THEME 4

Wireless Security

Telecommunication systems department

Lecturer: assistant professor Persikov Anatoliy Valentinovich

WLANS

In 1997, IEEE ratified the 802.11 standard for WLANs. The IEEE 802.11 standard supports three transmission methods, including radio transmission within the 2.4 GHz band. In 1999, IEEE ratified two amendments to the 802.11 standard – 802.11a and 802.11b – that define radio transmission methods, and WLAN equipment based on IEEE 802.11b quickly became the dominant wireless technology. IEEE 802.11b equipment transmits in the 2.4 GHz band, offering data rates of up to 11 Mbps. IEEE 802.11b was intended to provide performance, throughput, and security features comparable to wired LANs. In 2003, IEEE released the 802.11g amendment, which specifies a radio transmission method that uses the 2.4 GHz band and can support data rates of up to 54 Mbps. Additionally, IEEE 802.11g-compliant products are backward compatible with IEEE 802.11b-compliant products.

Standard	Max Data Rate	Typical Range	Band	Comments
802.11	2 Mbps	50-100 meters	2.4 GHz	
802.11a	54 Mbps	50-100 meters	5 GHz	Not compatible with 802.11b
802.11b	11 Mbps	50-100 meters	2.4 GHz	Equipment based on 802.11b has been the dominant WLAN technology
802.11g	54 Mbps	50-100 meters	2.4 GHz	Backward compatible with 802.11b

WI-FI ALLIANCE CERTIFICATION

While IEEE was examining the shortcomings of IEEE 802.11 security and starting to develop the 802.11i amendment, a non-profit industry consortium of WLAN equipment and software vendors called the Wi-Fi Alliance developed an interoperability certification program for WLAN products. The Wi-Fi Alliance felt it was necessary to create an interim solution that could be deployed using existing IEEE 802.11 hardware while IEEE worked on finalizing the 802.11i amendment. Accordingly, the Alliance created **Wi-Fi Protected Access** (**WPA**), which was published in October 2002; it is essentially a subset of the draft IEEE 802.11i requirements available at that time. The most significant difference between WPA and the IEEE 802.11i drafts is that WPA does not require support for Advanced Encryption Standard (AES), a strong encryption algorithm, because many existing IEEE 802.11 hardware components cannot support computationally intensive encryption without additional hardware components.

In conjunction with the ratification of the IEEE 802.11i amendment, the Wi-Fi Alliance introduced WPA2, its term for interoperable equipment that is capable of supporting IEEE 802.11i requirements. The Wi-Fi Alliance began testing IEEE 802.11i products for WPA2 certification shortly after the IEEE 802.11i amendment was finalized.

WIRELESS STANDARDS

In addition to the IEEE 802.11 and WPA standards, other wireless standards are also in use:

- **Wireless personal area networks (WPAN):** small-scale wireless networks that require little or no infrastructure. A WPAN is typically used by a few devices in a single room instead of connecting the devices with cables. For example, WPANs can provide print services or enable a wireless keyboard or mouse to communicate with a computer. Examples of WPAN standards include the following:
  - **IEEE 802.15.1 (Bluetooth).** This WPAN standard is designed for wireless networking between small portable devices. The original Bluetooth operated at 2.4 GHz and has a maximum data rate of approximately 720 kilobits per second (Kbps); Bluetooth 2.0 can reach 3 Mbps.
  - **IEEE 802.15.3 (High-Rate Ultrawideband; WiMedia, Wireless USB).** This is a low-cost, low power consumption WPAN standard that uses a wide range of GHz frequencies to avoid interference with other wireless transmissions. It can achieve data rates of up to 480 Mbps over short ranges and can support the full range of WPAN applications. One expected use of this technology is the ability to detect shapes through physical barriers such as walls and boxes, which could be useful for applications ranging from law enforcement to search and rescue operations.
  - ` This is a simple protocol for lightweight WPANs. It is most commonly used for monitoring and control products, such as climate control systems and building lighting.

WIRELESS STANDARDS

- **Wireless local area networks (WLAN).** IEEE 802.11 is the dominant WLAN standard, but others have also been defined. For example, the European Telecommunications Standards Institute (ETSI) has published the **High Performance Radio Local Area Network (HIPERLAN)** WLAN standard that transmits data in the 5 GHz band and operates at data rates of approximately 23.5 Mbps. However, HIPERLAN appears to have been supplanted by IEEE 802.11 in the commercial arena.
- **Wireless metropolitan area networks (WMAN):** networks that can provide connectivity to users located in multiple facilities that are generally within a few miles of each other. Many WMAN implementations provide wireless broadband access to customers in metropolitan areas. For example, IEEE 802.16e (better known as WiMAX) is a WMAN standard that transmits in the 10 to 66 GHz band range. An IEEE 802.16a addendum allows for large data transmissions with minimal interference. WiMAX provides throughput of up to 75 Mbps, with a range of up to 30 miles for fixed line-of-site communication. However, there is generally a tradeoff; 75 Mbps throughput is possible at half a mile, but at 30 miles the throughput is much lower.
- **Wireless wide area networks (WWAN):** networks that connect individuals and devices over large geographic areas, often globally. WWANs are typically used for cellular voice and data communications, as well as satellite communications.

IEEE 802.11 NETWORK COMPONENTS AND ARCHITECTURAL MODELS

IEEE 802.11 has two fundamental architectural components, as follows:

- **Station (STA).** A STA is a wireless endpoint device. Typical examples of STAs are laptop computers, personal digital assistants (PDA), mobile phones, and other consumer electronic devices with IEEE 802.11 capabilities.
- **Access Point (AP).** An AP logically connects STAs with a distribution system (DS), which is typically an organization's wired infrastructure. APs can also logically connect wireless STAs with each other without accessing a distribution system.

The IEEE 802.11 standard also defines the following two WLAN design structures or configurations:

- **Ad Hoc Mode.** The ad hoc mode does not use APs. Ad hoc mode is sometimes referred to as infrastructureless because only peer-to-peer STAs are involved in the communications.
- **Infrastructure Mode.** In infrastructure mode, an AP connects wireless STAs to each other or to a distribution system, typically a wired network. Infrastructure mode is the most commonly used mode for WLANs.

AD HOC MODE

This mode of operation, also known as **peer-to-peer mode**, is possible when two or more STAs are able to communicate directly to one another.

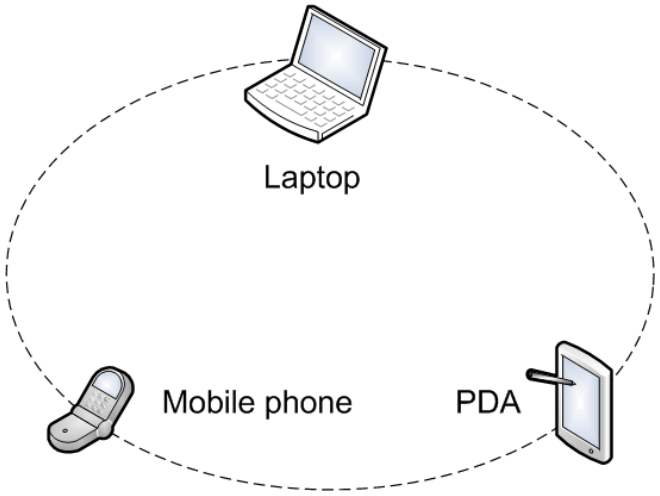

A set of STAs configured in this ad hoc manner is known as an **independent basic service set (IBSS)**.

Today, a STA is most often thought of as a simple laptop with an inexpensive network interface card (NIC) that provides wireless connectivity; however, many other types of devices could also be STAs. A fundamental property of IBSS is that it defines **no routing or forwarding**, so, based on the bare IEEE 802.11i spec, all the devices must be within radio range of one another.

INFRASTRUCTURE MODE

In infrastructure mode, an IEEE 802.11 WLAN comprises one or more **Basic Service Sets** (BSS), the basic building blocks of a WLAN. A BSS includes an AP and one or more STAs. The AP in a BSS connects the STAs to the DS. The DS is the means by which STAs can communicate with the organization's wired LANs and external networks such as the Internet.

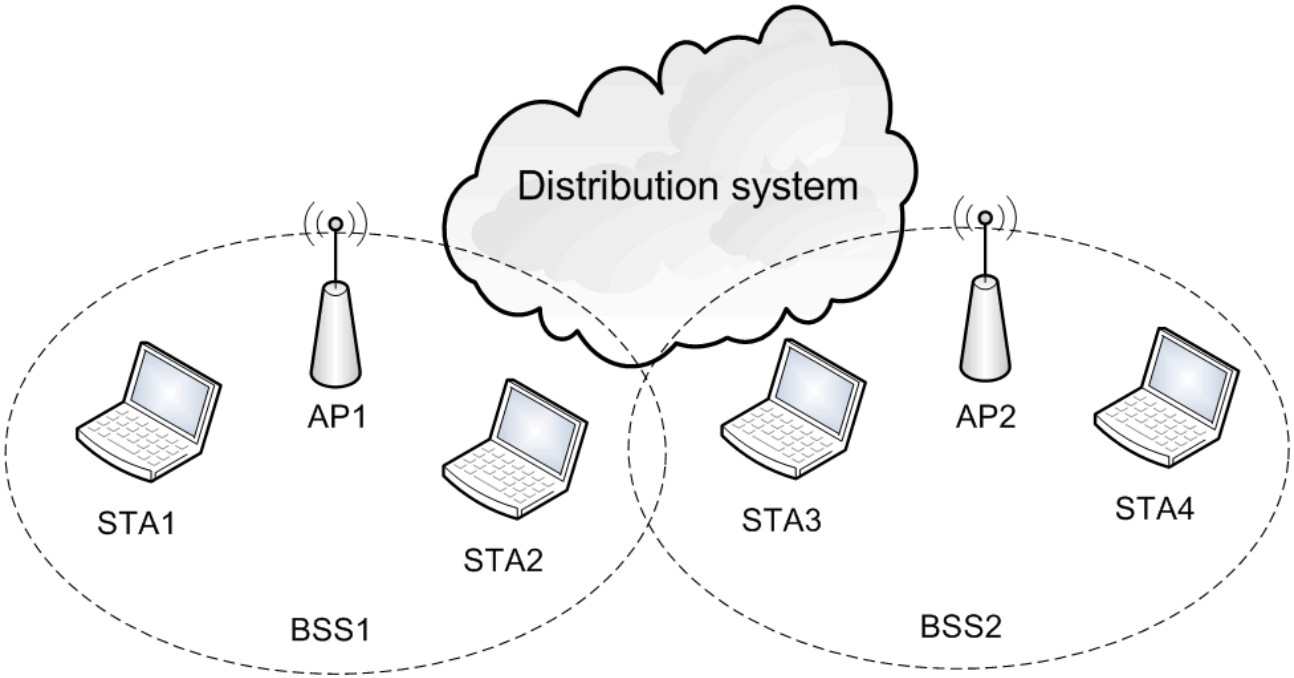

The DS and use of multiple BSSs and their associated APs allow for the creation of wireless networks of arbitrary size and complexity.

INFRASTRUCTURE MODE - EXTENDED SERVICE SET

In the IEEE 802.11 specification, this type of multi-BSS network is referred to as an extended service set (ESS). Figure conceptually depicts a network with both wired and wireless capabilities. It shows three APs with their corresponding BSSs, which comprise an ESS; the ESS is attached to the wired infrastructure. In turn, the wired infrastructure is connected through a perimeter firewall to the Internet. This architecture could permit various STAs, such as laptops and PDAs, to provide Internet connectivity for their users.

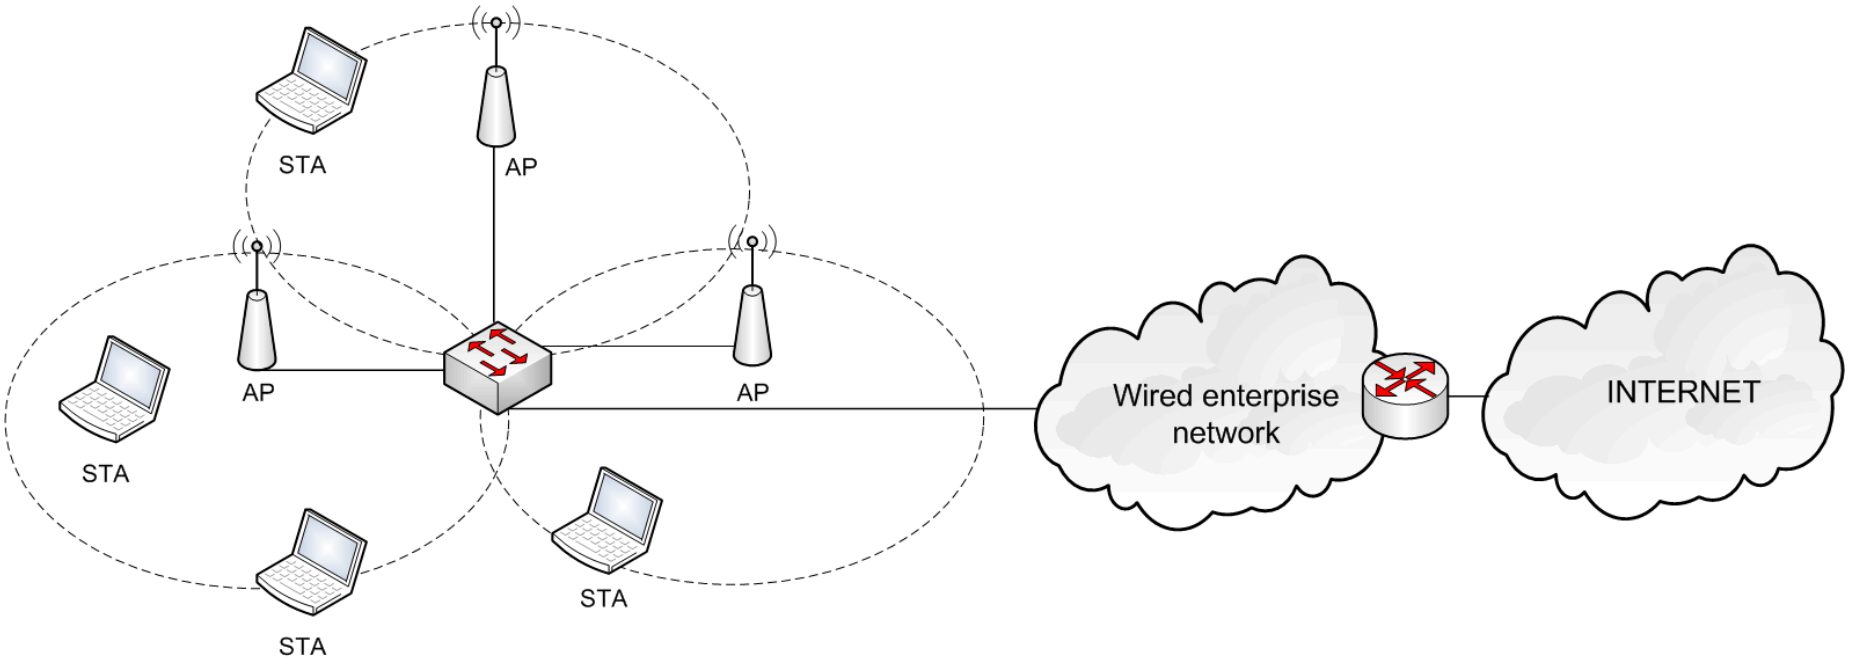

WLAN SECURITY CONCERNS

Like other wireless technologies, WLANs typically need to support several security objectives. This is intended to be accomplished through a combination of security features built into the wireless networking standard. The most common security objectives for WLANs are as follows:

- **Confidentiality –** ensure that communication cannot be read by unauthorized parties
- **Integrity –** detect any intentional or unintentional changes to data that occur in transit
- **Availability –** ensure that devices and individuals can access a network and its resources whenever needed
- **Access Control –** restrict the rights of devices or individuals to access a network or resources within a network.

The security objectives for wireless and wired LANs are the same, as are the major high-level categories of threats that they face.

MAJOR THREATS AGAINST LAN SECURITY

Threat Category	Description
Denial of Service	Attacker prevents or prohibits the normal use or management of networks or network devices.
Eavesdropping	Attacker passively monitors network communications for data, including authentication credentials.
Man-in-the-Middle	Attacker actively intercepts the path of communications between two legitimate parties, thereby obtaining authentication credentials and data. Attacker can then masquerade as a legitimate party. In the context of a WLAN, a man-in-the-middle attack can be achieved through a bogus or rogue AP.
Masquerading	Attacker impersonates an authorized user and gains certain privileges.
Message Modification	Attacker alters a legitimate message by deleting, adding to, changing, or reordering it.
Message Replay	Attacker passively monitors transmissions and retransmits messages, acting as if the attacker were a legitimate user.
Traffic Analysis	Attacker monitors transmissions to identify communication patterns and participants.

HISTORY OF PRE-RSN IEEE 802.11 SECURITY

Prior to the IEEE 802.11i amendment and its RSN framework, IEEE 802.11 had a number of serious security weaknesses. Many vendors have added proprietary features to their IEEE 802.11 implementations to compensate for security flaws in the standard, but proprietary features often prevent interoperability.

Access Control and Authentication

The original IEEE 802.11 specification defines two means to validate the identities of wireless devices attempting to gain access to a WLAN, open system authentication and shared key authentication; neither of these alternatives is secure. IEEE 802.11 implementations are required to support open system authentication; shared key authentication support is optional. Open system authentication is effectively a **null authentication mechanism** that does not provide true identity verification.

HISTORY OF PRE-RSN IEEE 802.11 SECURITY

In practice, a STA is authenticated to an AP simply by providing the following information:

- **Service Set Identifier (SSID) for the AP.** The SSID is a name assigned to a WLAN; it allows STAs to distinguish one WLAN from another. SSIDs are broadcast in plaintext in wireless communications, so an eavesdropper can easily learn the SSID for a WLAN. However, the SSID is not an access control feature, and was never intended to be used for that purpose.

- **Media Access Control (MAC) address for the STA.** A MAC address is a (hopefully) unique 48-bit value that is permanently assigned to a particular wireless network interface. Many implementations of IEEE 802.11 allow administrators to specify a list of authorized MAC addresses; the AP will permit devices with those MAC addresses only to use the WLAN. This is known as MAC address filtering. However, since the MAC address is not encrypted, it is simple to intercept traffic and identify MAC addresses that are allowed past the MAC filter. Unfortunately, almost all WLAN adapters allow applications to set the MAC address, so it is relatively trivial to spoof a MAC address, meaning attackers can gain unauthorized access easily.

HISTORY OF PRE-RSN IEEE 802.11 SECURITY

Additionally, the AP is not authenticated to the STA by open system authentication. Therefore, the STA has to trust that it is communicating to the real AP and not an impostor AP that is using the same SSID. Therefore, open system authentication does not provide reasonable assurance of any identities, and can be misused easily to gain unauthorized access to a WLAN or trick users.

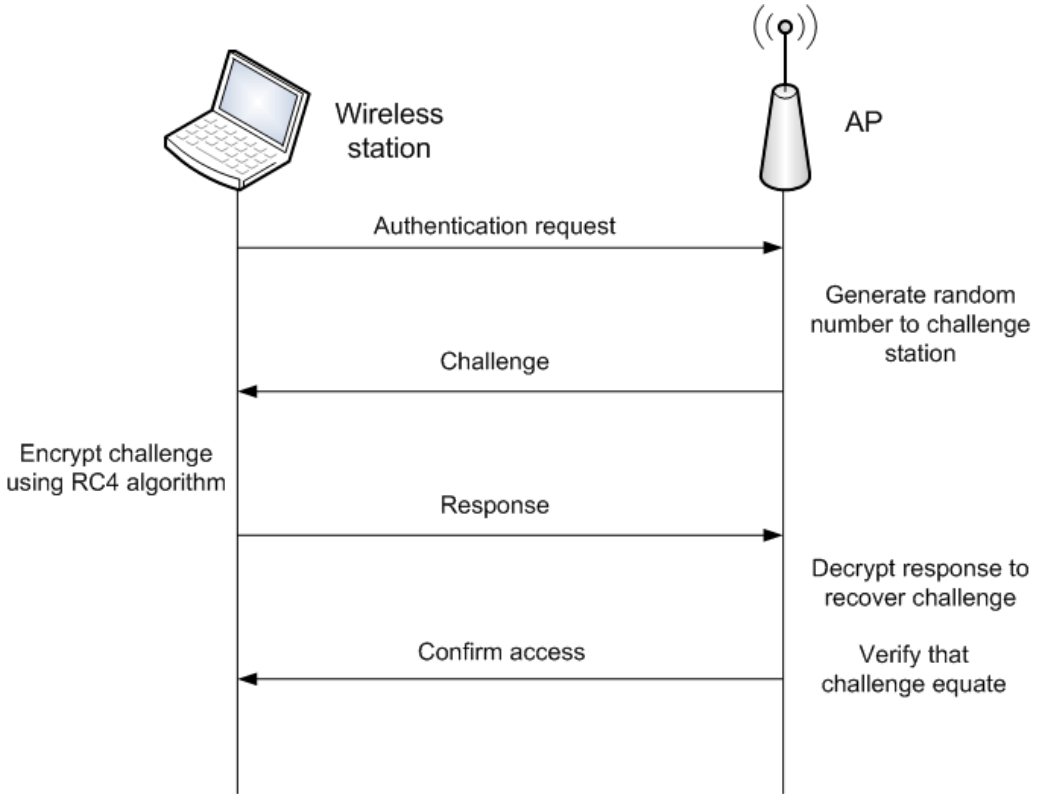

Shared key authentication was supposed to be more robust than open system authentication; in fact, it is equally insecure. As the name implies, shared key authentication is based on a secret cryptographic key known as a **Wired Equivalent Privacy** (WEP) key; this key is shared by legitimate STAs and APs. Shared key authentication uses a simple challenge-response scheme based on whether the STA seeking WLAN access knows the WEP key.

HISTORY OF PRE-RSN IEEE 802.11 SECURITY - ENCRYPTION

The **WEP protocol**, part of the IEEE 802.11 standard, **uses the RC4 stream cipher algorithm** to encrypt wireless communications, which protects their contents from disclosure to eavesdroppers. The standard for WEP specifies support for a **40-bit WEP key only**; however, many vendors offer non-standard extensions to WEP that support key lengths of **up to 128** or even **256** bits. WEP also uses a **24-bit value** known as an **initialization vector** (IV) as a seed value for initializing the cryptographic key stream. For example, a 104-bit WEP key with a 24-bit IV becomes a 128-bit RC4 key. Ideally, larger key sizes translate to stronger protection, but the cryptographic technique used by WEP has known flaws that are not mitigated by longer keys.

Most attacks against WEP encryption have been based on IV-related vulnerabilities. For example, the IV portion of the RC4 key is sent in cleartext, which allows an eavesdropper that monitors and analyzes a relatively small amount of network traffic to recover the key by taking advantage of the IV value knowledge, the relatively small 24-bit IV key space, and a weakness in the way WEP implements the RC4 algorithm.

HISTORY OF PRE-RSN IEEE 802.11 SECURITY - ENCRYPTION

Also, WEP does not specify precisely how the IVs should be set or changed; some products use a static, well-known IV value or reset to zero. If two messages have the same IV, and the plaintext of either message is known, it is relatively trivial for an attacker to determine the plaintext of the second message. In particular, because many messages contain common protocol headers or other easily guessable contents, it is often possible to identify the original plaintext contents with minimal effort. Even traffic from products that use sequentially increasing IV values is still susceptible to attack. There are less than 17 million possible IV values; on a busy WLAN, the entire IV space may be exhausted in a few hours. When the IV is chosen randomly, which represents the best possible generic IV selection algorithm, by the birthday paradox two IVs already have a 50% chance of colliding after about 2^{12} frames.

Another possible threat against confidentiality is network traffic analysis. Eavesdroppers might be able to gain information by monitoring which parties communicate at what times. Also, analyzing traffic patterns can aid in determining the content of communications; for example, short bursts of activity might be caused by terminal emulation or instant messaging, while steady streams of activity might be generated by video conferencing. More sophisticated analysis might be able to determine the operating systems in use based on the length of certain frames.

HISTORY OF PRE-RSN IEEE 802.11 SECURITY - DATA INTEGRITY

WEP performs data integrity checking for messages transmitted between STAs and APs. WEP is designed to reject any messages that have been changed in transit, such as by a man-in-the-middle attack. WEP data integrity is based on a simple encrypted checksum – a 32-bit cyclic redundancy check (CRC-32) computed on each payload prior to transmission. The payload and checksum are encrypted using the RC4 key stream and transmitted. The receiver decrypts them, recomputes the checksum on the received payload, and compares it with the transmitted checksum. If the checksums are not the same, the transmitted data frame has been altered in transit, and the frame is discarded.

Unfortunately, CRC-32 is subject to bit flipping attacks, which means that an attacker knows which CRC-32 bits will change when message bits are altered. WEP attempts to counter this problem by encrypting the CRC-32 to produce an **integrity check value** (ICV). The creators of WEP believed that an enciphered CRC-32 would be less subject to tampering. However, they did not realize that a property of stream ciphers such as WEP's RC4 is that bit flipping survives the encryption process—the same bits flip whether or not encryption is used. Therefore, the WEP ICV offers no additional protection against bit flipping.

HISTORY OF PRE-RSN IEEE 802.11 SECURITY

Integrity should be provided by a **cryptographic checksum** rather than a CRC. Also known as **keyed hashes** or **message authentication codes** (MAC), cryptographic checksums prevent bit flipping attacks because they are designed so that any change to the original message results in significant and unpredictable changes to the resulting checksum.

CRCs are generally more efficient computationally than cryptographic checksums, but are only designed to protect against random bit errors, not intentional forgeries, so they do not provide the same level of integrity protection.

Replay Protection

The cryptographic implementation provides no protection against replay attacks because it does not include features such as an incrementing counter, timestamp, or other temporal data that would make replayed traffic easily detectable.

HISTORY OF PRE-RSN IEEE 802.11 SECURITY - AVAILABILITY

Individuals who do not have physical access to the WLAN infrastructure can cause a **denial of service for the WLAN**.

One threat is known as **jamming**, which involves a device that emits electromagnetic energy on the WLAN's frequencies. The energy makes the frequencies unusable by the WLAN, causing a denial of service. Jamming can be performed intentionally by an attacker or unintentionally by a non-WLAN device transmitting on the same frequency.

Another threat against availability is **flooding**, which involves an attacker sending large numbers of messages to an AP at such a high rate that the AP cannot process them, or other STAs cannot access the channel, causing a partial or total denial of service. These threats are difficult to counter in any radio-based communications; thus, the IEEE 802.11 standard does not provide any defense against jamming or flooding.

Attackers can **establish rogue APs**; if STAs mistakenly attach to a rogue AP instead of a legitimate one, this could make the legitimate WLAN effectively unavailable to users. Although 802.11i **protects data frames**, **it does not offer protection to control or management frames**. An attacker can exploit the fact that management frames are not authenticated to deauthenticate a client or to disassociate a client from the network.

BRIEF OVERVIEW OF IEEE 802.11I SECURITY

The IEEE 802.11i standard is the sixth amendment to the baseline IEEE 802.11 standards. It includes many security enhancements that leverage mature and proven security technologies. For example, IEEE 802.11i references the **Extensible Authentication Protocol** (EAP) standard, which is a means for providing mutual authentication between STAs and the WLAN infrastructure, as well as performing automatic cryptographic key distribution. EAP is a standard developed by the Internet Engineering Task Force (IETF). IEEE 802.11i employs accepted cryptographic practices, such as generating **cryptographic checksums** through **hash message authentication codes** (HMAC).

The IEEE 802.11i specification introduces the concept of a Robust Security Network (RSN). An RSN is defined as a wireless security network that only allows the creation of Robust Security Network Associations (RSNA). An RSNA is a logical connection between communicating IEEE 802.11 entities established through the IEEE 802.11i key management scheme, called the 4-Way Handshake, which is a protocol that validates that both entities share a **pairwise master key** (PMK), synchronizes the installation of temporal keys, and confirms the selection and configuration of data confidentiality and integrity protocols.

BRIEF OVERVIEW OF IEEE 802.11I SECURITY

The entities obtain the PMK in one of two ways – either the PMK is already configured on each device, in which case it is called a **pre-shared key** (PSK), or it is distributed as a side effect of a successful EAP authentication instance, which is a component of IEEE 802.1X port-based access control. The PMK serves as the basis for the IEEE 802.11i data confidentiality and integrity protocols that provide enhanced security over the flawed WEP. Most large enterprise deployments of RSN technology will use IEEE 802.1X and EAP rather than PSKs because of the difficulty of managing PSKs on numerous devices. WLAN connections employing ad hoc mode, which typically involve only a few STAs, are more likely to use PSKs.

The IEEE 802.1X standard defines several terms related to authentication. The **authenticator** is an entity at one end of a point-to-point LAN segment that facilitates authentication of the entity attached to the other end of that link.

BRIEF OVERVIEW OF IEEE 802.11I SECURITY

The supplicant is the entity being authenticated. The STA may be viewed as a supplicant. The **authentication server** (AS) is an entity that provides an authentication service to an **authenticator**. This service determines from the credentials provided by the supplicant whether the supplicant is authorized to access the services provided by the authenticator.

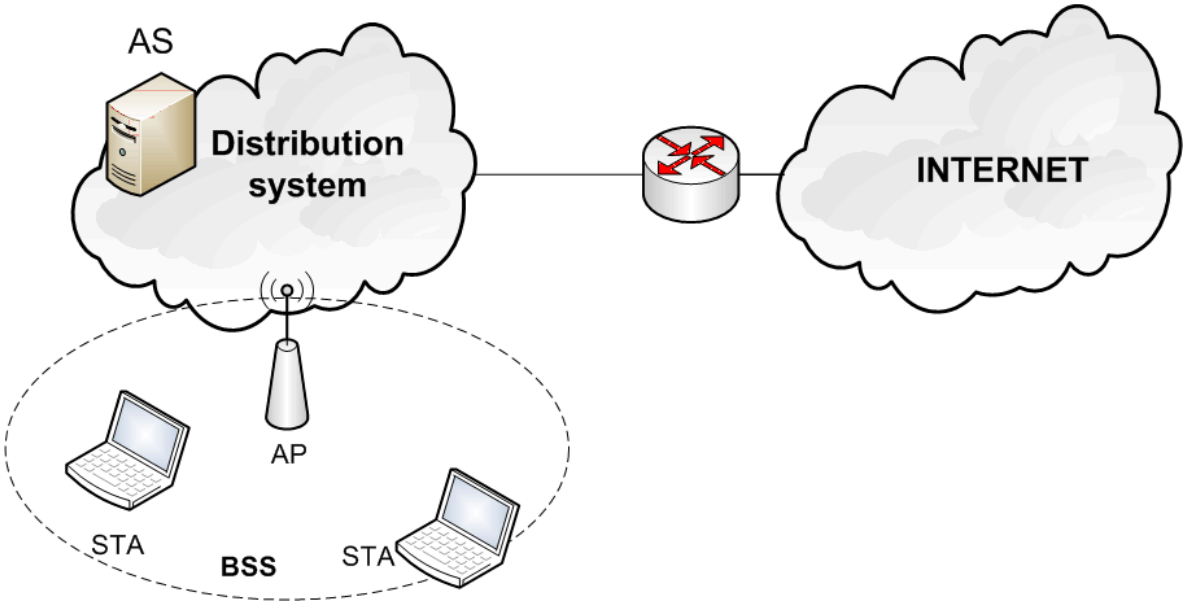

The AS provides these authentication services and delivers session keys to each AP in the wireless network; each STA either receives session keys from the AS or derives the session keys itself. The AS either authenticates the STA and AP itself, or provides information to the STA and AP so that they may authenticate each other.

BRIEF OVERVIEW OF IEEE 802.11I SECURITY

The AS typically lies inside the DS. When employing a solution based on the IEEE 802.11i standard, the AS most often used for authentication is an **Authentication, Authorization, and Accounting** (AAA) server that uses the **Remote Authentication Dial In User Service** (RADIUS) or **Diameter** protocol to transport authentication-related traffic. The **supplicant/authenticator model** is intrinsically a unilateral rather than mutual authentication model: the supplicant authenticates to the network. IEEE 802.11i combats this bias by requiring that the EAP method used provides mutual authentication.

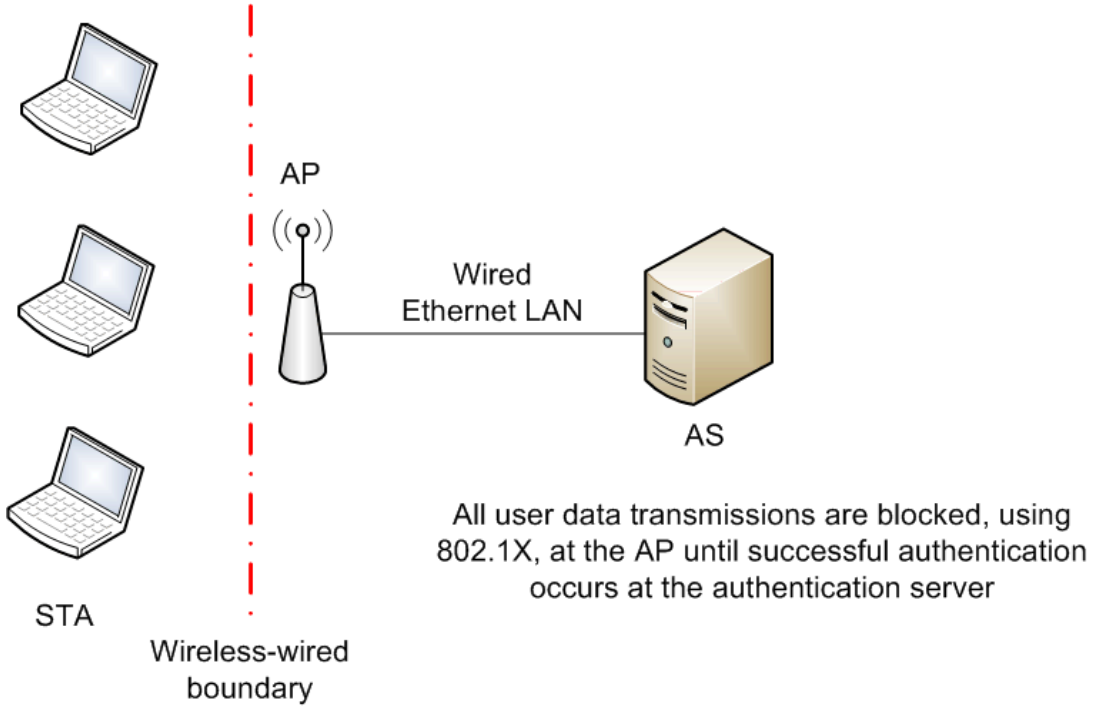

SECURITY FRAMEWORK FOR ROBUST SECURITY NETWORKS

With the addition of the IEEE 802.11i amendment in 2004, IEEE 802.11 offers two general classes of security capabilities for IEEE 802.11 WLANs. The first class, **pre-RSN security**, includes the legacy security capabilities developed in the original IEEE 802.11 specification: open system or shared key authentication for validating the identity of a wireless station, and WEP for the confidentiality protection of traffic. The second class of security capabilities includes a number of security mechanisms to create **RSNs**. An RSN includes security enhancements to address all the known flaws of WEP and provide robust protection for the wireless link, including data integrity and confidentiality.

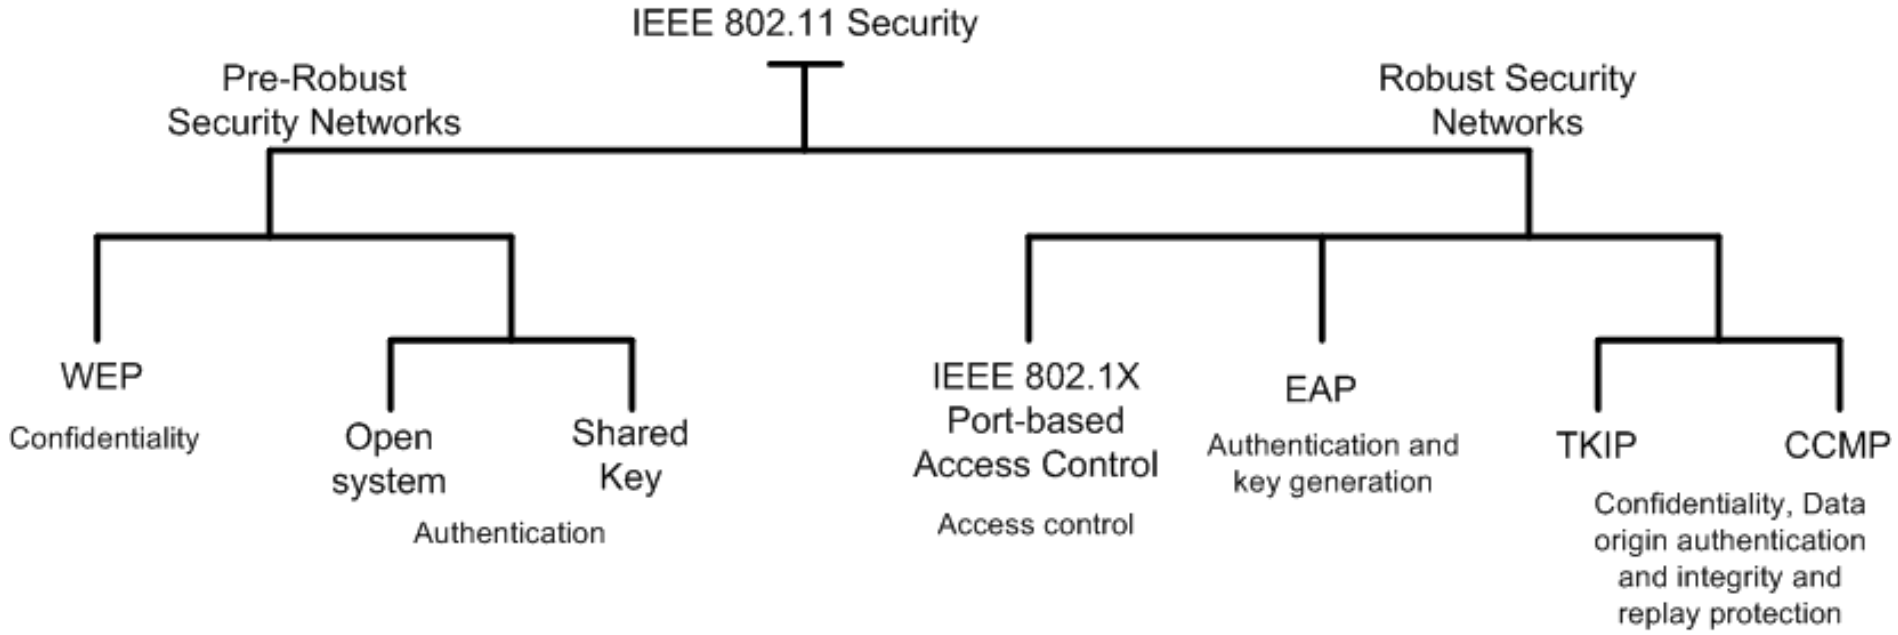

SECURITY FRAMEWORK FOR ROBUST SECURITY NETWORKS

At a high level, RSN includes IEEE 802.1X **port-based access control**, **key management techniques**, and the **TKIP** and **CCMP data confidentiality** and **integrity protocols**. These protocols allow for the creation of several diverse types of security networks because of the numerous configuration options. RSN security is at the link level only, providing protection for traffic between a wireless STA and its associated AP, or between one wireless STA and another wireless STA. It does not provide end-to-end application-level security, such as between a STA and an e-mail or Web server on the DS, because communication between these entities requires more than just one link.

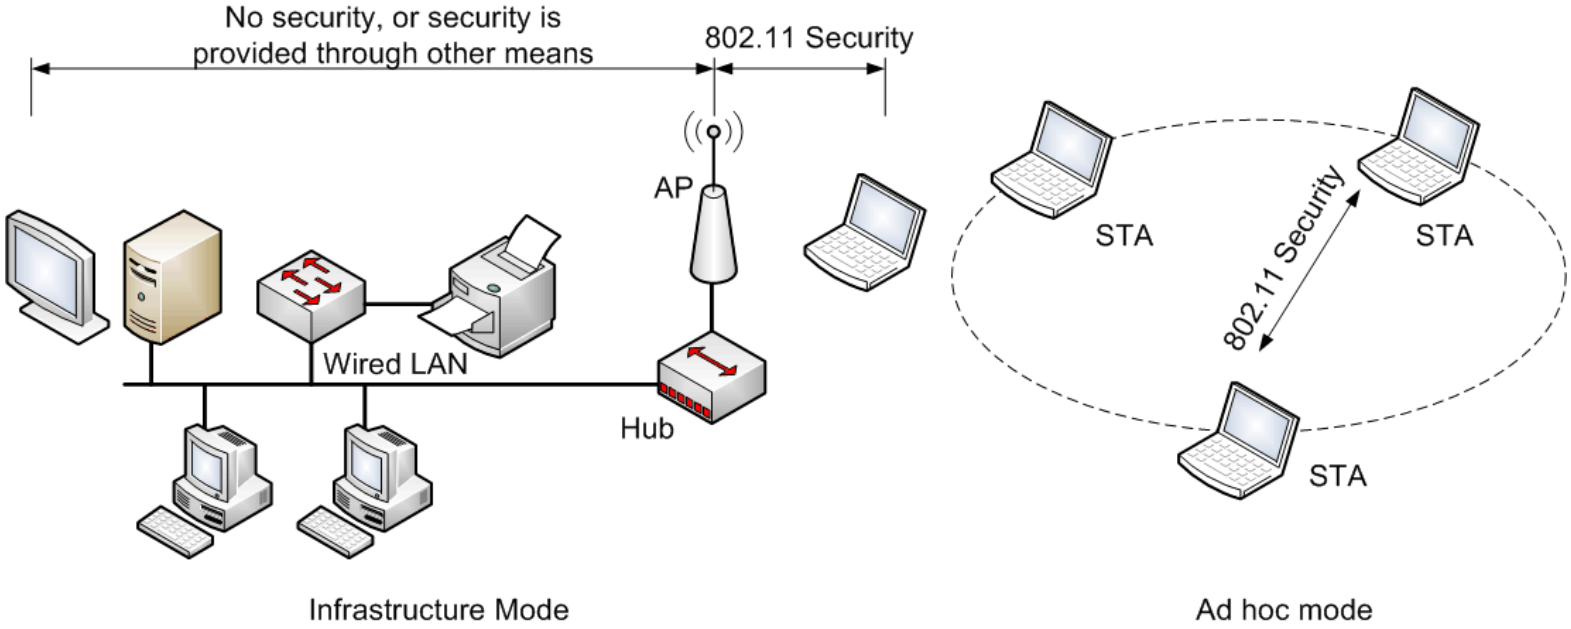

SECURITY FRAMEWORK FOR ROBUST SECURITY NETWORKS

To provide end-to-end security, organizations can implement network level security mechanisms such as **Transport Layer Security** (TLS) or IPsec. Also, RSN's security features apply only to the wireless portion of the overall network, not to communications on wired networks.

The IEEE 802.11i amendment defines an RSN as a wireless network that allows the creation of RSN Associations (RSNA) only.

An RSNA is a **security relationship** established by the IEEE 802.11i 4-Way Handshake. The 4-Way Handshake validates that the parties to the protocol instance both possess a **pairwise master key** (PMK), synchronizes the installation of temporal keys, and confirms the selection of cipher suites.

The PMK is the cornerstone for a number of security features absent from WEP. Complete robust security is considered to be possible only when all devices in the network use RSNAs. In practice, some networks have a mix of RSNAs and non-RSNA connections. A network that allows the creation of both pre-RSN associations (pre-RSNA) and RSNAs is referred to as a **Transition Security Network** (TSN). A TSN is intended to be an interim means to provide connectivity while an organization migrates to networks based exclusively on RSNAs.

SECURITY FRAMEWORK FOR ROBUST SECURITY NETWORKS

RSNAs enable the following security features for IEEE 802.11 WLANs:

- Enhanced user authentication mechanisms.
- Cryptographic key management.
- Data confidentiality.
- Data origin authentication and integrity.
- Replay protection.

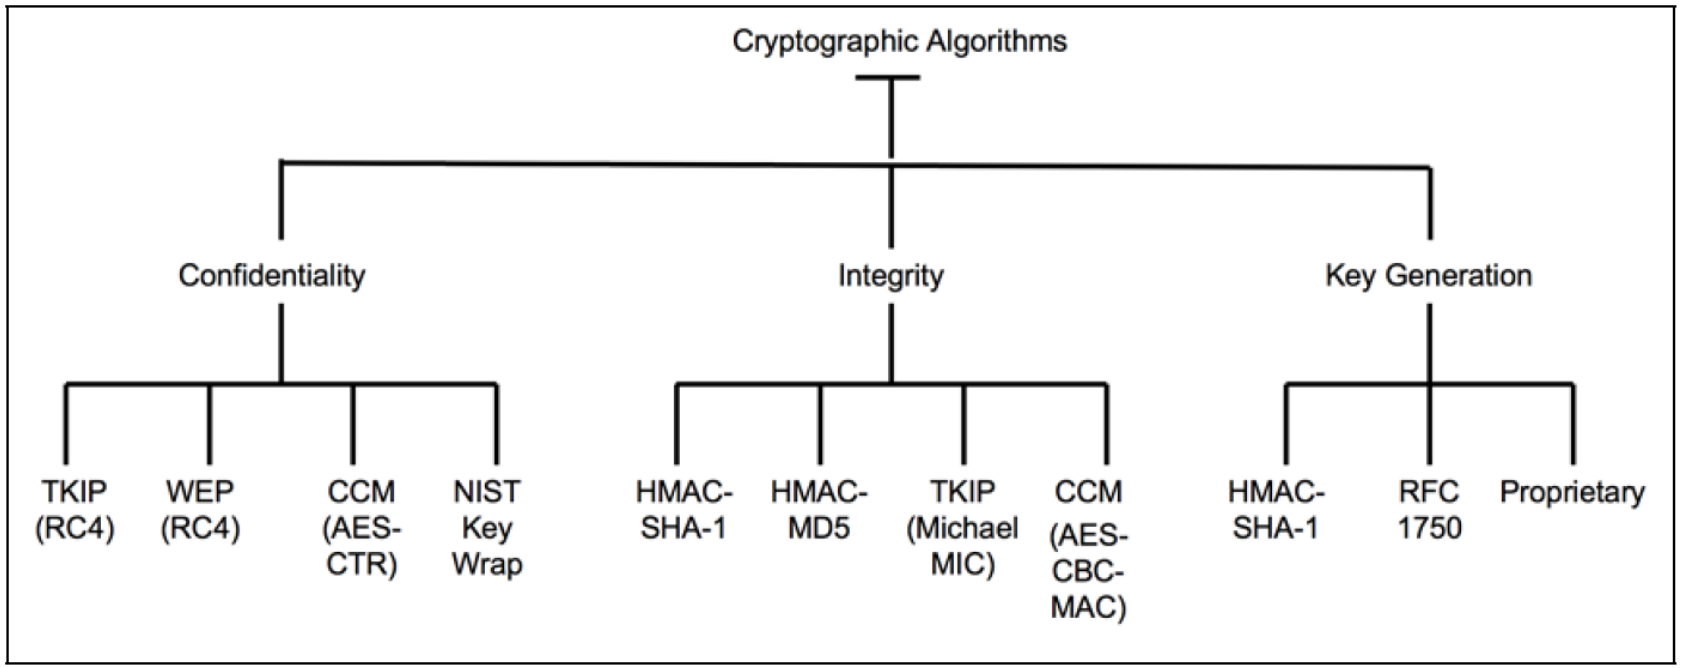

KEY HIERARCHIES AND KEY DISTRIBUTION AND MANAGEMENT

Fundamental to any cryptographic system are the cryptographic keys used in the transformation (enciphering or deciphering) processes. Since cryptography is the security foundation of IEEE 802.11 WLANs, the security of keys is particularly important.

Keys typically need to meet the following requirements:

- Randomly generated to reduce the probability that they can be determined by an adversary or that they will be reused.
- Changed frequently to reduce the possibility of discovery through sophisticated cryptanalysis.
- Protected while in storage, so that previous communications cannot be deciphered.
- Protected during transmission.
- Erased completely when no longer needed.

These requirements are related to the security service known as key management, which is defined as “the process of handling and controlling cryptographic keys and related material (such as initialization values) during their life cycle in a cryptographic system, including ordering, generating, distributing, storing, loading, escrowing, archiving, auditing, and destroying the material.”

KEY HIERARCHIES (PAIRWISE KEY HIERARCHY)

Pre-Shared Key (PSK). A PSK is a static key delivered to the AS and the STA through an out-of-band mechanism. The PSK must be put into place before establishing an association. The PSK may be generated and installed in any number of ways, including proprietary automated public-key cryptographic approaches, and manual means such as a USB device or a passphrase (which can be converted to a cryptographic key using one of a number of algorithms). If any of the PSKs are compromised, they must be re-distributed in the same way.

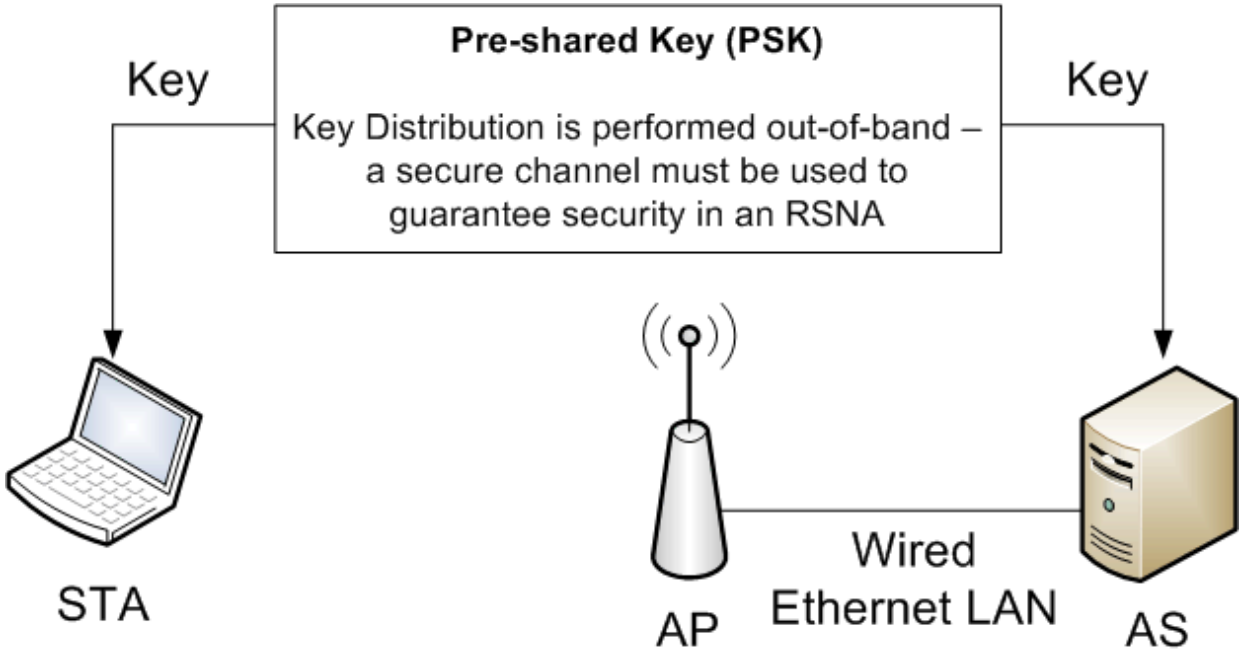

KEY HIERARCHIES (PAIRWISE KEY HIERARCHY)

The security of the WLAN is compromised if any of the PSKs does not possess sufficient cryptographic strength; the passphrase from which the PSK is generated must be a long and complex, possibly randomly generated. The IEEE 802.11 standard does not specify how PSKs are to be generated or distributed, so these decisions are left to implementers. As a result, organizations should review any PSK approach carefully for possible vulnerabilities and evaluate its performance implications.

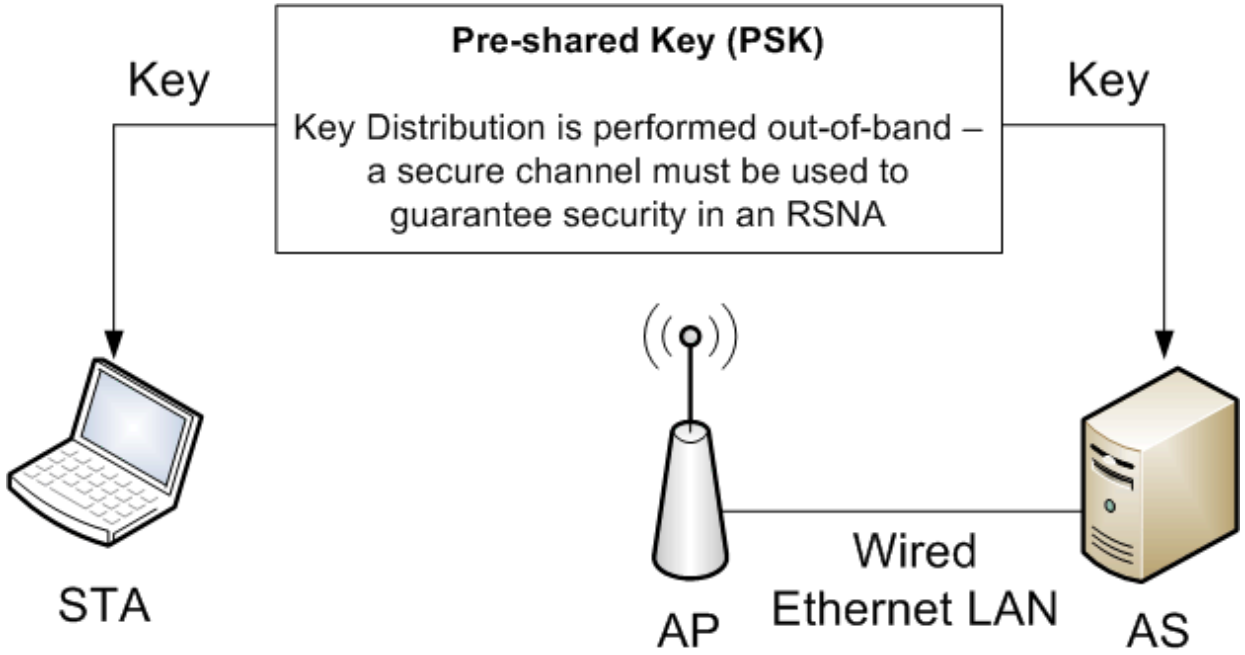

Distributing PSKs in a large network might be infeasible.

Due to client software limitations, a common practice is to assign a single PSK per SSID to enable roaming. In such a case, all users can decrypt the traffic of other users, even if the network is protected from outsiders.

KEY HIERARCHIES (GROUP KEY HIERARCHY)

The second key hierarchy defined by IEEE 802.11 is the **Group Key Hierarchy**, which consists of a single key, the **Group Temporal Key** (GTK). Unlike the PMK, which is generated using material from both supplicant and authenticator, the GTK is generated by the authenticator (AP) and transmitted to its associated STAs. Exactly how this GTK is generated is undefined and is likely to vary considerably in various vendor implementations, with possible implications for security. IEEE 802.11i, however, requires that its value is computationally indistinguishable from random.

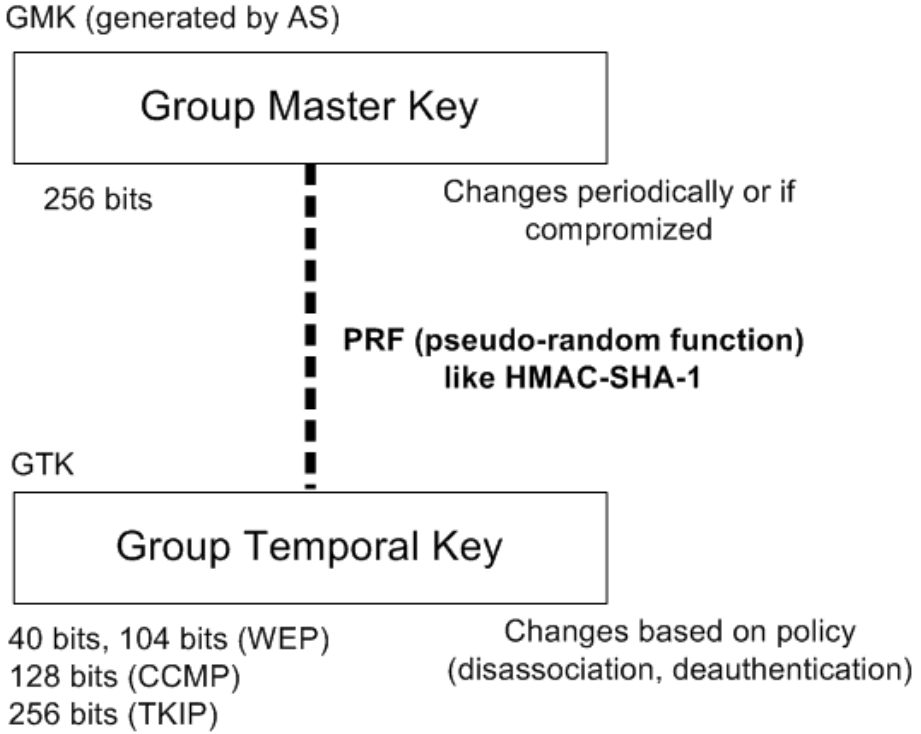

KEY HIERARCHIES (KEYS SUMMARY)

Abbr.	Name	Description / Purpose	Size (bits)	Type
AAA Key	Authentication, Accounting, and Authorization Key	Used to derive the PMK. Used with the IEEE 802.1X authentication and key management approach. Same as MSK.	≥ 256	Key generation key, root key
PSK	Pre-Shared Key	Becomes the PMK in pre-shared key environments.	256	Key generation key, root key
PMK	Pairwise Master Key	Used with other inputs to derive the PTK.	256	Key generation key
GMK	Group Master Key	Used with other inputs to derive the GTK.	128	Key generation key
PTK	Pairwise Transient Key	Derived from the PMK. Comprises the EAPOL-KCK, EAPOL-KEK, and TK and (for TKIP) the MIC key.	512 (TKIP) 384 (CCMP)	Composite key
TK	Temporal Key	Used with TKIP or CCMP to provide confidentiality and integrity protection for unicast user traffic.	256 (TKIP) 128 (CCMP)	Traffic key
GTK	Group Temporal Key	Derived from the GMK. Used to provide confidentiality and integrity protection for multicast/broadcast user traffic.	256 (TKIP) 128 (CCMP) 40, 104 (WEP)	Traffic key
MIC Key	Message Integrity Code Key	Used by TKIP's Michael MIC to provide integrity protection of messages.	64	Message integrity key
EAPOL-KCK	EAPOL-Key Confirmation Key	Used to provide integrity protection for key material distributed during the 4-Way Handshake.	128	Message integrity key
EAPOL-KEK	EAPOL-Key Encryption Key	Used to ensure the confidentiality of the GTK and other key material in the 4-Way Handshake.	128	Traffic key / key encryption key
WEP Key	Wired Equivalent Privacy Key	Used with WEP.	40, 104	Traffic key

TEMPORAL KEY INTEGRITY PROTOCOL (TKIP)

TKIP is a cipher suite for enhancing the WEP protocol on pre-RSN hardware without causing significant performance degradation. TKIP works within the processing constraints of first-generation STAs and APs, and therefore enables increased security without requiring hardware replacement.

TKIP provides the following fundamental security features for IEEE 802.11 WLANs:

- Confidentiality protection using the RC4 algorithm.
- Integrity protection against several types of attacks using the Michael message digest algorithm (through generation of a message integrity code).
- Replay prevention through a frame sequencing technique.
- Use of a new encryption key for each frame to prevent attacks such as the Fluhrer-Mantin-Shamir (FMS) attack, which can compromise WEP-based WLANs.
- Implementation of countermeasures whenever the STA or AP encounters a frame with a MIC error, which is a strong indication of an active attack.

TKIP ENCAPSULATION

Encapsulation is the process of generating the cryptographic payload (ciphertext) from the plaintext data. The plaintext data comprises user traffic and the source and destination MAC addresses. TKIP encapsulation builds upon the WEP encapsulation technique, modifying WEP with additional features through software, to bolster security without requiring hardware changes. TKIP uses three distinct keys: two integrity keys and an encryption key. The primary characteristics of TKIP encapsulation are presented briefly as follows:

- **Two 64-bit message integrity keys** are used with the Michael message digest algorithm to produce a **message integrity code** (MIC). One key is used to provide integrity protection for each half-duplex data channel between the STA and AP. The MIC is computed over the user data, source and destination addresses, and priority bits to provide data integrity. Due to design constraints, an attacker can use sophisticated methods to forge information without detection. Accordingly, TKIP decapsulation employs additional countermeasures, to partially mitigate the risk of these attacks.
- **A monotonically increasing TKIP Sequence Counter (TSC)** is assigned to each frame. The TSC provides protection against replay attacks. If frames do not arrive in order, the receiver simply drops them.

TKIP ENCAPSULATION

- A **two-phase cryptographic key-mixing process** occurs to produce a new key for every frame that is transmitted. The process takes a session Temporal Key along with the dynamically changing TSC to create a dynamic WEP key.
- The **original user frame, the computed 64-bit MIC, and the transmitter address** are encrypted using WEP (with RC4) and the per-frame WEP key. Dynamic key updates and other countermeasures provide additional security.

Although the destination and source addresses and priority and payload are used as inputs by the Michael algorithm, only the payload is encrypted. Because the frame header contains the source and destination addresses and priority, the MIC generated by Michael incorporates them. This prevents an adversary from modifying the frame header addresses to spoof the source or redirect the frame to an unauthorized destination. The TKIP encapsulation process also involves encrypting the MIC using WEP, which helps to hide information about the 64-bit MIC key.

TKIP DECAPSULATION AND COUNTERMEASURES

Decapsulation is the **process to recover the content of protected frames** – that is, to decrypt a received ciphertext packet. During decapsulation, various checks are performed on the frames. For example, if the TSC indicates a violation of proper frame sequencing (it should be monotonically increasing), the frame is discarded. Also, the MIC is recomputed and compared with the MIC in the packet; if they do not match, the frame is discarded and TKIP countermeasures are invoked, which serve as a TKIP safety net.

Although the Michael MIC offers increased message integrity protection in comparison with the legacy WEP and its use of an encrypted CRC, Michael is much weaker than what is usually required. Its objective is to provide reasonable levels of integrity assurance on pre-RSNA-compliant devices without requiring hardware upgrades. Michael is subject to a 2^{29} differential cryptanalysis attack, meaning an attacker could expect to create a forgery in about 2^{28} messages on average. Since the Michael MIC has known vulnerabilities, any failure of the message integrity check in TKIP represents a probable active attack. Therefore, TKIP employs additional countermeasures to help thwart these attacks.

TKIP DECAPSULATION AND COUNTERMEASURES

These countermeasures accomplish the following security goals:

- **Logging security events.** MIC failures during decapsulation at the STA or AP likely mean an active attack. These are to be logged, and a system or security administrator should investigate.
- **Limiting MIC failures.** A receiving STA or AP that detects two failures within a 60-second period disables reception for 60 seconds, not allowing any new associations for STAs using TKIP. This suspense mechanism thwarts an adversary's attempts at numerous attacks in a short period, limiting what an active attacker can learn about any Michael key. The countermeasures effectively limit the adversary to random guessing attacks.
- **Changing the PTK and GTK.** Temporal keys are erased and must be re-initialized.
- **Blocking the IEEE 802.1X ports.** If IEEE 802.1X authentication is used, the state machine is initialized, thereby blocking the controlled ports.

COUNTER MODE WITH CIPHER BLOCK CHAINING MAC PROTOCOL (CCMP)

CCMP is the second data confidentiality and integrity protocol that may be negotiated as a cipher suite for the protection of user traffic in an RSNA. Like TKIP, CCMP was developed to address all known inadequacies of WEP; however, CCMP was developed without the constraint of requiring the use of existing hardware. CCMP is considered the long-term solution for the creation of RSNs for WLANs. It is mandatory for RSN compliance.

CCMP is based on **CCM**, a generic authenticated encryption block cipher mode of AES. CCM is a mode of operation defined for any block cipher with a 128-bit block size. CCM combines two well-known and proven cryptographic techniques to achieve robust security. First, CCM uses **CTR for confidentiality** and **Cipher Block Chaining MAC** (CBC-MAC) for both authentication and integrity protection.

CCMP protects the integrity of both the packet data and portions of the IEEE 802.11 header. CCM for IEEE 802.11 employs a single 128-bit session key (TK) to protect the duplex data channel. The CCMP key space has size 2^{128} and uses a 48-bit packet number (PN) to construct a nonce to prevent replay attacks. The construction of the nonce allows the key to be used for both integrity and confidentiality without compromising either.

CCMP

As the long-term IEEE 802.11 WLAN solution for confidentiality and integrity, CCMP uses CCM, which was specifically designed to possess the following characteristics:

- A single cryptographic key for confidentiality and integrity to minimize complexity and maximize performance (minimize key scheduling time)
- Integrity protection of the packet header and packet payload, in addition to providing confidentiality of the payload
- Computation of some cryptographic parameters prior to the receipt of packets to enable fast comparisons when they arrive, which reduces latency
- Small footprint (hardware or software implementation size) to minimize costs
- Small security-related packet overhead (minimal data expansion due to cryptographic padding and integrity field, for instance)
- No encumbrance by any existing or pending patents.

CCMP ENCAPSULATION

CCMP encapsulation is the process of generating the cryptographic payload (ciphertext) from the plaintext data. The plaintext data comprises user traffic and a MAC header.

The primary steps of CCMP encapsulation are the following:

1) The packet number (PN) maintained for the session is incremented.
2) The PN and other portions of the address field are combined to form the nonce.
3) The identifier for the Temporal Key, or KeyID, and the PN are combined to form the CCMP header.
4) The frame header is used to construct the Additional Authentication Data (AAD). The AAD is a 22-byte or 28-byte parameter comprising several fields, including several addresses and the quality-of-service control field, that are used as additional input into the CCM authentication process.
5) The AAD, nonce, and plaintext data are provided as inputs to CCM along with the Temporal Key to encrypt the data.
6) The packet header, the CCM header, and the ciphertext data are concatenated to form the ciphertext (or encapsulated) packet.

CCMP ENCAPSULATION

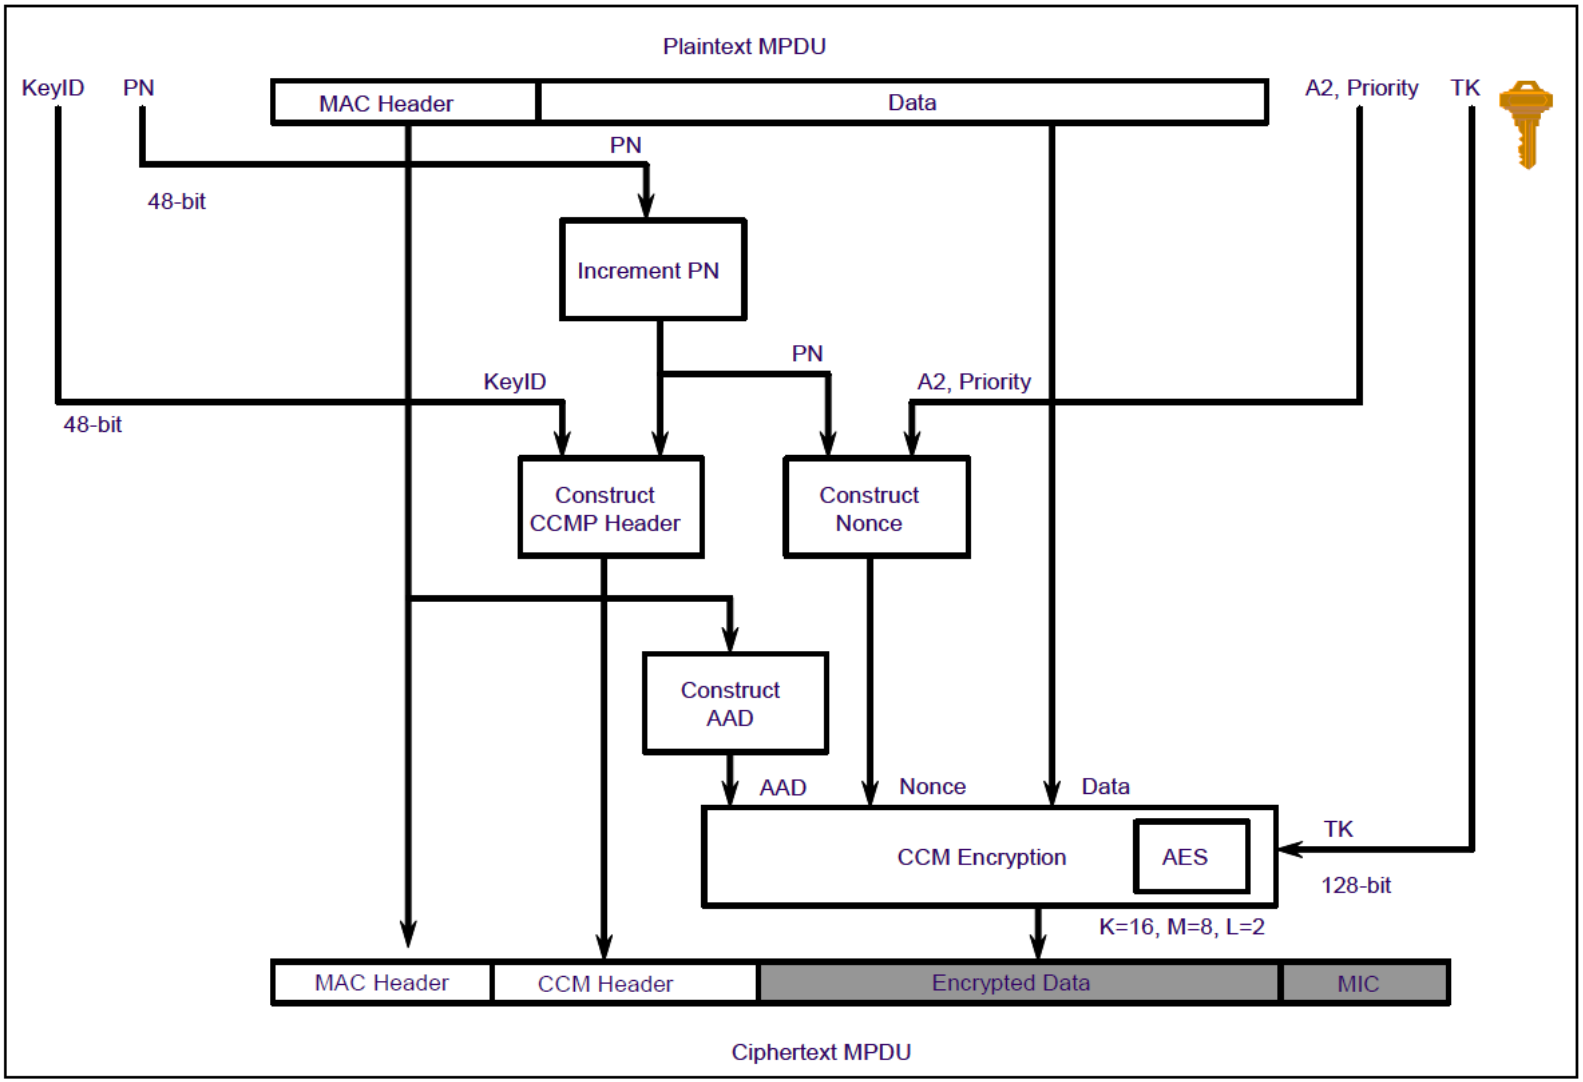

CCMP ENCAPSULATION

CCM is an “authenticate-and-encrypt” block cipher mode of AES. As such, it both encrypts and produces a MIC. As shown, the four inputs to the CCM processing are the following:

- 128-bit cryptographic key, TK
- 48-bit nonce (derived for a 48-bit packet number, PN)
- Additional Authentication Data (AAD)
- Variable length packet (frame body) with header.

CCM uses a new Temporal Key every session – with every new STA-AP association. Unlike TKIP, the use of AES at the core of CCM obviates the need to have per-packet keys. As a result, the two-phase key mixing functions of TKIP encapsulation are not present in the CCMP encapsulation.

CCMP DECAPSULATION

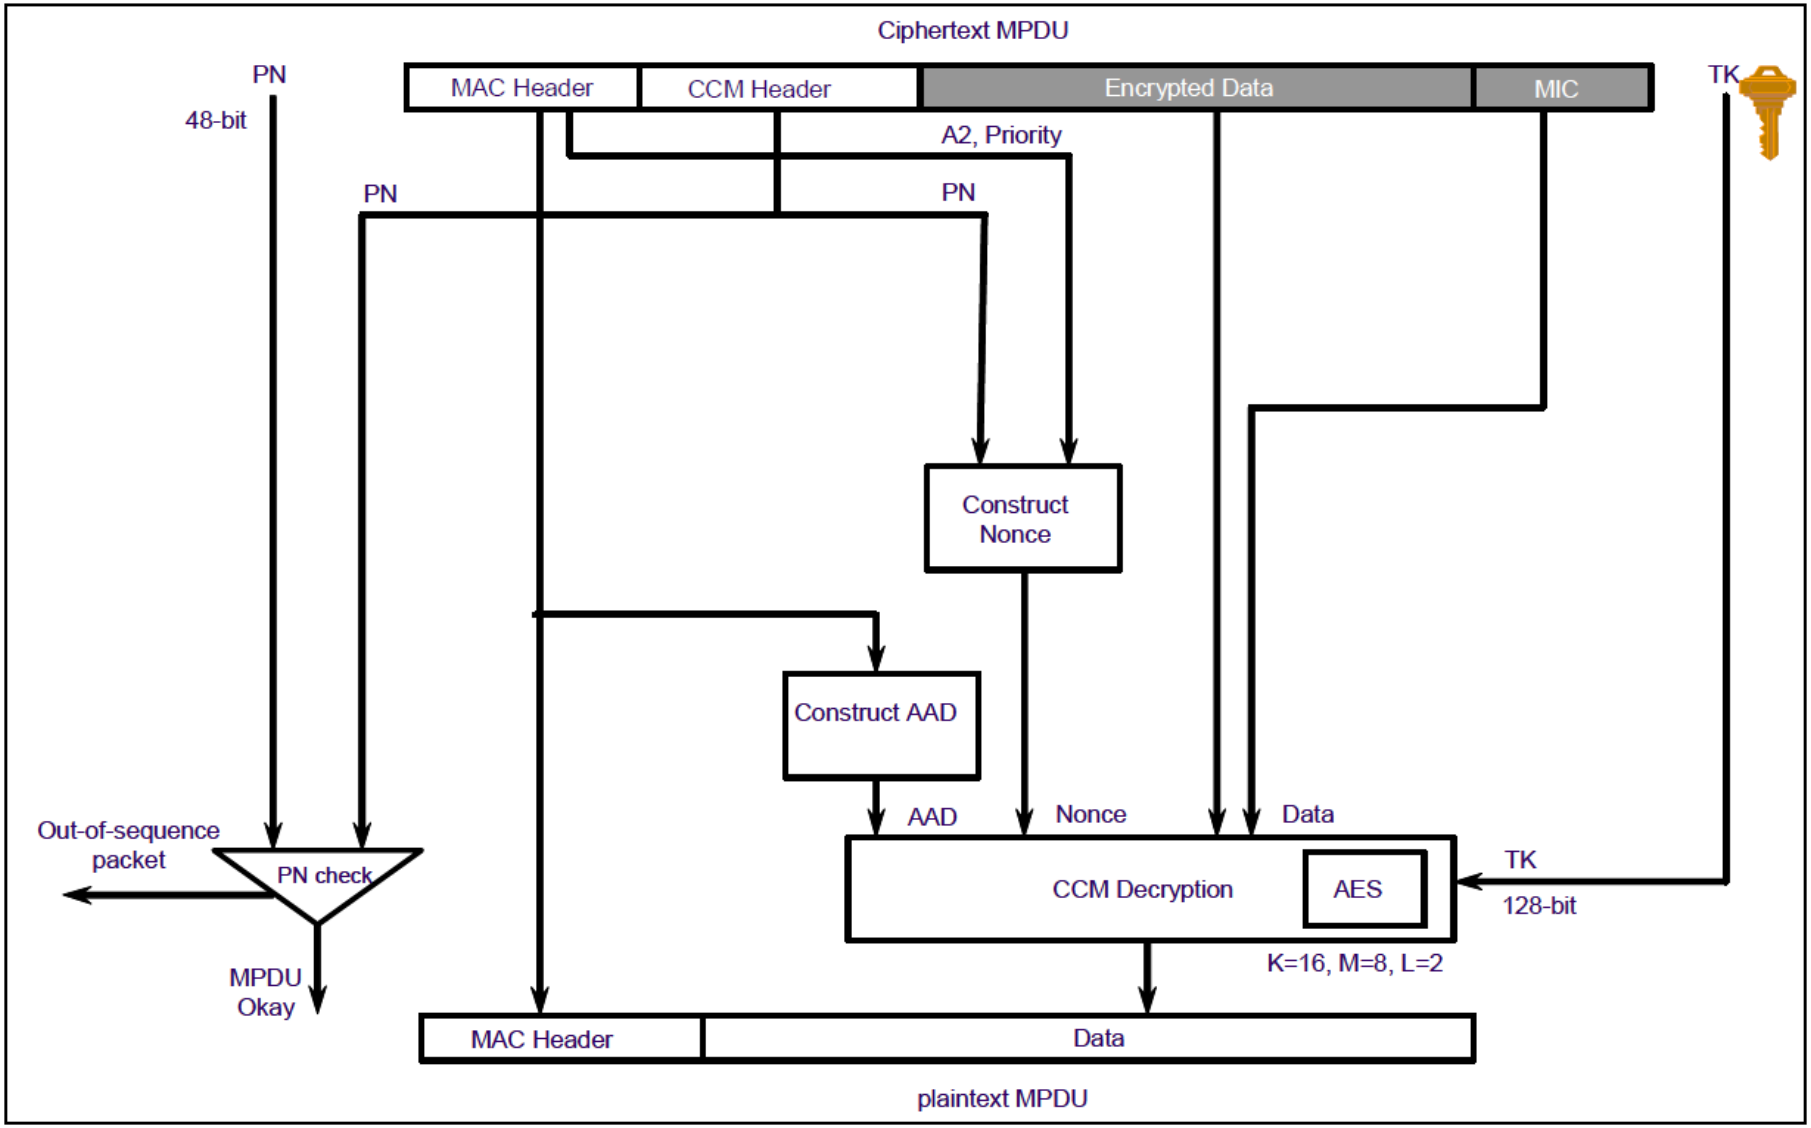

CCMP DECAPSULATION

CCMP decapsulation is used to recover and decrypt a transmitted frame:

1) The encrypted frame is parsed to re-construct the AAD and the nonce. The AAD is formed from the frame header.
2) The nonce is formed from the PN plus the A2 (transmit address) and Priority fields.
3) CCM uses the Temporal key, AAD, nonce, MIC, and encrypted payload to recover the plaintext data and to verify the MIC. If the MIC integrity check fails, CCM will not return the plaintext.
4) The received frame header and the plaintext data are concatenated to form the plaintext frame.
5) The PN in the frame is validated against the PN maintained for the session. If the PN received is not greater than the session PN, the frame is simply discarded; this check prevents replay attacks.

ROBUST SECURITY NETWORKS PRINCIPLES OF OPERATION

The IEEE 802.11 media access control (MAC) protocol supplies the functionality in WLANs that is required to provide reliable delivery of user data over the potentially noisy, unreliable wireless media. The IEEE 802.11 MAC protocol implements a frame exchange protocol in which the STA receiving a frame either returns an acknowledgement to the frame's source that the frame was received correctly, or notifies the source of an error. The frame exchange protocol is executed by each STA in the WLAN; every STA receives, decodes, and responds to information in the MAC header for every frame that it receives, with the exception of certain broadcast, multicast, and beacon frames.

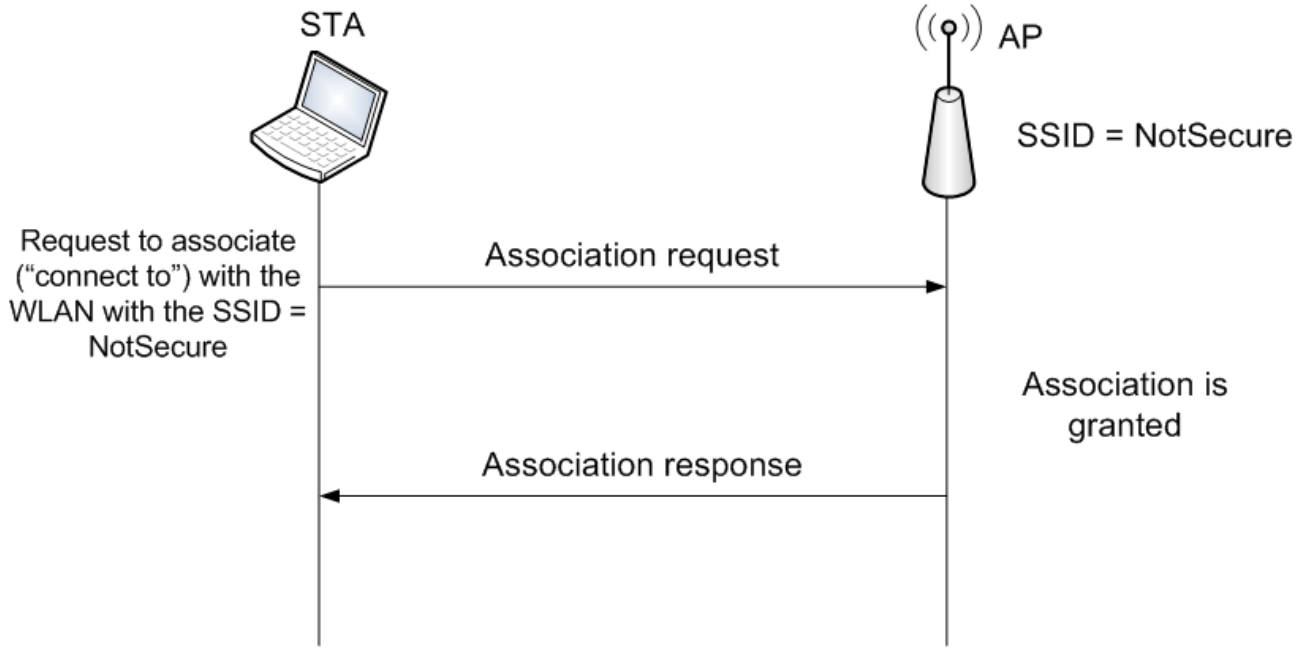

ROBUST SECURITY NETWORKS PRINCIPLES OF OPERATION

Typical two-frame flow for IEEE 802.11 WLAN communication that illustrates an Association Request and Response. First, the STA sends an Association Request frame to the AP, which is a request to connect to the WLAN with a Service Set Identifier (SSID) of “NotSecure”. The SSID is a text name assigned to the WLAN. The AP with the matching SSID then responds to the STA with either success or failure. If the response indicates success, the result is an association (not yet an RSNA) between the AP and STA. Association is a record-keeping procedure that allows the DS to keep track of STA location, so that frames from the DS are forwarded to the correct STAs.

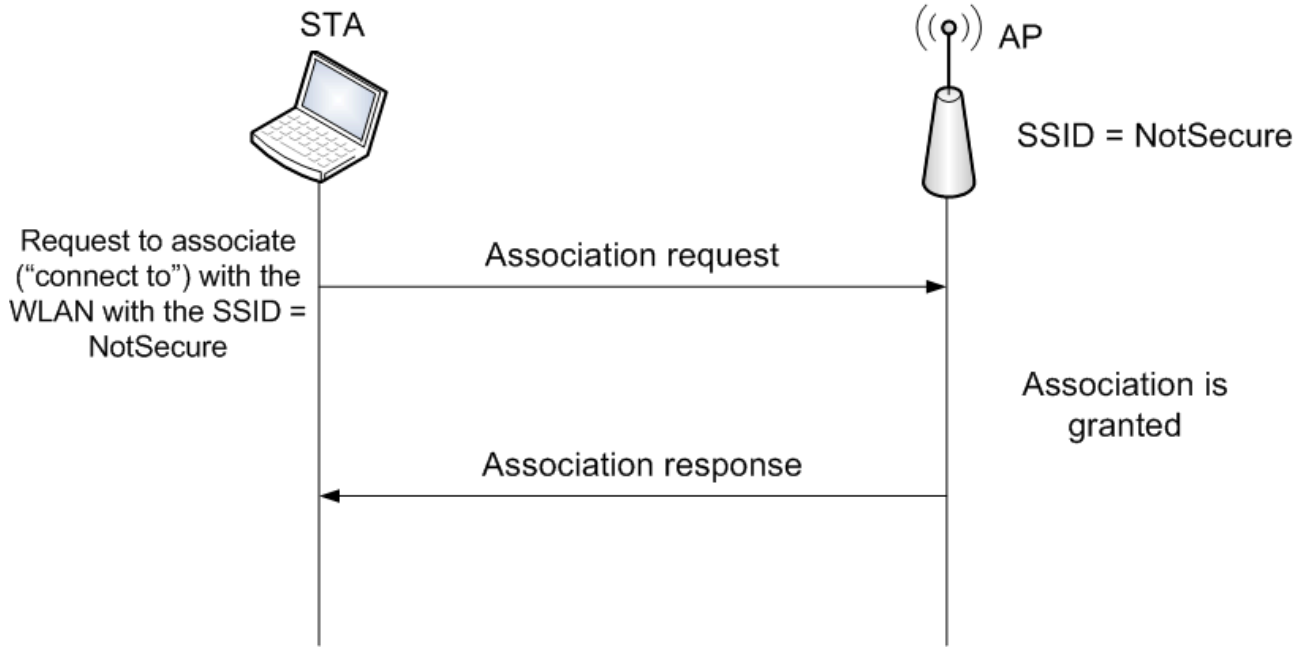

IEEE 802.11 FRAME TYPES

The IEEE 802.11 frame exchange protocol involves three types of frames, as follows:

- **Data Frame.** Data frames encapsulate packets from upper layer protocols, such as IP, which in turn might contain application data (e.g., e-mail, Web pages). Data frames allow for the delivery of the upper layer protocol packets to a STA or AP. RSNA security mechanisms protect these frames.
- **Management Frame.** Management frames carry the information necessary for managing the MAC. They provide the ability to perform management functions such as authenticating or associating (the wireless equivalent to connecting or registering). These frames can easily be forged, since IEEE 802.11i does not protect management frames. IEEE 802.11w is working on a standard to protect some management frames.
- **Control Frame.** Control frames are used for requesting and controlling access to the wireless media. An example of a control frame is the acknowledgement frame, which is used after data frames to ensure reliability. Its primary purpose is to alert the sender that the last frame was received correctly and there is no need to retransmit. This simple positive acknowledgement following each frame is expected, or the frame is considered lost. These frames can easily be forged, since IEEE 802.11i does not protect control frames.

IEEE 802.11 MANAGEMENT FRAMES

Frame Subtype	Description	Modified in IEEE 802.11i
Association Request	Used by a STA to request an association. The SSID is provided in this frame.	+
Association Response	Used to indicate the status (success or failure) of the Association Request.	
Reassociation Request	Used by a STA that has been associated with one BSS to request an association with another BSS with the same SSID. This frame includes the same information as the Association Request, with the addition of the current AP address.	+
Reassociation Response	Used to indicate the status (success or failure) of the Reassociation Request.	
Probe Request	Used by a STA to locate a WLAN quickly. This frame may be used to locate any WLAN or one with a particular SSID.	
Probe Response	Used by an AP to respond to a Probe Request. This frame contains essentially the same information as a beacon.	+
Beacon	Transmitted periodically by an AP to allow STAs to locate and identify a BSS.	+
Authentication	Used by an AP or STA to verify the identity of another STA.	
Deauthentication	Used by a STA to indicate termination of an authentication relationship.	
Disassociation	Used by a STA to indicate termination of an association.	
Announcement Traffic Indication Message (ATIM)	Used by a STA in an IBSS to notify other STAs that may have been operating in low power modes that it has data buffered and waiting to be delivered to the STA addressed in the ATIM frame.	

IEEE 802.11 MANAGEMENT FRAMES

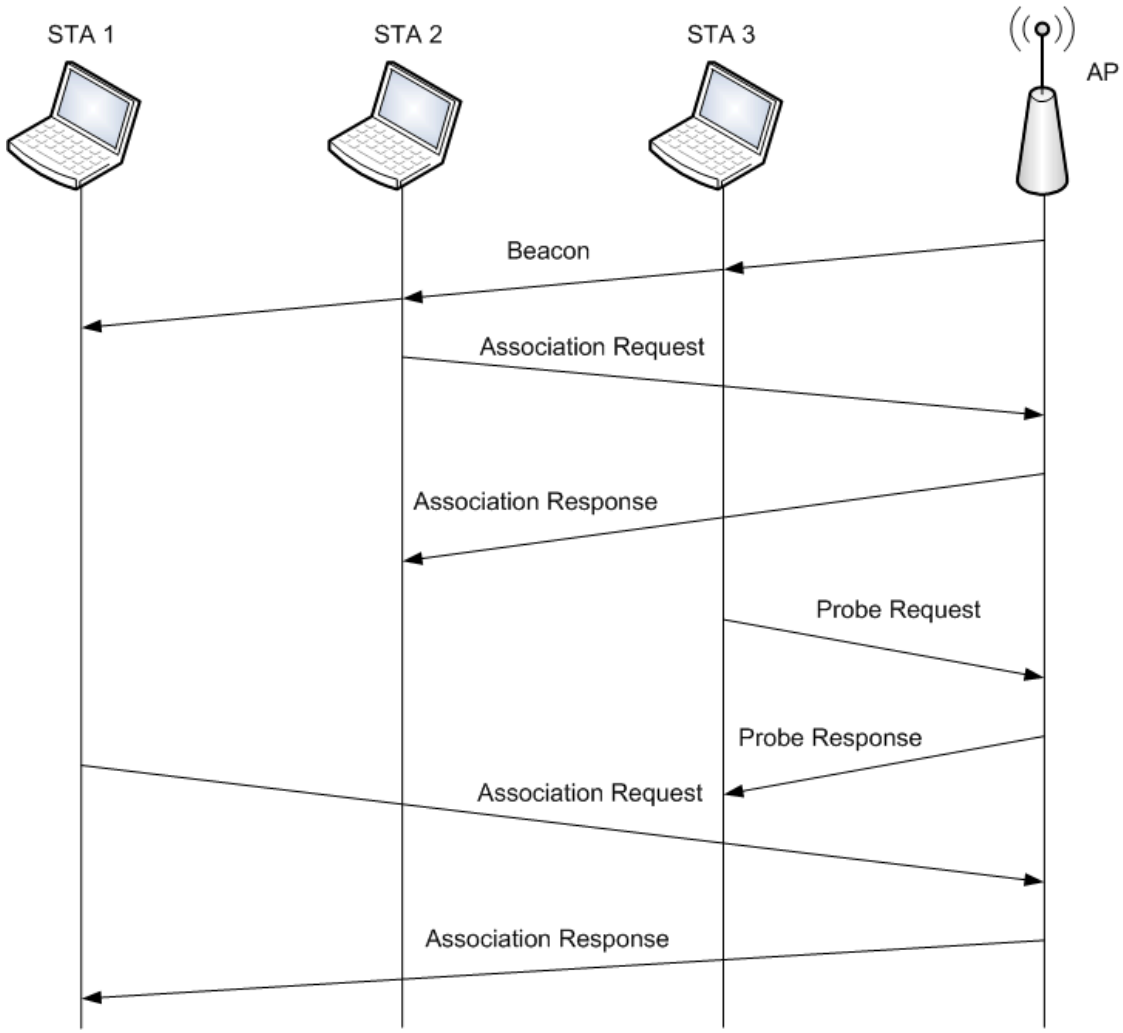

Figure illustrates the flow of management frames in a frame exchange between three STAs and an AP in a single infrastructure BSS. The AP periodically sends a Beacon frame, alerting all stations that the WLAN is operating in the area. After completing an IEEE 802.11 authentication exchange, the STAs are then able to connect to the AP by associating with it. STA1 and STA2 perform the Association Request-Response frame exchange with the AP to accomplish the STA registration for later frame delivery from the DS. STA3 joins the network after the beacon was transmitted. As a result, it sends a Probe Request frame – an active request for WLANs in the area – to determine the capabilities of the AP, and receives a Probe Response frame containing the requested information.

IEEE 802.11 DATA FRAME STRUCTURE

The data frame begins with a MAC header, which contains numerous fields for the transport of data in a WLAN. Most importantly, the header provides the MAC addresses of the source and destination, as well as the transmitter address, which identifies the address of the wireless network interface card that transmitted the frame onto the wireless medium, and the receiver address, which identifies the wireless station or group address that should process the frame.

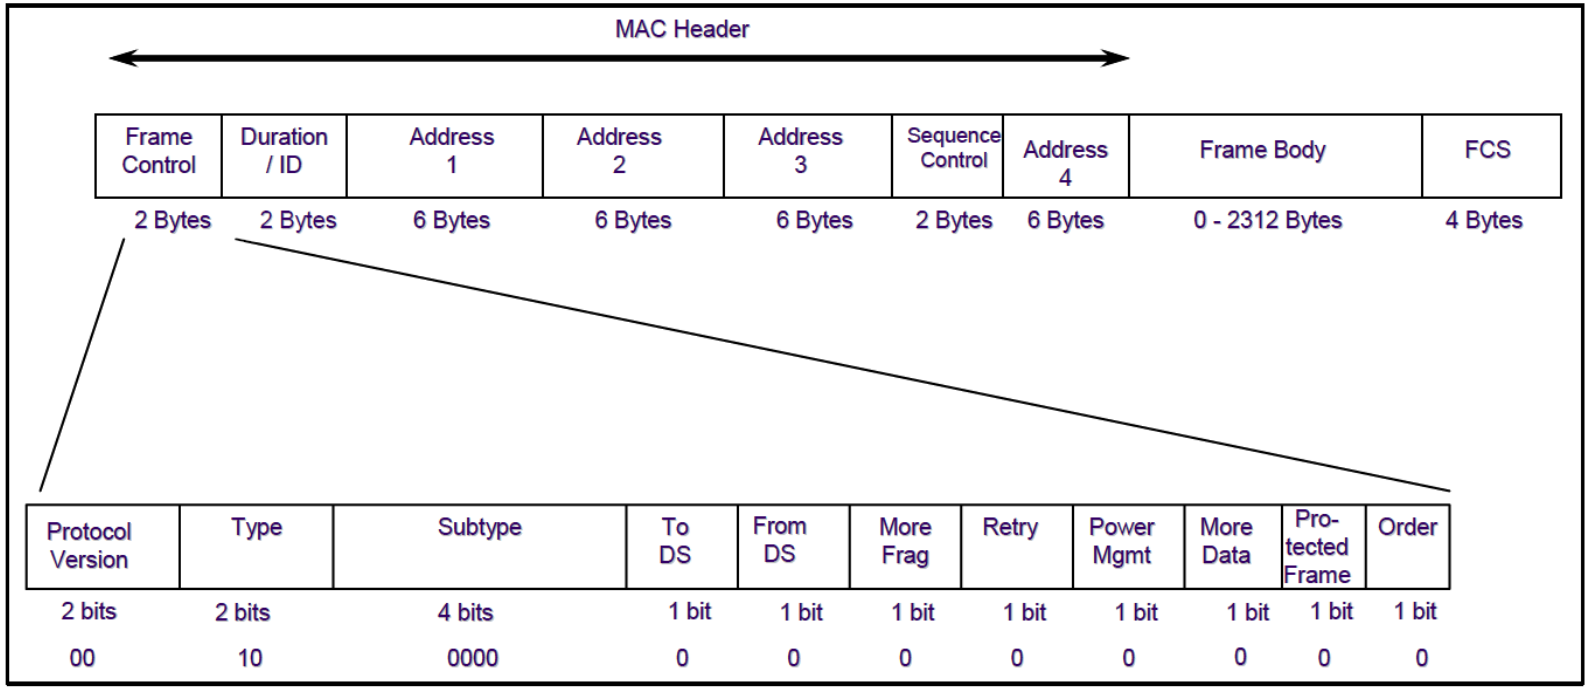

IEEE 802.11 DATA FRAME STRUCTURE

For example, when APs bridge wired LANs, a STA can send a message to a wired LAN end station connected to the AP, in which case the receiver address is the AP's address, and the destination address is the end station's address. Each STA and AP processes frames with a receiver address that matches its MAC address. Each AP also forwards frames to an attached LAN when a frame's destination address is different than the receiver address. In addition to the MAC addresses and other header fields, a data frame also contains a frame body, which is the encapsulated data from the higher layer protocol, and a **frame check sequence** (FCS), which is provided for error detection purposes.

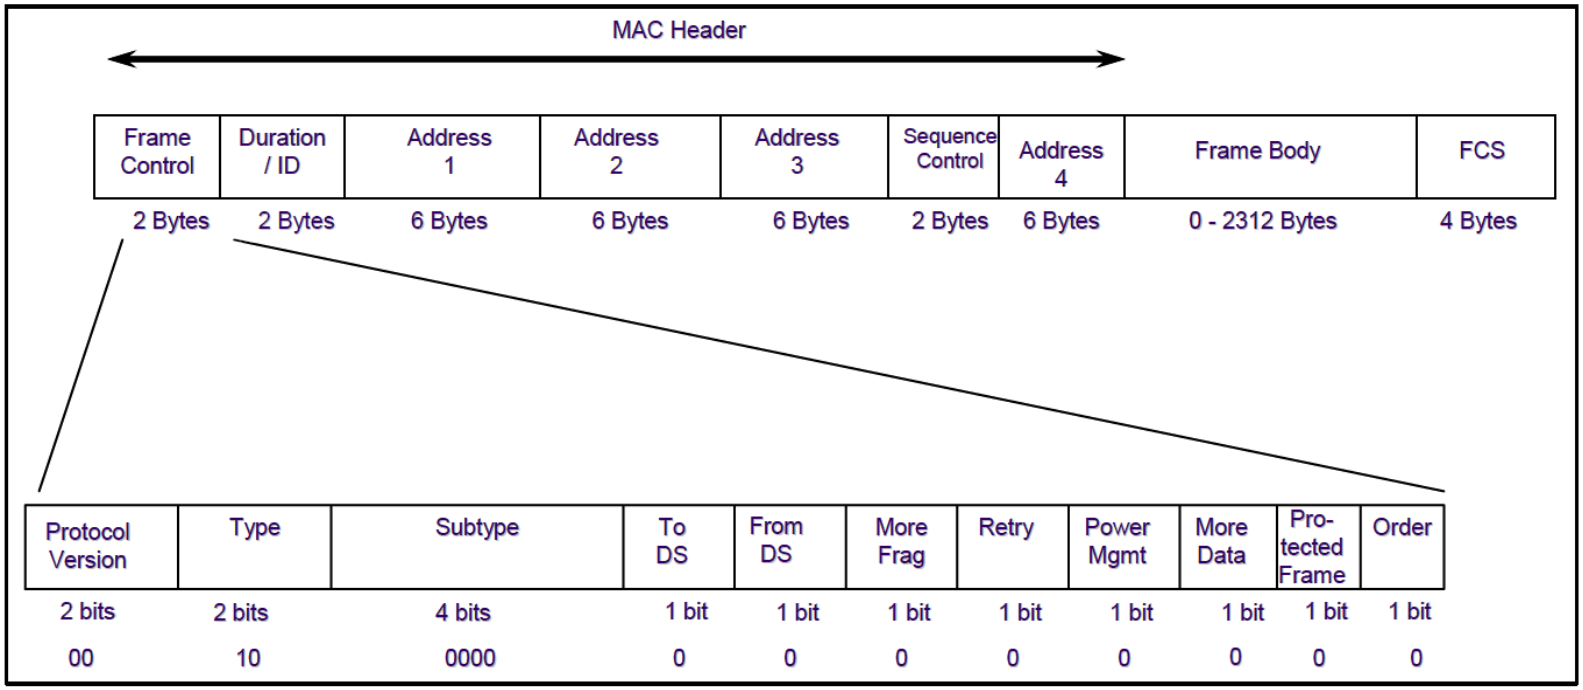

IEEE 802.11 DATA FRAME STRUCTURE

The following items briefly describe the frame body, FCS, and MAC header fields.

- **Frame Body.** This field, also called the Data field, holds a payload from a higher layer. The Frame Body field is variable in length, with a maximum size of 2312 octets.
- **FCS.** This field is used for error detection to detect random bit errors in the received frame. It contains the result of applying a 32-bit cyclic redundancy check (CRC-32) on the data. Because of this, the FCS is often called the CRC. The FCS calculation is performed on all data in the MAC header and frame body.
- **Frame Control Field.** This field defines a number of parameters for IEEE 802.11 operation. For example, it contains two bits used to identify the version of the IEEE 802.11 MAC. Another value within the field is the Protected Frame bit; if it is set to 1, the frame body is cryptographically protected using the negotiated ciphersuite (e.g., CCMP, TKIP, WEP). The Frame Control Field also indicates the frame type (e.g., management, control, data) and subtype (e.g., Association Request, Probe Response).
- **Duration/ID.** This field is used by a STA to retrieve frames buffered at an AP. The field identifies the remaining duration in the frame exchange between a STA and AP.
- **Sequence Control.** This field is used to allow a STA to identify received frames that are duplicates, and to assist it in reassembling fragmented frames.

IEEE 802.11 DATA FRAME STRUCTURE

– **Address Fields.** The MAC header for a data frame contains four distinct address fields, although in some cases not all fields contain relevant addresses. The address fields identify the original source address (SA) and final destination address (DA) in a frame exchange, as well as the receiver address (RA). Depending on the function of the frame, the address fields also identify either the transmitter address (TA) or the BSS identifier (BSSID), which is typically the address of the AP. The sequence of the addresses in the MAC header depends on two things: whether the transmitting station is in an IBSS or an infrastructure BSS, and whether the communicating stations are part of the DS.

Function	"To DS" Subfield	"From DS" Subfield	Address 1	Address 2	Address 3	Address 4
IBSS	0	0	RA = DA	SA	BSSID	N/A
Infrastructure BSS: From the AP	0	1	RA = DA	BSSID	SA	N/A
Infrastructure BSS: To the AP	1	0	RA = BSSID	SA	DA	N/A
Infrastructure BSS: Wireless DS (AP to AP)	1	1	RA	TA	DA	SA

PHASES OF IEEE 802.11 RSN OPERATION

The following items briefly describe each of the phases.

- **Phase 1: Discovery.** An AP uses Beacons and Probe Responses to advertise its IEEE 802.11i security policy. The STA uses these to identify an AP for a WLAN with which it wishes to communicate. The STA associates with the AP, which it uses to select the cipher suite and authentication mechanism when the Beacons and Probe Responses present a choice.
- **Phase 2: Authentication.** During this phase, the STA and AS prove their identities to each other. The AP blocks non-authentication traffic between the STA and AS until the authentication transaction is successful. The AP does not participate in the authentication transaction other than forwarding traffic between the STA and AS.
- **Phase 3: Key Generation and Distribution.** The AP and the STA perform several operations that cause cryptographic keys to be generated and placed on the AP and the STA. Frames are exchanged between the AP and STA only.
- **Phase 4: Protected Data Transfer.** Frames are exchanged between the STA and the end station through the AP. As denoted by the shading and the lock and key, secure data transfer occurs between the STA and the AP only; security is not provided end-to-end.

PHASES OF IEEE 802.11 RSN OPERATION

– **Phase 5: Connection Termination.** The AP and STA exchange frames. During this phase, the secure connection is torn down and the connection is restored to the original state.

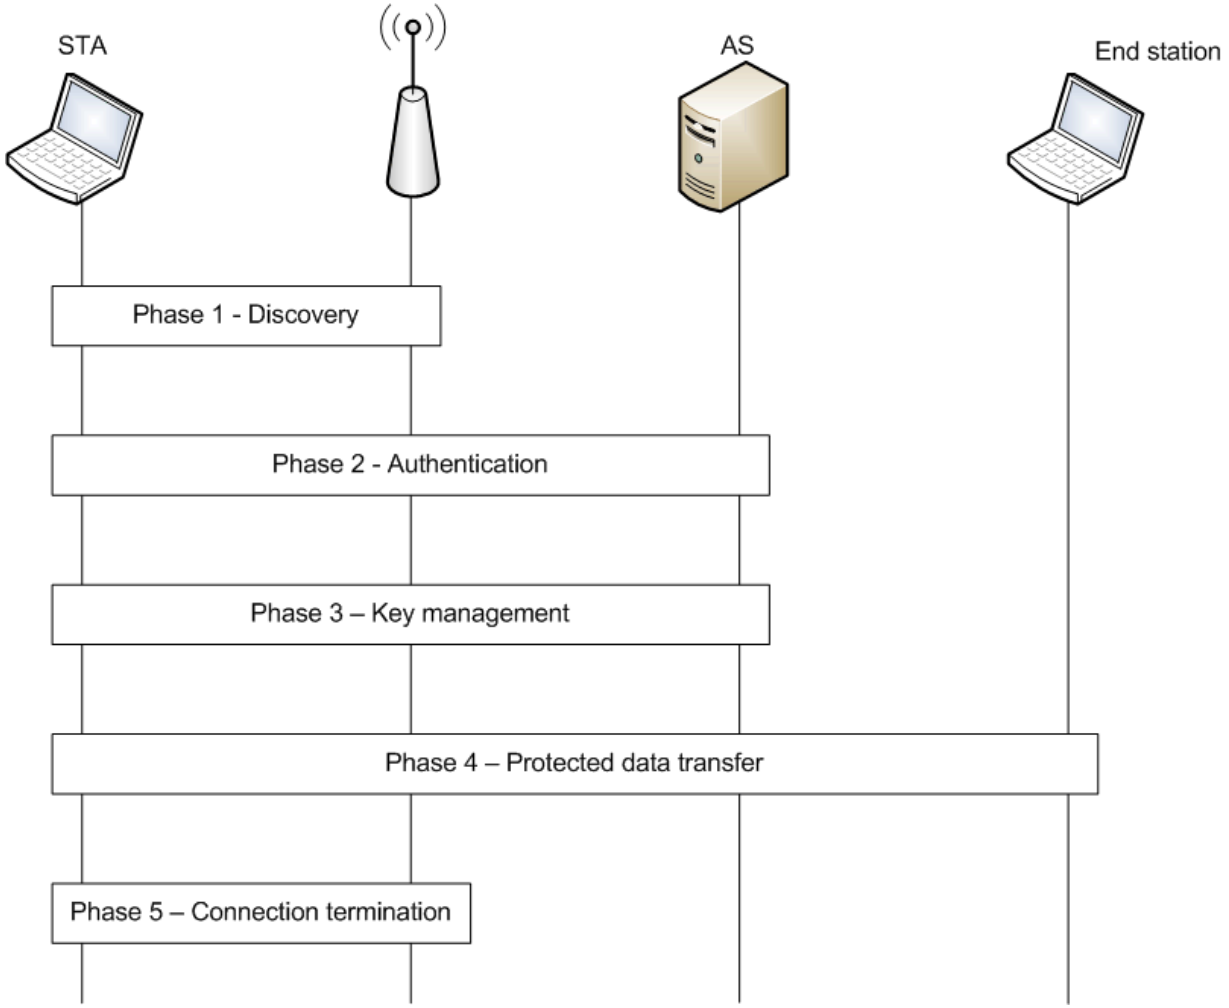

DISCOVERY PHASE

The discovery phase is the first phase in the process to establish RSNAs. During this phase, STAs discover the existence of a network with which to communicate. STAs locate and identify APs through the APs' periodic transmission of Beacon frames. A Beacon frame contains a timestamp, beacon interval, and capability information, which includes supported data rates and the SSID.

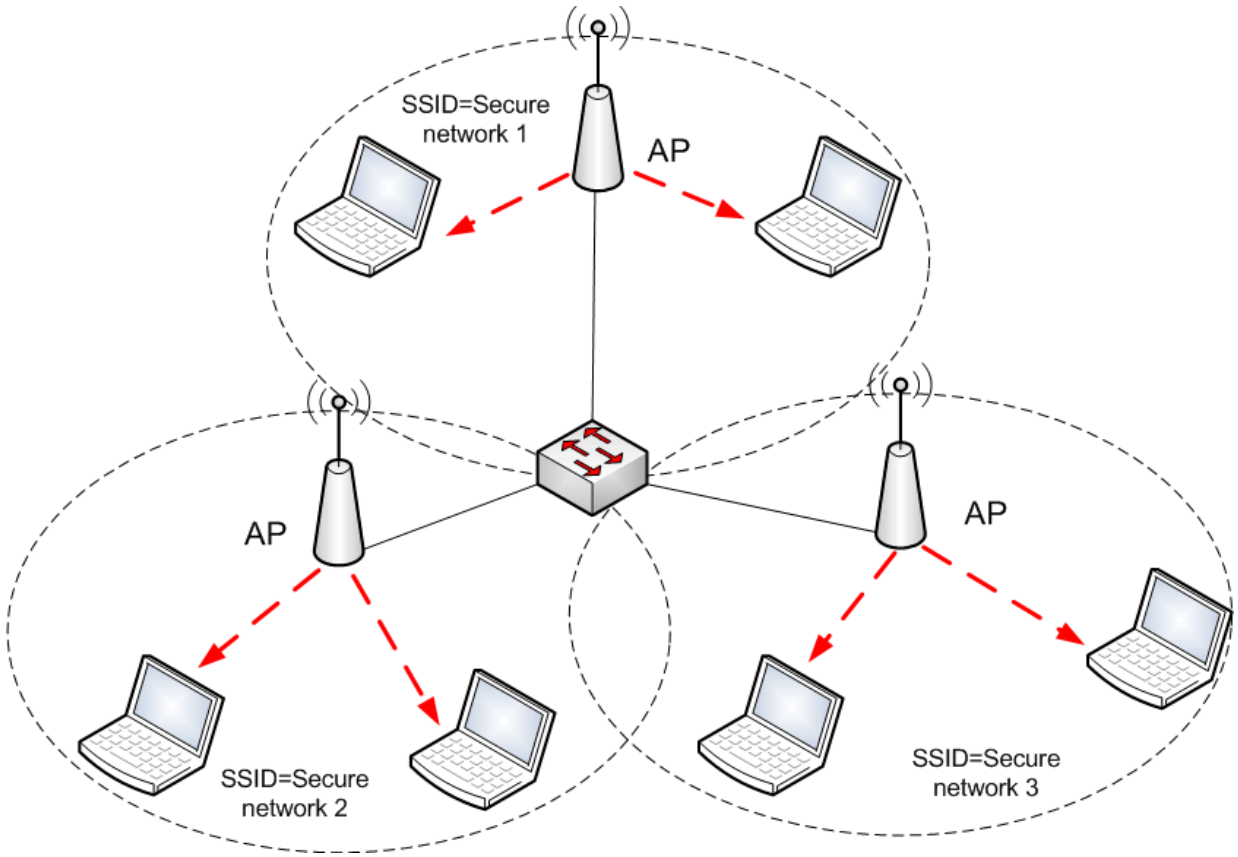

DISCOVERY PHASE

During the discovery phase, STAs and APs negotiate several things, including the SSID, supported data rates, and other technical operating parameters related to reliable communication, as well as a security policy. In general, 802.11i does not support extensive negotiation The AP describes the options that it supports, and only clients that are configured for compatible options will attempt to connect. Many APs and STAs can only store a single configuration at a time.

During the discovery phase, STAs and APs negotiate the following key security capabilities:

- Confidentiality and integrity protocols for protecting unicast traffic;
- Authentication method for mutual authentication of the AP and AS;
- Cryptographic key management approach;
- Pre-authentication capabilities.

Confidentiality and integrity protocols for protecting multicast/broadcast traffic are dictated by the AP, since all STAs in a multicast group must use the same cipher suite.

The possible cipher suites allowed by the IEEE 802.11i amendment are as follows:

- WEP, with either a 40-bit or 104-bit key;
- TKIP;
- CCMP which is the default choice according to the IEEE 802.11i standard;
- Vendor-specific methods (to allow for flexibility and expansion).

DISCOVERY PHASE

The possible AKM suites are as follows:

- **Mutual authentication and key management over IEEE 802.1X** or **using pairwise master key security association** (PMKSA) caching. Authentication is accomplished with an EAP method.
- **Pre-shared key.** No explicit authentication transaction takes place. If the PSK is unique for each STA, the STA and AP effectively authenticate each other by holding an identical pre-shared key, without which the data confidentiality and integrity services could not function properly.
- **Proprietary suites** developed by vendors; this allows for flexibility and expansion.

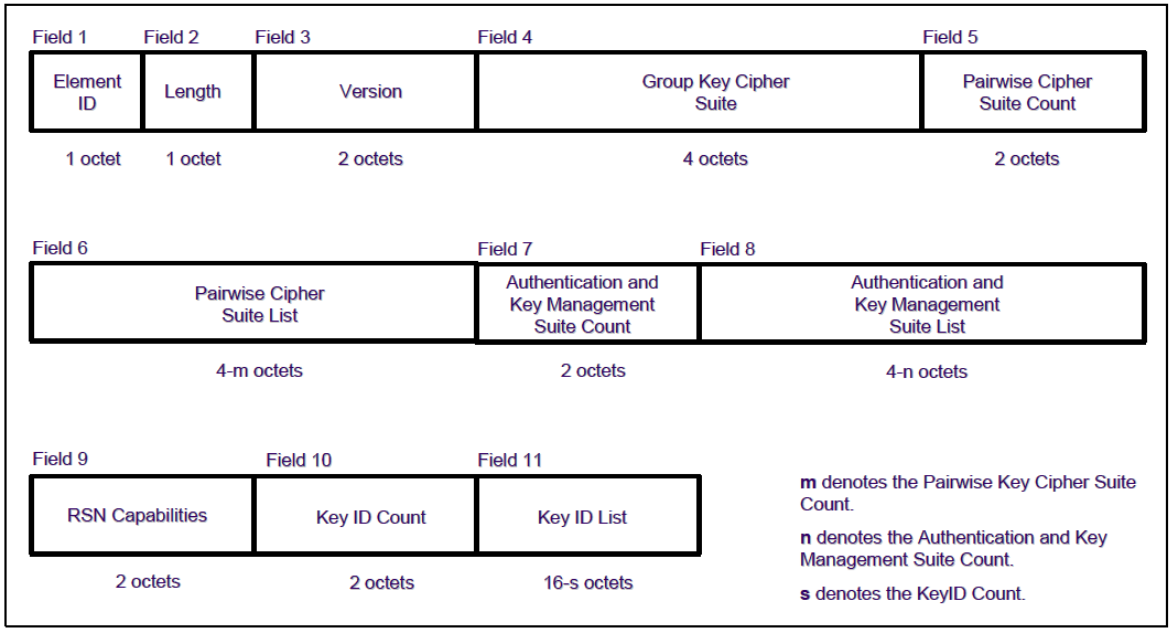

DISCOVERY PHASE FRAME FLOWS

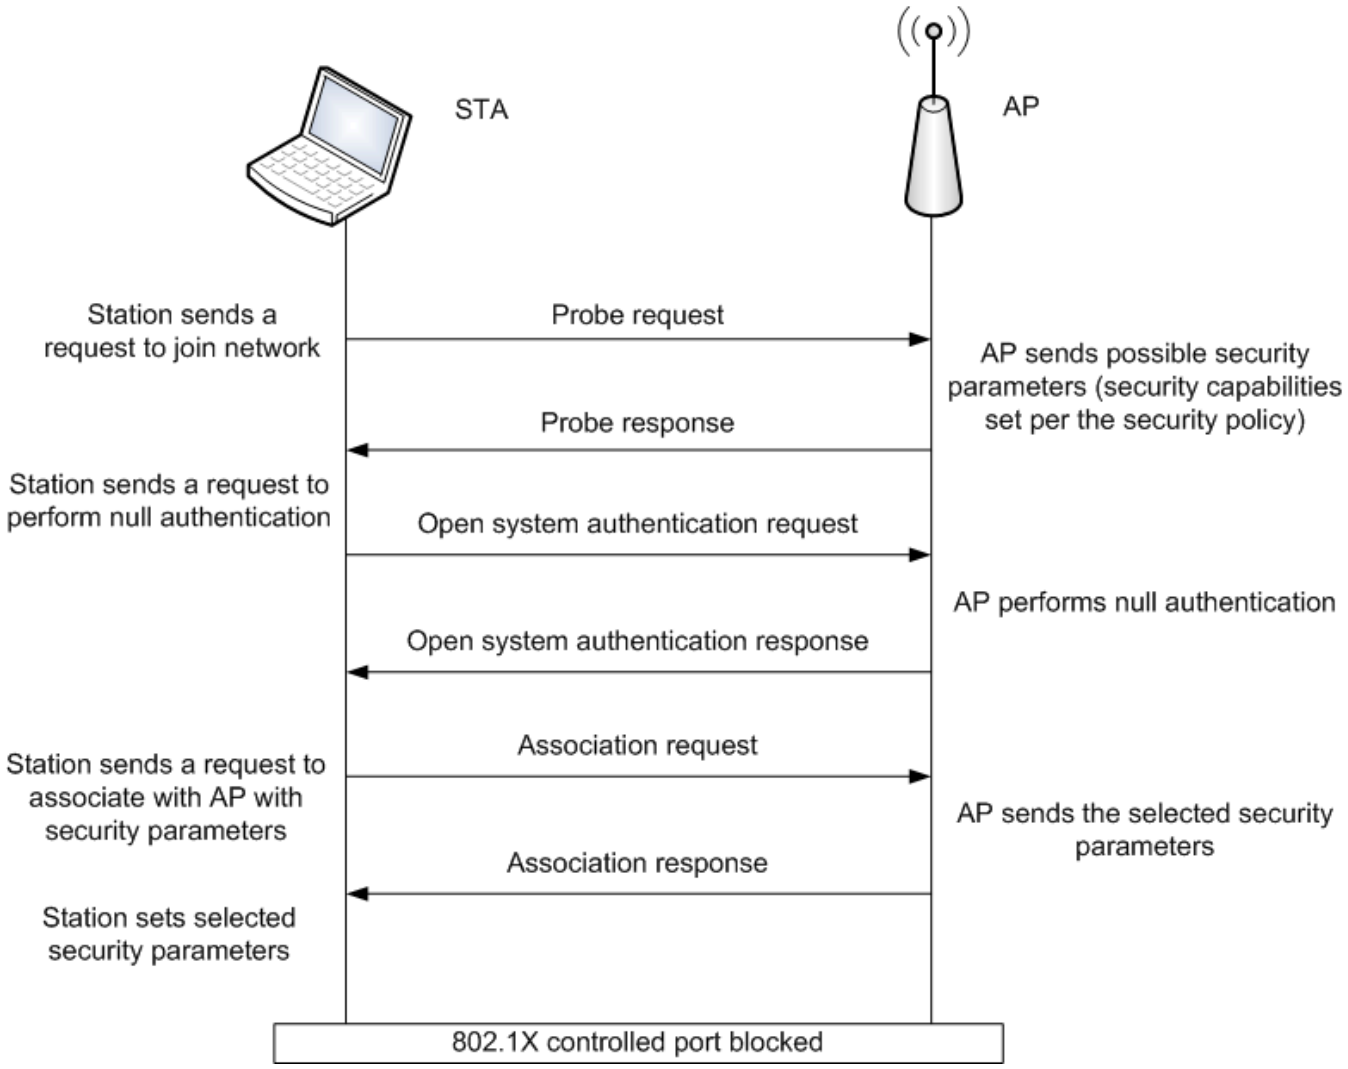

In this scenario, the STA sends a Probe Request frame to locate an AP in the area. The AP responds with its capabilities in the RSNIE field of the Probe Response frame; this includes all of its enabled encryption and authentication capabilities. When the STA receives the Probe Response frame, it performs open system authentication – null authentication – with the AP. The purpose of this frame sequence, which provides no security, is simply to maintain backward compatibility with the IEEE 802.11 state machine, as implemented in existing IEEE 802.11 hardware.

DISCOVERY PHASE FRAME FLOWS

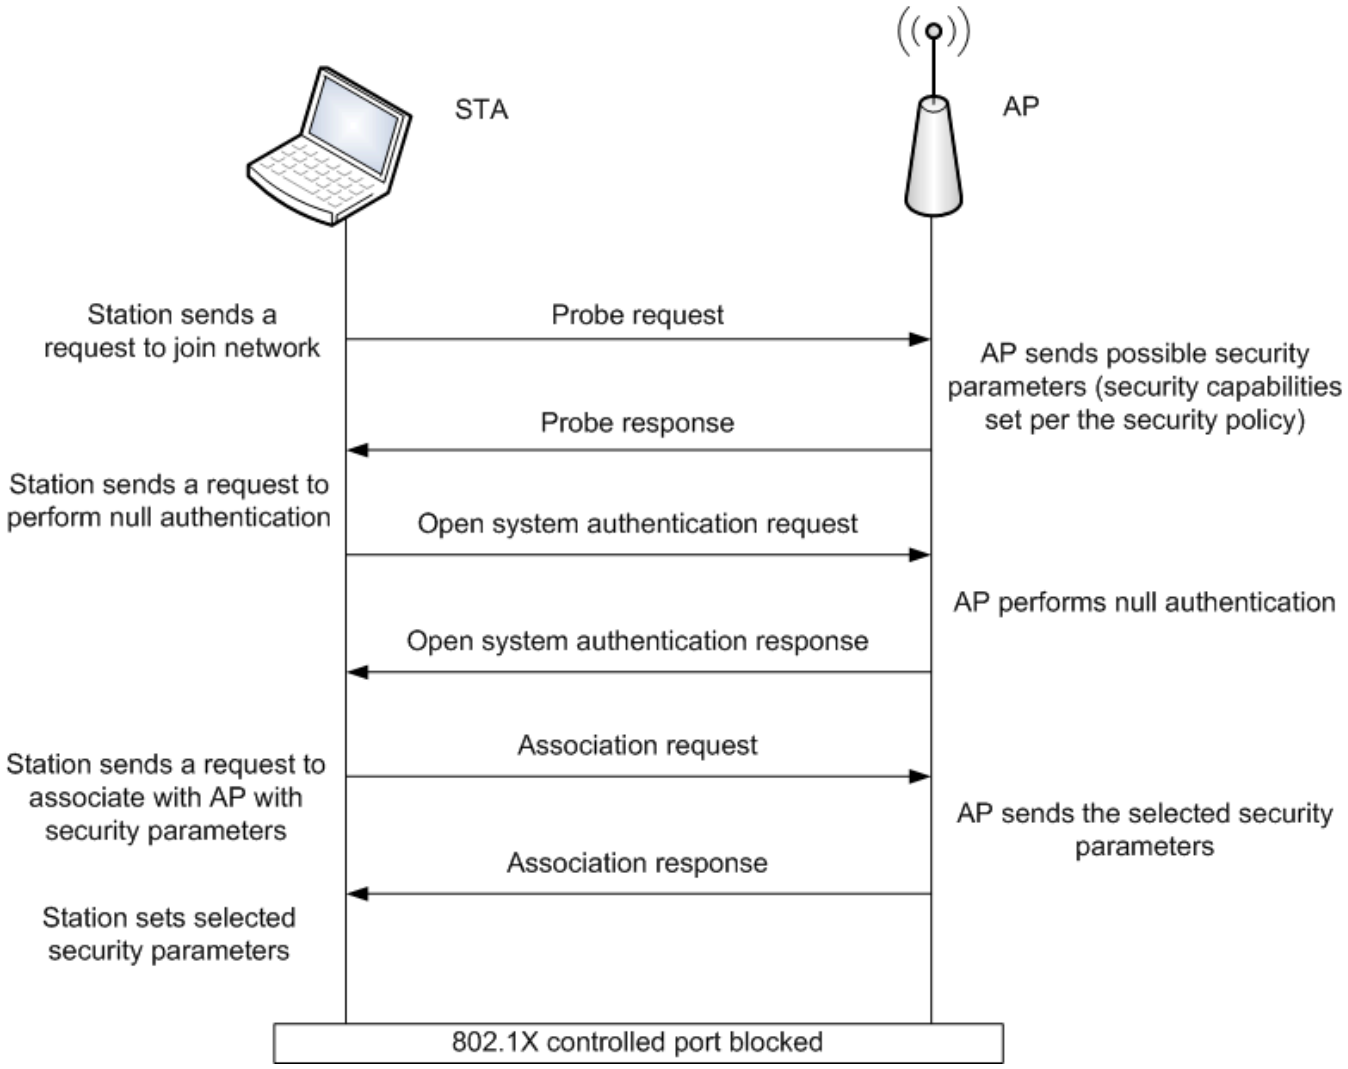

Following the authentication frame exchange, the STA then sends an Association Request frame to the AP. In this frame, the STA specifies one set of matching capabilities (one authentication and key management suite, one pair-wise cipher suite, and one group key cipher suite) from among those advertised by the AP. If there is no match in capabilities between the AP and the STA, the AP refuses the Associate Request. The STA blocks it too, in case it has associated with a rogue AP or someone is inserting frames illicitly on its channel.

DISCOVERY PHASE FRAME FLOWS

AP advertises its capabilities to the STA in a periodic Beacon frame, as set by the WLAN security policy. In this example, the AP informs the STA that it is capable of performing IEEE 802.1X authentication, CCMP for unicast traffic protection, and CCMP for broadcast traffic protection. The STA also indicates that it is capable of performing IEEE 802.1X authentication and CCMP for both types of traffic. At this point, the AP accepts the IEEE 802.1X and CCMP request, then self-configures these capabilities. It confirms its declaration with an Association Response frame, indicating the completion of security policy negotiation and the end of the discovery phase.

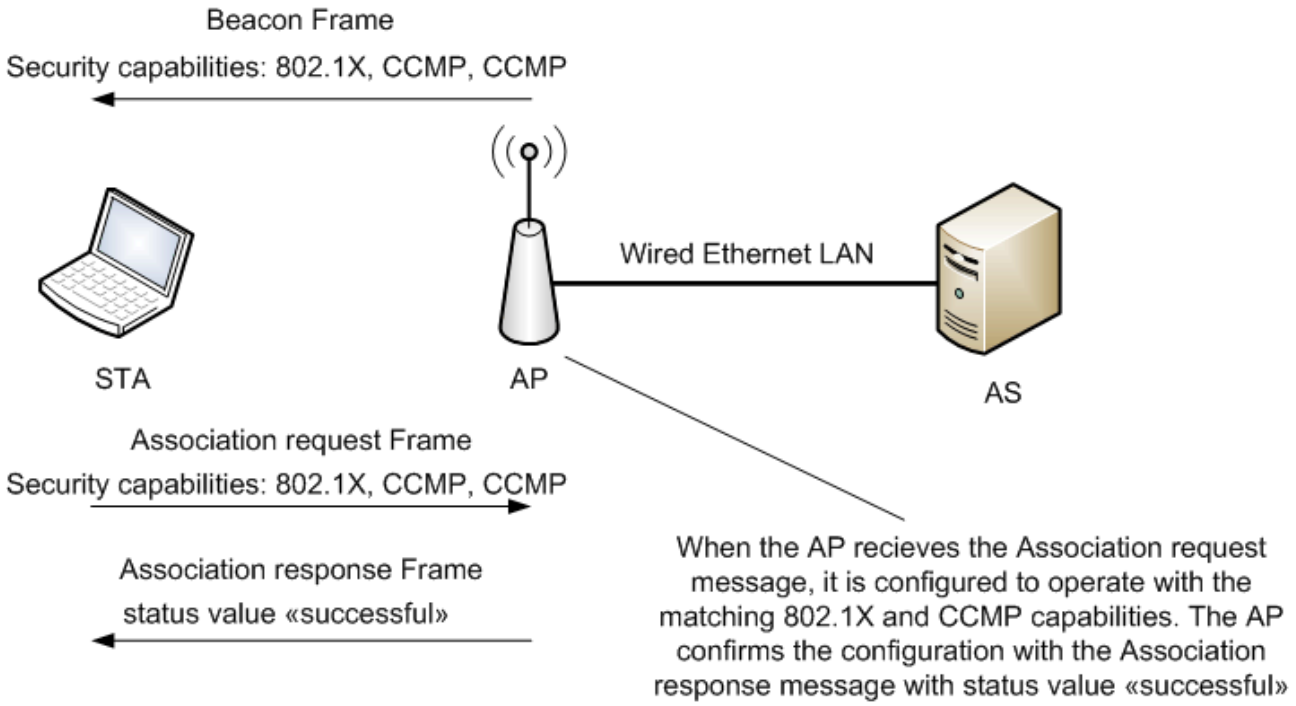

DISCOVERY PHASE FRAME FLOWS

During the discovery phase, a STA may decline to communicate with an AP or another STA that fails to disclose any of the following:

- Security policy in the Beacon or Probe Response frames
- Authorized SSID
- Authorized encryption and authentication cipher suites.

The IEEE 802.11 standard does not specify the manner in which these conditions are handled. This remains undefined and is left as a design choice to the RSN technology manufacturer. Manufacturers typically design their products so that these conditions are configurable via policy.

AUTHENTICATION PHASE

Upon successful completion of the discovery phase, the STA and AP enter the second phase in the establishment of an RSNA: the **authentication phase**. This phase provides the means for a STA to prove its identity to the WLAN. This security service is critical for preventing unauthorized access to network resources. In an infrastructure WLAN, authentication provides protection against unauthorized users in the DS, since the AP is the entry point into the ESS. Improper authentication can undermine all security measures in an enterprise. Mutual authentication also allows the WLAN to prove its identity to the STA, which allows the STA to validate positively that it is communicating with a legitimate WLAN, as opposed to an unauthorized or "rogue" WLAN.

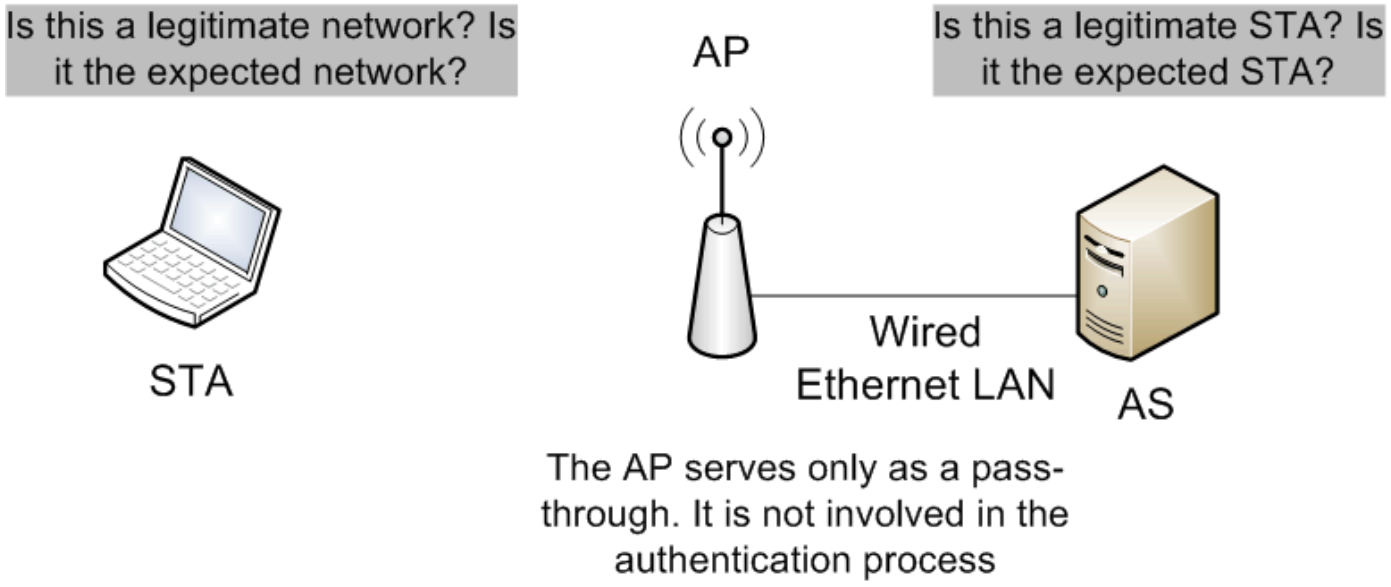

THE IEEE 802.1X FRAMEWORK: PORT-BASED ACCESS CONTROL

The IEEE 802.11 standard uses the IEEE 802.1X standard to provide mutual authentication between STAs and ASs. IEEE 802.1X is a general-purpose, extensible framework for authenticating users. The actual authentication mechanism incorporated into the framework is implemented by the STA and the AS using EAP. EAP provides a framework that allows the use of multiple methods for achieving authentication, including static passwords, dynamic passwords (e.g., one-time passwords, token generators), and public key cryptography certificates (on the AS only or on both the AS and STAs). Dozens of standard and proprietary EAP methods exist.

IEEE 802.1X authentication has three main components: a **client** (also known as a supplicant), an **authenticator**, and an **AS**. The authenticator simply passes authentication traffic between the client and AS. IEEE 802.1X controls the flow of data between the DS and STAs by use of a controlled/uncontrolled port model. EAP authentication occurs through the IEEE 802.1X **uncontrolled port** on the authenticator; non-EAP data frames are passed or blocked via the IEEE 802.1X **controlled port**, depending upon the success or failure of IEEE 802.1X authentication (which includes EAP). This model is known as port- based access control. Using this concept, IEEE 802.1X achieves the objective of blocking access for unauthorized parties in an IEEE 802.11 WLAN.

THE IEEE 802.1X FRAMEWORK: PORT-BASED ACCESS CONTROL

The authentication message flows between the client and the authenticator typically use the EAP over LAN (EAPOL) protocol. RADIUS is the protocol most commonly used to transport EAP messages between the authenticator and the AS. The steps in a typical successful IEEE 802.1X authentication exchange when RADIUS is used to support authentication-related traffic on the DS are as follows:

1. The supplicant (client) may start the exchange with an optional EAPOL-Start message.
2. The EAP exchange begins with the authenticator issuing an EAP-Request/Identity frame to the supplicant.
3. The supplicant replies with an EAP-Response/Identity frame, which the AP receives over the uncontrolled port. The packet is then encapsulated in RADIUS over EAP and passed on to the RADIUS server as a RADIUS-Access-Request packet.
4. The AAA server replies with a RADIUS-Access-Challenge packet, which is passed on to the supplicant as an EAP-Request. This request is of the appropriate authentication type and contains relevant challenge information.
5. The supplicant formulates an EAP-Response message and sends it to the authenticator. The response is translated by the authenticator into a Radius-Access-Request, with the response to the challenge as a data field. Steps 4 and 5 may be repeated multiple times, depending on the EAP method in use. For TLS tunneling methods, it is common for authentication to require 10-20 round trips.

THE IEEE 802.1X FRAMEWORK: PORT-BASED ACCESS CONTROL

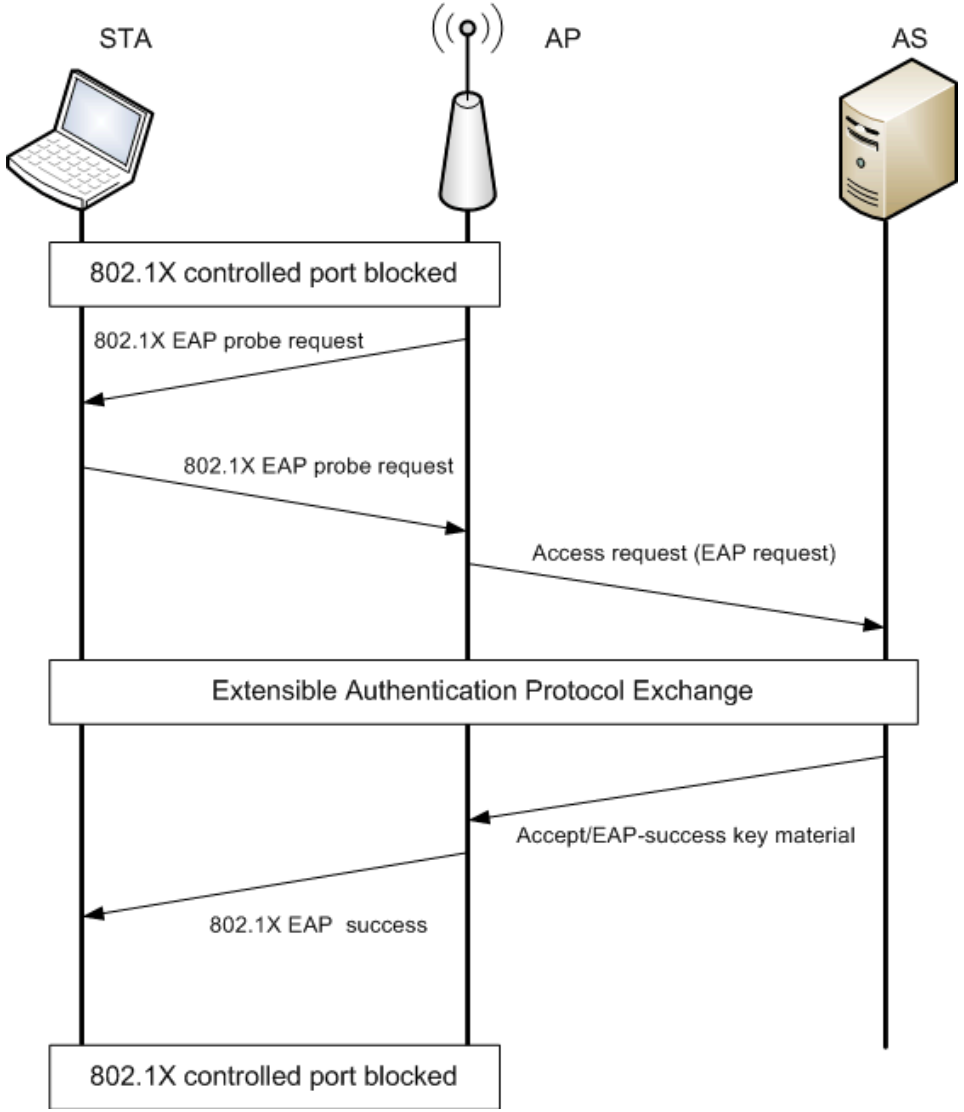

6. The AAA server grants access with a Radius-Access-Accept packet. The authenticator issues an EAP-Success frame. (Some protocols require confirmation of the EAP success inside the TLS tunnel for authenticity validation.) The controlled port is authorized, and the user may begin to access the network.

7. During the termination phase, when the supplicant is finished accessing the network, it may send an optional EAPOL-Logoff message to restore the controlled port to an unauthorized state.

After the seven-step authentication process has been completed, the AAA key is installed in the STA and the AS.

AUTHENTICATION WITH THE PSK

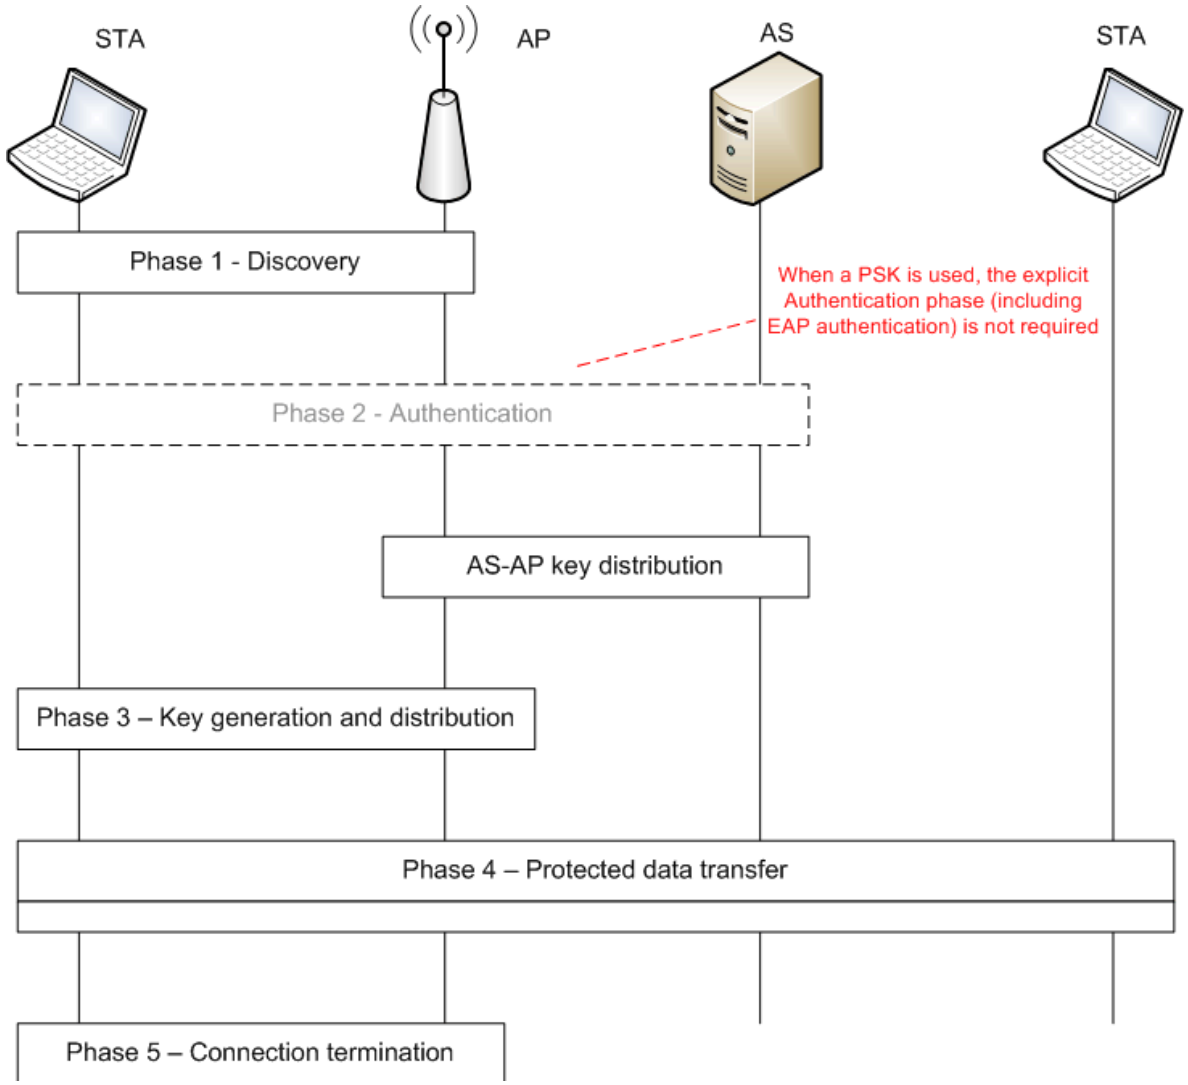

Typically, the authentication phase provides mutual authentication of a STA and an AS in an RSNA and delivers the Master Session Key to the AP and, sometimes, to the STA. However, in an RSNA that has negotiated the PSK AKM during the discovery phase, the authentication phase is not required, because the shared key has already been distributed and installed in an out-of-band manner that has implicitly provided authentication.

AS TO AP CONNECTIONS

Although the details of the communications interface between the AS and the AP are outside the scope of the IEEE 802.11i amendment, the amendment does contain several requirements for the interface to ensure that the security of an RSN is not compromised. Specifically, the communication link between the AS and AP must provide the following:

- Robust, mutual authentication between the AS and AP
- An end-to-end channel between the AS and the authenticator for the mutual authentication
- The ability to transfer the cryptographic key generated by the AS to the AP securely. The AS to AP communication must provide confidentiality and integrity, and the AS must prevent key compromise during storage.

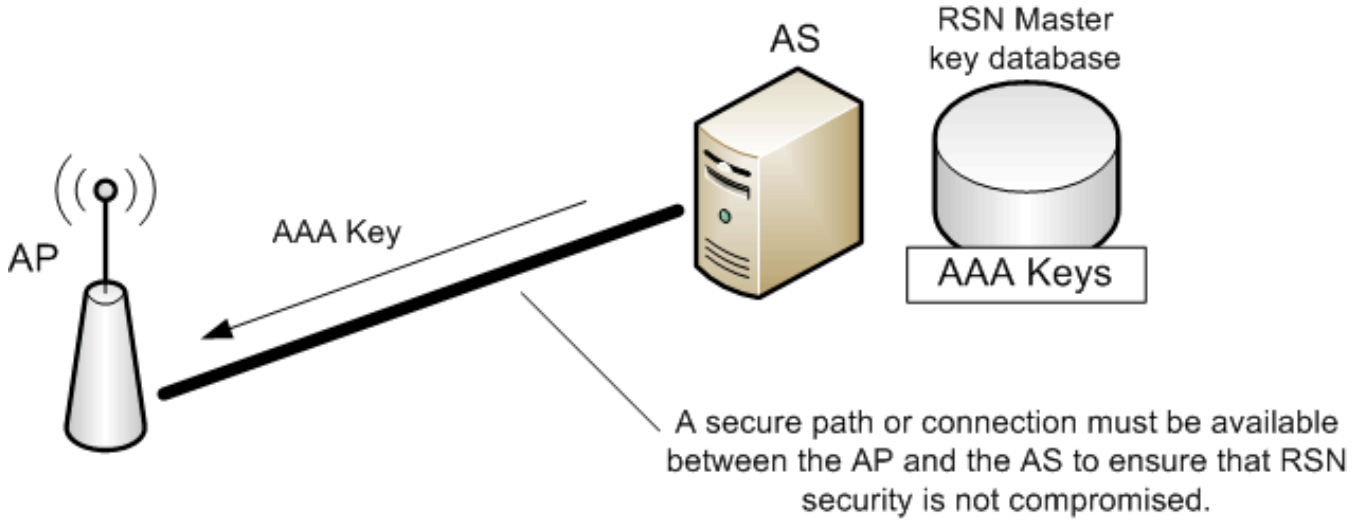

AS TO AP CONNECTIONS

The AS is a critical component of overall RSN security. The IEEE 802.11 standard assumes the following with respect to the AS:

- It does not expose or compromise the PMK (a subset of the AAA key) to other entities besides the AP.
- It does not masquerade as a STA to the AP.
- It does not masquerade as an AP to the STA.

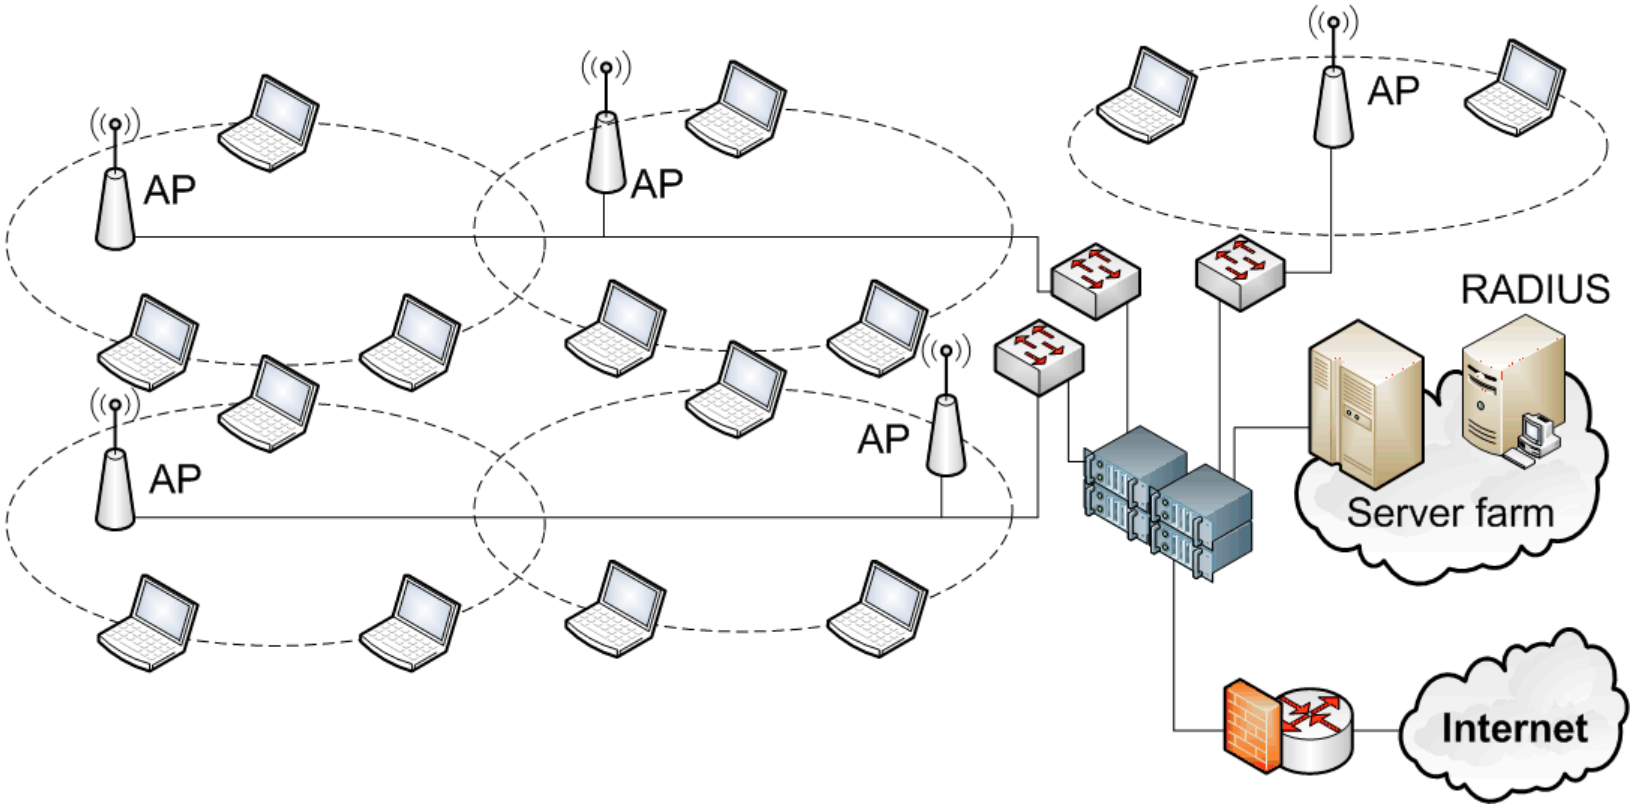

KEY GENERATION AND DISTRIBUTION

Following the successful completion of the authentication phase, the STA and AP perform a series of functions that position cryptographic keys in both entities. This phase is called the **key generation** and **distribution** (KGD) **phase**. It provides the final step in authentication and allows the STA and AP to derive keys that make secure data transfer possible. The KGD phase has several purposes, including the following:

- Confirming the existence of the Pairwise Master Key (PMK)
- Ensuring the security association keys are new
- Deriving and synchronizing the installation of traffic encryption keys (temporal keys) in the AP and STA
- Distributing a group key for multicast and broadcast traffic protection
- Confirming the cipher suite selection.

KEY GENERATION AND DISTRIBUTION

The KGD phase includes two types of handshakes: a 4-Way Handshake and a Group Handshake. The Group Handshake is necessary only when STAs participate in multicast or broadcast traffic. Both types of handshakes employ the following fundamental security features:

- Message integrity checking, to protect against tampering and to validate the source of traffic;
- Message encryption, to protect against unauthorized disclosure of data.

The confidentiality and integrity algorithms used for both handshakes are configurable to either of the following:

- RC4 Encryption with HMAC-MD5. RC4 is the well-known stream cipher that forms the basis of WEP. RC4 uses the 128-bit EAPOL-KEK derived from the PTK using the PRF.
- AES Key Wrap with HMAC-SHA-1-128. The AES Key Wrap was designed specifically to encrypt keying material (cryptographic keys). The key wrap parses data into n blocks of 64-bits and “wraps” (encrypts) the key contents. The key wrap uses the AES codebook mode along with the EAPOL-KEK derived from the PTK.

Both RC4 and the AES Key Wrap use the HMAC along with the EAPOL-KCK derived from the PTK using the PRF to provide integrity during the 4-Way Handshake.

4-WAY HANDSHAKE

The KGD phase begins with the 4-Way Handshake. During the handshake, four frames are exchanged between the STA and the AP. To generate data for the frames and verify data received in frames, both the STA and the AP perform several computations. At the successful conclusion of the 4-Way Handshake, the AP and STA have been mutually authenticated. At that point, the IEEE 802.1X controlled ports are opened to allow the flow of frames for data traffic.

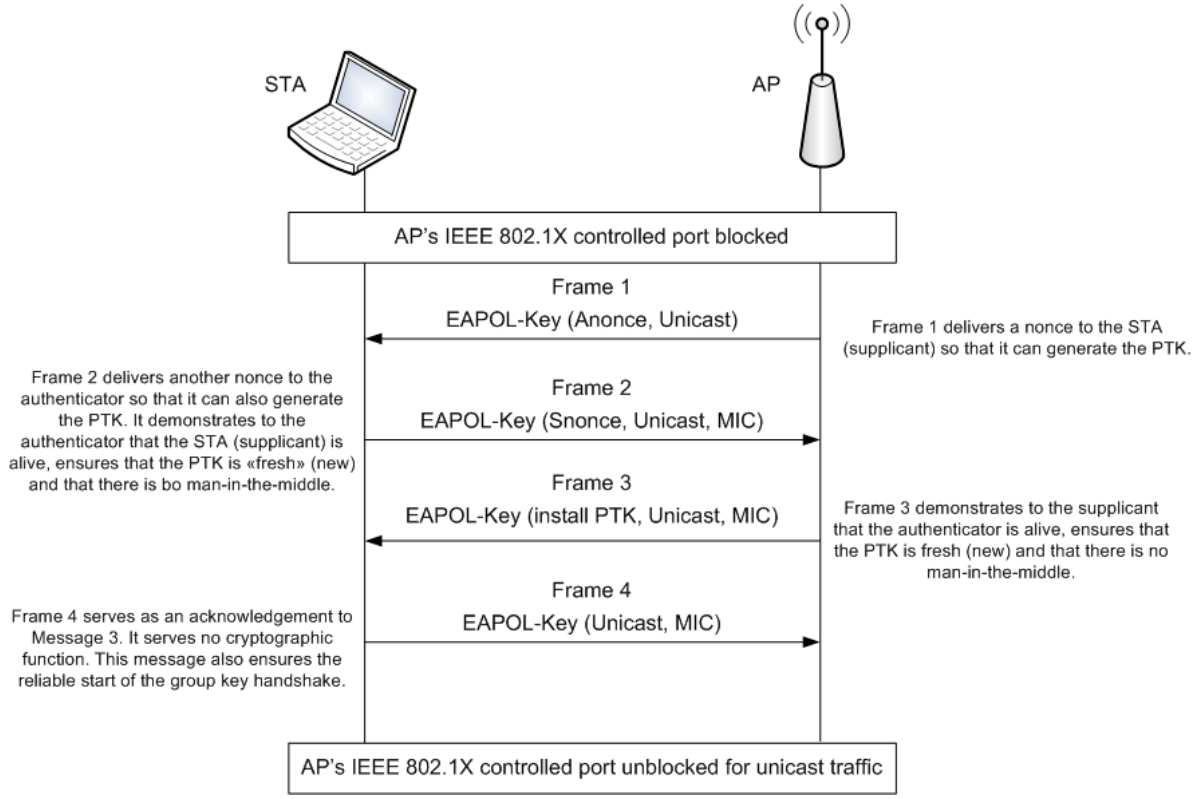

GROUP KEY HANDSHAKE

The Group Key Handshake is used by the AP to send a new GTK to a STA. It may occur immediately after the 4-Way Handshake or upon STA initiation. It is necessary to support multicast or broadcast traffic.

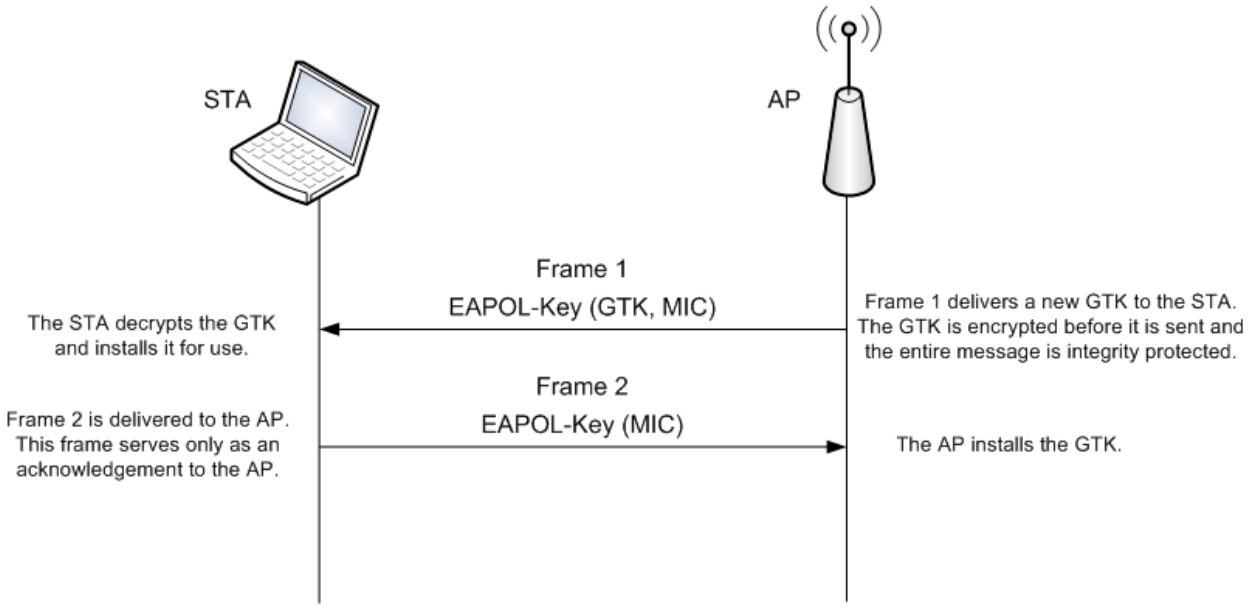

After the Group Key Handshake is complete, the AP and STA are ready for operation. The Group Key Handshake also may be used to distribute subsequent GTKs.

GROUP KEY HANDSHAKE

The AP can use the handshake to update the GTK in STAs under the following conditions:

- on disassociation or deauthentication of a STA;
- upon occurrence of an event on the STA that triggers an update, such as a configuration change in the STA's local security policy.

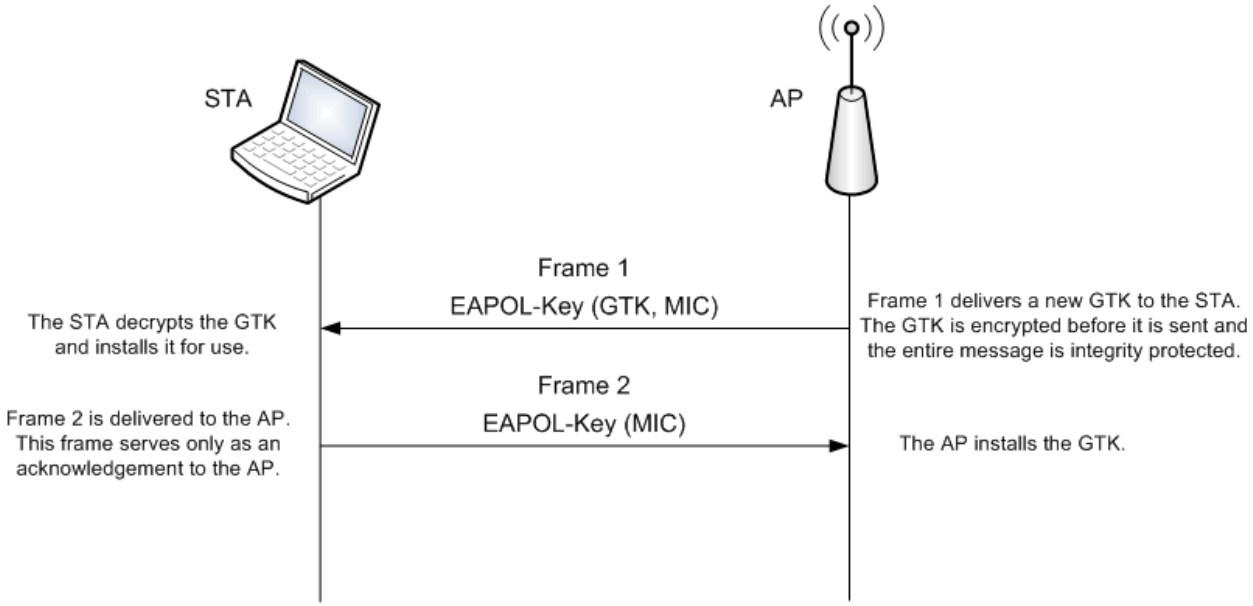

The GMK used to derive the GTK may be updated in the AP at a time interval configured into the system. Periodic updating of the GMK may be included in the security policy for the WLAN. Organizations should update the GMK to prevent exposure of subsequent traffic between STAs and the AP, if the GMK is ever compromised.

PROTECTED DATA EXCHANGE

The fourth phase in the operation of an RSN is the protected data exchange phase. Before this phase, the AP and STA have already done the following:

- Become associated and negotiated a security policy (discovery phase)
- Mutually authenticated using EAP and derived a Master Session Key using the uncontrolled IEEE 802.1X port, or implicitly authenticated through a previously installed preshared key (authentication phase)
- Generated, distributed and confirmed the session keys through the 4-Way Handshake (KGD phase)
- Derived a pairwise transient key and unblocked the IEEE 802.1X ports (KGD phase).

These actions have prepared the AP and STA to communicate securely. During the protected data exchange phase, the AP and STA may now share data securely. The traffic between the AP and STA is protected using the data confidentiality and integrity algorithms chosen during the discovery phase. IEEE 802.11i supports three methods of data transfer: unicast, multicast, and broadcast.

PROTECTED DATA EXCHANGE

For RSNs, unicast (also called “directed”) is the type of data transfer used most often during the protected data exchange phase. Unicast data transfer can occur when a unique association exists between the AP and the STA and a pairwise transient key is used for the protection of the traffic. Protections afforded unicast frames include encryption, integrity protection, and replay protection. Additionally, because data forgery is a major security concern in WLANs, unicast frames are equipped with a data origin authentication mechanism that prevents masquerading attacks. The mechanism allows a STA to confirm whether or not a received data frame originated from the claimed STA.

The broadcast and multicast data transfer mechanisms (also called “group”) allow for common data to be transferred to multiple devices efficiently. Communication between the AP and the STAs is protected using CCMP. Unique Group Key Handshakes with each STA insert the GTK used with CCMP to protect the data exchanges. Because all STAs share the same GTK, a single breach of the GTK affects all STAs.

CONNECTION TERMINATION

The fifth and final phase in the operation of an RSNA is the connection termination phase. During this phase, the association between the STA and the AP is deleted, and the wireless connection is terminated. This phase provides the elegant teardown of a connection and a restoration to an initialized state.

During the connection termination phase, the following events occur:

- The AP deauthenticates the STA.
- The security associations, used internally by the AP to keep track of associations between STAs and APs, are deleted.
- The temporal keys used for encrypting and protecting the integrity of data traffic are deleted.
- The IEEE 802.1X controlled port returns to a blocked state so that user traffic cannot pass.

CONNECTION TERMINATION

The connection termination phase may be entered in several ways, including the following:

- Radio communication between the STA and AP is lost (e.g., STA moves out of range).
- The 4-Way Handshake or Group Key Handshake times out during execution.
- The RSNIE check during the 4-Way Handshake fails.
- The user powers down the STA or disables the NIC.
- The security policy indicates a termination of the connection (implementation-specific).

This phase restores the AP and STA to an initialized state. If further communication is subsequently required, then these devices begin anew at the discovery phase with the re-discovery of the available resources and capabilities.