Underexposed risks of public Wi-Fi hotspots

This thesis outlines the consideration that should be given by Wi-Fi hotspot users to the significant and as yet unresolved security risk such use presents. It promotes awareness and outlines current thinking in addressing the issue

by Daan Stakenburg and Jason Crampton

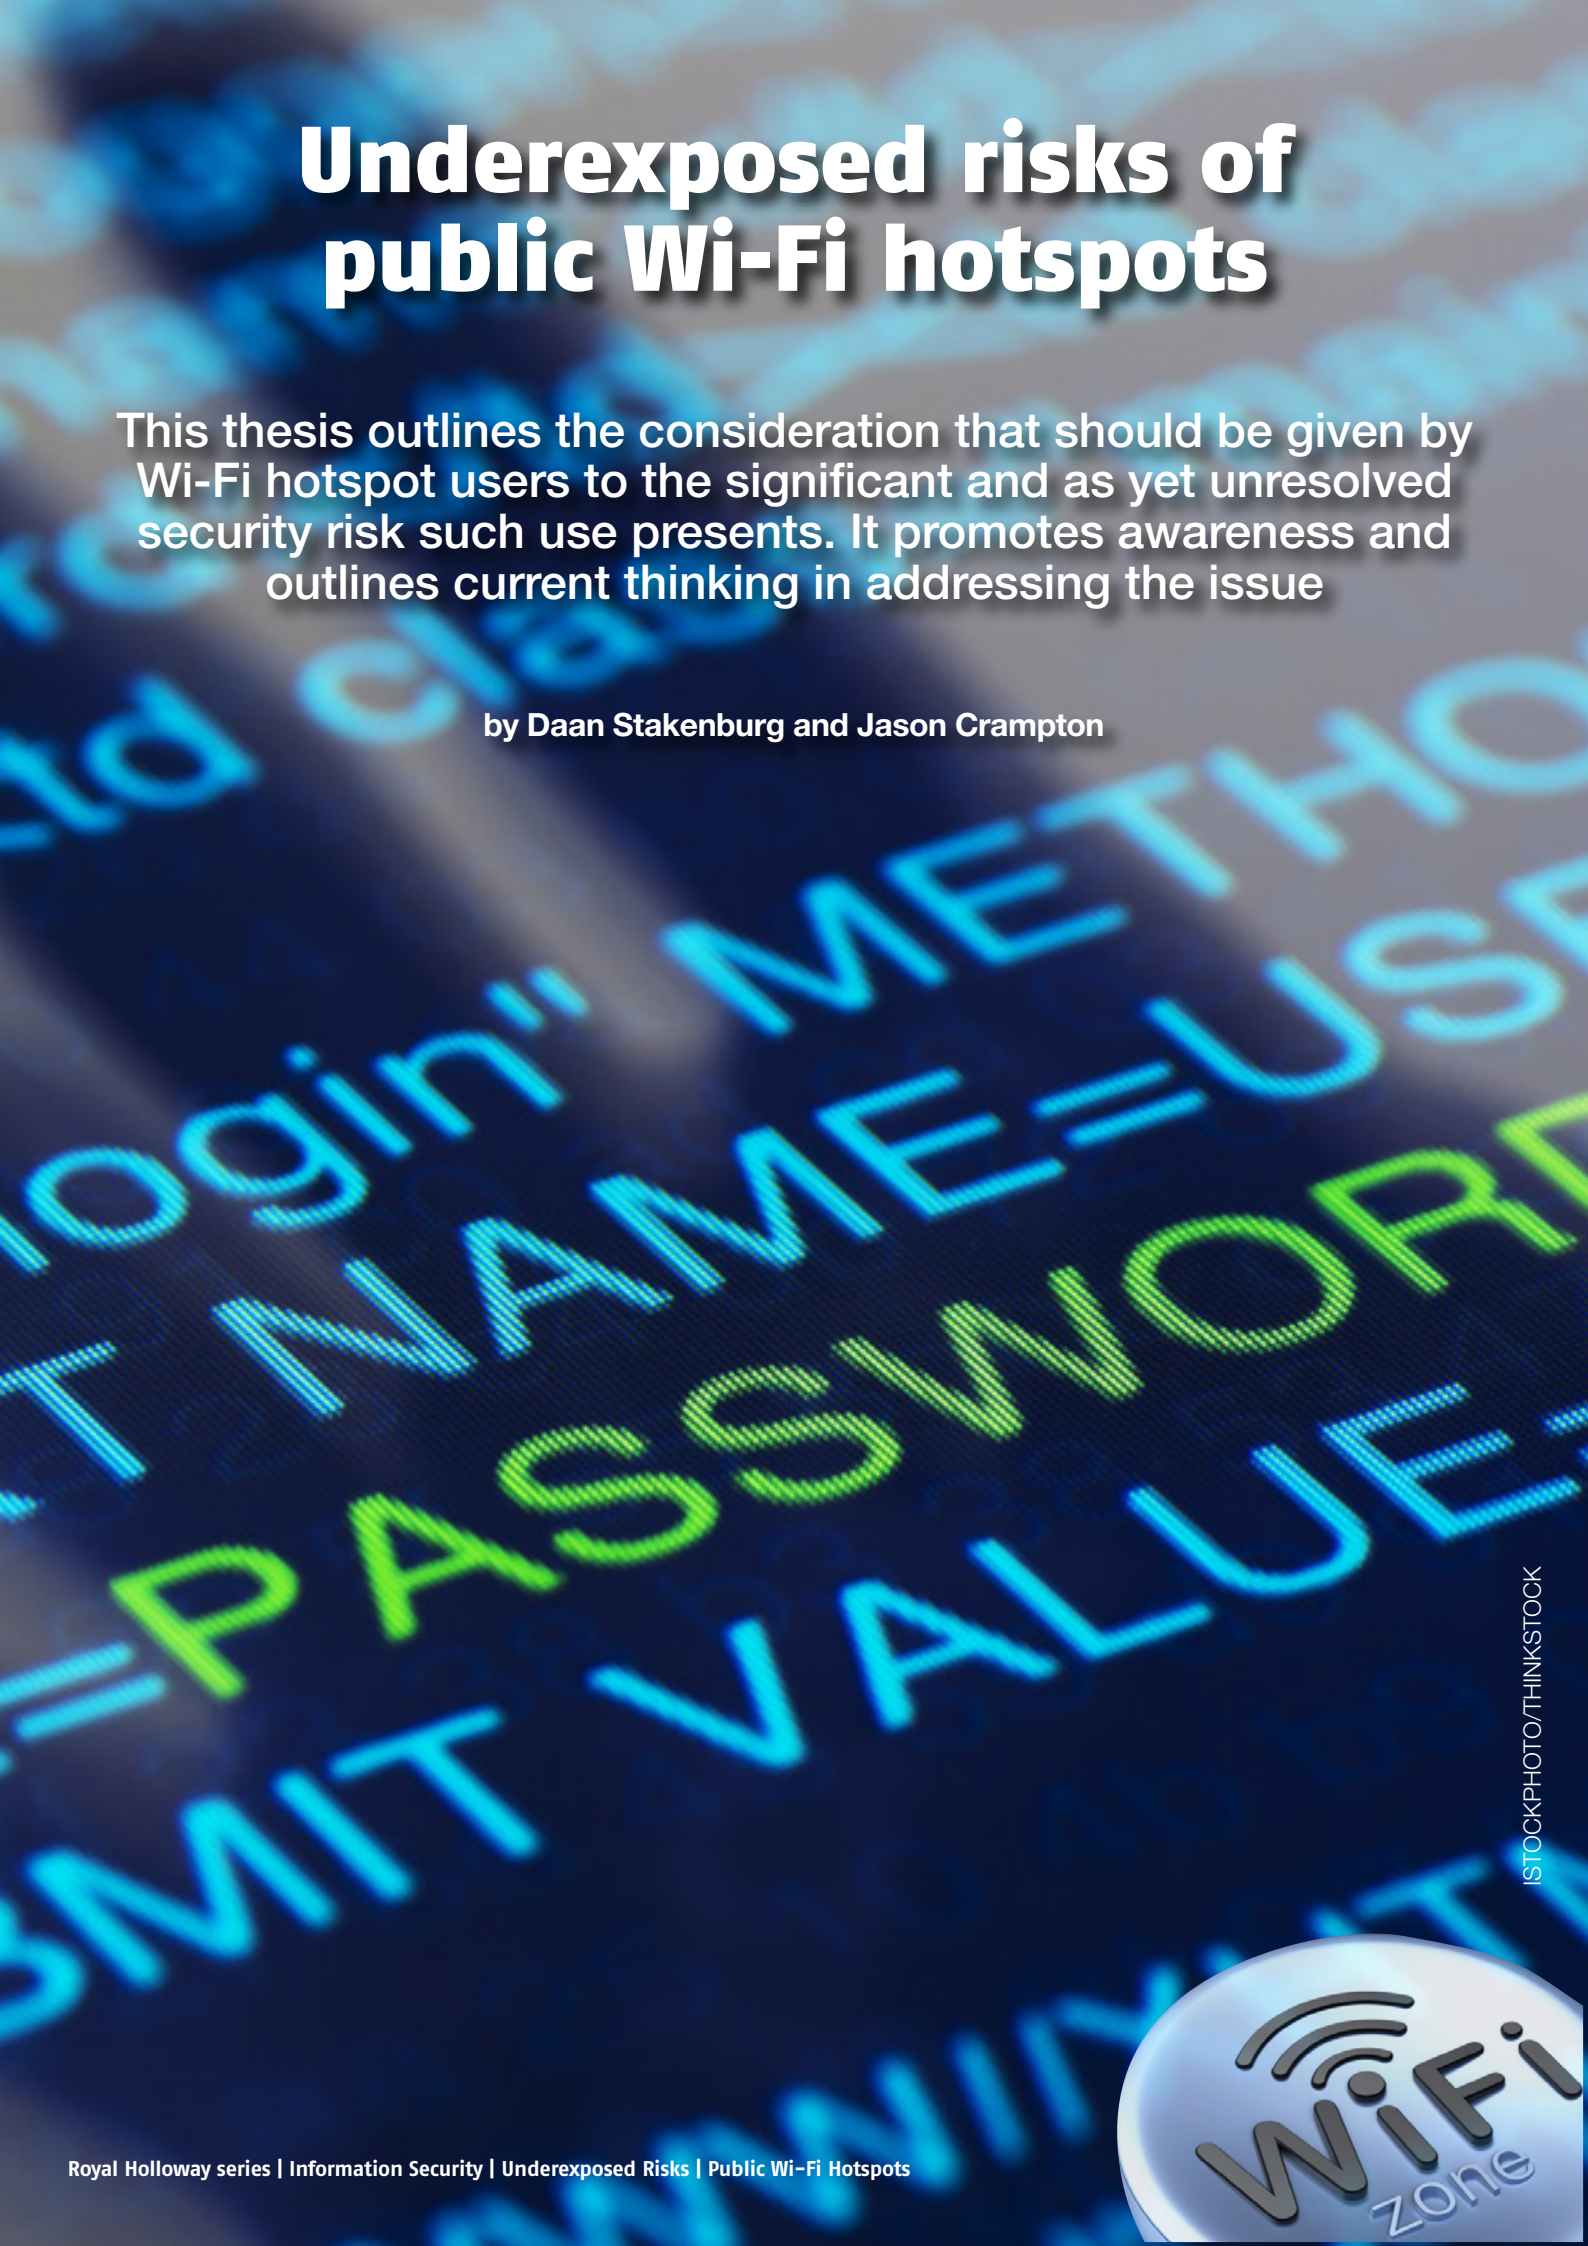

ISTOCKPHOTO/THINKSTOCK

Raising awareness of risks associated with the use of public Wi-Fi hotspots

All is not always as it appears when users access public Wi-Fi networks via seemingly authentic and trustworthy providers. This thesis seeks to raise awareness of the underexposed risks for identity and data theft by exploring the status quo and potential developments for minimising those risks

by Daan Stakenburg and Jason Crampton

With the increasing use of smartphones and the deployment of 4G or long-term evolution (LTE) mobile networks, it is easy to forget the other high-speed internet access method that is widely used by roaming mobile users: publicly accessible wireless networks, also known as Wi-Fi hotspots. Hotspots lack the standard security measures seen in other wireless networks, and if applications do not implement security features, or implement them badly, these hotspots can, in some cases, result in very insecure combinations that may lead to identity theft and information loss.

Hotspots are very common in airports, hotels and coffee shops. They are also starting to appear on public transport, such as trains and buses, in supermarkets and in other, less obvious, establishments. The reason for their success is due to the fact that they are relatively easy to set up and provide the network owner with an additional direct or indirect stream of revenue. BT, for example, established around 500,000 Wi-Fi hotspots in London for the 2012 Olympics.

Although there is nothing wrong with providing or using such networks, it is important to recognise that such networks could be deployed by people with malicious intent.

The problem

What's in a name? Everything! The main threat to users of hotspots is that it is difficult to corroborate the identity of a public wireless access point. Each wireless access point announces a so-called service set identifier (SSID) to identify the wireless network. These SSIDs are the names that are listed when a mobile device searches for nearby wireless networks. Once successfully connected, the SSID is usually saved by the device to facilitate subsequent connections to the same network.

Anyone can buy and configure a wireless access point and choose the name it announces. This could be related to the name of an establishment such as DonutsAndBagels, but it could also be the name of a wireless network that is commonly used, such as the wireless network name of a local telecom provider or the name of a restaurant chain.

Herein lies the problem. When visiting a shop called Donuts & Bagels, one would assume that the wireless network announcing the name DonutsAndBagels in the shop's vicinity is associated with the shop, but this assumption may be unjustified. By deceiving users into thinking they are connecting to a genuine access point, an attacker could perform man-in-the-middle attacks, potentially allowing him to gain access to sensitive information.

Becoming the man-in-the middle does not have to be the attacker's primary objective. There are other ways of abusing the inability to properly authenticate a nearby access point and SSID.

Example of a captive portal

Wi-Fi roaming: it's not a bug, it's a feature

As and when necessary, most networked client devices will connect to other access points announcing a known SSID – an understandable feature when you consider the fact that a wireless network may contain more than one wireless access point. This allows a client device to jump or roam between access points without the need for user interaction. But the same happens when there is a rogue access point advertising the SSID of a wireless network to which the device has connected before.

An attacker simply has to find the name of a network to which the laptop has previously connected

Consider the following scenario when a teleworker returns to the office with his dual-homed laptop (*Figure 1*). A dual-homed system is a device that can connect into two different networks. In this case the laptop would be connected to the corporate network with a cable (e.g. through its docking station) while having a wireless network interface that is continuously listening for the availability of any previously configured wireless networks it can connect to.

While the teleworker is sitting at his desk with a firewall protecting the corporate network from the threats on the internet, an adversary tries to use the laptop as a proverbial back door to gain access to the corporate network, systems and data. By setting up a nearby access point with a previously used SSID, the attacker could trick the laptop into connecting to this rogue access point. This may allow him to gain access to the laptop and use it as a stepping stone to other parts of that corporate network.

This may sound unlikely, but an attacker simply has to find the name of a network to which the laptop has previously connected. With so many commonly known hotspot networks in use today, there is a good chance an attacker will be able to find one to which the user previously connected.

This type of attack can be mitigated through the use of a properly configured host firewall running on the laptop. But that firewall could allow certain traffic and/or the system may have software vulnerabilities. It is far better to enable the network auto-switch feature that some laptop suppliers provide. This auto-switch feature ensures that only one network interface can be active at a time. The setting can be enabled in the BIOS of the laptop, but is not always enabled by default.

A wireless security service (WSS) such as rogue access point detection (RAPD) would not necessarily work. They are usually configured to detect rogue access points that announce the same SSID as the corporate wireless network.

In our scenario, the corporate office may not even have a wireless network. RAPD cannot provide a complete solution in this case. One might wonder whether there is a security feature that would allow devices to confirm the identity of a nearby hotspot. The short answer is: no, not really.

Figure 1: Dual homed attack

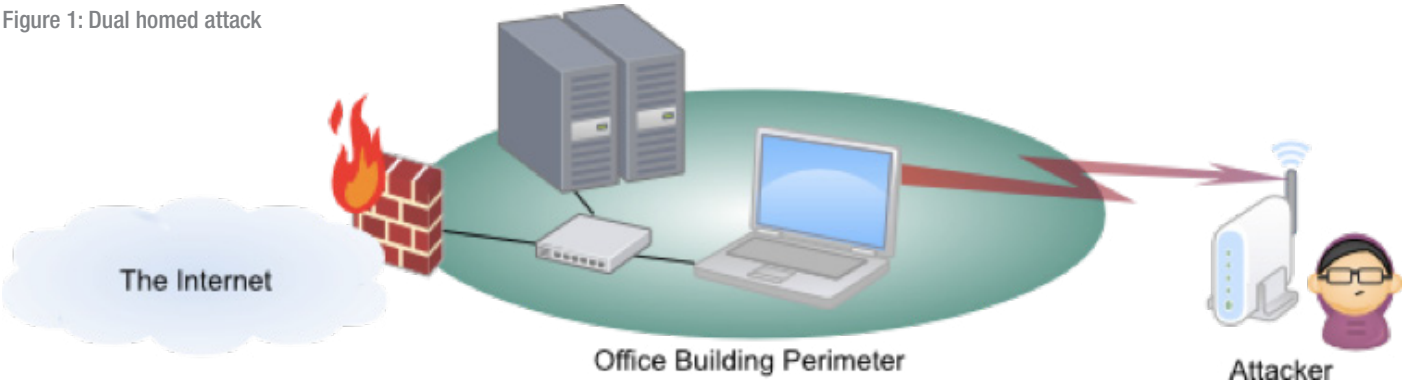

In physical networking the identity of the network is not fully confirmed either, but a physical network has physical boundaries. It is normally quite safe to assume that a network wall socket in an office building is connected to the company's network. Wireless networks do not have these kinds of boundaries, and thus require their identity to be confirmed.

Current solutions

Hotspot authentication: Hotspots currently use a technique called universal access method (UAM) and use a captive portal to authenticate/authorise customers. As mentioned earlier, access to the wireless network is open to any device in the vicinity of the access point. Only when you want to use the internet service provided by the underlying network are you requested to provide some form of credentials or payment.

The captive portal does not provide any credible form of authentication to the user of a wireless network, as an attacker could create an exact copy of the genuine captive portal.

An attacker could, just like the genuine portal, even add an X.509 certificate to his captive portal to increase its apparent trustworthiness. The certificate only authenticates the host name of the website being accessed, which could be any host name, including ones available to an attacker.

The validity of the certificate may be confirmed by the client, but this only confirms the certificate authority (CA) signed the certificate for a particular website; the CA signing the certificate may not be aware how the certificate will be used. It should be clear that even though the two actions may appear to be the same, authenticating the captive portal is not the same as authenticating the wireless network.

By creating a portal of his own, an attacker could, for example, harvest credit card information. This may not be of interest to an attacker at a small establishment in a rural area, but has great potential at a public place, such as an airport, with thousands of travellers passing through each day. Users should be wary of this possible attack and only use keys or codes purchased in advance.

Pre-shared key: Small wireless network set-ups use a so-called pre-shared key (PSK). Users requiring access to the network are given this PSK in advance. This PSK and associated SSID would normally provide a good level of authentication because there is a high probability that the combination is unique, and there is a small chance that the PSK is known to an attacker.

There is, however, only one PSK for all users. Using a PSK would prevent the hotspot provider from uniquely identifying each user. This makes it impossible for the provider to distinguish one user from the other, thereby restricting the provider's ability to charge each user separately. The PSK should be changed regularly to prevent former users from gaining unauthorised access. But that could only take place when there are no active users, as changing the PSK would instantly prevent all users from accessing the network.

Using a PSK would also imply a level of trust between its users. That is clearly an unreasonable assumption for hotspots in a public area such as a railway station. It would allow an attacker to set up a hotspot with the same SSID and PSK, thereby defeating the brittle level of authentication taking place. Using a PSK to authenticate hotspots is therefore not an option.

Extensible Authentication Protocol (EAP): Larger networks use a protocol standard called 802.1X, or "EAP over LAN", to authenticate users. The latter

Small wireless network set-ups use a so-called pre-shared key. This should be changed regularly to prevent former users from gaining unauthorised access

could be read as: LAN, not wireless LAN, because of the following reasons: Two protocols using this standard, called protected EAP (PEAP) and tunnelled transport layer security (TTLS), allow the use of username and password for authentication.

To confirm the identity and authorisation of a connecting client, an access point would allow the client to send its credentials to an external authentication server. Before the client transmits its credentials, it would want to know if the authentication server is a trusted entity. The identity of the authentication server is confirmed with a digital certificate. But that certificate only confirms the identity of the authentication server, not the association with either the SSID or the access point.

To confirm the identity and authorisation of a connecting client, an access point would allow the client to send its credentials to an external authentication server

This leaves an attacker with a way to set up a rogue 802.1X-enabled network (*Figure 2*) with a known SSID of his own. When trying to authenticate such a rogue access point, users may receive a warning about the authentication server certificate presented not being trusted. This is because an 802.1X wireless network is usually pre-configured on a client with a setting that only trusts a specific CA to have signed the certificate presented by the authentication server.

Using a public CA would in this case be worse as it would allow the attacker to get a signed, and thus trusted, certificate for his rogue authentication server and may result in no warning to the user at all. Even if the user were warned of certificate issues, it could easily be ignored.

Once the user accepts the rogue server, running an altered version of the authentication service, would allow an attacker to retrieve the transmitted credentials from users while accepting any password submitted to prevent users from becoming suspicious. These credentials could then be used to access the legitimate wireless network and maybe even systems residing on the underlying network.

These implementations of 802.1X are not entirely robust methods for authenticating hotspots for the reasons mentioned above. It would also require the client configuration of credentials, the names of trusted authentication servers and signing certificate authorities to be distributed in advance. Local distribution of these settings would require hotspot providers to invest and set up nearby vending machines or counters.

It is worth noting that administrators running 802.1X-based networks in their enterprise should deploy their configurations to managed clients with restrictions on possible authentication servers and certificates, and not use public CAs to sign the authentication server certificates to prevent the above attacks from taking place.

Figure 2: Spoofing an 802.1X wireless network

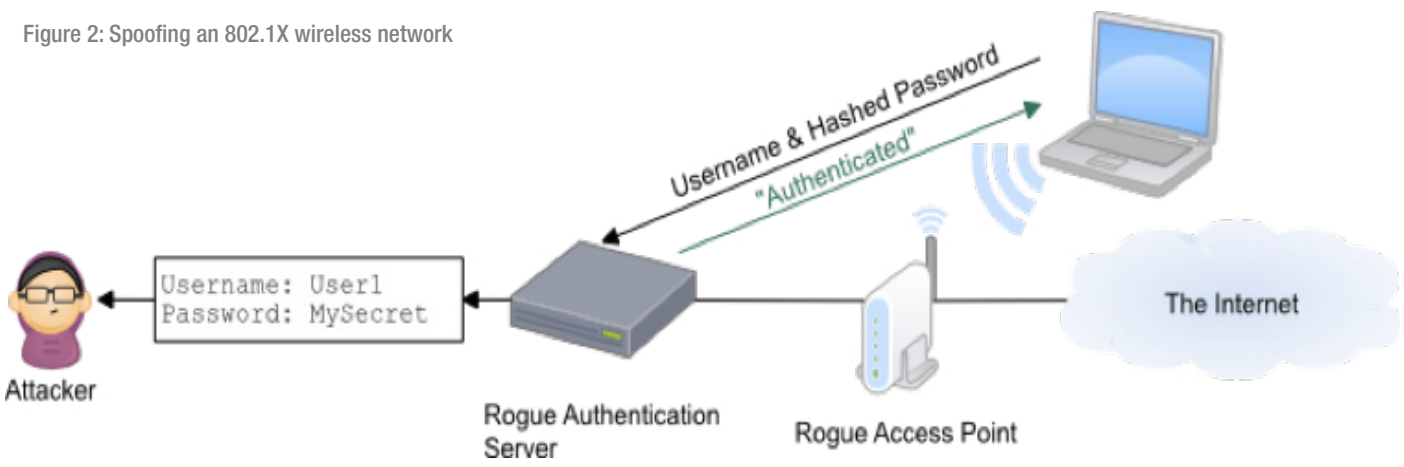

The Future

Open Wireless Network Authentication Protocol (OWNAP): A solution could be sought in the way an internet browser authenticates a website using X.509 certificates. Each access point, or at least each wireless network, would need to have such a certificate in place.

A client trying to establish a connection would perform the required authentication by verifying that the name in the certificate matches both the SSID and the access point hardware address. The client would even be able to establish a secret key to encrypt all further communications.

This would, however, require the SSID to be globally unique. It is something that could be achieved by registering a new top-level domain (TLD), say Wi-Fi, and dedicating this TLD to SSID registration alone. Then the registration of SSIDs would provide CAs signing the SSID certificate requests with a familiar process for verifying the owner of the SSID – CAs are used to doing this when processing certificate renewal requests for the regular, fully qualified domain names used with, for example, SSL/TLS.

However, deploying certificates to all access points would introduce substantial overheads as X.509 certificates expire and need to be replaced or, in some cases, revoked. It is therefore not expected that the proposed OWNAP protocol would ever reach such maturity.

Hotspot 2.0: The Wi-Fi Alliance has launched an alternative to captive portal authentication under the name Passpoint or Hotspot 2.0. It is based on the new IEEE standard 802.11u.

In regular wireless networks, devices use active probing techniques or listen for access point broadcasts to discover nearby wireless networks. The discovery process of 802.11u is still the same, but it allows devices to query a Passpoint-enabled hotspot for more information. This includes queries on the ability to roam via the hotspot, similar to how a mobile phone roams on guest mobile networks.

But, just like mobile networks, service providers must establish mutual roaming agreements that cover credential validation and billing before roaming can take place.

Once implemented Passpoint would allow the mobile service provider to charge customers for both mobile and Wi-Fi roaming. Charging is, of course, only possible when the device is identified and thus requires any connecting device to be authenticated. Passpoint includes the following three techniques to accommodate this requirement:

- EAP-SIM, which uses the GSM authentication triplet for credentials;
- EAP-AKA, which uses the UMTS authentication quintet for credentials. This method is more secure as it includes an encrypted sequence number that is incremented for each authentication event. It allows the device to confirm that the home network has actually produced the challenge;
- Software agent for legacy devices such as laptops using credentials (e.g. a username and password).

The pros and cons of Hotspot 2.0

Passpoint will increase transparency and ease of the use of secure authenticated wireless networks and will not require the user to submit credit card details when he/she tries to gain access to the internet. This is a big security advantage.

The Wi-Fi Alliance has launched an alternative to captive portal authentication, under the name Passpoint or Hotspot 2.0, based on the new IEEE standard 802.11u

There are, however, some things to recognise, as follows:

- The three techniques would need to implement the same certificate and server restrictions mentioned in the EAP section, as the SSID and access point are still not authenticated;
- An attacker could set up an access point with the same SSID as a Passpoint access point and prevent access to the genuine access points. Though user devices would not connect automatically to it, users can still be social engineered into connecting and providing sensitive information: something along the lines of "Our system is down. Your credit card will not be charged but please provide it as a method of identification";
- Though interesting for mobile telecom companies, others providers may not want to upgrade to Passpoint as it will force them to share the revenue generated. Mobile telecom providers may not allow the necessary second set of credentials to be stored on the SIM;
- Even if a user has a Passpoint-enabled device, he or she may prefer to use another nearby wireless network that is free of charge instead of paying roaming charges. That network may use captive portal and not provide proper protection;
- Tablets are being sold without a slot for a SIM card, possibly preventing these systems from using Passpoint access points. Their access to the wireless interface may be restricted to built-in apps;
- Legacy devices will still require a set of credentials in advance, possibly limiting their "plug and surf" capability. The intention is to leave captive portals in place for legacy systems

Network entry and exit points are easy points of attack and should therefore be properly authenticated

Conclusion

People generally have a good sense of whether or not they can trust other people. These senses are part of our social skills that have been developing since the beginning of mankind. We rely on them to protect us from harm. We cannot do the same with technology as it requires a level of technical understanding available to very few users.

With most preventative and detective security measures focusing on client authentication and blocking unauthorised access, it is difficult for devices to confirm the identity of a hotspot and establish a secure channel of communication. Other higher level protocols, such as HTTPS, could prevent eavesdropping, but those are not implemented all the time either. If implemented, these higher level protocols would only provide a single layer of security, where multiple layers would be better because a security warning may be considered a glitch and ignored by users.

This article has shown that network entry and exit points are easy points of attack and should therefore be properly authenticated. With Passpoint, a first step is made towards enhancing the security measures on the user's side. ■

About the authors

Daan Stakenburg is an Information Security Architect at a global technology company listed on the London Stock Exchange. With almost 15 years of IT experience in a variety of multinational organisations, he acts as an internal consultant advising departments and project teams on a multitude of network and information security matters

Jason Crampton has a BSc in Mathematics (University of Manchester) and a PhD in Computer Science (University of London). He has worked in the Information Security Group at Royal Holloway, University of London since 2002, where he is a Professor of Information Security. His research focuses on access control in multi-user computer systems