



802.11n: Look Before You Leap

Key considerations for moving to 802.11n



The WLAN world is changing

With the introduction of 802.11n, the field of wireless local area networks (WLANs) is undergoing a paradigm shift as phenomenal as the beginning of WLANs themselves. The extraordinary data rates that the final 802.11n standard will support solidify the realization of an all-wireless enterprise network. Rich multimedia applications will be seamlessly streamed to every point in a facility at better performance than possible with legacy 802.11a/b/g technologies.

But the question remains, "How do you plan a high performance 802.11n network?" Despite continual attention from WLAN equipment vendors and endless discussion about the benefits of 802.11n, no one has adequately explained how 802.11n network planning works and, perhaps more importantly, how it doesn't work! "How can I reap the maximum benefit from a clean-slate 802.11n implementation?", "What happens if I simply rip-and-replace access points (APs) when migrating my existing network to 802.11n?", "How would I best pursue a phased-migration to 802.11n?" In this paper, we will explain the 802.11n basics necessary for answering these questions, and provide guidance to help you choose the best strategy for your organization.

How 802.11n affects the three C's of network planning

Network planning should always be approached from the perspective of Context, Coverage, and Capacity (the three C's of network planning). The technological advances of the 802.11n standard affect all three of these. With respect to Context, network planners must consider how interference from new 40MHz channels and the complex site-specific dependencies of Multiple Input Multiple Output (MIMO) technology affect channel planning and AP placement. When considering Coverage, designers must understand coverage differences between 802.11n and legacy systems and correctly define coverage requirements based on network demands. Finally, Capacity is improved due to the 802.11n standard's increased transmit data rates and MAC layer efficiencies. However, these capacity improvements can only be fully utilized when the network client distribution is properly planned.

Context

The context of the environment in which a WLAN is deployed is critical. Interference can be caused by neighboring APs or other wireless transmitters broadcasting within the same frequency band.

This form of wireless congestion results in dropped packets, slower networks, and reduced capacity. In addition to traditional co-channel and adjacent-channel interferences, 802.11n 5GHz band deployments must also consider potential interference from radar systems. Context also includes the site-specific structure of the environment containing the WLAN, which dramatically affects the performance of 802.11n's MIMO technology.

Interference and channel planning

Providing a more than two-times improvement in data rate performance over the 20MHz of frequency bandwidth used per channel in 802.11a/b/g, the new 40MHz channels of 802.11n are a must for truly high performance wireless networks. Unfortunately, bigger channels mean greater potential for interference and reduced spectrum for channel planning.

In the United States, if a 40MHz channel is used in the 2.4GHz band, only one other non-overlapping 20MHz channel is available. The result is a greater likelihood for adjacent-channel interference in the 2.4GHz band. Since channel planning in 2.4GHz was already a difficult task with only three non-overlapping channels in 802.11a/b/g, the use of 40MHz channels is not recommended for 2.4GHz deployments utilizing 802.11n.

Fortunately, the 5GHz band frees 802.11n users from the tight spectrum constraints of the 2.4GHz band. In the United States, the 5GHz band allows for 11 non-overlapping 40MHz channels if the AP is fully compliant with the dynamic frequency selection (DFS) restrictions (more is mentioned on DFS in the next section). The abundance of non-overlapping 40MHz channels in the 5GHz band allows an 802.11n deployment to take advantage of this performance gain and is the recommended deployment strategy for high performance WLAN networks. See Table 3 for 802.11n channel overlap.

The effect of radar avoidance on 5GHz band channel planning (DFS)

As stated previously, to achieve the maximum number of non-overlapping 40MHz channels in the 5GHz band an AP must be fully DFS compliant. As defined in section 15 of the FCC rules and regulations (47 CFR§15), this means that if a device detects in-band interference from a nearby radar system it must immediately stop all transmission within that band for 30 minutes and switch to another, non-interfering channel. Clearly, compliance with this federal regulation will have an effect on 5GHz channel planning since it requires the AP channel change dynamically.

Though the inclusion of DFS into a 5GHz 802.11n deployment may cause issues, the best practices for planning in a DFS band (5.25-5.35GHz and 5.47-5.725GHz) do not change much from planning in a non-DFS band. The first step in the deployment process should involve a site-survey to determine whether any radar interferers exist in the environment. Secondly, the network channel plan should be designed to avoid operation on any channels where DFS has been detected. Lastly, since the DFS standard requires the operating channel change dynamically, empty channels should be made available for device utilization if radar interference is detected. A good rule of thumb is to provide at least one unused channel in a non-DFS band.

Site-specific effects of MIMO

MIMO is inherently site-specific

In legacy systems, interference caused by reflections and diffractions of the transmitted signal (called multipath) was viewed as a hindrance to system performance and was compensated for by including large fade margins in the system design to improve signal quality in areas with heavy multipath interference. In contrast, in MIMO systems multipath is the cornerstone of improving system performance! By utilizing complex signal processing, a MIMO system is capable of sending multiple data streams at the same time. Effectively, this means the received signal strength (RSSI) alone is no longer sufficient for predicting system performance. Given the site-specific nature of MIMO, the use of site-specific planning and management tools for 802.11n networks is highly recommended.

Contrasting performance of MIMO in dense office vs. long hallway environments

What is the best environment for optimal MIMO performance? To use an academic term, multipath rich environments are the best scenarios for MIMO performance. A multipath rich environment is one in which the received signal is evenly distributed among multiple different paths from the transmitter to the receiver. The amount of difference between individual paths is a basic metric for the richness of the multipath. To best explain this, consider two frequently occurring deployment scenarios: a complex office building and a long, straight hallway.

In an office building, an AP is usually centralized within the desired coverage area. The room with the AP is usually surrounded by other rooms, which may be connected by short winding hallways. In general, the environment is dense with obstacles to the signal path (typically walls) and there are very few, if any, Line-of-Sight (LOS) reception paths which aren't in the room containing the AP. The complexity of this environment generates many different paths for the transmitted signal and MIMO systems will perform very well. This is the preferred MIMO deployment scenario.

Changes in coverage from Legacy to 802.11n

Transmit Data Rate	Difference in Coverage from 802.11a/b/g to 802.11n
Minimum (1Mbps for 2.4GHz, 6Mbps for 5GHz)	802.11a/b/g clients will see very similar range when compared to existing networks
	802.11n clients will see very similar range when compared to existing networks
802.11a/g Maximum (54 Mbps)	802.11a/g clients will see some coverage improvement
	Largest increase in coverage area when using 802.11n clients

Table 1: Coverage improvements of 802.11n networks compared to 802.11a/b/g networks.

In a long, straight hallway scenario, the dominant path of the received signal is strongly LOS, and the main multipath contributions come from reflections of the signal along the walls of the hallway. In this environment, a network designer can expect their MIMO performance to drop substantially as the distance between the AP and the receiver increases down the hallway. Multipath components in this scenario are fairly similar, and therefore the environment is not multipath rich and the MIMO performance gain (while still available) is not as large as the complex office scenario. With legacy hardware, placing an AP to maximize LOS coverage down a hallway was an accepted best-practice; however, it is a hindrance to 802.11n performance and is not an optimal deployment scenario.

Coverage

Coverage differences between legacy 802.11a/b/g systems and new 802.11n deployments are an area plagued by much confusion. To truly assess the difference in coverage between 802.11a/b/g and 802.11n hardware, the term “coverage” must be clearly defined. Throughout this paper, coverage will be defined as “communication at a specified minimum transmit data rate at a given location.” “Range”, on the other hand, is defined

as “communication at the minimum supported transmit data rate for an AP at a given location.” The following discussion outlines the truth regarding 802.11n coverage differences and is summarized in Table 1.

Fundamentally, 802.11n radios are still bound by the same governmental power output regulations (the Effective, Isotropic Radiated Power, or EIRP) as those in the 802.11a/b/g standards. This means in an “apples to apples” comparison, a signal transmitted by an 802.11n access point will go no farther than a signal transmitted by legacy hardware. Despite a lack of transmit range improvement, the best 802.11n implementations will leverage the increased number of antennas on an 802.11n AP for increased receiver diversity gains. This allows the AP to hear fainter signals and effectively increases the “visible” coverage area, and thereby reduce hidden node problems¹.

Since a transmitted signal from an 802.11n AP travels no further than a legacy AP, the transmit data rate performance at a given RSSI becomes a vital metric for indicating the differences between 802.11n and legacy coverage. The transmit data rate indicates the speed at which an individual data packet is wirelessly transmitted over the air. With respect to coverage differences between 802.11n

¹ Hidden node problems cause interference issues in WLAN networks when devices with overlapping coverage “talk over each other” when communicating to other devices. The problem is minimized when all devices on a network can “hear” all other devices on the network and therefore know when it is their turn to communicate.

and legacy networks, it is noteworthy that the coverage area of 802.11n at a transmit data rate of 54Mbps is greater than the 54Mbps coverage area of 802.11a/g. This observation, combined with the lack of range improvement for 802.11n, outlines the main truth of 802.11n versus 802.11a/b/g coverage differences. Specifically, coverage improvement is greater at higher transmit data rates and the improvement is less at lower data rates. Further, it is important to note that the use of 40MHz channels will additionally improve the coverage area of an 802.11n AP at higher data rates, but it also will not increase the transmit range of an 802.11n device.

Capacity

Due to the increased transmit data rates of the 802.11n standard the capacity of 802.11n APs may be greater than that of legacy APs. However, this improvement is only available when 11n clients are associated to the 11n AP and within the coverage area of the higher transmit data rates of the 802.11n standard. Since legacy clients on the network now have the potential to actually decrease overall system performance, client distribution planning is a major factor for high-performance 802.11n deployments.

Mixed networks

A tremendously important feature of 802.11n is that it is fully backward compatible to 802.11a/b/g networks and devices. While this provides a smooth path for migrating existing wireless networks to 802.11n, it also means that the 802.11n network

must sacrifice some of its performance when legacy devices are transmitting over the network. The three main cases of client distribution and their effects on network performance are indicated in Table 2.

To help mitigate the effect of legacy clients on a high-performance 802.11n network, it is important to note that 802.11n can operate in both the 2.4GHz and 5GHz bands. The 5GHz band has long been left relatively empty by the mediocre adoption of 802.11a networks, but this is an ideal scenario for the new 802.11n standard. Since so few 802.11a clients exist in the 5GHz band, “n-only” deployment scenarios can be carried out with relative ease in this space without having to worry about the network being bogged down by legacy clients. Deploying 802.11n in the 5GHz band is the recommended deployment scenario for n-only, high performance WLANs.

Targeted Upgrades to the Wired Network

Since the transmit data rates of the 802.11n standard have increased significantly, for the first time it’s possible that a wireless network could routinely out-perform a 100-BaseT network. What results is a need to intelligently upgrade wired network infrastructure to support gigabit Ethernet on backhaul connections for 802.11n APs only as necessary. When a well-planned wireless network can deliver more than 100 Mbps and advanced meshing technology can decrease the demands on the wireless network in general, upgrading the entire wired network to support gigabit Ethernet can be an unnecessary, and costly luxury.

802.11n Network performance with various client distributions

Client Distribution	Performance in terms of throughput
802.11a/b/g clients only	802.11a/b/g clients will perform as well or slightly better than if the network was 802.11a/b/g
802.11a/b/g and n clients	802.11a/b/g clients will perform as well or slightly better than if the network was 802.11a/b/g
	802.11n clients will perform better than 802.11a/b/g clients, but not as well as if it was only 802.11n clients
802.11n clients only	Best performance for 802.11n clients

Table 2: Relative throughput performance of an 802.11n network under various client distributions

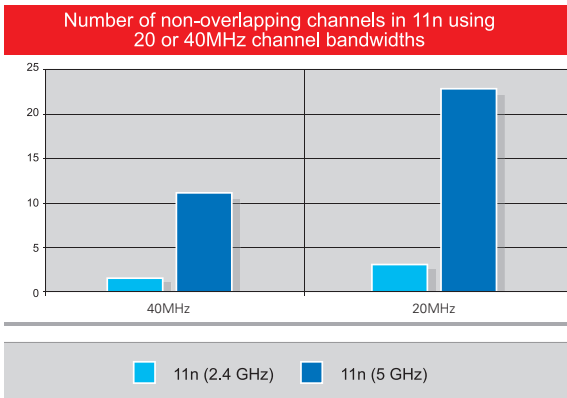


Table 3: 802.11n Channel Overlap

Recommendations for 11n network migration

Until this point, this paper has primarily focused on how the basics of 802.11n technology relate to the three C's of network planning, Coverage, Context, and Capacity. Now, we will consider three strategies to migrate a legacy network to 802.11n in a way that considers the three C's, namely Clean-Slate, Rip-and-Replace, and Phased Migration.

Clean-slate design

Clean-slate network migration is based on the removal of existing network infrastructure and the subsequent deployment of new network infrastructure into locations that are chosen based on maximization of network performance. Despite the increased costs of performing a clean-slate design for an 802.11n network, the clean-slate design migration strategy provides the advantage of a wireless network that can be specifically designed from the ground up to take full advantage of 802.11n's unique qualities. The clean-slate design is discussed first, since it is the model for an optimally performing 802.11n deployment.

To minimize the adverse effect of legacy clients on the network, use the 5GHz band for 802.11n clients and high bandwidth applications. The 5GHz n-only

client segregation strategy has been highly-touted by many industry professionals for good reason. This methodology adapts into any migration strategy, since a dual band 802.11n AP can support both the 2.4GHz and 5GHz bands simultaneously, while maintaining n-only segregation. Assuming the vast majority of legacy clients are operating in the 2.4GHz band, 40MHz channels can be used in the 5GHz band without much trouble. Further, as 802.11n clients are introduced into the 2.4GHz band it should not cause too much trouble for the 11n network, since there will likely be a higher capacity for 11n clients than for legacy clients. However, due to the coexistence of legacy clients, 11n clients in the 2.4GHz band cannot operate at the same performance as in the 5GHz band.

Additionally, AP placement for 802.11n networks should consider the effects of multipath on MIMO system performance. APs should be placed in locations which both maximize NLOS paths to the receiver and minimize the signal path loss. This results in a design decision depending on the environment the network is being deployed. For example, in a building with thick, highly attenuating walls, it may still be more beneficial to place APs in hallways despite its detrimental effect on MIMO system performance gain since the signal will not propagate far enough via the NLOS paths for a MIMO system to effectively utilize the rich multipath present in the NLOS link. Due to MIMO link design complications it is recommended a network planning and simulation tool be used during AP placement since it drastically reduces the required calculation time.

Ultimately, a clean-slate 802.11n design should produce a highly optimized network maximizing performance and minimizing hardware costs. In addition to client distribution planning and AP placement issues, upgrades to the wired backhaul for the wireless network still need to be included, but this is the case for all 802.11n deployments and migrations. While a clean-slate migration may be the most expensive option from a planning perspective, it is also the most likely to achieve a robust, reliable network prepared to meet the current network requirements and scale as those requirements increase in the future.

Rip-and-replace migration

Rip-and-replace migration involves the one-to-one replacement of existing APs for new 802.11n APs. The rip-and-replace migration strategy is often discussed within 802.11n circles. The attractiveness of this strategy is that it does not require any additional cabling or installation costs beyond simply exchanging existing APs for new 802.11n units. Unfortunately, the simplicity of the rip-and-replace migration strategy also faces drawbacks, since the segregation of 802.11n clients into the 5GHz band may not provide sufficient coverage for some legacy systems. Additionally, legacy network design practices can conflict with best practices for maximizing MIMO system performance.

The problem with 5 GHz segregation within a 802.11n network is existing networks designed for maximum coverage at 2.4 GHz may have nodes too far apart to satisfy the coverage requirements of a 5GHz deployment. Since networks designed on the transmit range to maximize an AP's coverage area are common among 802.11b/g deployments, this can be a persistent problem for rip-and-replace migrations wishing to provide 5GHz segregation. This is because, for the same transmit power, signals do not travel as far in the 5 GHz band as they do in the 2.4 GHz band. Since first-round 802.11n hardware does not have an increased range over existing 5 GHz APs, a network which was designed based on the maximum range at 2.4 GHz will likely have numerous coverage holes in its 5GHz 802.11n network.

As discussed previously, MIMO system performance is maximized in multipath rich environments. In rip-and-replace migrations (from an existing 802.11a/b/g network), it is likely there are several APs positioned to maximize coverage in hallways, as this was a widely accepted best-practice for legacy network deployments. A rip-and-replace migration of these "long hallway" APs reduces performance gains of the 802.11n migration.

Further, the issues presented above are for a network where it is assumed performance requirements have not increased from its initial deployment. If, as most likely the case, the performance requirements have increased, a rip-and-replace strategy may be impossible since increases in AP density may be necessary even for an 802.11n network. In contrast, an existing high-performance 802.11a/b/g network may be able to migrate to 802.11n with fewer APs, since there are known coverage improvements for 802.11n at higher data rates.

In the end, the problems associated with a rip-and-replace migration strategy simply highlight that planning will always be an integral step for any change to a wireless network. As such, even when the change is as simple as a one-for-one device replacement the best practice for network planning is still to assess your network performance requirements, assess the capability of your hardware and plan carefully.

Phased migration


The main goal of a phased migration is to supplement an existing 802.11a/b/g network with 802.11n APs in order to satisfy the demands of the existing network without incurring the cost associated with a full network migration. Phased migrations are useful if a portion of your facility is already adopting 11n clients or if high bandwidth applications like large file transfers are needed in specific locations.

The first step of any network migration should be to identify how the existing network is failing to satisfy current demand. This is of particular importance for a phased migration, since the stated goal is to meet the demands of specific portions of the existing network. A comprehensive network management system is very beneficial in this initial, requirements gathering, stage of the network migration.

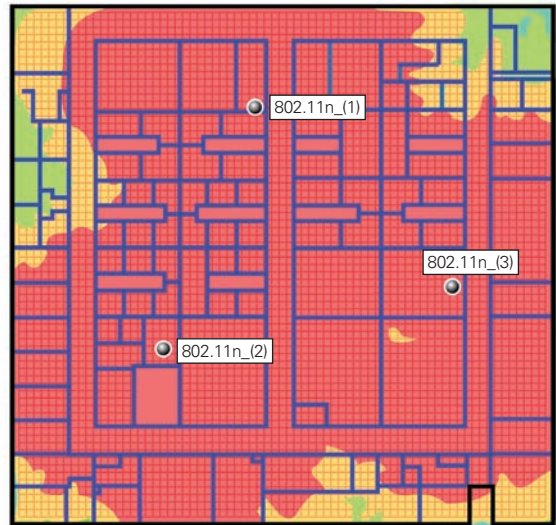
Once performance requirements have been determined, if an existing legacy client network is already operating at maximum capacity, simply replacing existing APs with 802.11n APs is not likely to satisfy capacity requirements since legacy clients on the migrated 802.11n network still only operate at legacy speeds. As mentioned previously, if a portion of the network is already adopting 802.11n clients and capacity requirements are being met, then a one-to-one exchange of a legacy AP for 802.11n hardware would be a prudent, short-term option. Otherwise, Motorola recommends AP density be increased by providing additional 802.11n APs into the capacity-strapped region of the network. These new APs should be given priority in terms of coverage area per AP via judicious power-planning since the end goal of a phased migration is to have the 802.11n network as the only network.

Unless network requirements can be met through a rip-and-replace of a network's subset, it is recommended that new 802.11n APs be placed independently with respect to existing legacy AP locations. This may also necessitate small adjustments to the existing AP placements, but it will allow the phased migration to maximize the


Using LANPlanner to Define Network Requirements




Original Network



Clean-Slate




Rip-and-Replace



Phased Migration

Transmit Data Rates



MIMO system performance of 802.11n hardware and makes channel planning for the final hybrid network much simpler.

How can Motorola LANPlanner® assist in the 802.11n planning process?

As various impediments to 802.11n planning have been discussed throughout this paper, it has become increasingly apparent that properly planning an 802.11n network is a daunting task and is quite difficult without the use of site-specific network modeling techniques. Fortunately, Motorola has tools to assist network planners in this process. Whether the project requirements involve a clean-slate design, a rip-and-replace network migration, or simply a pre-deployment survey to assist in requirements gathering for a phased migration, LANPlanner 11.0 has the capability to meet the demands of 802.11n.

Site-specific performance predictions

802.11a/b/g planning has always been a computationally intense challenge. Now that 802.11n includes additional non-intuitive multipath factors, site-specific modeling is even more important. As previously discussed, 802.11n performance will change depending on the environment surrounding the AP. LANPlanner 11.0 supports 802.11n planning, and has the capability to include site-specific effects into its calculations of the transmit data rate. These predictions are a must when planning new 802.11n deployments or migrating 802.11a/b/g networks to support 802.11n. Motorola recommends any company deploying an 802.11n network simulate their deployment plans in LANPlanner prior to making an upgrade decision.

Legacy-to-11n Migration Wizard

Many industry discussions have indicated that the most likely scenario for first-round 802.11n deployments will be migrations of existing 802.11a/b/g networks to support 802.11n. Since this is the case, LANPlanner 11.0 provides a Network Migration Wizard for this very purpose. Based on user-defined constraints for the migration such as region of interest, client distribution, and migration strategy, LANPlanner provides a suggested baseline deployment for the specific needs of each environment. Network planners can easily input the performance requirements of their network and then simulate multiple migration scenarios to

efficiently determine not only the best migration path for their network, but also the costs associated with the upgrade.

Comparison of 802.11a/b/g vs. 802.11n

With LANPlanner 11.0, 802.11n performance improvements are no longer left to complicated spreadsheets and non-intuitive numbers. When planning an 802.11n network, users are able to clearly visualize the improvements they are receiving from their 802.11n deployment within a site-specific heatmap. Combined with the exceptional, automated reporting capability of LANPlanner, the benefits of 802.11n migration can be effectively communicated with ease.

802.11n enabled survey capability

Due to the channel planning issues surrounding 40MHz channels, DFS compliance in the 5GHz band, and the effect of client distribution on network performance, pre-deployment surveying for 802.11n is even more important than it was for 802.11a/b/g networks. As such, the LANPlanner survey tool Motorola SiteScanner™, provides native support for 802.11n surveys using widely available, industry standard 802.11n WLAN cards.

In addition to pre-deployment surveys, this capability allows customers to verify the performance of their 802.11n APs once they have been deployed. Since characterization of MIMO system performance requires client association and data transfer from the AP, the AP Performance mode of SiteScanner has been improved to seamlessly facilitate these new requirements.

Automated tuning of performance predictions

As always, LANPlanner provides users with powerful tools for optimizing their LANPlanner network planning models. Performance data collected from SiteScanner measurement surveys can be used to close the feedback loop for 802.11n performance predictions. Through optimization, the user can achieve the maximum possible prediction accuracy for their 802.11n AP model.

LANPlanner — part of the 802.11n solution

LANPlanner is part of Motorola's comprehensive 802.11n solution, which also includes Motorola's AP-7131 802.11n access point and the RFS6000 and RFS700 802.11n ready WLAN switches.



MOTOROLA

motorola.com

Part number WP-802.11nPNM. Printed in USA 09/08. MOTOROLA and the Stylized M Logo are registered in the US Patent & Trademark Office. All other product or service names are the property of their respective owners. ©2008 Motorola, Inc. All rights reserved. For system, product or services availability and specific information within your country, please contact your local Motorola office or Business Partner. Specifications are subject to change without notice.