# Introduction

Mobility is a key focus for many enterprises as they see the business benefits of giving staff real-time access to network resources wherever they may happen to be. However, building a scalable and manageable platform for mobilising a workforce is no easy task. Where companies once adopted a tactical approach to mobility projects, they now take a more strategic approach and are looking to Service Providers to build a platform for mobilising their entire workforce.

## Business Drivers for Mobility

Specific business drivers include:

- Wire-speed wireless connections, allowing a new array of applications (CRM, ERP, media-rich web applications, etc.) to be mobilised.
- An explosion in the kinds of device that have built in wireless connectivity, from traditional laptops and PDA's to equipment such as printers and projectors.
- A move beyond email into new collaborative working practices, allowing staff to provide input to projects by updating resources over a network.
- Adoption of 'Fixed Mobile Convergence' technologies that reduce costs by seamlessly switching client devices from wide-area to local-area connections.

## The Service Provider's Challenge

The move to mobility is a significant opportunity for the Service Provider market to expand their portfolios and generate new revenue streams. However, in spite of all the work and investment that has gone on in the last decade to develop wireless broadband network access technologies, there are still remarkably few solutions that address the needs of the end-user and meet the requirements of the Service Provider.

The Service Provider's business model relies on being able to supply its customers with high quality services at low marginal cost. After the fall out from the tech market crash, Service Providers have invested heavily in building a single networking platform that is robust and flexible enough to provide economies of scale while being able to meet each individual customer's needs. The foundation of this platform is Multi-Protocol Label Switching (MPLS), which we discuss briefly below.

Easily available bandwidth makes it relatively simple to provide connectivity and access to an increasing array of applications. However, mobilising end-users remains stubbornly difficult. Providing high quality in-building wireless network access to users at low marginal cost is problematic for many reasons. We will examine how Wireless Local Area Networking (WLAN) equipment has developed over the last decade and how each successive generation has been designed to provide better integration to the Enterprise network. While each generation represents a significant improvement on the last, selecting the infrastructure that provides the best fit with the Service Provider's business model is vital. This document outlines a number of areas for consideration when making choices about the infrastructure to deploy.

## MPLS as the Backbone of Service Delivery

MPLS gives Service Providers the ability to tailor their offer to meet the specific needs of each customer from a single network platform. It was designed to combine the reliability and Quality of Service (QoS) associated with circuit-switched networks with the efficiency and scalability of packet-switched networks. As its name suggests, MPLS can transport data from many different kinds of networking solutions (Frame Relay, ATM, IP, PPP, etc.) by encapsulating packets and adding a label that keeps different data types separate. This label designates the eventual destination of each packet, the type of packet and the QoS requirements, among other things. Inherently, MPLS provides a great deal of visibility of traffic flows across a network and allows for seamless services such as traffic management and real-time monitoring, allowing Service Providers to offer transparent and individualised SLAs to their customers.

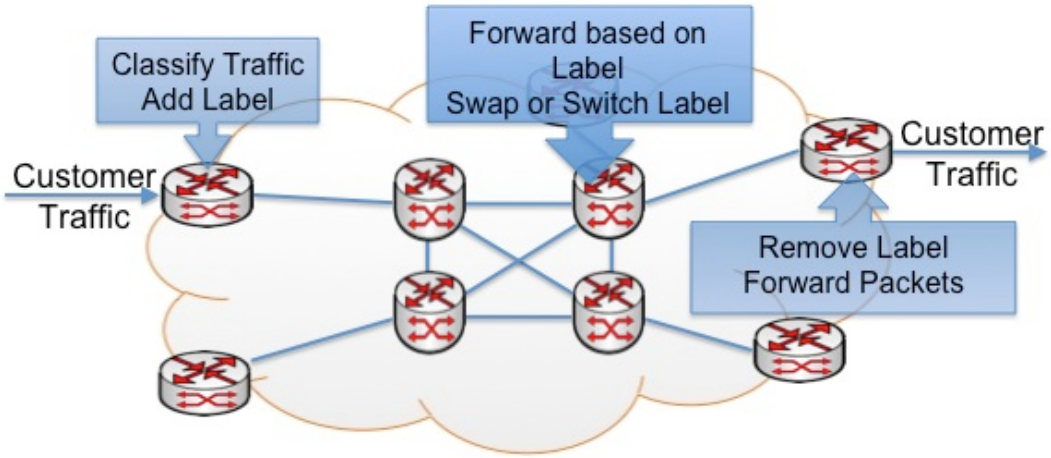


The diagram above shows how networks are interconnected using MPLS. Essentially, data is taken from the customer's network (in this case it is an ATM network, but it could just as well be an IP network), has a label appended to it and is transmitted across the network to its destination. On reaching the destination network, the label is removed and the data is returned to a state that the receiving system can understand (again, in this case ATM). As the data travels across the network, each intermediate router only needs to inspect the MPLS label in order to determine what to do with the packet. This makes the core of the network incredibly efficient and allows powerful but easily deployable control of traffic.

Below we see how traffic is classified as it travels through the core of an MPLS system. In trying to deploy mobility as part of their overall portfolio, Service Providers needs to be able to map different classes of mobile user efficiently onto this structure. Furthermore, in order to meet the basic requirements of their business models, the wireless infrastructure deployed must provide extremely high levels of reliability, be easy to deploy and manage, and provide the flexibility required for adding new applications without the need for extensive reconfiguration and, crucially, site visits.

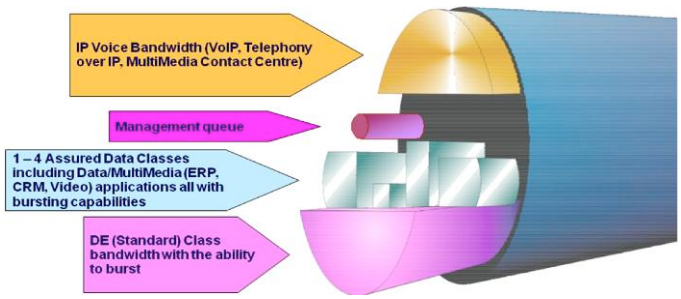

# Wireless Generations

The Institute of Electrical and Electronic Engineers (IEEE) ratified the first 802.11 standard, popularly known as WiFi, in 1997, and from this point forward the story of the WLAN has been one of continuous development.

The IEEE continues to build on the underlying architecture to make the technology appeal to a broader audience, offering increased data rates, enhanced quality of service and industry standard security mechanisms. From a product and deployment perspective, vendors have continuously re-modelled their solutions to become more closely embedded into the fabric of corporate networks, to be easy to manage, to react to changes in the radio environment and to quickly identify performance issues and security threats.

There have been four distinct phases in the development of WLAN infrastructure. In the following pages we offer a brief overview of the advancements offered at each stage.

# 1st and 2nd Generation Infrastructure

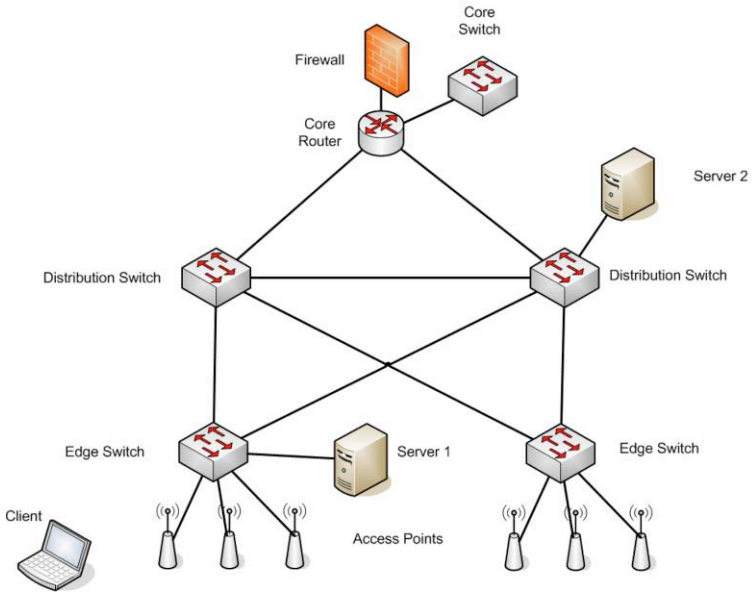


1st Generation WLAN hardware provided and extension to the network by means of a simple layer 2 bridge, translating 802.11 frames for transmission on an Ethernet network. They also provided basic security services, such as encryption, access control and authentication.

2nd Generation equipment essentially provided the same services, but allowed centralised control and configuration, thus simplifying the installation and commissioning process. MIBs became available so that AP's could become better integrated with network management systems, providing network managers with proactive monitoring and management.

| **Features** |
| --- |
| Extends network connectivity through basic layer 2 bridge |
| Security provided by Access Control Lists, Wireless Encryption, User Authentication |
| Extremely lmited QoS capabilities |
| Centralised control and management of AP's (2nd Generation Only) |
| **Benefits** |
| Very simple network design process |

## 3rd Generation Infrastructure

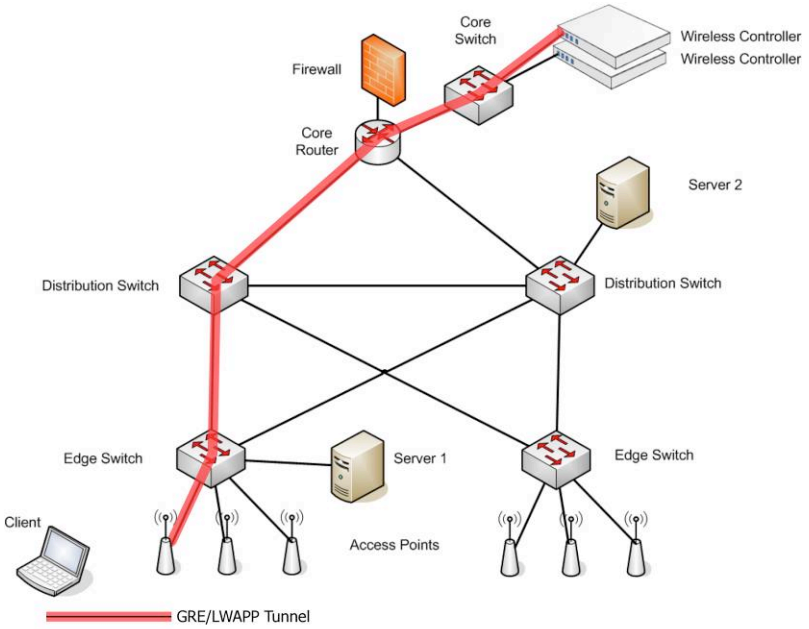


As the 802.11 standards developed it became apparent that stand-alone AP's would soon be unable to process high-level security and QoS requirements for clients. Many manufacturers took a view that a new architecture was needed, pulling all of the intelligence off the AP and instead using central controllers to manage these new services. For the first time, network managers were able to implement standards-based QoS and robust security mechanisms for mobile users.

This new architecture also allowed each AP to run multiple WLAN's simultaneously, segregating users and application flows. This increased the flexibility of the systems and allowed for new services to be deployed, such as mobile VoIP and 'hotspot' access.

| **Features** |
| --- |
| Provide multiple WLAN's to separate users and application flows |
| Traffic routing capabilities to provide clients with Layer 3 roaming |
| Traffic management capabilities including standards based QoS |
| Real-time management of radio environment |
| **Benefits** |
| Reduced cost of ownership |
| Greater control of mobile users |
| Flexible framework for adding new services, such as VoIP or guest internet access |

# 4th Generation Infrastructure

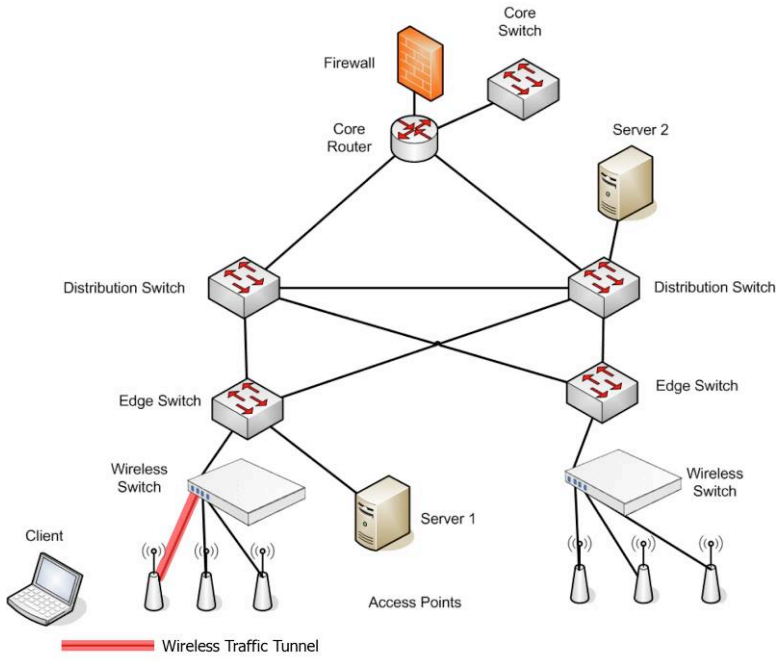


4th Generation systems use controllers located at the edge of the network communicating directly with 'ultra-thin' AP's. These AP's contain multiple radios that provide blanket coverage on all channels simultaneously. Each blanket of coverage appears to the client as a single cell, so that devices can move seamlessly without the need for roaming.

These slight changes in the architecture developed for 3rd generation systems give network owners all of the flexibility and control they require in a more manageable and efficient package.

| **Features** |
| --- |
| Efficient use of radio spectrum by providing blanket coverage on multiple channels |
| Physical and logical separation of traffic from different user groups |
| **Benefits** |
| Traffic takes direct route to network resources |
| End-to-end QoS maps directly onto MPLS architecture |
| Adding new services requires minimal configuration |
| Transparent traffic flows simplifies service delivery and SLA management |

# Practical Considerations

While all WLAN equipment provides basic support for mobilising users and applications, Service Provider's requirements are more complex. In order to deliver Mobility as a Service, wireless infrastructure must:

- Offer a stable platform for service delivery
- Give complete centralised control
- Provide consistent high performance
- Provide a secure environment
- Allow easy addition of new services
- Offer the segregation of user groups
- Integrate with MPLS platform
- Give a transparent view of traffic for monitoring, troubleshooting and billing purposes

With these criteria in mind, let us now examine how each generation of infrastructure stacks up.

## Radio Spectrum Design and Use

Radio spectrum is a scarce resource. There are many systems making use of the airwaves, from radio and television services to mobile telephony, building alarm systems to RFiD tags. WiFi has been designated for use in unlicensed radio spectrum with strict limitations placed on the power that can be transmitted. The diagram to the right shows the available channels that can be used in WiFi systems, with each channel sharing almost 80% of the spectrum with those around it.

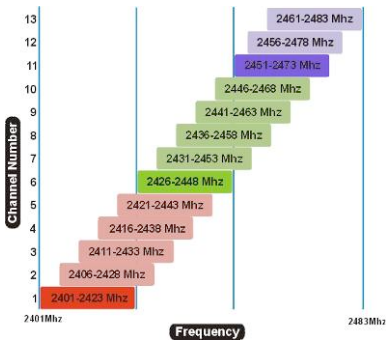


Channel availability for Wi-Fi systems in 2.4GHz radio spectrum

WiFi is a cellular system, much like the mobile phone network, where an Access Point/Port (AP) transmitting and receiving on a single fixed channel controls each cell. The 802.11 protocol insists that each node on the network must be sure that the channel they are transmitting on is clear before they can send any data. Therefore it is important that adjacent AP's use channels that do not overlap, otherwise transmissions in one cell will disrupt connections in those nearby. The bold colours in the diagram show that there are only three non-overlapping channels in the available spectrum. This makes designing the radio spectrum in a WiFi system an extremely challenging task.

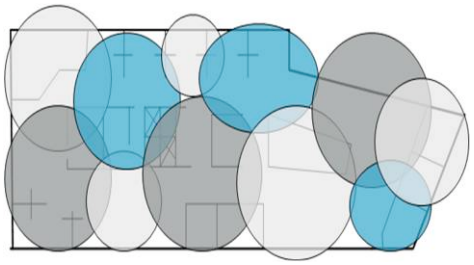

Traditional 'vertical' channel planning to avoid interference between cells.

To ensure that the system is as efficient as possible, many companies recommend that a site survey be conducted prior to installation. This is a laborious and expensive task, involving many dozens of measurements being taken to ensure correct placement of AP's. Incorrect planning, however, can cause the system to become unusable once installed. To overcome these issues, many 3rd Generation WiFi systems can automatically configure channel settings to avoid conflict, with many being able to self-heal by adjusting transmission power in the event of an AP failure. While this takes into account many potential issues it is far from a perfect solution. 4th Generation use multiple independent radios in each AP, thereby providing blanket coverage on all channels in all locations, doing away with the need for an RF design process at all.

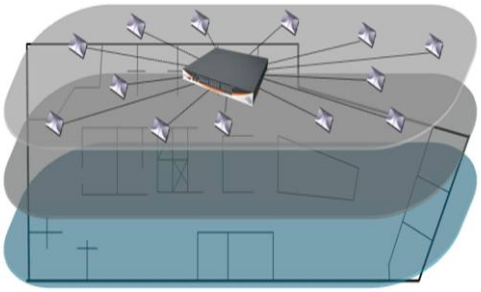

4th Generation 'horizontal' RF environment gives blanket coverage with no cross-cellular interference

## Roaming Performance

Wireless connections are by their nature unreliable. Radio signals are affected by their environment, reflecting off metallic surfaces but being absorbed by other materials such as wood and paper. As devices move throughout a facility, the rate at which they can communicate with an AP varies, as the strength of the signal they receive alters depending on distance. Therefore, clients must make intelligent decisions about whether and when to roam between cells to maintain the underlying quality of their service levels.

The algorithms that clients use to make roaming decisions have never been standardised. Generally, users are left at the mercy of component manufacturers with regards to how their device will use the WLAN. So whether a device roams before scaling back the rate of communication, or whether it clings on to a connection until it becomes completely untenable is uncertain until it is in use. From a network administrator's point of view, this is a nightmare scenario. The air is a shared medium and so one device running on a slow connection will disrupt communications for all those around it.

Since the 802.11 standard beefed up security and added prioritisation mechnisms, these roaming decisions have become increasingly complex. Now with extra authentication handshaking and registering for relevant QoS, the number of roaming messages has increased dramatically. This measurably slows the hand-off process to a degree that can significantly harm VoIP communications. The significance of these changes is illustrated to the right. The extended process shows a client using EAP authentication and registering for HCCA Scheduled Access

While roaming between AP's is common to all 1st, 2nd and 3rd generation infrastructure, 4th Generation WLAN's provide a single cell of coverage so that clients do not need to roam. Roaming is a process that requires a significant amount of management traffic between client and AP. Doing away with roaming reduces this traffic markedly, so that more air time is available for clients to send data. 4th Generation WLAN's therefore inherently provide greater performance and reliability than previous generations, which is especially important for low-latency traffic such as voice.

To overcome the issues in earlier generations, the IEEE recently added an addendum to the 802.11 standard to speed up and standardise the roaming process. However, implementation of the new process has so far been slow.

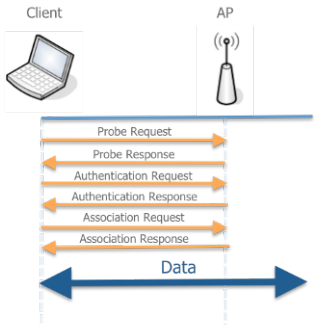


Original 802.11 connection process

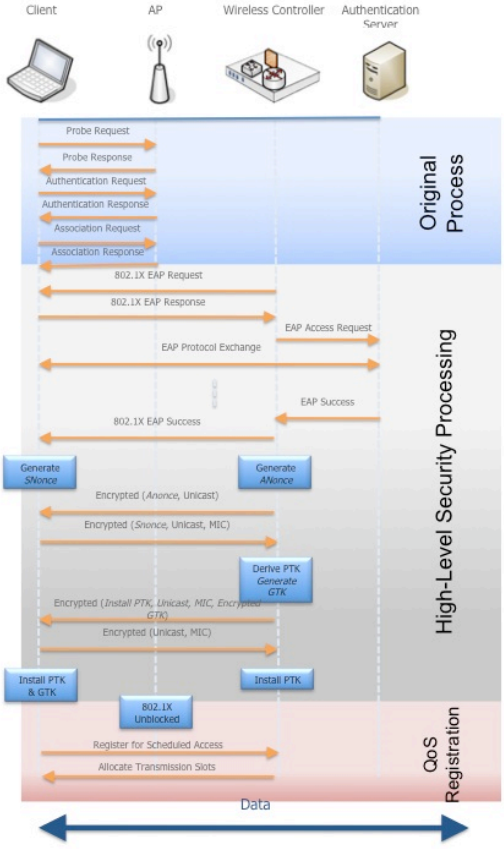


Extended 802.11 connection process