LANCOM™ Techpaper
WPA and 802.11i

Introduction

The WLAN standards WPA and 802.11i are currently redeeming the reputation of WLAN security, an issue which has recently been under attack. The processes incorporated into the original standard proved insuficient in practice. This lack led on the one hand to a series of proprietary extensions of the standard, like „CKIP" from Cisco, or „KeyGuard" from Symbol Technologies, and on the other hand to solutions which offered the required security on higher protocol layers with tools like PPTP or IPSec. All these processes are quite functional, but they introduce limitations, for instance those relative to interoperability or data transmission rates.

In the recently released standard 802.11i, the IEEE Committee has redefined the topic „WLAN and security" from the ground up. The result is a set of standardised methods that enable the construction of secure and manufacturer-independent WLANs in line with current standards.

On the way from the original WEP of the 802.11 standard to 802.11i, a whole series of concepts have arisen that have tended to increase confusion and insecurity among the users. This document should help to explain the concepts and the processes used, in chronological order of their development.

Some basic concepts

Even though one constantly hears the blanket term ‚Security' when talking about computer networks, it is still important for the coming exposition to differentiate a little more closely between the requirements it actually entails. The first point in security is access security:

- Here, a protective mechanism is involved which allows access to the network only to authorised users.
- On the other hand, however, it must also be ensured that the client is connected to the precise desired access point, and not with some other access point with the same name which has been smuggled in by some nefarious third party. Such an authentication can be provided, for example, using certificates or passwords.
- Once access is provided, one would like to ensure that data packets reach the receiver without any falsification, that is, that no-one can change the packets or insert other data into the communication path. The manipulation of data packets themselves cannot be prevented, but changed packets can indeed be identified using suitable checksum processes, and then discarded.

Quite separate from access security is confidentiality, that is, unauthorised third parties must not be able to read the data traffic. To this end, the data are encrypted. This sort of encryption process is exemplified by DES, AES, RC4, or Blowfish. Along with encryption, of course, there must also be a corresponding decryption on the receiving end, generally with the same key (also-called symmetric encryption process). The problem naturally then arises, how the sender can give the key to the receiver for the first time—a simple transmission could very easily be read by a third party, who could then easily decrypt the data traffic.

In the simplest case, this problem is left to the user, that is, one simply assumes that the user can make the key known at both ends of the connection. In this case, one speaks of pre-shared keys, or ‚PSK'.

More sophisticated processes come into play when the use of pre-shared keys is impractical, for instance in an HTTP connection built over SSL—in this case, the user can't retrieve a key from a remote web server quite so easily. In this case, socalled assymetric encryption

LANCOM™ Techpaper
WPA and 802.11i

methods such as RSA can be used, that is, to decrypt the data, a different key is used than the one used to encrypt it. Such methods are, however, much slower than symmetric encryption methods, which leads to a two-phase solution: one side possesses an asymmetric key pair and transmits the encryption key to the other side, generally as a part of a certificate. The other side chooses an arbitrary symmetric key, and encrypts this symmetric key with the asymmetric key previously received. The owner of the asymmetric key pair can now decrypt it, but a potential eavesdropper cannot—the aim of the secure key exchange is achieved.

In the following sections, we will see these methods again, sometimes in modified form.

WEP

WEP is an abbreviation for Wired Equivalent Privacy. The primary goal of WEP is the confidentiality of data. In contrast to signals which are transmitted over cables, radio waves spread out in all directions—even into the street in front of the house and other places where they really aren't desired. The problem of undesired interception is particularly obvious in wireless data transmission, even though it can also arise in larger installations with wired networks—however, access to cables is far more easily restricted than is the case with radio waves.

During the development of the WLAN security standard, the IEEE Committee did not intend to develop a „perfect" encryption method. Such high security encryption methods are, for instance, required and also used in electronic banking—in this case, however, the applications themselves use high-quality encryption methods, and it would be unnecessary to repeat this effort at the radio transmission level. With the new security standards, only those applications which normally work without encryption in wired LANs should be provided with sufficient security against eavesdropping by unauthorised third parties.

Figure 1 shows the process of WEP encryption—decryption runs in precisely the opposite manner. WEP is therefore a symmetrical encryption method. WEP uses RC4 algorithm as its basic encryption technology, a process already well-known in other areas and considered highly secure. RC4 uses a key between 8 and 2048 bits in length, which is used to generate a pseudo-random series of bytes using a predetermined process. The data packet is then XOR'd byte by byte with this byte stream. The receiver simply repeats this process with the same key and thus with the same sequence, in order to retrieve the original data packet—a double application of the XOR operation with the same values cancels out. The advantage of RC4 is that the operations

- generation of the byte sequence from the key
- XOR operation on the data stream

on the sending and receiving sides are identical—so the hardware need only be built into the WLAN card once, and then can be used for both transmission and receiving. Since the data in the WLAN are transmitted half-duplex only, simultaneous transmission and receiving will never occur. However, RC4 has one serious disadvantage:

one may only use a particular RC4 key once for a single packet! If the same RC4 key is used for two different data packets, then a potential eavesdropper is able to take the two packets and XOR them together. This operation doesn't result in clear text, but the pseudo random sequence, and thus the encryption, cancels out, and one has the XOR combination of wo clear text packets. If one already knows the contents of one of the two packets, then the clear text of the other is easily determined.

LANCOM™ Techpaper
WPA and 802.11i

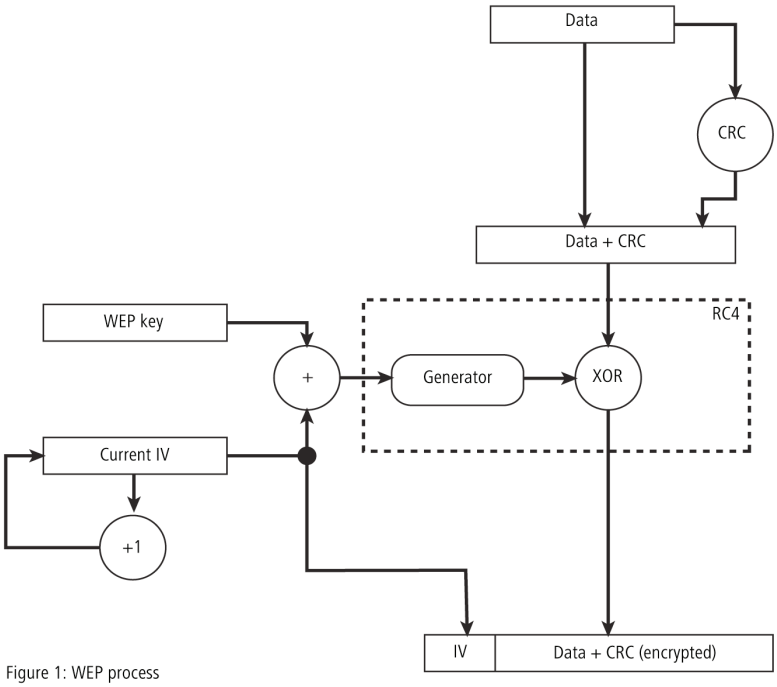

Figure 1: WEP process

Thus WEP does not directly use the key entered by the user for the RC4 algorithm, but rather combines it with a so-called Initial Vector (IV) to arrive at the actual RC4 key. This IV is automatically changed from packet to packet by the sender, generally by simple incrementation, and is transmitted along with the encrypted packet. The receiver uses the IV included in the packet in order to recon struct the RC4 key actually used for the packet.

WEP also calculates a CRC checksum for the unencrypted packet and appends it to the packet before it is RC4-encrypted. The receiver can check this CRC checksum after decryption and determine whether the decryption was faulty—for example, due to an incorrect WEP key. In this way, WEP also happens to offer a certain degree of access security, since an intruder without knowledge of the WEP key can only generate „defective" packets, which will automatically be filtered out by the WLAN card. This additional IV explains some of the confusion one sees about the key length in WEP—since larger key lengths sound more secure, the 24 bits of the IV sound nice when added to the actual key length, although the user can of course only configure the left-over portion.

The IEEE standard originally foresaw a relatively short key length of 40 bits, which was probably oriented towards the then-existing US export restrictions on strong cryptography— this variant is usually called WEP64 in brochures. Most WLAN cards today support a variant in which the user an configure a 104-bit key, which results in a 128 bit long RC4 key—correspondingly, this is often called WEP128. More seldom are key lengths of 128 bits (WEP152) or 232 bits (WEP 256).

As explained above, RC4 can in principle work with key lengths up to 2048 bits, which would correspond to WEP keys of up to 2024 bits. In the practice, key lengths reach a simple limit at which the user can manage to enter the columns of digits without making a mistake. Since WEP is a pure PSK method, the keys must be entered identically on both sides of the connection.

LANCOM Systems

LANCOM™ Techpaper
WPA and 802.11i

The IEEE standard provides no mechanism to distribute WEP keys in a WLAN automatically. Some manufacturers have, for instance, attempted to simplify entry for users by requiring entry not of the WEP key itself, but rather a passphrase (a sort of overly long password) from which the key can be calculated. However, this procedure varies from manufacturer to manufacturer so that the same passphrase or different manufacturers might lead to different WEP keys—besides, users have a tendency to choose passwords which are relatively easy to guess, so that the resulting keys are usually weaker than 40 or 104 bits (the current IEEE standards, for instance, assume that a typical password has a strength of about 2.5 bits per character.)

The IEEE standard specifies that up to four different WEP keys can exist in one WLAN. The sender encodes the number of the WEP key used in the encrypted packet along with the IV, so that the receiver can use the appropriate key. The idea behind this was that old keys in a WLAN could gradually be exchanged for new keys, in that stations which had not yet received the new key could still use an old key during a transition period.

Based on WEP, the 802.11 standard also defines a Challenge-Response procedure for authentication of clients. The access point sends a clear-text packet which contains a 128-byte long challenge, which the client encrypts and sends back with WEP. If the access point can successfully decrypt this answer (that is, the CRC is correct) and the result is the originally transmitted challenge, it can assume that the client has a correct WEP key and thus is authorised for access.

Unfortunately, this process provides a potential attacker with 128 bytes of clear text and the corresponding encrypted text, which offer scope for crypto analysis. Furthermore, many clients don't implement this variant, so that this process, called hared Key, is seldom used—instead, processes started after the WLAN registration are used for authentication, such as 802.1x (see below).

While the WEP process theoretically sounds good up to now, in practice there are unfortunately serious flaws which significantly reduce the advantages—regardless of the WEP key length used. These weaknesses really should have been found by closer analysis at the time when WEP was being defined. Unfortunately, no cryptology experts participated in the WEP definition process, so these flaws only became obvious once the WEP process was massively implemented thanks to the market success of 802-11b WLAN cards (earlier 2MB designs often included no encryption at all—WEP is an optional function in the 802.11 standard).

The chief weakness of WEP is the IV length, which is far too short. As already mentioned, the reuse of a key in RC4 is a serious security loophole—but it occurs in WEP at least every 16 million packets, when the IV counter overflows from 0xffffff to zero. An 11MB WLAN can achieve a net data rate of around 5MB/sec; with a maximum packet length of 1500 bytes, that comes to about 400 packets per second at full throttle. After about 11 hours, the IV counter would theoretically overflow, and an eavesdropper receives the information needed to ‚crack' the WEP key. In practice, the attacker will actually receive this information much sooner.

Mathematical analyses of RC4 have shown that for certain values of the RC4 key, conclusions may be drawn about the first values of the pseudorandom sequence it generates—thus about the bytes with which the beginning of the packet are encrypted. This property of RC4 can be relatively easily avoided, for instance by discarding the first bytes of the pseudo random byte sequence and only using the „later" bytes for encryption, and this is often done nowadays when RC4 is used.

LANCOM Systems

LANCOM™ Techpaper
WPA and 802.11i

But when this discovery was first made WEP in its described form was already part of the IEEE standard and indelibly incorporated into the hardware of the widely distributed WLAN cards. Very unfortunately, these „weak" values of RC4 keys can be recognised by particular values in the first bytes of the RC4 key, and in WEP that happens in the IV in each packet—which is transmitted in clear text. Once this connection was discovered, specialised sniffer tools quickly appeared on the Internet, which watched for packets with these ‚weak IVs', and thus only had to process a fraction of the total traffic. Depending on the amount of data being transferred in a WLAN, such tools can crack the encryption in a fraction of the time mentioned above. With longer WEP keys (such as 104 instead of 40 bits) this may take a little longer, but the time required for cracking grows at best linearly with the key length, not exponentially, as is usually the case.

Unfortunately the CRC checksums contained in the packets also haven't lived up to expectations. Ways were found to change encrypted packets under certain conditions even without knowledge of the WEP key in such a way that the CRC is still valid after decryption on the receiving end. So WEP therefore cannot guarantee that a packet hasn't been changed on the way from sender to receiver.

These weaknesses unfortunately degraded WEP to an encryption scheme which at best could be used to protect a home network against ‚accidental eavesdroppers.' These discoveries gave rise to much controversy, gave WLAN the reputation of being unsafe technology, and forced manufacturers to action. WLAN is, however, a standardised technology, and better standards don't come into being from one day to the next—which is why there were a few intermediate steps to a secure solution, which at least blunted the worst of WEP's design flaws.

WEPplus

As explained in the previous section, the use of ‚weak' IV values was the problem which weakened the WEP process most. Only a few weeks after the publication, tools like ‚WEPCrack' and ‚AirSnort' appeared on the Internet, which could automatically crack an arbitrary WLAN connection within a few hours. With this, WEP was essentially worthless.

A first ‚quick shot' to secure WLANs against this kind of program was the simple notion that the weak IV values are known, and that they could simply be skipped during encryption—since the IV used is after all transmitted in the packet, this procedure would be completely compatible with WLAN cards which didn't understand this extension, dubbed WEPplus. A true improvement in security would naturally only result once all partners in the WLAN were using this method. In a network equipped with WEPplus, a potential attacker again has the chore of listening to the entire data traffic, waiting for IV repetitions—simply waiting for the few packets with weak IVs is no longer an option. This raised the bar for an attacker again, particularly if one didn't simply set the IV counter to zero when initialising a WLAN card, but rather initialised with a random value:

the IV counter at an access point only starts to count when the firststation logs in and starts transmitting data. If the access point and station each initialised their IV counters to zero, packets with identical IV values occur immediately after the connection is made. By initialisation to a random value, the collision can at least be delayed by an average of 223 packets, that is, half the space of possible IVs — with more than one station in a WLAN, this value is naturally reduced.

LANCOM Systems

LANCOM™ Techpaper
WPA and 802.11i

WEPplus is thus technically only a slight improvement—but it did serve to calm the user base enough to make WEP acceptable again, at least for home use (as long as a new key was configured often enough.) For use in a professional environment, of course, that didn't suffice.

EAP and 802.1x

Obviously, an ‚add-on' like WEPplus can't eliminate the basic problem of too-short IVs, without changing the format of packets on the WLAN, thus rendering all existing WLAN cards incompatible. There is, however, a possibility of solving several of our problems with one central change: no longer use the formerly fixed WEP key, but to negotiate them dynamically instead. As the process to be used for this purpose, the Extensible Authentication Protocol has emerged. As the name suggests, the original purpose of EAP is authentication, that is, the regulated access to a WLAN—the possibility of installing a valid WEP key for the next session is more or less a byproduct. Figure 2 shows the basic process of a session secured by EAP. In the first phase, the client registers with the access point as usual, and enters the state in which it can now send and receive over the access point in normal WEP or WEPplus—but not with EAP, because in this state the client still doesn't have a key to secure its data traffic from eavesdropping. Instead, the client is in an ‚intermediate state' from the point of view of the access point, in which only particular packets from the client are forwarded, and these are only directed to an authentication server. These packets implement EAP/ 802.1x as already mentioned, which can easily be distinguished from other protocols due to its Ethernet type 0x888e. The access point packages these packets in RADIUS queries and sends them on to the authentication server. The access point converts the replies coming from the RADIUS server back into EAP packets, and sends them back to the client.
The access point is thus a sort of middle man between client and server. it doesn't have to check the contents of these packets, it just has to check that no other data traffic to or from the client can occur.

This process has two advantages:

- The implementation effort in the access point is low. While the client and the server are usually PCs with high levels of resources, access points are devices which are limited both in memory and in computing power.
- New processes for authentication require no firmware upgrade on the access point.

Over this tunnel through the access point, the client and server authenticate one another, that is, the server checks the client's access privilege to the network, and the client checks that it is talking to the right network. „Wild" access points set up by hackers can be recognised in this way. A whole series of authentication processes exist which can be used in this tunnel. A current process (and one supported by Windows XP) is for instance TLS, in which server and client exchange certificates; another is TTLS, in which only the server supplies a certificate—the client is uthenticated using only a username and password. After the authentication phase, a secure tunnel even without WEP encryption has been set up, in which the access point is connected in the next step. For this, the RADIUS server sends the socalled ‚Master Secret', a session key calculated during the negotiation, to the access point. Although the LAN behind the access point in this scenario can be viewed as secure, this transmission is also encrypted. With this session key, the access point now

LANCOM™ Techpaper
WPA and 802.11i

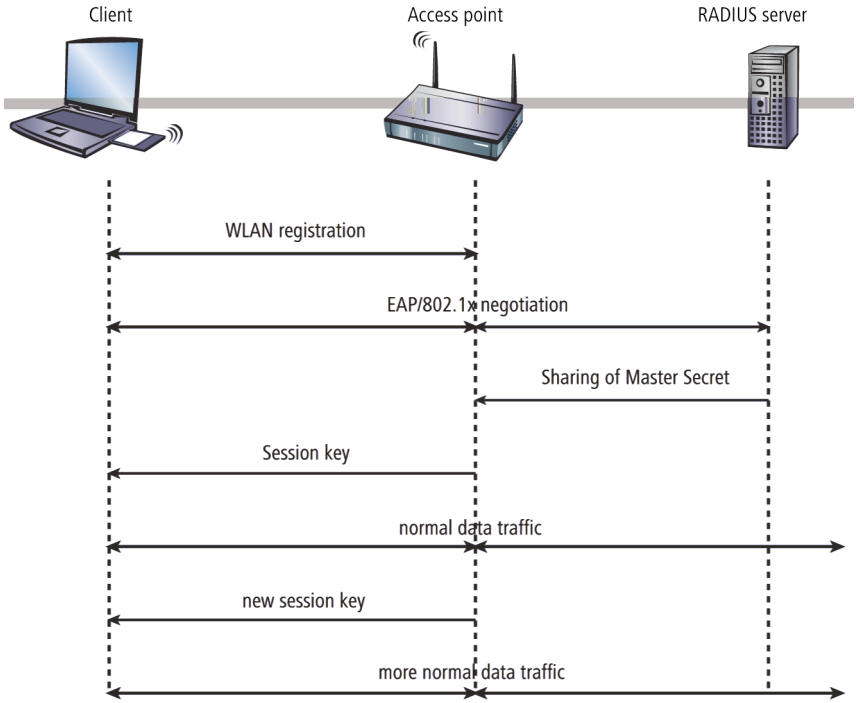

Figure 2: Schematic process of a WLAN session with EAP/802.1x

takes over the tunnel and can use it to provide the actual WEP key to the client.

Depending on the capabilities of the access point hardware, this can be a true session key (that is, a WEP key which will only be used for data packets between the access point and precisely this client), or a so-called group key, which the access point will use for communication with multiple clients. Classical WEP hardware can usually handle only group keys, these being the four mentioned in the chapter on WEP.

The particular advantage of this procedure is that the access point can regularly change the WEP key over the EAP tunnel, that is, it can perform a so-called rekeying. In this way, WEP keys can be replaced by new ones long before they run the risk of being cracked due to IV collisions. A common ‚use time' for such WEP keys might be 5 minutes. Further advantages of this procedure include its simple implementation in the access point, with little extension to existing hardware.

The disadvantage of the procedure is its complexity. The maintenance of the central RADIUS server and the certificates stored there is generally only possible in large installations with a separate IT department—it is less suitable for use in the home or in smaller companies. Furthermore, a minimum set of procedures has not been established which a client or a server must support. Thus scenarios are quite thinkable in which a client and a server cannot establish an EAP tunnel, because the sets of procedures they support don't match. These practical hurdles have thus limited EAP/802.1x to professional use so far—the home user must simply make do with WEPplus, or address security problems on the applications level.

LANCOM™ Techpaper
WPA and 802.11i

TKIP and WPA

As should be clear from the last section, the WEP algorithm is flawed and insecure in principle; the measures taken so far were largely either ‚quick fixes' with limited improvement, or so complicated that they were basically impractical for home use or smaller installations. The IEEE started a Task Group after the discovery of the problems with WEP which addressed the definition of better security mechanisms, and which should eventually result in the IEEE 802.11i standard. The composition and ratification of such a standard, however, generally takes several years. In the meantime, market pressure had grown to the point where the industry could no longer wait for the finalisation of 802.11i. Under the auspices of Microsoft, therefore, the WiFi Alliance defined the Wifi Protected Access (WPA) ‚standard'. The WiFi Alliance is an association of WLAN manufacturers which promotes the manufacturer-independent function of WLAN products and, for example, awards the Wifi logo.

In the definition of standards, and 802.11i is no exception, the basic mechanisms are generally known fairly quickly. The publication of the standard mostly takes such a long time because of the fine details. These details are often important only for rare applications. WPA thus took the pragmatic route of extracting the parts of the 802.11i proposal which were already clear and important for the market, and packing them into their own standard. These details include:

- TKIP and Michael as replacement for WEP
- A standardised handshake procedure between client and access point for determination/ transmission of the session key.
- A simplified procedure for deriving the Master Secret mentioned in the last section, which can be performed without a RADIUS server.
- Negotiation of encryption procedure between access point and client.

TKIP

TKIP stands for Temporal Key Integrity Protocol. As the name suggests, it involves an intermediate solution for temporary use until a truly strong encryption procedure is introduced, but which deals with the problems of WEP, never the less. One design requirement was therefore that the new encryption procedure should be implementable on existing WEP/RC4 hardware with a reasonable effort. When TKIP was defined, it was already foreseeable that it would be used well into the era of 54/108Mbit LANs, and a purely softwarebased encryption would be associated with too high a speed penalty on most systems. In the ‚block diagram' of TKIP (Figure 3), therefore, there are many components of WEP to be seen, which generally exist in hardware in WEP cards and thus can effectively be used for TKIP. As components already familiar from WEP, one recognises the RC4 engine used for the actual encryption and decryption, as well as the CRC module. As a new component (green), however, besides the CRC, the unencrypted package also has a so-called Michael-MIC attached. This is a hash algorithm developed especially for WLAN, which was designed so that it can be computed on older WLAN hardware with reasonable overhead. Since in contrast to the CRC a second key (the Michael key) must be agreed in this hash, it can neither be calculated nor used to falsify a data packet without detection by the receiver. This is only remains true if an attacker doesn't break the Michael hash with brute force techniques.

LANCOM™ Techpaper
WPA and 802.11i

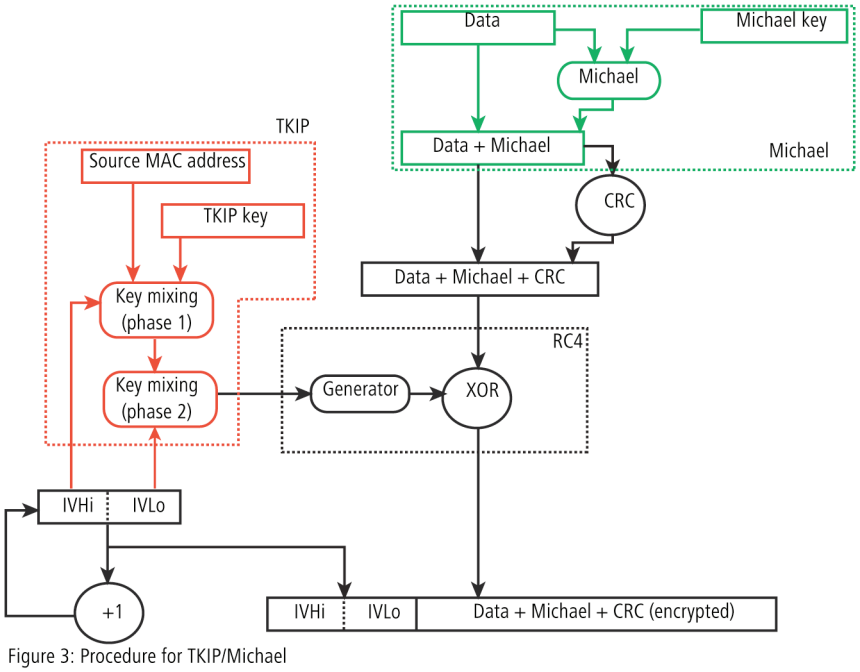

Figure 3: Procedure for TKIP/Michael

Due to the requirement of high run-time efficiency, Michael makes a few compromises: although a 64-bit key is used, the effective strength of Michael is only about 40 bits. This was still seen as sufficient, since a potential attacker would have to break the TKIP components in the first place in order to generate data packets which would get past the CRC check of the WEP/RC4 components.

TKIP (red) takes care of the calculation of the actual key for the RC4 engine. In contrast to WEP, the actual key and the IV contained in the packet are never used directly as the RC4 key, but rather it runs through two so-called key mixing phases along with the IV—so an attacker can draw no direct conclusions about the RC4 key from the IV contained in clear text, which solves the problem of ‚weak' IVs in WEP (the key mixing itself is designed so that weak RC4 keys can never occur). Furthermore, the internally incremented IV transmitted in clear text in the packet is 48 bits long instead of 24 - so a sender can now transmit some 280 trillion packets before the 128-bit TKIP key must be changed. Even in a modern WLAN with a net 108 Mbps, which achieves a net rate of around 50 Mbps, using the same assump tions made above for WEP, this would correspond to about 2000 years. It must still be noted that the IV is split into two parts for reasons of optimisation: a 16-bit low part and a 32-bit high part. The background for this is that the key mixing proceeds in two phases, as shown in the illustration:

- For the first (computationally intensive) phase, only the upper part is needed, so it only needs to be performed once for every 65,536 packets.
- The second, relatively simple phase of the key mixing uses the result of the first phase along with the low part of the IV (which changes with each packet) in order to create the actual RC4 key.

LANCOM Systems

LANCOM™ Techpaper
WPA and 802.11i

In contrast to WEP, it is additionally determined in TKIP that the IVs to be used from packet to packet must increase in a strictly monotone manner, so the receiver only has to perform phase 1 for every 65,536 received packets. The decryption part of TKIP checks this sequentiality and discards packets which contain an already-used IV, which prevents replay attacks. As a further detail, TKIP also mixes the MAC address of the sender into the first phase. This ensures that the use of identical IVs by different senders cannot lead to identical RC4 keys and thus again to attack possibilities. As mentioned above, the Michael hash does not represent a particularly tough cryptographic hurdle: if the attacker can break the TKIP key or get encrypted packets past the CRC check via modifications similar to those for WEP, then not many barriers remain. For this reason, WPA defines countermeasures if a WLAN card detects more than two Michael errors per minute: both the client and the access point break data transfer off for one minute, afterwards renegotiating TKIP and Michael keys.

The key handshake

In the discussion of 802.1x it was already noted that EAP/802.1x provides a possibility to inform the client at the outset of a session of the key valid for it. WPA now places that on a standardised basis, and considers the session-key option offered by modern access points that, in addition to the four ‚global' keys, assigns each registered client with a session key that is used exclusively with data packets to or from that client. If you take another look at the procedure shown in Figure 2, the newly defined key handshake replaces the phase in which the access point transmits the WEP key to the client after receiving the Master Secret from the RADIUS server. The key handshake breaks down into two phases: first the pairwise key handshake, then the group key handshake (Figure 4). As you can see, the handshake consists of pairs of packets which each consist in turn of a ‚query' of the access point and a ‚confirmation' of the client. The first pair serves mostly for the client and access point to exchange the specific random values (socalled nonces) to be used for this negotiation. The Master Secret already known to both sides is now mixed with these nonces and after a predetermined hash procedure, further keys are generated, on the one hand to take care of securing further exchanges, and on the other to be used as a pairwise key for this station. Since the Master Secret isn't used directly, it can be reused later for any necessary renegotiations, since it can then be mixed with new random value and thus will deliver different eys. In the second pair, the access point instructs the client to install the calculated TKIP session key, and as soon as the client confirms this, the access point does the same. This concludes the pairwise handshake, and as a result it is now possible to exchange data between client and access point via TKIP. The client still can't be ‚approved', however, because the access point must still transmit a further key—the group key, which it uses to transmit broadcast and multicast packets simultaneously to all stations. This must be determined unilaterally by the access point, and it is simply transmitted to the station, which confirms receipt. Since at this point a pairwise key is already installed on both sides, both of these packets are already encrypted. After a successful group key handshake, the access point can finally release the client for normal data transfer. The access point is free to perform a rekeying again during the session using the same type of packets. In principle, the client may also request rekeying from the access point. WPA also takes the case of older WLAN hardware into account, in which the access point does not support pairwise keys, but only group keys.

LANCOM™ Techpaper
WPA and 802.11i

The first phase of the handshake in this case proceeds exactly as before, but doesn't result in the installation of a pairwise key—the group key handshake simply proceeds in clear text, but an encryption in the EAP packets themselves prevents an attacker from simply reading the keys.

WPA with passphrase

The handshake described in the previous section runs strictly under WPA, i.e. the user will never have to define any TKIP or Michael keys. In environments in which no RADIUS server is available to provide master secrets (for instance in smaller companies or home networks), WPA therefore provides the PSK method besides authentication using a RADIUS server; here, the user must enter a passphrase of 8 to 32 characters on the access point and on all stations, from which the master secret is calculated along with the SSID used using a hash procedure. The master secret is therefore constant in such a PSK network; the nonces ensure, however, that different TKIP keys still result.

In a PSK network—similar to classical WEP—both access security and confidentiality depend on the passphrase not being divulged to unauthorised people. As long as this is the case, WPA-PSK provides enormously improved security against break-ins and eavesdropping over any WEP variant. For larger installations in which such a passphrase would have to be made known to too large a user community for it to be kept secret, EAP/802.11i is used in combination with the key handshake described here.

LANCOM Systems managed to close this potential loophole by inventing the LEPS feature (LANCOM Enhanced Passphrase Security).

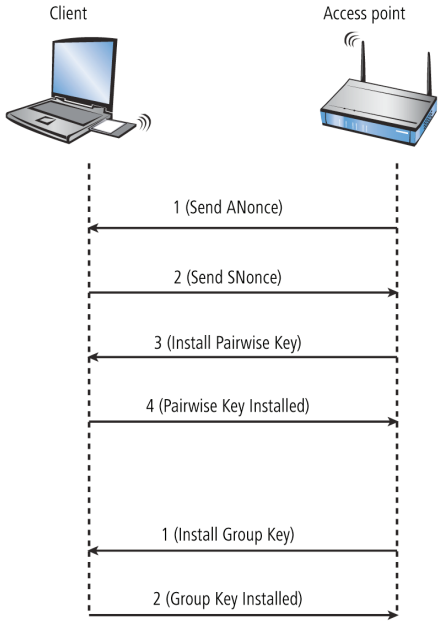

Figure 4: Key handshake in WPA

Without the need to set up a complicated and extensive server infrastructure for 802.1x, every clients MAC address is assotiated in the ACL (Access Contzrol List) to an individual passphrase. It is not necessary anymore to use company-wide one similar passphrase which eliminates the risk of giving the passphrase to unauthorised persons.

Negotiation of the encryption method

The original WEP definition only specified a fixed key length, so that only a single bit was required in the registration packets from the station and access point to show whether encryption should be used or not. This became insufficient the moment WEP was used with key lengths other than 40 bits—the user just had to take care that not only the same value but that the same length was defined as well. WPA provides a mechanism with which client and access point can

LANCOM Systems

LANCOM™ Techpaper
WPA and 802.11i

agree on the encryption and authentication procedures to be used. For this purpose, a new info element was defined which can contain the following:

- The encryption method to be used for broadcasts in this network (also the type of group key). Each client wanting to register in a WPA-WLAN must support this procedure. Here, besides TKIP, WEP is also still allowed, in order to support mixed WEP/ WPA networks—in a pure WPA network, TKIP will be selected.
- A list of encryption methods which the access point provides for the pairwise key—here, WEP is explicitly disallowed.
- A list of authentication methods a client may use to show itself to the WLAN as authorised for access—possible methods are currently EAP/ 802.1x or PSK.

The access point broadcasts such an element with its beacons, so that clients know whether this network is suitable for them or not. When registering at the access point, the client sends another such packet, in which it gives the desired type of pairwise key as well as its authentication scheme. The access point then starts either the EAP/802.1x negotiation, or starts directly with the key handshake. Since neither beacons nor registration packets are cryptographically secured, it is possible that a third party might interfere in this exchange and force the client and/or the access point down onto a weaker method than the one actually desired. Both the access point and the client are therefore required to exchange these info elements again during the key handshake, and if the element received doesn't match the one from the registration, they immediately break the connection. As mentioned, the original WPA standard specifies only TKIP/Michael as an improved ncryption method. With the further developmentof the 802.11i standard, the AES/ CCM method described below was added.

In a WPA network it is now possible for some clients to communicate with the access point using TKIP, while other clients use AES.

AES and 802.11i

In mid-2004, the long awaited 802.11i standard was approved by the IEEE, which should put the entire security concept of the WLAN on a new basis—which is to be expected, since errors as serious as those encountered during the introduction of WEP are unlikely to occur with 802.11i. As mentioned in the last section, WPA has already implemented a whole series of concepts from 802.11i—so in this section we will only describe the components which are new compared to WPA.

AES

The most obvious extension is the introduction of a new encryption process, namely AES-CCM. As the name already hints, this encryption scheme is based on DES's successor AES, in contrast to WEP and TKIP, which are both based on RC4. Since only the newest generation of WLAN chips contain AES hardware, 802.11i continues to define TKIP, but with the opposite prerequisites: any 802.11i-compliant hardware must support AES, while TKIP is optional—in WPA that was exactly the other way around. Due to the widespread adoption of non-AES-compatible hardware, however, it is to be expected that every AES-capable WLAN card will still support WEP and TKIP. WLAN devices will, however, probably provide configuration options which prevent use of TKIP—many agencies in the USA consider TKIP insufficiently secure, which due to the comparatively weak Michael hash is fairly well justified. The suffix CCM denotes the way in which AES is used in WLAN packets. The process is actually quite complicated,

LANCOM™ Techpaper
WPA and 802.11i

for which reason CCM is only sensibly implemented in hardware—software-based implementations are possible, but would result in significant speed penalties due to the processors commonly used in access points In contrast to TKIP, AES only requires a 128-bit key, with which both the encryption and protection against undetected changes to packets is achieved. Furthermore, CCM is fully symmetric, i.e. the same key is used in both communications directions—a compliant TKIP implementation, on the other hand, requires the use of different Michael keys in the send and receive directions, so that CCM is significantly simpler in use than TKIP. Occasionally one finds other AES variants in older publications or drafts of the 802.11i standard, namely AES-OCB or WRAP. In these variants, AES was used in a different form, which was dropped in favor of CCM in the final standard. WRAP is nowadays meaningless.

Similar to TKIP, CCM uses a 48-bit Initial Vector in each packet—an IV repetition is impossible in practice. As in TKIP, the receiver notes the last IV used and discards packets with an IV which is equal to or less than the comparison value.

Pre-authentication and PMK caching

As mentioned earlier, the delay in publishing standards is usually due to the details. In the case of 802.11i, there were two details which should particularly help with the use of WLAN for speech connection (VoIP) in enterprise networks. Especially in connection with WLAN-based wireless telephony, quick roaming (switching from one access point to another without lengthy interruptions) is of special significance. In telephone conversations, interruptions of 100 milliseconds are irritating, but the full authentication process over 802.11x, including the subsequent key negotiation with the access point, could take significantly longer. For this reason, the so-called PMK caching was introduced as a first measure. The PMK, of course, serves as the basis for key negotiation in an 802.1x authentication for both client and access point. In VoIP environments it is possible that a user moves back and forth among a relatively small number of access points. Thus it may happen that a client switches back to an access point in which it was already registered earlier. In this case it wouldn't be sensible to repeat the entire 802.1x authentication again. For this reason, the access point can provide the PMK with a code, the so-called PMKID, which it transmits to the client. Upon a new registration, the client uses the PMKID to ask whether this PMK is still stored. If yes, the 802.1x phase can be skipped and only the exchange of six short packets is required before the connection is restored. This optimisation is unnecessary if the PMK in a WLAN is calculated from a passphrase as this applies everywhere and is known

A second measure allows for some acceleration even in the case of first-time registration, but it requires a little care on the part of the client. The client must already detect a degrading connection to the access point during operation and select a new access point while it is still in communication with the old access point. In this case it has the opportunity to perform the 802,1x negotiation with the new access point over the old one, which again reduces the „dead time" by the time required for the 802.1x negotiation.

LANCOM™ Techpaper
WPA and 802.11i

Summary

After the security loopholes in WEP encryption became public knowledge, the presentation of short-term solutions such as WEPplus and the intermediate steps like WPA, the IEEE committee has now presented the new WLAN security standard 802.11i. The TKIP procedure used by WPA is based on the older RC4 algorithm, the foundation of WEP. AES is the first important and conclusive step towards a truly secure encryption system. 802.11i/AES have confined the practical and theoretical security loopholes in previous methods to history. The AES procedure provides security on a level that satisfies the Federal Information Standards (FIPS) 140-2 specifications that are required by many public authorities. LANCOM Systems equips its 54Mbps products with the Atheros chip set featuring a hardware AES accelerator. This guarantees the highest possible level of encryption without performance loss. The user-friendly preshared key procedure (entry of a passphrase of 8-63 characters in length) makes 802.11i quick and easy for anybody to set up. Professional infrastructures with a larger number of users can make use of 802.1x and RADIUS servers. In combination with further options such as Multi-SSID and VLAN tagging, it is possible to provide highly secure networks for multiple user groups and with different levels of security.

- VLAN tagging is available as of LCOS version 3.32.
- Multi-SSID is available as of LCOS 3.42.
- LANCOM Systems provides the PSK procedure with the LCOS version 3.50.
- 802.1x is supported since LCOS version 3.52.
- Easy authentication by the individually assigned passphrase to each MAC address (LEPS) with LCOS version 4.0