



**WHITE PAPER**



## **802.11n for Enterprise Wireless LANs**

Copyright © 2010, Extricom Ltd. All rights reserved. Extricom, TrueReuse, UltraThin and the Extricom logo are trademarks of Extricom Ltd.

No part of this publication may be reproduced, stored in a retrieval system or transmitted, in any form or by any means, photocopying, recording or otherwise, without prior written consent of Extricom Ltd. No patent liability is assumed with respect to the use of the information contained herein.

While every precaution has been taken in the preparation of this publication, Extricom Ltd. assumes no responsibility for errors or omissions. The information contained in this publication and features described herein are subject to change without notice. Extricom Ltd. reserves the right at any time and without notice, to make changes in the product.

--

# Table of Contents

<b>EXECUTIVE SUMMARY .....</b>	<b>1</b>
<b>EVOLUTION OF THE IEEE 802.11 STANDARD .....</b>	<b>2</b>
NEW FEATURES OF 802.11N .....	3
<i>Physical Layer Improvements</i> .....	3
Higher Modulation Rates – from 54 Mbps to 58.5 Mbps .....	3
Improved Forward Error Correction – from 58.5 Mbps to 65 Mbps .....	3
Short Guard Interval – from 65 Mbps to 72.2 Mbps .....	3
Channel Bonding, 40 MHz Channels – from 72.2 Mbps to 150 Mbps .....	3
MIMO – from 150Mbps to 600Mbps .....	4
Maximal Ratio Combining (MRC) .....	4
<i>802.11n MAC Layer Improvements</i> .....	5
Aggregation of Packets .....	5
Block Acknowledge .....	7
Lower Overhead: Reduced Interframe Space .....	7
Power Save Modes .....	7
Backwards Compatibility .....	8
BENEFITS OF 802.11N.....	8
<i>Improved Performance - High Throughput</i> .....	8
<i>Improved Range and Coverage</i> .....	9
<i>Flawless Backwards Compatibility</i> .....	9
<i>More Capacity</i> .....	9
<b>THE REAL-WORLD 802.11N CHALLENGE .....</b>	<b>10</b>
CHANNEL BONDING .....	10
UNPREDICTABLE COVERAGE.....	12
CO-CHANNEL INTERFERENCE .....	13
PLANNING AND DESIGN.....	15
<i>What Applications Will Be Supported?</i> .....	16
<i>What Is the Deployment Model?</i> .....	17
<i>What Are the Site Characteristics?</i> .....	18
<i>Transmission Power Control</i> .....	18
<i>Rate versus Range</i> .....	19
<i>Client Mixed Mode and Backward Compatibility</i> .....	20
<i>Client Variations</i> .....	20
<i>Unbalanced Link</i> .....	21
<i>Frame Aggregation</i> .....	22
<i>MIMO Power Consumption</i> .....	22
<i>Wired Network Infrastructure</i> .....	23
<i>Power Save (PSMP)</i> .....	23
<i>Roaming</i> .....	24
Intelligent Roaming Mechanisms .....	24
Re-Establishing the Connection .....	24
<i>Edge Users</i> .....	25
<i>802.11k</i> .....	26
<b>THE EXTRICOM ADVANTAGE .....</b>	<b>28</b>
HOW EXTRICOM WORKS.....	28
WHY EXTRICOM WORKS BETTER FOR 802.11N .....	29

## Executive Summary

The 802.11 standard that started as single access point solution (“hot spot”), dramatically evolved, reaching speeds higher than wired Ethernet and providing infrastructure for crowded enterprise sites and organizations. With the ratification of the 802.11n standard, new technologies are used to increase performance, capacity, and coverage. On the client side, laptops, smart phones and other devices are available to be used in a variety of applications.

While enterprises are seeking more performance, the underlying fundamentals of wireless LAN (WLAN) have remained the same. This paper will review what additions to the standard are included in 802.11n, along with the challenges of implementing it in a real-world environment.

While deployment of WLAN in a home-like environment for a single AP or even a few APs may be straightforward, an enterprise deployment with a large AP count, dense client population, critical applications requirements and client roaming can become problematic. It can even reach the point that the wireless infrastructure “melts down” and can no longer effectively support the goals of the enterprise.

The situation is further complicated by the mixture of client types, each one with different algorithms for roaming and power save, the lack of available spectrum channels, single and aggregated client throughput requirements, dynamic coverage patterns formed by multi path beams, and more. Given the complexities of the picture, it’s no surprise that enterprises have looked for alternatives to conventional WLAN architectures.

One alternative leading approach is the Channel Blanket WLAN architecture from Extricom. While the Extricom WLAN offers multiple and varied benefits, its value in an 802.11n implementation boils down to the following:

- Full performance with concurrent support for legacy 802.11b/g or 802.11a devices
- Immunity to channel shortages and the unpredictability of MIMO coverage
- Supports a mix of 20MHz and 40MHz channels, in either band, for all client types, from the same physical infrastructure. This provides the ability to segregate client devices or segment radios per service
- Higher aggregated capacity by available channel numbers per radios
- No need for rate or range compromise

While the promise of 802.11n is clear, making that promise a reality in the enterprise is not as simple as it may seem. The solution chosen must incorporate a practical, scalable fulfillment path, and an Extricom WLAN is an excellent choice in this regard.

## Evolution of the IEEE 802.11 Standard

The Institute of Electrical and Electronic Engineers (IEEE) is a professional organization of electrical engineers that has developed a variety of standards. The IEEE 802 committee is responsible for networking standards that define the operation of the Medium Access Control (MAC) and physical (PHY) layers.

Ethernet (IEEE 802.3) and Token Ring (IEEE 802.5) are important wired LAN standards developed by the IEEE 802 committee. The IEEE 802.11 project was formed in the early 1990s to develop wireless LAN standards. The original 802.11 standard was published in 1997 and defined two different radio-based PHYs that offered data rates of 1 and 2 Mbps.

IEEE wireless standards have continually evolved, primarily driven by end-user demands and needs. In recent times, emerging technologies, especially video-based applications like videoconferencing or video streaming, have increasingly required higher speed and range than was available in the 802.11a/b/g standards.

802.11g promised support for such applications, but even though so-called “g” products have a theoretical maximum data speed of 54 Mbps, in real-world conditions data transmission speeds are reduced by 50% or more, thereby falling short of the applications’ requirements.

	802.11	802.11b	802.11a	802.11g	802.11n Draft 2	802.11n
<b>Ratified</b>	1997	1999	1999	2003	2007	2009
<b>Operating Band</b>	2.4GHz	2.4GHz	5GHz	2.4GHz	2.4GHz 5GHz	2.4GHz 5GHz
<b>Maximum Data Rate</b>	2 Mbps	11 Mbps	54 Mbps	54 Mbps	300 Mbps	600 Mbps
<b>Channel Width</b>	20 MHz	20 MHz	20 MHz	20 MHz	20 MHz 40 MHz	20 MHz 40 MHz

**Table 1:802.11 Evolutions**

The 802.11n standard was designed to address this need for speed, and 802.11n is the only Wi-Fi technology today with the bandwidth to support multiple HDTV streams at 20 Mbps each. But this performance also allows for implementing some long-standing goals of Wi-Fi networks, for example, combining wireless multimedia, voice/VoIP, data, video, and gaming in residential applications. Another objective was to achieve throughput, Quality of Service, and security levels that would compare favorably with those of Ethernet, as this is increasingly a business requirement for enterprise-grade, campus, and municipal networks.

IEEE 802.11n is a huge step forward for the standard and is a major improvement in almost every aspect of Wi-Fi. The current 802.11n products that are available today offer a variety of data rates up to 300 Mbps and ultimately will support up to 600 Mbps. In addition, 802.11n offers excellent backwards compatibility with previous versions of the standard.

## New Features of 802.11n

To achieve the goal of higher throughput goals, 802.11n defines a number of improvements at both the Physical and Data Link layers. The maximum performance is only possible when all of the performance enhancement mechanisms are employed by both the sender and receiver.

The enhancements of 802.11n are implemented in the context of maintaining compatibility with legacy based on earlier versions of the standard – 802.11a, 802.11b and 802.11g. 802.11n systems share the same spectrum and must coexist and interoperate with systems based on previous versions of the standard.

### *Physical Layer Improvements*

#### Higher Modulation Rates – from 54 Mbps to 58.5 Mbps

The 802.11n has continued the improvement in data transmission that 802.11a and g pioneered. These two standards leveraged orthogonal frequency division multiplexing (OFDM) to increase data rates. Under OFDM, a single 20-MHz radio channel was divided into 48 “sub-carriers” or smaller channels that could transmit data independently of each other.

It was this innovation that allowed maximum data rates to be increased from 11 Mbps of 802.11b to 54 Mbps of 802.11a and 802.11g using the same 20 MHz channel. 802.11n also uses OFDM data transmission, but with significant differences. The number of sub-carriers has increased from 48 to 52 in a 20-MHz channel.

A further innovation in 802.11n is that each spatial stream can use a different type of modulation, i.e. QPSK, 16QAM, or 64QAM, vastly increasing the number of data rates available. Although this theoretically greatly increases the operational flexibility of WLAN, in real-world conditions it may be difficult to harness it to any real effect.

#### Improved Forward Error Correction – from 58.5 Mbps to 65 Mbps

Forward Error Correction (FEC) is a system of control whereby the sender adds redundant data to allow the receiver to detect and correct errors. The original 3/4 coding rate is improved to 5/6 and that yields an improvement to 65 Mbps in a 20 MHz channel.

#### Short Guard Interval – from 65 Mbps to 72.2 Mbps

A technically arcane but significant performance improvement is the introduction of a Short Guard Interval or SGI feature with 802.11n. This is a 50% reduction in the “guard interval” (or non-transmission period) between symbols, the base unit of a wireless transmission (equivalent to a bit of data), from 800 to 400 nanoseconds. SGI is an optional feature and offers potential improvements in bandwidth, compared to a full guard interval when both the sending and receiving stations are able to support SGI.

#### Channel Bonding, 40 MHz Channels – from 72.2 Mbps to 150 Mbps

One of the major changes introduced by 802.11n is the ability to bond two 20 MHz channels into one channel of 40 MHz width. Doubling the channel width provides up to double the bandwidth to an individual user but comes at the expense of using twice the available spectrum for each channel. In the 2.4GHz band, only one independent 40 MHz channel is possible.




Figure 1: Channel Bonding in 2.4GHz

### MIMO – from 150Mbps to 600Mbps

The most notable new technology introduced in 802.11n is called multiple input, multiple output (MIMO). MIMO uses several antennas, spatial multiplexing, and sophisticated signal processing to transmit multiple data streams.

Instead of sending and receiving a single stream of data, MIMO can simultaneously transmit up to four spatial streams. The transmitting and receiving stations each have multiple RF chains with multiple antennas. The 802.11n mandates at least two and up to four spatial streams. Below is an example of two streams of data being transmitted and received by a station with three RF chains for receive. Each of these “spatial” streams arrives at the receiver with different amplitude (signal strength) and phase.




Figure 2: MIMO

Thanks to spatial multiplexing, multi-path interference (wireless signals bouncing off walls and other structural elements and arriving at the receiver at slightly different times) and fading are no longer an obstacle to wireless transmissions. In fact with MIMO, multi-path now is an enabler of significantly increased performance. The end-user experiences much faster data transmission, as well as significant increases in the range over which data can be transmitted.

### Maximal Ratio Combining (MRC)

While 802.11g access points have two antennas, they only have one radio receiver. A diversity switch allows the single receiver to briefly sample the signal quality on the two antennas, after

which it picks one to receive each packet. The signal from one antenna is used while the other antenna's signal is simply discarded.

In contrast, 802.11n systems have a full, dedicated receiver for each antenna. They operate simultaneously to capture up to four signals and add them together. By adjusting the weight and phase of each signal as they are added, noise received on multiple antennas can actually be cancelled out, resulting in the best possible signal-to-noise ratio for the combined signal; hence the name "Maximal Ratio Combining".

The enhanced throughput and link reliability from MRC can be enjoyed by any client device – it is fully backward compatible with 802.11a and 802.11g clients since no changes are required on the client transmitters.




Figure 3: MRC

## 802.11n MAC Layer Improvements

The 802.11n PHY improvements increase raw data rate, MAC layer improvements lead to more efficient operation and result in more of the raw data rate actually being delivered as useful throughput.


### Aggregation of Packets

The 802.11n standard includes changes to how packets are aggregated, with the intent of improving performance, especially on busy wireless networks.

The 802.11 MAC layer protocol has a large amount of overhead compared to wired network protocols, especially in the inter-frame spaces and control frames such as acknowledgements. When data is transmitted in short packets, the protocol overhead can actually be longer than the entire data frame. A further MAC layer issue is contention for the air and collisions, which reduces the maximum effective throughput of 802.11. The 802.11n standard introduces changes in the MAC layer to address the throughput impact from fixed overhead and contention losses.




Each 802.11 frame has fixed overhead in the radio preamble and MAC frame fields. Even where a high data rate is possible, the fixed overhead restricts actual throughput (illustrated in Figure 4).



**Figure 4: Frame Overhead**

The 802.11n standard introduces frame aggregation, a major improvement designed to address the MAC layer shortcomings of WLAN. Frame aggregation, in simple terms, puts two or more frames together into a single transmission. 802.11n includes two methods for frame aggregation: MAC Service Data Units (MSDU) aggregation and Message Protocol Data Unit (MPDU) aggregation. Both aggregation methods reduce the overhead to only a single radio preamble and MAC header for each frame transmission (Figure 5).



**Figure 5: Packet Aggregation**

By sending multiple frames as a single transmission, the amount of potential collisions, as well as time lost to back-off, is greatly reduced. To compensate for the larger aggregated frame size, 802.11n also has augmented the maximum frame size from 4 KB to 64 KB.

## Block Acknowledge

Rather than sending an individual acknowledge following each data frame, 802.11n introduces the technique of confirming a burst of up to 64 frames with a single Block ACK frame. The Block ACK even contains a bitmap to selectively acknowledge individual frames of a burst.

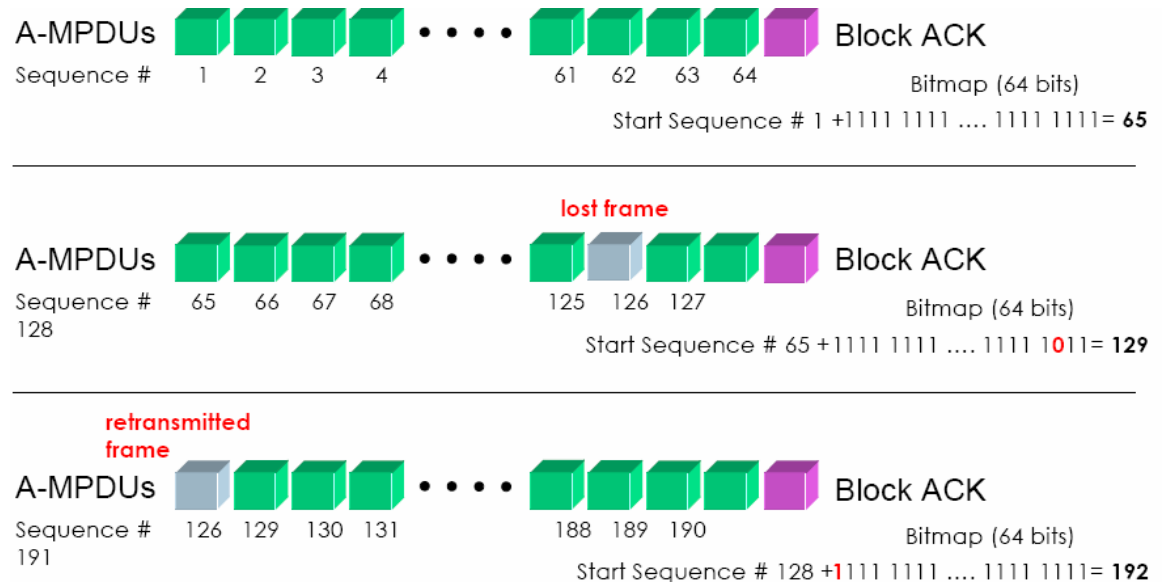


Figure 6: Block ACK

## Lower Overhead: Reduced Interframe Space

802.11n provides a mechanism to reduce the overhead involved with transmitting a stream of frames to different destinations. This mechanism reduces the interframe space between receiving a frame, typically an acknowledgement frame, and sending a subsequent frame.

The 802.11e extension for quality of service added the ability for a single transmitter to send a burst of frames during a single, timed transmit opportunity. During the transmit opportunity, the sender does not need to perform any random backoff between transmissions, separating its frames by the smallest allowable interframe space, the short interframe space (SIFS).

802.11n improves on this mechanism, reducing the overhead between frames, by specifying an even smaller interframe space, called the reduced interframe space (RIFS). RIFS cuts down further on the dead time between frames, increasing the amount of time in the transmit opportunity that is occupied by sending frames.

## Power Save Modes

The **spatial multiplexing (SM) power save** mode allows an 802.11n client to power down all but one of its radios. The static SM power save mode has the client turn off all but a single radio, becoming essentially equivalent to an 802.11a or an 802.11g client. The client's access point is notified that the client is now operating in the static single-radio mode, requiring the access point to send only a single spatial stream to this client until the client notifies the access point that its additional radios are again enabled and operating.

The dynamic SM power save mode also turns off all but one of the client's radios. But in this mode of operation, the client can rapidly enable its additional radios when it receives a frame

that is addressed to it. The client can immediately return to the low power state by disabling its additional radios immediately after its frame reception is complete.

The **Power Save Multi-Poll (PSMP)** mode extends the Automatic Power Save Delivery (APSD) mechanism defined in 802.11e. PSMP provides the same delivery-enable and trigger concepts, extending the ability of the client to schedule the frames that it transmits as the trigger for delivering the downlink frames.

This scheduling mechanism reduces the contention between clients and between the client and the access point. Reducing contention also reduces the time the client spends in backoff and reduces the number of times a frame must be transmitted before it is delivered successfully. This dramatically improves power conservation in the clients. PSMP is also a dynamic method that immediately adjusts to changes in traffic demand by the clients using it.

### Backwards Compatibility

Compatibility with existing 802.11a, b, and g devices is a critical issue addressed in 802.11n. Just as 802.11g provides a protection mode for operation with 802.11b devices, 802.11n has a number of mechanisms to provide backward compatibility with 802.11a, b, and g devices, allowing these devices to understand the information necessary to ensure that 802.11n devices can operate in the same area.

Like 802.11g, 802.11n transmits a signal that can't be decoded by devices built to an earlier standard. To avoid absolute chaos from massive interference and collisions, 802.11n operating in mixed mode transmits a radio preamble and signal field that can be decoded by 802.11a and 802.11g radios.

This provides enough information to the legacy radios to allow them to indicate that there is another transmission on the air and how long that transmission will last. Following the legacy preamble and signal field, the 802.11n device sends the remaining information using 802.11n rates and its multiple spatial streams, including an 802.11n preamble and signal field.

In addition to the legacy preamble and signal field, it can also be necessary to use additional protection mechanisms provided by 802.11g to allow the MAC in legacy devices to correctly determine when it is allowed to transmit and when it must perform backoff before transmission. The mechanism provided by 802.11g and utilized by 802.11n when either 802.11g or 802.11a devices are present is the CTS-to-self mechanism.

CTS-to-self allows the 802.11n device to transmit a short CTS frame, addressed to itself, that includes the timing information necessary to be communicated to the neighboring legacy MACs that will protect the 802.11n transmission that will follow. The CTS frame must be transmitted using one of the legacy data rates that a legacy device will be able to receive and decode.

## Benefits of 802.11n

### *Improved Performance - High Throughput*

The greatly enhanced 802.11n PHY introduces a number of options for improving the communications link between stations by utilizing more of antennas, transmitting simultaneous spatial streams, improved modulation method and forward error correction (FEC) coding options. When all of the performance features are utilized, 802.11n certainly

lives up to its charter of higher throughput. The new 802.11n MAC protocol is more efficient than previous 802.11 MACs, so it delivers a higher percentage of the raw data rate as useful throughput.

### *Improved Range and Coverage*

Many of the 802.11n mechanisms that are designed to increase the data rate also improve coverage. Multiple antennas with MRC and better RF receive chains lead to better coverage. 802.11 systems will still offer the rate versus range tradeoff. The highest data rates are available at shorter range. Greater range can be achieved by lowering the data rate. For 802.11n, a larger area will support the high data rates, and there will be less variability within the coverage area in most environments.

### *Flawless Backwards Compatibility*

As with any network standard in the 802 architecture, 802.11n is compatible with previous generations of the standard. Clients based on 802.11a/b/g will interoperate with 802.11n access points and operate at their full speed on 802.11n networks. 802.11n clients are able to operate on legacy 802.11a/b/g networks. 802.11n clients work with legacy infrastructure at the highest possible performance of the legacy standard and likely at increased range.

The 802.11n standard defines three types of networks to address backwards compatibility:

- **Legacy:** A network in which all stations use PHY and MAC features of the legacy 802.11a/b/g standards. 802.11n clients participate at the performance of 802.11a/b/g.
- **Mixed:** A network with 802.11n access points configured to support a mixture of legacy 802.11a/b/g as well as native 802.11n MAC and PHY features. Protocol mechanisms have been defined to allow clients based on the older standards to recognize and handle 802.11n traffic. This is the most general configuration, but the 802.11n clients operate at reduced performance in order to coexist with legacy clients.
- **Greenfield:** Stations are all 802.11n, communicating solely to other 802.11n stations. This is the highest performing configuration, but legacy clients will not be able to join the network or operate successfully on the same channel.

### *More Capacity*

More capacity is possible with 802.11n simply because of the higher data rate per channel. 802.11n also allows more possible channels. The extended range provided by 802.11n overcomes many of the coverage limitations of 5GHz 802.11a networks. The combined performance benefits of 802.11n enable more practical enterprise deployments at 5GHz. By utilizing both 2.4GHz and 5GHz channels, 802.11n will be able to operate effectively across many more channels and therefore deliver much higher capacity in a given area.


## The Real-World 802.11n Challenge

802.11n brings significant benefits to enterprises that have made WLAN part of their business model. But at the same time, it has a number of well-documented technical challenges which are important to understand.

### Channel Bonding

The problem lies with the wireless topology that has traditionally been the de facto standard: a cell-based WLAN. The 2.4GHz band supports 802.11b/g devices, and this band is the norm in most corporate implementations today.

In 2.4GHz, only three non-overlapping channels are available to plan the cells of the system, and most enterprise Wi-Fi implementations depend on a careful balancing of the signals of multiple APs. Bonding two channels together uses two-thirds of the spectrum for a single channel, leaving only one other 20MHz channel. This makes cell planning impossible.



**Figure 7: Channel bonding in 2.4GHz**

In the 5GHz band, there are theoretically 11 non-overlapping 40MHz channels available for 802.11n channel bonding. In reality, the number is much lower due to local regulation which reserves some 5GHz channels for governmental use and other technical factors; in some cases, as few as three 5GHz channels may be available if channels are bonded.

With so few channels, all the challenges of cell-planning previously experienced in 2.4GHz also start to appear in 5GHz. In indoor environments, the range of a 5GHz signal is approximately 40 to 60% of an equivalent 2.4GHz signal, since it is less capable of penetrating walls. 2.4GHz band is the default standard for many applications and devices.

Although in theory, there are 23 non-overlapping 20MHz channels in the 5GHz band there are several issues in utilizing them:

- Not all client devices are dual-band capable
- Varying power limitations on each channel result in unequal cell coverage. Transmit Power in channel bonding is split between channels in constant spectral density

- Channel bundle of 40MHz is allowed only in the following combinations, reducing the non-overlapping channels in 40MHz to 11

Frequency Band	Channel ID
UNII – 1 (lower)	36-40 or 44-48
UNII – 2 (middle)	52-56 or 60-64
UNII – 2 extended (high)	100-104, 108-112, 116-120, 124-128, and 132-136
UNII – 3 (upper)	149-153 or 157-161

While eleven non-overlapping channels might seem sufficient, regulatory mandates require the Dynamic Frequency Selection (DFS) mechanism to be applied on specific frequency bands. In order to avoid RF collisions with radars, a radar pulse detection function is used. Once there's a case of dual channel occupancy the Access Point has to stop transmission and check for alternative channel availability. This exercise may take minutes and affect adjacent Access Point channel configuration, resulting in lack of service for clients. The following table describes the channels requiring DFS per regulatory domain:

Frequency Band	Channel ID	Frequency (GHz)	FCC/US & Canada	ETSI/Europe	MIC/Japan
UNII – 1 (lower)	34	5.170			√
	36	5.180	√	√	
	38	5.190			√
	40	5.200	√	√	
	42	5.210			√
	44	5.220	√	√	
	46	5.230			√
	48	5.240	√	√	√
UNII – 2 (middle)	52	5.260	√	√	√
	56	5.280	√	√	√
	60	5.300	√	√	√
	64	5.320	√	√	√
UNII – 2 extended (high)	100	5.500	√	√	√
	104	5.520	√	√	√
	108	5.540	√	√	√
	112	5.560	√	√	√
	116	5.580	√	√	√
	120	5.600	√	√	√
	124	5.620	√	√	√
	128	5.640	√	√	√
	132	5.660	√	√	√
	136	5.680	√	√	√


Frequency Band	Channel ID	Frequency (GHz)	FCC/US & Canada	ETSI/Europe	MIC/Japan
	140	5.700	√	√	√
UNII – 3 (upper)	149	5.745	√		
	153	5.765	√		
	157	5.785	√		
	161	5.805	√		
ISM	165	5.825	√		

Requires DFS

Other issues that directly impact 802.11n implementations include the fact that channel bonding implies a performance tradeoff, providing less coverage in return for greater transmission strength. In addition, the receiver is more susceptible to interference because the input filter is wider and less selective, the “edge user” effect can be more pronounced with 802.11n than with 802.11b/g since the rate span is twice as wide.

## Unpredictable Coverage

The coverage of an 802.11n AP is more irregular than the coverage of legacy APs based on 802.11 a/b/g. This can lead to more coverage holes and possibly higher co-channel interference. Multipath is very sensitive to changes in the environment changes (for example, large groups of people in one area, closed doors, etc.) since it leverages RF reflections. An additional element of complexity is introduced by client antenna capabilities (for example, 1X1), resulting in a different coverage map for each client type.




**Figure 8: 802.11a/b/g vs. 802.11n Coverage Pattern**



**Figure 9: Cell Planning with 3 Non-Overlapping Channels—802.11b/g vs. 802.11n**

## Co-Channel Interference

Large open spaces and dynamic environments create unique challenges in terms of radio coverage, particularly for high speed access. Radio coverage when an area is empty may be completely different to when it is full of people. The multiple AP cell based wireless solution shown below provides a graphic illustration of these issues




**Figure 10: 300 Mbps Coverage Area**

The above diagram shows the 300 Mbps coverage area of nine APs. The same nine APs, transmitting at the same power, provide a very different coverage pattern (as shown below) if a speed of 54Mbps or above is considered. The actual coverage of the APs is the same, but as clients can utilize a weaker signal to run at 54Mbps they can be further from the AP and run at this speed.

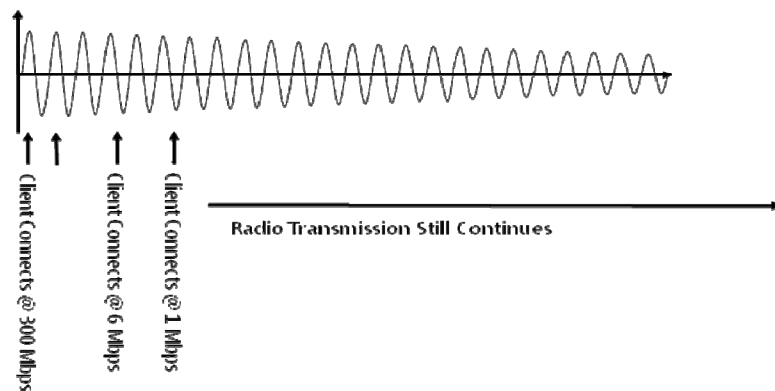
The areas of overlap seen below would be further exaggerated if speeds of 11Mbps or 1Mbps were considered and it should be remembered that a signal will exist even beyond coverage at these speeds, though the signal will become so weak as to be unusable yet still sufficient to cause interference.





**Figure 11: 54 Mbps Coverage Area**

As can be seen, the larger range at slower speeds can cause significant interference between APs on the same channel. This can drastically reduce the performance of cell-based systems, since each AP has to compete with its neighbors and clients transmitting on the same channel for air time.



Client devices in this environment would also have to decide on which individual AP to talk to and would need to roam between APs as they moved around the environment. These roaming events tend to take place at very low data rates meaning that potentially 25% of the APs available would be visible to the client and thus make the client roams a more complex process, with potential for significant impact on the quality of service.

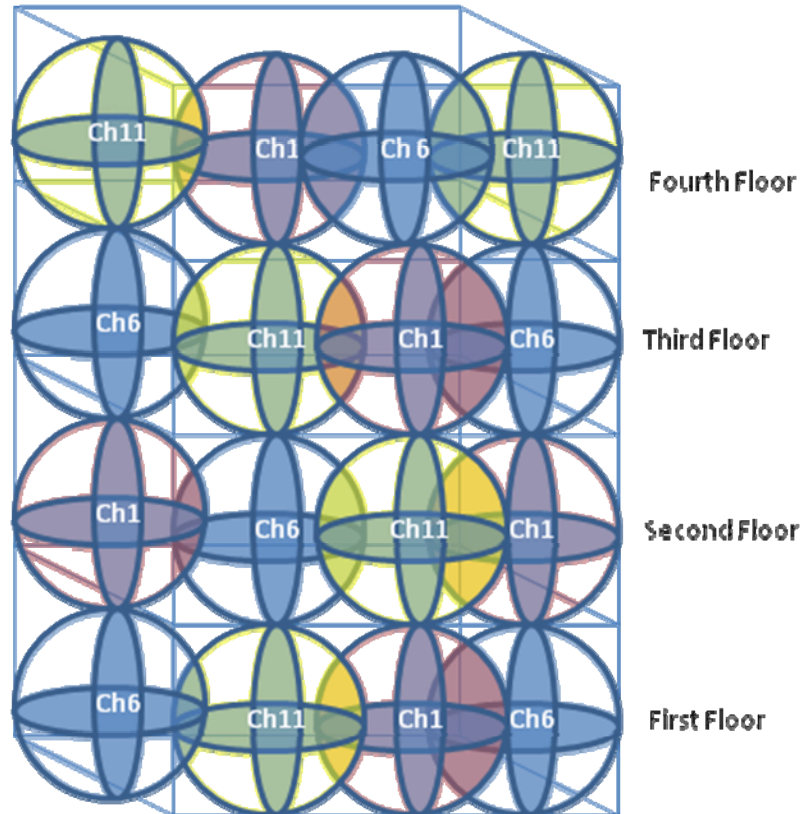
Tests show that the carrier-to-interference ratio (CIR) for 802.11b CCK modulation must be at least 8dB, which means that the carrier must be 2.5 times closer than any co-channel interference to be able to read the data packet; that is, the client must be 2.5 times closer to one AP than to the next one on the same channel to be able to reuse the frequency channel without interference.

For 802.11g or 802.11a modulation, the CIR has to be 12dB-26dB depending on the rate (12Mbps/54Mbps respectively). 802.11n which uses the same OFDM/QAM modulation will require at least the same amount of energy in order to transmit at a higher rate.

Practically speaking 802.11n has great range, but its interference range is greater too, so 11n will generate more co-channel interference than 802.11a/b/g. Interference from neighboring WLAN equipment will result in low packet rates and packet drops, negating the frame aggregation benefits. The higher number of sub-carriers increases the possibility that there will

be coverage holes. Another important factor to take into account is that wireless clients will also produce a great deal of co-channel interference.

The above examples illustrate the coverage area in two-dimensional terms, but real-world deployments are 3-dimensional. An example of vertical inter-floor co-channel interference in a deployment with four APs per floor is shown below.



**Figure 12: 3D Cell Plan**

Transmissions from nearby cells on the same RF channel can prevent local devices from transmitting, even when the intended recipient of the transmission might not be in an area of interference. If the unwanted signal has comparable signal strength relative to the wanted frame, errors will cause a discard, and the lack of an acknowledgement will cause the sender to re-transmit the frame, impacting throughput.

## Planning and Design

Currently, planning and design constitutes a significant phase of the deployment for conventional cell-based WLANs. But are the baseline assumptions and insights derived from 802.11a/b/g deployments still applicable to 802.11n infrastructure and clients?

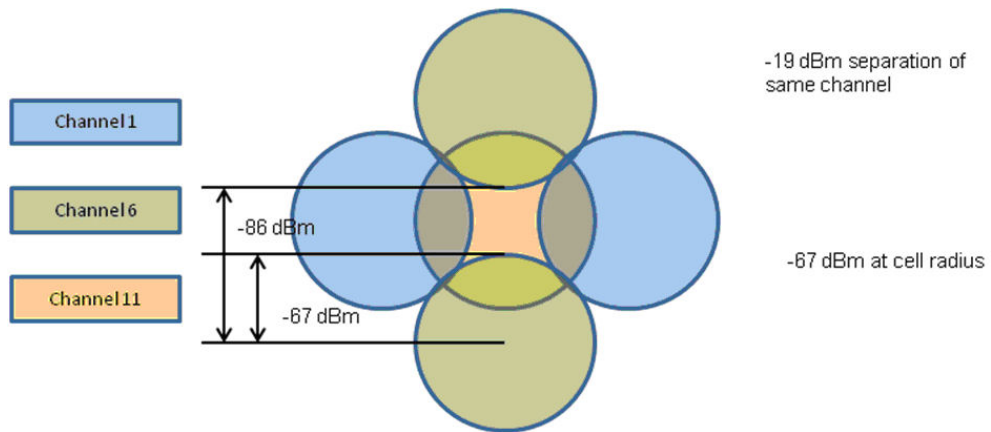
When planning any WLAN deployment, some key factors need to be addressed. Some factors will be based on assumptions, such as the percentage of the users that will use legacy devices and whether the wired network will have enough bandwidth capacity. Other issues are more definitive, such as which bands will be used or what the required power scheme is. The

following sections review some of the critical questions involved in planning an 802.11n deployment.

### *What Applications Will Be Supported?*


For WLANs that will be supporting VoIP, it is recommended as part of the design that minimum power of -67 dBm be present at the cell boundary. This results in cell sizes that are smaller than those used in data WLAN designs.

The -67 dBm threshold is a general recommendation for achieving a packet error of one percent at the highest data rate desired for a given channel, which requires an SNR value of 25 dB or greater (local noise conditions impact this requirement). The -67 dBm signal strength measurement has been used by 802.11b phone vendors for a number of years, and tests indicate that this same rule of thumb measurement also works well for 802.11g and 802.11a clients. It leads to about 15% cell overlap.



**Figure 13: Cell Overlap**

While this design criteria was ideal for 802.11b, new handset devices with 802.11n chipsets are coming to market, such as the iPhone 3GS and Nexus One. In an office or education environment, these 802.11n clients will be used to access email and VoIP applications, so the challenge becomes how to provide ubiquitous WLAN coverage while maintaining the required channel separation. In the dynamically changing MIMO cell environment, this will be difficult as the graphic below indicates.



**Figure 14: MIMO Cell Overlap**

Another emerging application is voice and data convergence from a single client, such as a laptop running Microsoft Office Communications Server (OCS) for VoIP and downloading data files in parallel, or a laptop providing Skype for VoIP calls and sending emails at the same time. While WMM is required to prioritize voice over data, the same WLAN link is used for both applications.

### ***What Is the Deployment Model?***

The coverage-based model is for low-bandwidth deployments, where low roaming wireless signal coverage is required for applications such as a scanner solution or guest access. A coverage-based deployment might consist of APs placed roughly 100 to 200 feet apart running at 50-75% of power. The deployment would consist of an AP installation base with clients associating at greater distances and at lower traffic rates. This deployment will probably use only legacy client types.

The capacity-based model, however, is for dense deployments with high traffic rates due to the number of users and the types of applications they'll be accessing. For example if the requirement is a "desk-top like" experience for employee laptops, where the employee may be running multiple applications simultaneously, the site would require a capacity-based deployment. The deployment would include a dense AP installation base with clients associating at closer distances and at higher traffic rates.

A capacity-based deployment might consist of APs placed roughly 45 to 60 feet apart running at 25-50% or 50-75% of power when using 802.11g or 802.11a. As 802.11n provides higher throughput, it is fair to assume it will be used to support the capacity-based deployment model. At present, green-field deployments consisting of 802.11n clients only are very rare. Assumptions on AP placement and transmission power will have to be field-proven, but the

amount of variables may make it difficult to formulate general deployment rules that will cover a range of scenarios.

### ***What Are the Site Characteristics?***


For the WLAN deployment to be effective, wireless signal has to penetrate a range of obstacles. So the physical environment of the deployment becomes an important success factor: is the site an open space, or are there large amounts of wood or concrete present?

While there is a significant and well-tested body of knowledge on what the effective range of 802.11a/b/g clients will be in such environments, how will multipath affect these characteristics? Also, metal and glass elements such as cubicles, steel doors, or glass window, will have different RF attenuation characteristics, and this will affect Wi-Fi signal penetration.

### ***Transmission Power Control***

The 802.11 standard stipulates use of Clear Channel Assessment (CCA) before a node transmits to ensure that the frequency is clear for transmission. CCA measures the amount of energy in the channel without regard to packets, transmission speed, or even source of energy. It is actually CCA reach that will determine the collision domain size, and not the cell size at the highest data rate. The CCA mechanism is notoriously unreliable, however, and may provide both false positives and false negatives. The CCA may indicate that the channel is clear when, in fact, it is not (false negative). This will result in co-channel interference. On the other hand, the CCA may indicate that the channel is not clear when, in fact, it could be reused.

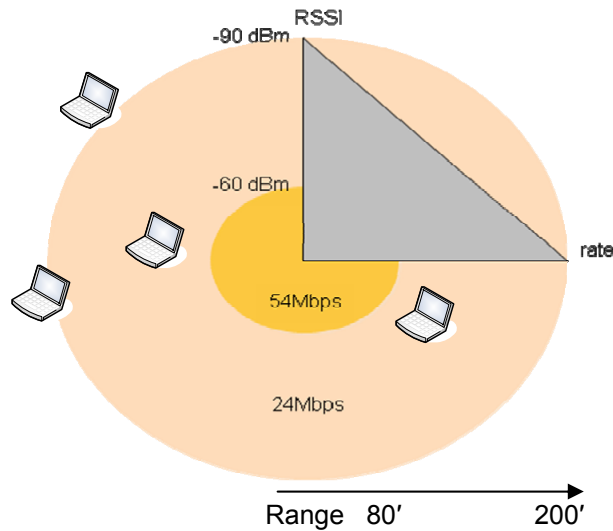
In order to reduce energy expenditure and compensate for this phenomenon, Transmission Power Control is provided to reduce a radio's "pollution area". So for example in the 5GHz band, the 802.11a maximum transmission power levels vary on different channels by as much as 6 dB. This means that when using the maximum allowed transmission power throughout a site that allows all channels, there will not be equal cell coverage on all channels. It also means that if dynamic channel selection is used, the cell coverage edge may change based on the channel number. As a general guideline, cell transmit power on all APs should not exceed the maximum or desired transmit power of the client device. If the client maximum or set transmit power is 10 dBm, WLAN infrastructure vendors usually recommend that the maximum transmit power on the AP be set to an equal level or, if that's not possible, the next higher transmit power level. Equal transmit power is recommended to avoid a one-way link. The AP generally has better receiver sensitivity and diversity support than the client, so it should be able to receive the slightly lower strength uplink signal.



**Figure 15: TPC – 30mW vs. 5mW**

### *Rate versus Range*

In real-world deployment conditions, clients are dispersed in varying ranges from the center of cell coverage.



**Figure 16: Rate vs. Range Considerations**

Cell planning does not increase capacity as actually the slowest user link determines the aggregated data rate because of distance or poor signal reception. It may lead to the “Edge User” effect, in which users who are the farthest from the AP consume a disproportionate amount of the air time, slowing down the entire network. In addition, as the number of users increases, the chances of collisions in the air increase rapidly, thus reducing aggregate throughput.

802.11n high throughput rates (Modulation Coding Schemes) can use OFDM for PHY with 64QAM. Such settings are extremely sensitive to the range between the client and AP, so in practice these rates can be achieved only at a very short range, typically several feet. Since

Equivalent isotropically radiated power (EIRP) is capped by regulation, range reduces speed making mobility even a bigger issue.


### *Client Mixed Mode and Backward Compatibility*

For the foreseeable future, 802.11n will need to operate in the presence of legacy 802.11a/b/g devices. This mixed-mode operation will continue until all the devices in an area have been upgraded or replaced with 802.11n devices. The mixed-mode protection mechanism for 802.11n is similar to the protection mechanism of 802.11g.

This additional legacy preamble and signal field, as well as the CTS-to-self, means more overhead on every 802.11n transmission. This reduces the benefits of all the 802.11n improvements, resulting in significantly lower effective throughput for 802.11n devices in mixed environments.

If there are legacy devices on the same channel on any nearby APs, that will cause protection mechanisms to be invoked as well. Protection mechanisms will in all likelihood be in use in the 2.4GHz band (802.11b and 802.11g) until nearly every legacy device has disappeared. This is because there are too few channels available in that band to effectively overlay pure 802.11n WLANs in the same areas as legacy 2.4GHz WLANs.

Co-existence of 802.11a/b/g with 802.11n on a channel will slow down the throughput dramatically. In the case where all clients need to transfer the same data, air time will be utilized according to the following approximate following dispersal pattern.



**Figure 17: Mixed Client Medium Share**

Since the 802.11 standard allows equal access to all devices, uplink transmission on the channel (i.e., from the clients to the AP) may impact downlink transmission on the channel, thus severely reducing downlink performance as well.

### *Client Variations*

Wi-Fi client devices are very diverse. Each has a Wi-Fi driver that was created with a different goal in mind. Each client device has one or more antennas; some are equipped with Omni antennas and others with directional antennas.


Some legacy clients support only one band, while new 802.11n-compliant clients support two band. Some clients support all MAC layer features, while others only support a subset of these features. All of these factors make it hard to predict how a specific client or a group of clients will behave, especially when mixed with differing clients in the same location.

In addition, clients consume more power when working with 802.11n MIMO antennas. Some clients, such as barcode readers, location tags, etc. are designed to operate for long periods of time between recharging, so companies may not want to move to MIMO to support these types of clients.

In particular, VoIP handsets will probably use a 1X1 antenna configuration with a single stream., whereas laptops will probably use 3 data streams and a 3X3 antenna configuration. This will result in a different AP coverage pattern for each client type.

### *Unbalanced Link*

While 802.11g APs have two antennas, they only have one radio receiver. A diversity switch allows the single receiver to briefly sample the signal quality on the two antennas, after which it picks one to receive each packet. The signal from one antenna is used while the other antenna's signal is simply discarded. In contrast, MRC systems have a full, dedicated receiver for each antenna. They operate simultaneously to capture three signals and add them together. By adjusting the weight and phase of each signal as they are added, noise received on multiple antennas can actually be cancelled out, resulting in the best possible signal-to-noise ratio for the combined signal; hence the term "maximal-ratio combining". The enhanced throughput and link reliability derived from MRC can be leveraged by any client device - it is fully backward compatible with 802.11g clients since no changes are required on the client transmitters.



**Figure 18: MRC**

MRC helps address multiple challenges, but coverage is often limited by the uplink - the link from the client device to the AP. While the downlink signal can be improved by increasing the transmit power of the access points, this is rarely practical for client devices, where battery life, size, and cost often limit transmitter power.



In addition, APs often have higher gain antennas and are mounted in a physically elevated position. This can result in a phenomenon in which the AP is able to receive the client transmissions at a high RSSI while the downlink transmission is not received by the client, making the link unbalanced. In this case, incorrect link advertisement (e.g. data rate) may be selected or there is no transmission acknowledgement (ACK) from the client, triggering retransmission.

### *Frame Aggregation*

The MAC layer enhancement of Frame Aggregation introduced with 802.11n has some specific limitations that one should be aware of, as follows:

- All aggregated frames have to be addressed to the same mobile client or AP, a possible disadvantage in an environment where clients de-associate and re-associate frequently.
- Frame Aggregation requires that all aggregated frames be ready to transmit at the same time, potentially creating delays in transmission.
- Frame size can be truncated in a highly mobile environment due to channel coherence time issues.
- Reduced inter frame space (RIFS) is restricted to green field deployments, i.e. deployments where there are no legacy 802.11a/b/g devices in the area.
- The form of Selective ACK is not supported by all clients. Most of the retry failures are in the form of the entire aggregation failing, not just part of it.
- A possible Denial of Service vulnerability is an attacker can manipulate the block acknowledgment process by transmitting a surreptitious Add Block Acknowledgment (ADDDBA) frame to the recipient, spoofing the source of the victim. Advertising a window of sequence numbers that is not currently in use by the victim, the attacker can cause the receiver to drop all frames from the victim. While the ADDDBA frame is a type of management action frame, this frame is not protected with management frame protection introduced in IEEE 802.11w.
- MSDU aggregation is limited in its frame size to 8kb making it less efficient than MPDU.

### *MIMO Power Consumption*

Multi-radio APs with MIMO draw more power than the 802.3af PoE of 15.4W standard can provide. In addition, the more powerful transmissions from MIMO devices have the potential to be a lot less “neighbor friendly” to legacy 802.11a/b/g systems.

The alternative is to reduce transmission rates, and/or number of antennas, and/or transmission power. This limits APs to what’s called a 2x2 antenna configuration, which breaks a data stream into two sub-streams for transmitting and receiving. To run both radios at the same time in a 3X3 antenna configuration, however, requires higher wattage.

One might mistakenly assume there is no difference in throughput, just in the reliability of the signal. However testing of a 2x3, 2 stream system, against a 2x2, 2 stream system with 20MHz channel width shows a 20% improvement in the uplink throughput from 20-200 foot range. A real-world test with the same configuration using distances of 12-42 feet shows similar results.

The 802.3at standard provides double the power of the original 802.3af standard. At 30 watts, 802.3at provides sufficient power to operate an 802.11n access point. When migrating to 802.11n, therefore, must involve some consideration of how to provide the power to support this new PoE standard.

There are at least two alternatives. The first is to provide power directly from the Ethernet switch to which the AP is connected. This may involve upgrading the Ethernet switch to one that supplies the new PoE standard. The second alternative is to use an inline power injector, which is installed in the wiring closet, along with the Ethernet switches, and inserted into the Ethernet cable between the switch and the AP.

When deployed using PoE, the power drawn from the power sourcing equipment will be higher by some amount depending on the length of the interconnecting cable. This additional power may be 1.5W to 2.45W.

### *Wired Network Infrastructure*

The increased bandwidth of 802.11n APs may move the network bottleneck. In legacy WLAN systems, the bottleneck was always in the wireless access edge. 802.11n may create enough bandwidth to move the bottleneck back to the wired network. Gigabit Ethernet backbones may no longer be sufficient, as potentially each radio can generate about 150Mbps of traffic load for the wired Ethernet.

New wired and wireless architectures may be required to get the most out of 802.11n in the enterprise. A large, centralized WLAN controller at the core may not be the best architecture as routing at network edge is more efficient.

### *Power Save (PSMP)*

PSMP trades the overhead of listening to the PSMP beacon (and potentially listening through some voice packets destined for other nodes) for less contention on the medium when a large number of VoIP devices are trying to operate at the same time. Several factors cause PSMP to have a significant amount of overhead, so this mode only pays off if the number of clients is high:

- The PSMP beacon has all other devices to go silent after they hear it, and it carries a schedule for when the AP is going to send downlink packets to the PSMP-capable devices—different VoIP phones, for example. The AP also uses the beacon to schedule these devices' uplinks. A device that has to wake up and receive the PSMP beacon is pure overhead.
- Depending on how many other wireless devices are in the network and the timing of the activities, the device may not have time to transition into the sleep mode, wake back up, and receive a packet. It may be that the gap in time is short enough that the device must stay awake while other devices receive their packets. If so, that time is also wasted overhead.
- As devices move around, the supportable data rate increases and decreases. As a result, PSMP devices need to be conservative and use a lower data rate because failed packets are a burden on the system. If a packet fails, on the other hand, the system has to make up for that failure at a much later time. The system has to allocate a new time slot to retransmit the packet, because the whole idea is to pack transmissions right next to each other. There

is no time allowed for retransmitting failed packets. In contrast, with APSD, a packet might fail, but it can be retransmitted immediately, without centralized coordination. This makes the penalty for a failed packet much lower, and allows APSD devices to use higher data rates more aggressively, reducing power consumption.

## *Roaming*


Roaming or handoff is a key requirement and has a large influence on WLAN design. A mobile client will roam from one AP to another when it moves out of the AP's BSS range, or when the current operating channel conditions deteriorate.

This process of switching from one AP to the other can impact applications running on the mobile client, including, in the worst case, terminating the application session and requiring it to be restarted after the handover is completed. In other instances, the application may either experience a temporary outage or, at best, a short delay in transfer of packets.

### Intelligent Roaming Mechanisms

To maintain application continuity without compromising the key aspects of the connection, such as security and power save mode, requires an intelligent roaming mechanism. Clients use proprietary background scan methods to discover APs. During a VoIP call, the client has about 20 ms between packets to carry out this background scan. It may, therefore, have to repeat the scan in case all desired channels are not scanned during the 20 ms window.

Scanning also affects battery consumption on the client. Some clients are designed to “stick” as long as possible to an AP by shifting the rate and increasing transmission power. Depending on client dispersal and AP deployment, this may not always lead to battery consumption; it may, in fact, be better for the client to scan and change channels.



Once it has built a list of APs, the wireless client can roam when conditions dictate. The decision to roam is the responsibility of the client, and factors such as Received Signal Strength Indicator (RSSI), signal-to-noise ratio, and frequency of packet re-tries are taken into account by the client when deciding whether or not to roam.

### Re-Establishing the Connection

The next important task during roaming is establishing a connection with the new AP. Depending on the AP's security settings, this may involve the 802.1x/Extensible Authentication Protocol (EAP) mechanism. 802.1x/EAP can be a time-consuming process,

based on the EAP type in use. The simplest 802.1x authentication can take a minimum of 100 ms depending on network conditions. In WPA/WPA2 Pre-Shared Key (PSK) security mode, the re association process requires an additional 4-way handshake, but not 802.1x/EAP. This handshake takes additional time, but typically is completed within 50ms.

The WPA/WPA2 Enterprise security mode involves the most time-consuming 802.1x/EAP authentication, together with the four-way handshake mechanism, during re-association. In order to meet this requirement, the IEEE standard proposes a pre-authentication technique in which the client completes the 802.1x/EAP process with the new AP through the distribution system before it decides to roam to this AP. During handoff, therefore, the client would only need to use four-way handshake to obtain the PTK from the new AP, and as a result, would only take as much time to roam as in the PSK case.

Pre-authentication is not scalable because it requires every AP to remember the PMSKA (Pair wise Master Key Security Association – the context resulting from a successful 802.1X exchange) for all clients that could possibly roam to it. Each client has to authenticate to each AP. Clients have to guess which AP they may hand off to, and therefore how many APs to pre-authenticate with.

Opportunistic Key Caching (OKC), sometimes called Proactive Key Caching (PKC), is a method whereby the client authenticates to the network for the first time using 802.1X/EAP, the result is calculated per the AP (PMKIDs). When the client wishes to roam from one AP to another, a new key is calculated based on the previous key, BSSID of the new AP and the station (supplicant) MAC address. When the infrastructure uses a controller for MAC functions, the controller has access to all keys and can forward them among nearby APs for future building of PMKIDs.

IEEE	802.11i	802.11d	802.11h	802.11e	802.11k	802.11r
802.11n						
802.11a						
802.11g						
802.11b						

**Table 2: 802.11 IEEE standards**

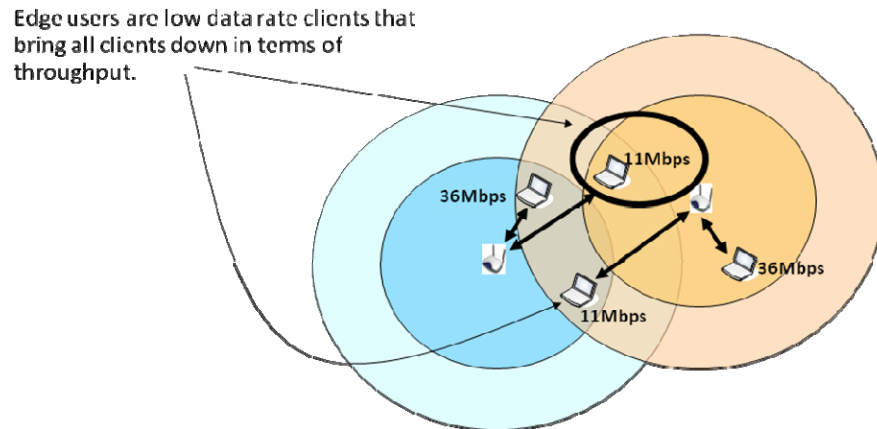
An AP announces its support for fast BSS transition in beacons, probe responses and re-association responses. The Fast Transition (FT) over the air is part of the association request/response frame exchange.

802.11r merely generalizes OKC. However, neither 802.11r nor OKC addresses how the handoff decision is made, only how quickly it can occur. The client will still make the roaming decision, usually based on RSSI. Mobile clients have to scan across more than 30 channels (in the US) in multiple bands just to identify which APs might serve them. If the chosen AP doesn't have the resources to handle the new association, the client is then forced to repeat the scanning, authentication and association processes.

### *Edge Users*

Remote clients in the outskirts of the AP's coverage area will generally use a lower data rate than clients closer to the AP. As the air medium is a shared resource among clients, a low-rate associated client will consume more air time (either uplink or downlink), thus reducing the

overall aggregate throughput of the cell. Fairness algorithms that tend to balance medium sharing among clients do not factor in this. Since roaming decisions are different for each client, it may happen that two adjacent clients residing in overlapped coverage zones between two APs will use channels and different rates. 802.11k tries to remedy this situation, but requires client-side support.




**Figure 19: The Edge User Phenomenon**

## 802.11k

The Neighbor Reports mechanism allows a client to locate its neighboring APs, thereby reducing the client scan effort. The intent is to make the handoff decision faster, but Neighbor Reports don't effectively guide the clients and 802.11k doesn't specify what a client is supposed to do with the information once it receives it.

In addition, it's not clear what a neighbor means – the AP closest to the client or the AP closest to the AP from which the client requests the Neighbor Report? One approach is to have the Neighbor Report replace client scanning but there may be a case where the list of neighboring APs is either empty or confusing; in this instance, the client should ignore it.

802.11k becomes less effective as the number of channels in a deployment increases. As more channels are added, the wireless environment becomes more complicated: in 802.11n if using the 5GHz band with DFS channels, a cell-planned design will attempt to have a very small number of APs (optimally one) that are configured to a specific channel available at any given physical location. The Neighbor Report from an AP can include six neighboring APs. Assuming these surrounding APs are in-band the required DFS, then the client is required to not trust the Neighbor Report's indication that an AP is on the channel. What that means is that the client can tune to the channel, but it is forced to wait for the next beacon before it can start association (which requires transmission). That can lead to a wait of up to 100ms per channel, since active probing isn't allowed.



**Figure 20: Neighbor Report Roaming Issues**

The figure above illustrates this issue. It's likely that only one of the six APs will be in range of the station because it moved away from the center AP towards the green circle, as represented by the arrow. The client can pick only one of the neighboring channels and has a low chance of choosing a channel with an AP in range. In the worst case, it will have to wait up to 100ms on each of the six channels before it finds its neighbor. This adds up to a possible 600ms of passive scan time (the scan order proceeding clockwise from the red AP), and if the neighbor AP is full (without available resources), 802.11k guidance is not useful. There is still no way to make handoff predictable and each station has to figure out when and how to perform it. Both 802.11r and 802.11k perform best in scenarios where they are needed the least.

## The Extricom Advantage

For all of the technical challenges 802.11n presents, it seems inevitable that it will feature prominently in Enterprise wireless strategies in the years to come. All of the major wireless vendors have announced 802.11n products, and are actively touting their features – though not always how they resolve the more thorny issues presented by 802.11n. Clearly 802.11n requires a highly adaptable Wi-Fi architecture, but the cell-planning approach that dominates most network designs is not known for flexibility in the face of technical challenges.

Extricom's uniquely innovative Channel Blanket architecture for WLAN constitutes an alternative approach, one that will provide Enterprises with significant advantages when implementing 802.11n. Virtual solutions may try to mimic blanket advantages through software logic, yet may fail in practice due to high performance and real-time computation required.

## How Extricom Works

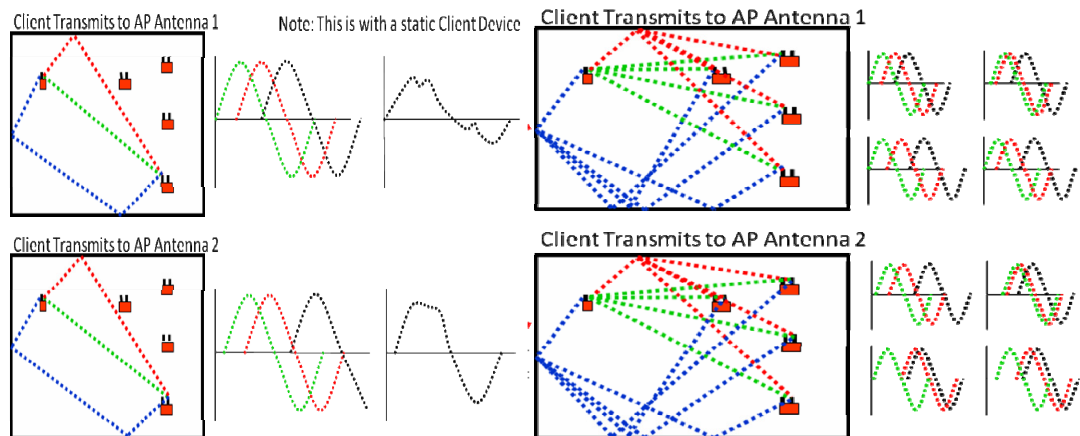
The Extricom solution is based on a centralized WLAN architecture, in which the switch makes all of the decisions for packet delivery on the wireless network. In this configuration, the access points (APs) simply function as radios, with no software, storage capability, MAC or IP address. Even the basics of connecting are different: clients associate directly with the switch, not with the AP. The AP acts as an "RF conduit" to rapidly funnel traffic between the clients and the switch.

Centralization of the Wi-Fi environment enables enterprises to deploy 802.11a/b/g/n channels at *every* AP, creating overlapping "Channel Blankets". The channel's bandwidth is delivered across the blanket's service area (i.e. the combined coverage of all APs connected to the switch), with interference-free operation and consistent capacity throughout.

Within each Channel Blanket, the switch avoids co-channel interference by permitting multiple APs to simultaneously transmit on the same channel, but only if they won't interfere with each other. In summary, the Extricom solution eliminates the traditional performance limitations caused by RF cell planning, co-channel interference, edge users, rate adaptation, mixed b/g/n devices, and frequent AP-to-AP handoffs.

To better understand the mechanics of the Channel Blanket, it is useful to detail how the uplink and downlink functions in the Extricom system. Since the client is associated to the switch rather than any single AP, the client does not sense how many APs exist in the Channel Blanket. The client only knows that the channel it is using is continuously available over the Channel Blanket's range.

Since the APs transmit and receive on the same channels, multiple APs will usually be able to receive a client's uplink transmission (illustrated in Figure 4 on the right hand side as opposed to a single AP on the left). As a result, the switch receives multiple copies of the same client transmission. The switch will then dynamically, and on a packet-by-packet basis, decide which packet copy it prefers, and select it while discarding the others.



**Figure 21: Uplink Diversity**


As the client moves throughout the blanket, different APs will be in the best position to serve the client at different times. The switch always uses the uplink and downlink path that is optimal to serve the client. While this is going on “behind the scenes,” the client never experiences an AP-to-AP handoff (i.e. de-association and re-association), resulting in seamless mobility.

## Why Extricom Works Better for 802.11n

As detailed in “Channel Bonding” beginning on p. 10, the effect of channel bonding in the 2.4GHz spectrum is to reduce the number of available channels to one 40MHz channel and one 20MHz channel, making conventional cell-planning impossible for 802.11n.

Extricom’s Channel Blanket architecture inherently has no requirement for cell-planning because all APs transmit and receive on the same channel, thus it’s capable for example of a 40MHz in 2.4GHz deployment. A single set of APs enables deployment of multiple high-data-rate Channel Blankets with overlapping coverage, resulting in multiplied aggregate capacity. From the perspective of 802.11n, separate Channel Blankets also offer the unique ability to physically separate 802.11n traffic from 802.11a/b/g traffic.

For this reason, Extricom is capable of supporting full-performance 802.11n clients and permanent co-existence with full-performance 802.11b/g clients, all in the same AP deployment.



**Figure 22: 802.11n on a Separate Channel**

An alternative to 2.4GHz that some major wireless industry players have proposed is to operate 802.11n with 20MHz channel width and 802.11n with 40MHz channel width in the 5GHz band, thus leveraging its supposed larger number of channels to make channel bonding work effectively. There are some problems with this approach as well, namely:



- **Reduced channels:** Although there are up to 11 non-overlapping 40MHz channels in the 5GHz band, regulatory factors such as DFS2 can reduce this to as little as three channels. For cell-based WLANs, this will once again create the complexity of conventional channel planning, as historically seen in the 2.4GHz band. This means losing the very flexibility that makes 5GHz an attractive proposition in the first place.
- **MIMO Variability:** The variability inherent to MIMO makes for highly unpredictable or even physically disconnected coverage areas, making 802.11n cell-planning extremely difficult.




Figure 23: MIMO Variability's Effect on Cell-Planning

Extricom is the only technology immune to cell-induced problems and variability of MIMO coverage. In fact, because MIMO improvements are opportunistic, the Extricom architecture actually improves MIMO performance by providing the “best” MIMO path on a packet-by-packet basis to a wireless client, thus adding robustness and stability to a highly variable technology. MIMO doesn't impact the cell plan, because there's *no cell plan* to impact!





Figure 24: Using Channel Blankets to Take Maximum Advantage of MIMO

- **Complete Choice of Clients:** An added drawback to operating 802.11n exclusively in the 5GHz band means that all equipment deployed by the Enterprise must be compatible with 802.11a/n. But the simple fact of the matter is that many legacy wireless clients only operate in 802.11 b/g mode; swapping this equipment out could prove immensely expensive. What's more, handhelds such as PDAs, phones, and tablets may not use 802.11n at all due to the MIMO antenna form factor and battery drain.
- **Client Separation:** Extricom's topology allows enterprises to separate legacy 802.11b/g devices from applications requiring the speeds of 802.11n, ensuring maximum flexibility. There is no need to forklift out legacy wireless equipment base; instead, devices can be gradually phased out to suit the organization's operational needs.


What about the other strategy advocated by those who design cell-based WLANs? The one that says to build a dual-band WLAN using two-radio APs that operate in 2.4GHz and 5GHz at the same time, leaving b/g clients in 2.4GHz and placing 802.11n clients only in 5GHz?

Since the same physical set of APs is to be used to drive two cell plans, one in each frequency band, this strategy assumes that the placement of APs can be set to simultaneously maximize and optimize the performance of a 3-channel 2.4GHz cell plan and a 3-channel 5GHz cell plan (worst case, without DFS2 channels available). The figure below shows a simplified view of this approach. Since the traditional challenges of cell planning are further aggravated by MIMO coverage variability, how difficult and risky will this approach be?



**Figure 25: The Challenge of a Dual-Band Cell Plan from the Same Set of APs**

The alternative is to deploy a multi-channel, multi-band, multi-mode WLAN using the Extricom Channel Blanket topology. The figure below shows what this would look like. The only design task is to decide how many APs to place, wherever needed to reach the desired target bandwidth. Link robustness, mobility, Quality of Service, capacity are simply intrinsic. And there are no cross-constraints created between bands, channels, or modes.



**Figure 26: The Extricom Multi-Layer, Multi-Band and Mode WLAN**

Leveraging the Extricom WLAN approach, the organization retains full choice of bands and devices, today and in the future. Existing a/b/g clients permanently and peacefully co-exist with the new 802.11n devices, and neither undermines the performance of the other. Deployment is radically easier and more straight-forward than any alternative, ensuring that the introduction of 802.11n really delivers its promised benefits, with the least project and technology risk.

The Extricom WLAN elegantly accomplishes both goals of technical capability and real-world implementation effectiveness. This solution delivers the performance promise of the 802.11n standard in a comprehensive communication system.