.brighttalk.com/webcast/5522/279895/designing-with-dfs-channels>
- A Practical Introduction to DFS (Blog post):
<https://www.adriangranados.com/blog/practical-intro-dfs>
- Dynamic Frequency Selection Parts 1 to 3 (Jennifer Jabbusch articles) :
<https://www.networkcomputing.com/wireless/dynamic-frequency-selection-why-its-critical-80211ac/1439499958>
- (ETSI) EN 301 893 standard:
http://www.etsi.org/deliver/etsi_en/301800_301899/301893/01.07.01_60/en_301893v010701p.pdf
- Wikipedia WLAN Channels:
https://en.wikipedia.org/wiki/List_of_WLAN_channels
- UK 5GHz channel summary (WiFiNigel PDF):
<http://bowdennetworks.co.uk/downloads/5GHz%20Spectrum%20Usage%20UK%20-%202017%20-%20v1.pdf>

Shop Related Products

[Ads by Amazon](#)



[Network Security Through Data Analysis: Building Situational Awar...](#)

~~\$44.95~~ \$40.00

Bestseller

(13)

DEAL OF THE DAY

ENDS IN



[Basic Well Log Analysis, 2nd Edition](#)

\$59.00

Bestseller

(14)

DEAL OF THE DAY

ENDS IN



[Universal Laptop Mini PCI-E Wireless WiFi Internal Antenna ipx...](#)

\$7.98

Bestseller

(38)

DEAL OF THE DAY

ENDS IN



[Dueling with Kings: High Stakes, Killer Sharks, and the Get-Rich P...](#)

\$11.99

Bestseller

(126)

DEAL OF THE DAY

ENDS IN

[Ads by Amazon](#)

Popular posts from this blog

[Microsoft NPS as a RADIUS Server for WiFi Networks: Dynamic VLAN Assignment](#)

– [March 18, 2014](#)



[Read more](#)

[Microsoft NPS as a RADIUS Server for WiFi Networks: Self Signed Certificate](#)

– [March 15, 2014](#)



[Read more](#)

WiFi Nigel (CWNE #135)



Archive



Labels



Profile

[About Me](#)

IT Blog Awards



Follow Me

- [Twitter](#)
- [YouTube](#)
- [GitHub](#)

Podcast




Wi-Fi For Beginners Podcast

Useful Links / Wi-Fi Blogs

- [My YouTube Channel](#)

- [My GitHub Repositories](#)
- [My LinkedIn Profile](#)
- [Wi-Fi For Beginners Podcast](#)
- [Revolution WiFi](#)
- [Devin Aiken](#)
- [Eddie Forero](#)
- [CWNP Blog](#)
- [Ekahau Blog](#)
- [Wireless LAN Pros](#)
- [My 802.11](#)



Viewed using [Just Read](#)

