


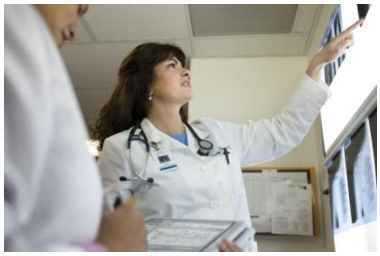

Reciprocal Wireless Access

Business Drivers

In many public sector organisations there are requirements to share buildings and workspace. In the United Kingdom, this is particularly true for NHS trusts, teaching hospitals, and universities in close proximity of each other. There is increasing pressure from the government to make cost savings as detailed in the Comprehensive Spending Review of November 2010.

The government signalled the need for unprecedented cultural and business change, requiring organisations to merge operations, share services, and introduce new ways of working. ICT infrastructure was identified as one of the key services that must be shared.

On March 30, 2011, the Cabinet Office published a report on the new government ICT strategy. The strategy recommends an approach to efficiency and cost savings based on deployment of common ICT infrastructure and the use of ICT to enable and deliver change.

This paper defines a solution that takes advantage of existing Cisco® wired and wireless infrastructure to increase location-independent working capabilities across public sector estate. There are a number of benefits to be achieved by adopting this approach. These include cost savings and productivity gains. For example, cost savings are made because we are leveraging an existing investment in the ICT infrastructure. There are no separate overlay networks to deploy. In addition, users will experience productivity gains because they can access their network resources seamlessly without any client reconfiguration.

Use Case Scenarios

Scenario 1

University students studying for a medical degree may well spend time on the hospital campus in workplace experience initiatives. They will require access to their email, files and applications on their home network in the university.

Scenario 2

A consultant from the hospital may spend time on the university campus presenting lectures to medical students and require access to applications, email and file shares on the hospital network.

Scenario 3

Community nurses and social care workers can use their nearest shared office to work from, regardless of whether that is a primary care trust office, hospital campus, or community centre.

Architectural Overview

The Cisco Unified Wireless Network architecture centralises WLAN configuration and control onto the Cisco Wireless LAN Controller (WLC). The Cisco Unified Wireless Network simplifies operational management by collapsing large numbers of managed access points into a single managed system comprised of the WLAN controller(s) and its corresponding, joined access points.

In the Cisco Unified Wireless Network architecture (see Figure 1), access points are "lightweight", meaning that they cannot act independently of a WLC. Access points are "zero-touch" deployed and no individual configuration of access points is required. The access points learn the IP address of one or more WLCs via a controller discovery algorithm and then establish a trust relationship with a controller via a "join" process. Once the trust relationship is established, the WLC will push firmware to the access point if necessary, as well as a configuration.

Once joined to a controller, the access points are also lightweight in the sense that they handle a subset of 802.11 MAC functionality. Typically, this subset includes only real-time 802.11 MAC functionality, with the controller handling all non-real-time 802.11 MAC processing. This architecture enables support for seamless mobility and a number of advanced features in an elegant and scalable way.

Figure 1. Cisco Unified Wireless Network

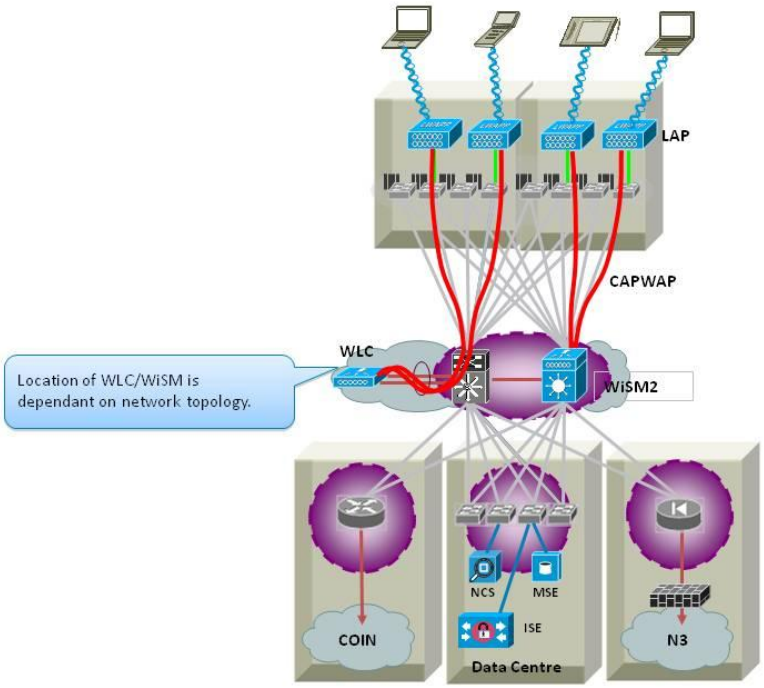

Building on this foundation, we can use **mobility anchoring** to create a solution for sharing wireless infrastructure across organisations.

Mobility Anchoring

Under normal roaming conditions, client devices join a wireless LAN and are anchored to the first controller that they contact. If a client roams to a different subnet, the controller to which the client roamed sets up a foreign session for the client with the anchor controller. However, when you use the auto-anchor mobility feature, you can specify a controller or set of controllers as the anchor points for clients on a wireless LAN.

As an example, if we take a client that normally resides on the university campus, and connect this client to the wireless on the NHS trust, that client will be anchored to a controller on the university campus. This is called the anchor controller and all traffic is tunneled back from the wireless controllers in the NHS trust to this anchor controller in the university.

In Figure 2, users from the university can connect their Wi-Fi devices to the university service set identifier (SSID) that is being broadcast in the hospital. Traffic on this WLAN is tunneled via Ether over IP (EoIP) back to the university Anchor WLC and then switched on to the university LAN via an 802.1q trunk on to the core switch.

Figure 2. Mobility Anchoring

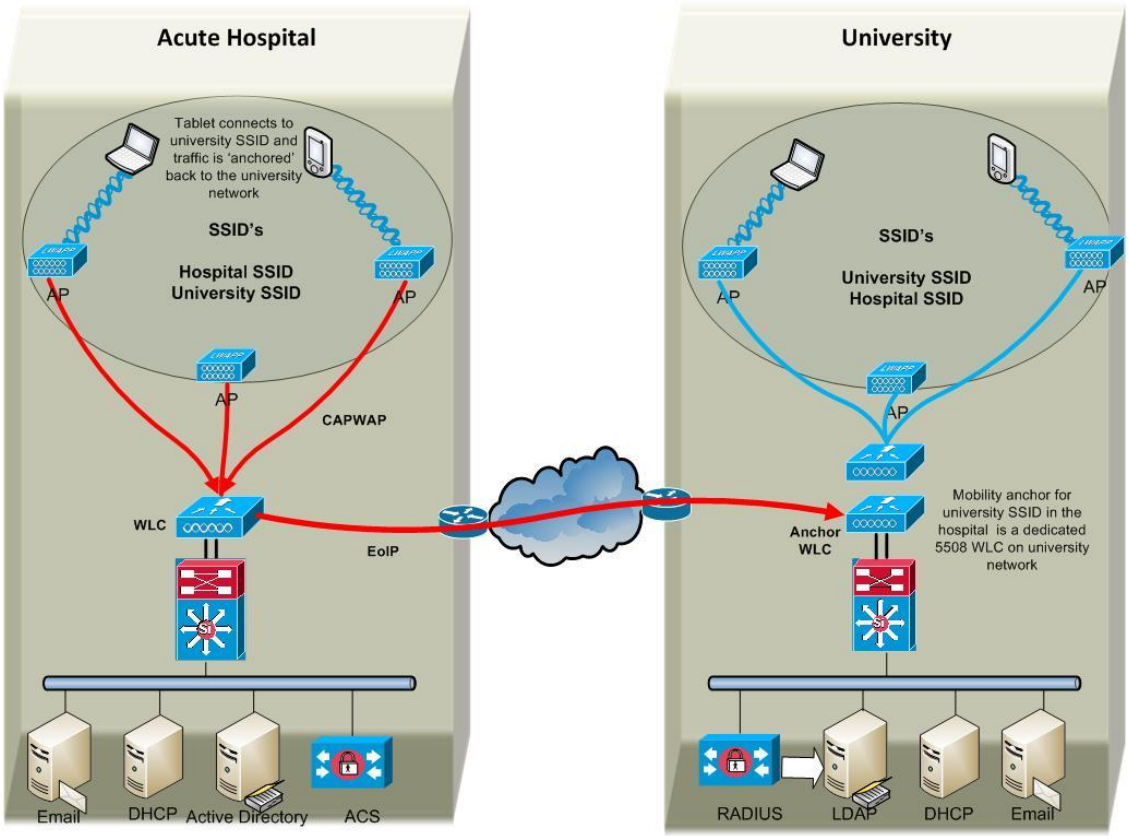

This is also a reciprocal arrangement in that when hospital users spending time in the university associate to the hospital SSID, this traffic will be tunneled back to an anchor controller on the hospital LAN, as depicted in Figure 3.

Figure 3. Reciprocal Mobility Anchoring

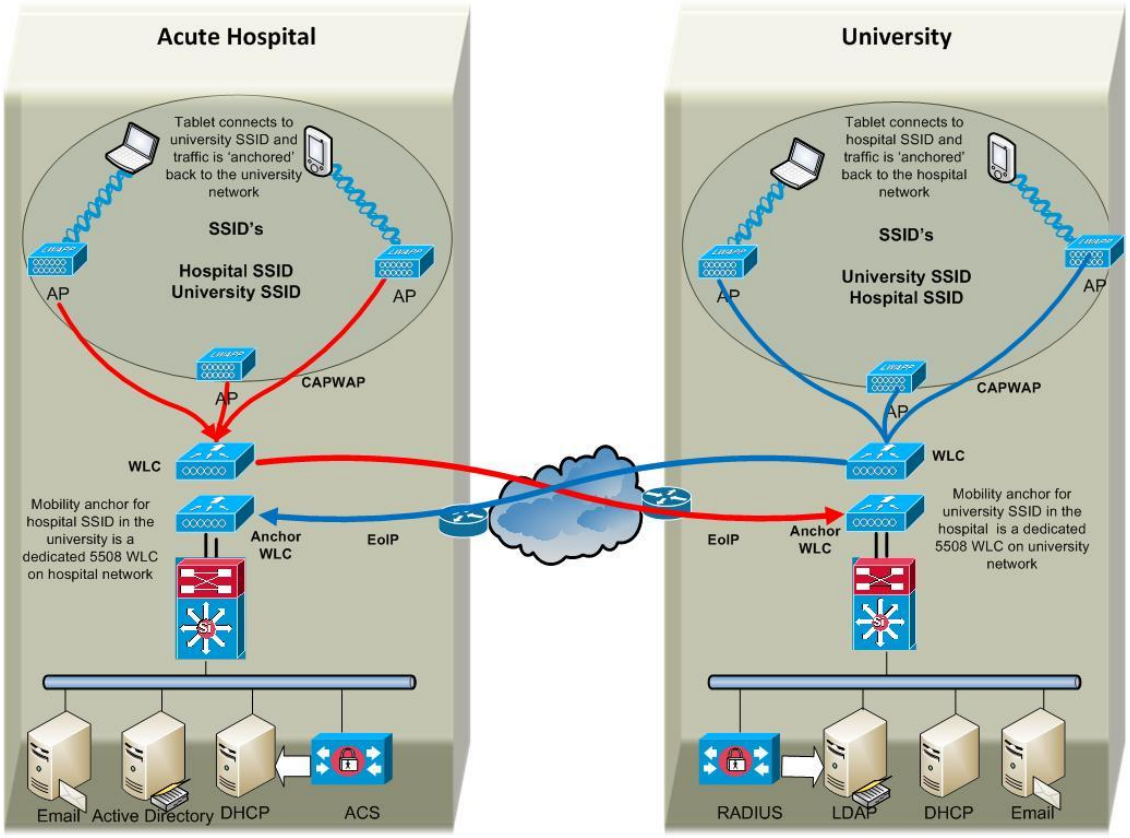

Path Isolation

The anchor controller is usually located in the ”home” network. Other internal WLAN controllers from where the traffic originates are located in the ”foreign” network enterprise LAN. An EoIP tunnel is established between the foreign controller and the anchor controller in order to help ensure **path isolation** of visitor traffic. Path isolation is a critical security management feature for visitor access. It helps to ensure that security and quality of service (QoS) policies can be separate, and are differentiated between visitor traffic and internal traffic.

An important feature of the Cisco Unified Wireless Network architecture is the ability to use an EoIP tunnel to statically map one or more provisioned WLANs (that is, SSIDs) to a specific anchor controller within the network. All traffic—both to and from a mapped WLAN—traverses a static EoIP tunnel that is established between a remote controller and the anchor controller.

Using this technique, all associated visitor traffic can be transported transparently across the enterprise network to an anchor controller that resides in the home network.

The EoIP tunnel carries the visitor traffic from the internal WLAN controller to the anchor controller in the clear with no encryption. It is important to point out that the primary requirement for securely carrying guest traffic is path isolation. Security features such as confidentiality or integrity, delivered by technologies such as IPsec, are not required and offer little to no additional risk mitigation when compared to EoIP tunnels. If a guest user wishes to ensure confidentiality of their traffic, they will simply rely on solutions such as IPsec VPN clients or application level encryption in the form of Secure Sockets Layer/ Transport Layer Security (SSL/TLS). It could be argued that in fact not encrypting traffic delivers a level of visibility to the NHS trust to be able to monitor visitor user activity and intervene if malicious activity is detected.

There are no known attacks against EoIP tunnels that would allow a user to break-out and start attacking the internal infrastructure. Of course, misconfiguration of the WLAN could lead to compromise. However, this attack vector is mitigated not through technology, but through comprehensive and robust operating and change-control procedures.

Security Policy

Security policy dictates how a user is authenticated and what that user is authorised to access. This policy may depend on who the user is, which device the user has, where they are accessing the network from, and what time they are accessing the network. A policy can then be applied based on these criteria.

Security policies and implementations will vary across organisations to some degree. For example, an NHS trust may use 802.1x/Protected Extensible Authentication Protocol (PEAP) to authenticate users via RADIUS in to Microsoft Active Directory. A university may use EAP-TLS and not use Active Directory at all.

Therefore, it is important to be able to configure the appropriate security settings on the WLAN that is being offered to visitors so that network access is consistent and seamless wherever they are in the shared estate (Figure 4).

Figure 4. Security Applied to a Consultant Visiting the University

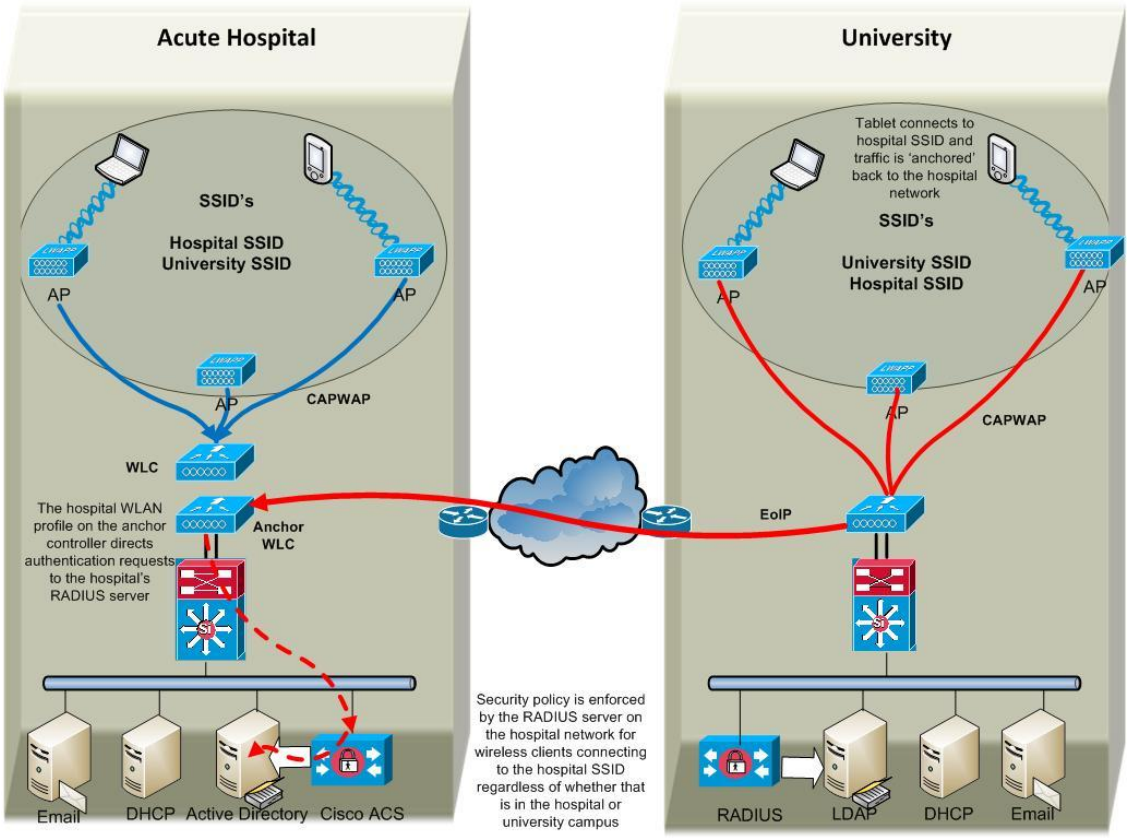

In Figure 4, a consultant from the hospital is spending some time in the university and needs to access his email and file share in the hospital data centre. The hospital has a security policy that mandates the use of 802.1x/PEAP for authentication and Wi-Fi Protected Access 2 (WPA2) for encryption over the wireless network. This policy may be different to what the university has and we don't want to have to reconfigure the consultant's iPad to access a university SSID. Alternatively, the consultant could associate to a Guest SSID, but this will require him to have a VPN client set up on his iPAD and create a VPN session via the Internet.

The simplest solution is to use the same SSID that he would normally connect to in the hospital. Therefore, we configure and broadcast the hospital SSID on the university access points. This WLAN is configured with the same security parameters as the WLAN in the Hospital, so that it is using 802.1x/PEAP with WPA2 and the RADIUS servers that reside in the hospital.

When the consultant connects to the hospital SSID in the university, access to the network is not allowed until successful authentication is completed. This is achieved by sending the authentication request to the hospital's RADIUS server and applying policies from here. The result is access for the consultant that is as seamless as if he were on his own network.

Mobility Groups and Roaming

A WLAN client must be able to maintain its association seamlessly from one access point to another, securely, and with as little latency as possible. These mobility requirements are completely supported by the Cisco Unified Wireless Network architecture.

A wireless client roams when it moves its 802.11 association from one access point to another access point. Wireless client devices initiate roaming based on the internal roaming algorithms programmed into the client radio firmware. Typically, a client's roaming logic is triggered by crossing a received signal strength indicator (RSSI) or signal-to-noise ratio (SNR) threshold that causes the client to look for a better signal from a new access point. Device roaming behavior and performance differ by vendor, so it is wise to characterise device roaming and look to device configuration best practices from the client vendor.

WLAN clients are always reauthenticated by the system in some way on a roam; this is always necessary to protect against session spoofing and replay attacks. Normally, the reauthentication requires a full authentication transaction. In the case of 802.1X authentication, a full EAP reauthentication and rekeying will be required. However, the Cisco Unified Wireless Network supports two methods of fast secure roaming that short-cut the reauthentication process while maintaining security: Cisco Centralized Key Management (CKM) and Proactive Key Caching (PKC). While no special client software is required for roaming, Cisco CKM and PKC do require supplicant support.

From a roaming client's perspective, it is only roaming between access points. However, there is considerably more going on in the background on the controllers that is not visible to the client. When a wireless client associates and authenticates to an access point, the controller that this access point is joined to places an entry for the client in its client database. This entry includes the client’s MAC and IP addresses, security context and associations, QoS context, WLAN and associated access point. The WLC uses this information to forward frames and manage traffic to and from the wireless client. When a client roams, this client database information must be updated and possibly copied or moved to another controller.

Cisco Unified Wireless Network client roaming is categorised in three ways:

Intra-controller roaming
Inter-controller Layer 2 roaming
Inter-controller Layer 3 roaming

Figure 5 illustrates the first case, intra-controller roaming.

Figure 5: Layer 2 – Intra Controller Roaming

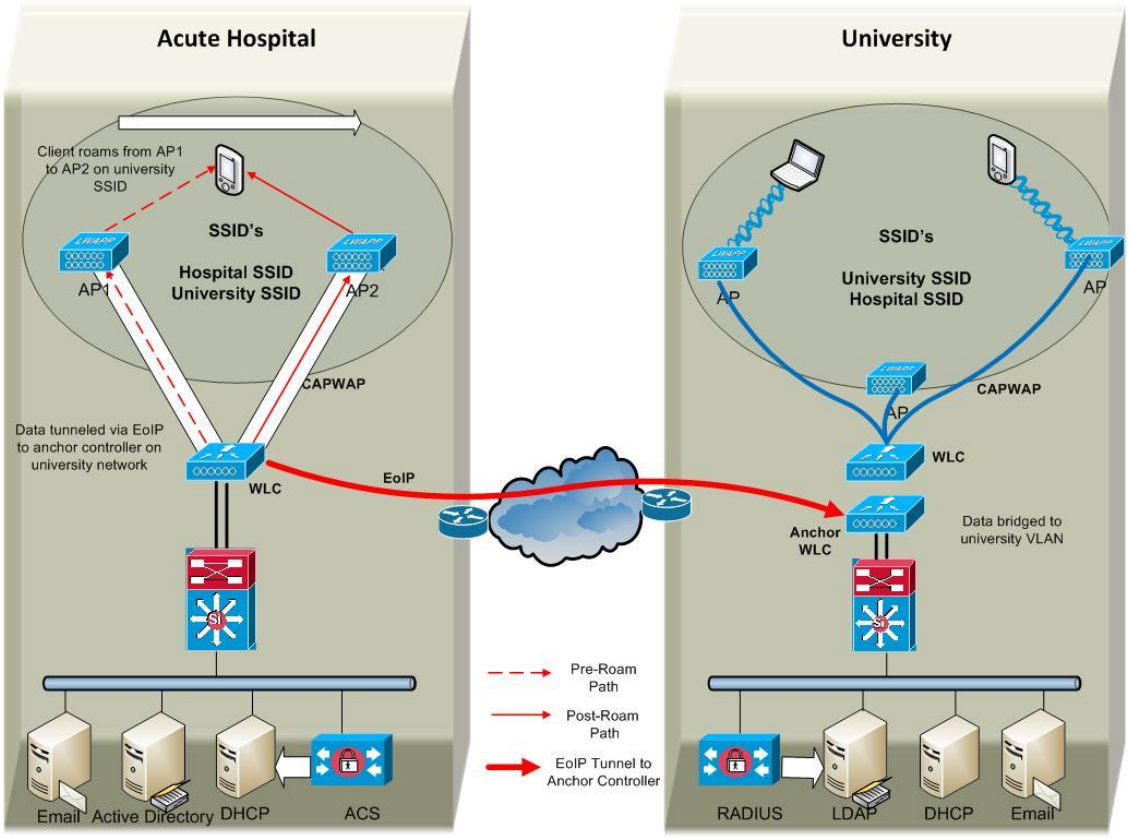

When the wireless client moves its association from one access point to another, the WLC simply updates the client database with the new associated access point. The client is reauthenticated to establish a new security context.

Consider what happens when a client roams from an access point joined to one WLC and an access point joined to a different WLC. Figure 6 illustrates an inter-controller roam in the event of a Layer 2 roam.

Figure 6: Layer 2 Inter-controller Roam

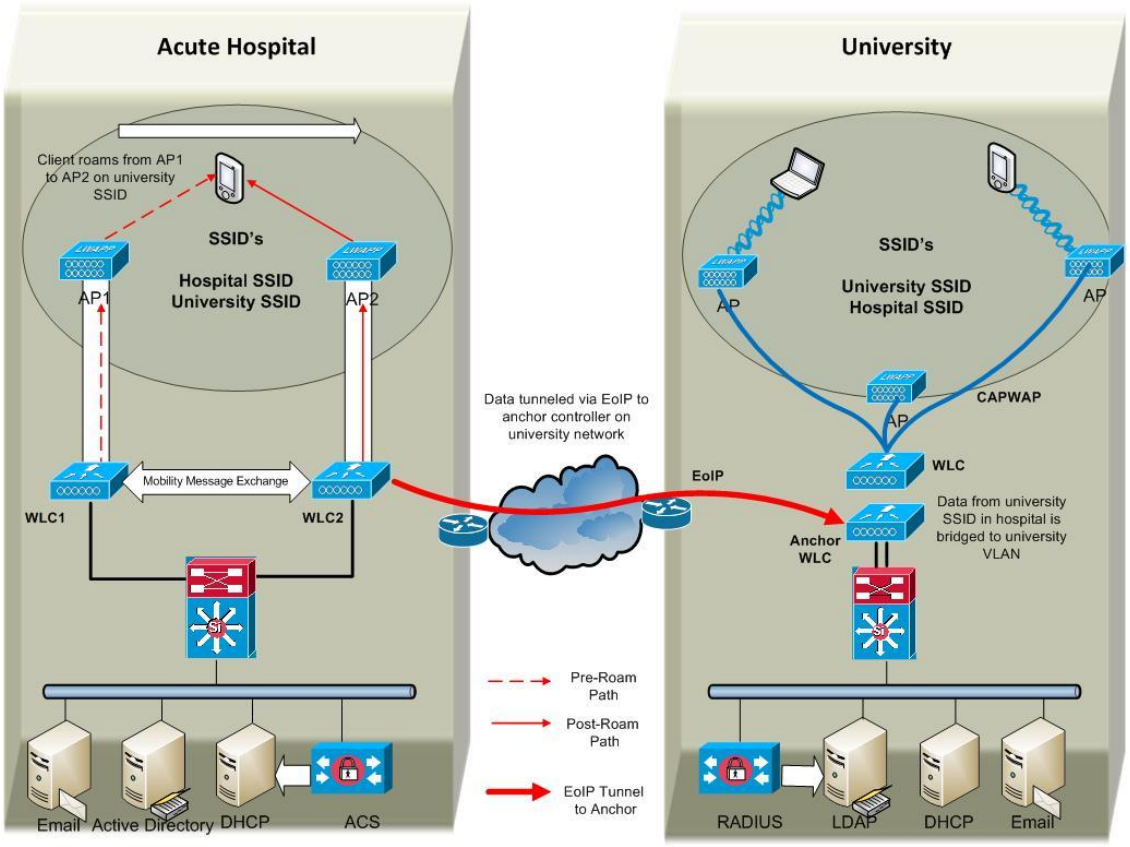

The reciprocal wireless access solution with mobility anchoring supports Layer 2, intra- and inter-controller roaming within the same mobility group.

Layer 3 inter-controller roaming is not supported with mobility anchoring. However, for the shared wireless access solution, it is not required, as all traffic is tunneled back to the anchor controller and is always bridged on to the same VLAN.

Cisco Centralised Key Management

Cisco Centralized Key Management (CKM) shortcuts the reauthentication process by caching a set of master key materials in the controller and on the client device. When the device roams, the master keying material is used both to authenticate the client and negotiate a new Pairwise Transient Key (PTK) for the new session.

To implement Cisco CKM requires support on the client devices and supplicant. Cisco CKM is supported for all EAP types in Cisco Compatible Extensions Version 4. However, you will need to check with your device vendor for Cisco CKM support.

Proactive Key Caching

Proactive Key Caching (PKC) is an extension of an optional component of the 802.11i security specification. With PKC, the controller and client cache Pairwise Master Keys (PMK) after the initial authentication. When the client device roams, it embeds a unique number called the PMKID in its reassocation request. If the controller can match the PMKID using its cached PMKID, the controller and client can skip EAP authentication, going directly to the four-way handshake to derive a new PTK.

PKC is extremely fast, but like Cisco CKM, requires client and supplicant support. PKC also is only supported with WPA2 and Advanced Encryption Standard (AES) encryption. Again, check with your device and supplicant vendors for PKC support. PKC is sometimes called Opportunistic Key Caching (OKC).

Latency Expectations

Latency in roaming times will vary depending on various factors such the client device itself and RF conditions. Therefore, we need to take these factors in to account when evaluating vendor claims:

> Client devices initiate roaming. The algorithms that clients use to decide when to roam vary based on vendor implementation. Some clients are inherently more “sticky” than others.
> The algorithms that clients use to select a new access point once they've decided to roam vary by vendor implementation. Some clients take longer than others to select a new access point.
> Reauthentication times are affected by factors like network latency and RADIUS server performance.
> The RF environment makes a difference; busy RF environments affect roaming frequency and times.
> Access point placement, co-channel interference levels, and supported data rates affect client roam times.

These factors have implications for roaming times in your environment, so it is wise to test your client devices in your environment. In most environments, roam times well under 100 milliseconds are achievable with Cisco CKM and PKC. These roam times are more than adequate to support a properly installed voice over WLAN (VoWLAN) implementation.

Roaming between two different mobility groups is not supported. For example, the University and the Hospital would be in different mobility groups; therefore, a client’s connection would drop and reassociate as it moved between these two networks.

QoS Considerations

QoS on a pervasive WLAN is much more than simply prioritising one type of packet over another. WLAN traffic is nondeterministic; channel access is based on a binary back-off algorithm defined by the IEEE 802.11 standard (CSMA/CA) and is by nature variable, based on the number of clients accessing the network. Mobility makes this challenge more difficult.

All WLAN traffic that passes between the access point and the wireless LAN controller is encapsulated using User Datagram Protocol (UDP) CAPWAP. CAPWAP encapsulation maintains the Layer 3 marking in the original packet. Once the CAPWAP packet is de-encapsulated at the access point or wireless LAN controller, the original Layer 3 marking is again used by QoS mechanisms in the network infrastructure. With this capability enabled in CAPWAP and the Cisco Unified Wireless Network infrastructure, the network can achieve end-to-end QoS for latency sensitive traffic, over the air and across the wired network.

Conclusion

This document has explained how the network can be used as the platform to increase location-independent working capabilities across public sector organisations. In addition, this document has highlighted an approach to efficiency and cost savings based on the deployment of a common ICT infrastructure and the use of ICT to enable and deliver change.

Cisco advocates that organisations take an architectural approach to business strategy allowing technology to become an enabler of business and service transformation. The network is pervasive and is therefore the most logical place to create the platform whereby projects can be delivered on a flexible, intelligent infrastructure that supports the organisations need to be agile and at the same time enabling it to re-design services without being constrained by technology.

Appendix A - References

UK Government ICT Strategy resources

http://www.cabinetoffice.gov.uk/resource-library/uk-government-ict-strategy-resources

HM Treasury Spending Review

http://www.hm-treasury.gov.uk/spend_index.htm

Enterprise Mobility Design Guide

http://www.cisco.com/en/US/docs/solutions/Enterprise/Mobility/emob41dg/emob41dg-wrapper.html

Important Notice

"The guidance provided in this document is of a generic nature and cannot be specific to your organisation or operations. Please contact your Cisco partner or Account Manager to discuss your specific requirements. The guidance is provided in good faith based upon reference materials sourced from Public Sector organisations up to the date of publication. Errors and omissions are accepted. No warranty is given or implied."