



# Authentication and Remote Access: Account Policies

Slide 1



Principles of Computer Security, Fifth Edition


## Authentication and Remote Access



Account Policies

Copyright © 2018 by McGraw-Hill Education. All rights reserved.

Howdy! In this video, we discuss account policies.



Principles of Computer Security, Fifth Edition

## Account policies


- Good set of policies guide security professionals in daily tasks
- Policies needed for a wide range of elements
  - Naming conventions to operating rules, such as audit frequency and other specifics
- Having issues resolved as a matter of policy enables security professionals to go about the task of verifying and monitoring systems
  - Avoids adjudication of policy type issues with each user case

Copyright © 2018 by McGraw-Hill Education. All rights reserved.

A good set of policies guide security professionals in daily tasks.

Policies are needed for a wide range of elements, from naming conventions to operating rules, such as audit frequency and other specifics.

Having these issues resolved as a matter of policy enables security professionals to go about the task of verifying and monitoring systems rather than trying to adjudicate policy type issues with each user case.



Principles of Computer Security, Fifth Edition

## Account policy Enforcement

- Passwords: primary method of account policy enforcement
- Foundation of a solid account policy:
  - Each user ID is traceable to a single person's activity
  - No sharing of passwords and credentials
- Passwords need to be
  - managed to provide appropriate levels of protection
    - strong enough to resist attack
    - not too difficult for users to remember.
- Password policy ensures necessary steps taken to enact a secure password solution

Copyright © 2018 by McGraw-Hill Education. All rights reserved.

Passwords are the primary method of account policy enforcement.


The foundation of a solid account policy includes that each user ID is traceable to a single person's activity and that passwords and credentials are not shared (which supports user IDs being traceable to a single person's activities).

Most account policies state that the owner of the account is responsible for the actions of the account, putting the onus on the user to keep their password strong and secret.

Passwords need to be managed to provide appropriate levels of protection.

They need to be strong enough to resist attack, and yet not too difficult for users to remember.

A password policy ensures that the necessary steps are taken to enact a secure password solution both by users and by the password infrastructure system.



Principles of Computer Security, Fifth Edition

## Credential Management


- **Credential management:**
  - Processes, services, and software used to store, manage, and log the use of user credentials
- **Credential management solutions:**
  - Typically aimed at assisting end users manage their growing set of passwords
- **Credential management products**
  - Provide secure means of storing user credentials
  - Make credentials available across a wide range of platforms

Copyright © 2018 by McGraw-Hill Education. All rights reserved.

**Credential management** is made up of processes, services, and software used to store, manage, and log the use of user credentials.

Credential management solutions are typically aimed at assisting end users manage their growing set of passwords.

Credential management products provide secure means of storing user credentials and make credentials available across a wide range of platforms.



Principles of Computer Security, Fifth Edition

## Account Maintenance

- **Account maintenance** is the routine screening of all attributes for an account.
- Best practice: perform in accordance with risk associated with the profile.
- Account maintenance is a joint responsibility

Copyright © 2018 by McGraw-Hill Education. All rights reserved.

**Account maintenance** is the routine screening of all attributes for an account.

The best practice is to perform account maintenance in accordance with the risk associated with the account.

System administrators, and other privileged accounts, need greater scrutiny than normal users.


Shared accounts, such as guest accounts, also require scrutiny to ensure they are not abused.

The higher the risk, the more often and thorough the screening should be.

It is also important to review which accounts have particular high-risk permissions.

Account maintenance is a joint responsibility.

The job of determining who has what access is actually one that belongs to the business, not the security group.



Principles of Computer Security, Fifth Edition

### Usage auditing and review

- Examination of logs to determine user activity.
- Important: reviewing access control logs for root level accounts
- Root-level changes in a system tend to be significant changes
  - Approved changes in advance in production environment.
- Compare all root-level activity against approved changes
  - Assists in the detection of activity that is unauthorized.

Copyright © 2018 by McGraw-Hill Education. All rights reserved.

Usage auditing and Review is an examination of logs to determine user activity.


Reviewing access control logs for root level accounts is an important element of securing access control methods.

Due to their power and potential for misuse, administrative or root-level accounts should be closely monitored

A strong configuration management environment includes control of access to production systems by users who can change the environment.

Root-level changes in a system tend to be significant changes and so need to be approved in advance for the production environment.

A comparison of all root-level activity against approved changes will assist in the detection of unauthorized activity.




Principles of Computer Security, Fifth Edition

## Time-of-Day Restrictions

- **Time-of-day** restrictions limit when a user can log in, when certain resources can be accessed, and so on.
- From a security perspective, time of day restrictions can be very useful.
- Time of day restrictions can also serve as a mechanism to enforce internal controls of critical or sensitive resources.

Copyright © 2018 by McGraw-Hill Education. All rights reserved.

Time-of-day restrictions limit when a user can log in, when certain resources can be accessed, and so on. From a security perspective, time of day restrictions can be very useful as they support the principle of least privilege. Time of day restrictions can also serve as a mechanism to enforce internal controls of critical or sensitive resources.



Principles of Computer Security, Fifth Edition

## Account Expiration

- Setting an end time for an account's validity.
- Organizations must define whether accounts are deleted or disabled when no longer needed.
  - Deleted – removed from the system permanently.
  - Disabled – marked as unusable temporarily.

Copyright © 2018 by McGraw-Hill Education. All rights reserved.

Operating systems allow administrators to specify the length of time an account is valid and when the “account expires” or is disabled.


This is a great method for controlling temporary accounts, or accounts for contractors or contract employees. For these accounts, the administrator can specify an expiration date; when the date is reached, the account automatically becomes locked out and cannot be logged into without administrator intervention. A related action can be taken with accounts that never expire: they can automatically be marked “inactive” and locked out if they have been unused for a specified number of days. Account expiration is similar to password expiration, in that it limits the time window of potential compromise. When an account has expired, it cannot be used unless the expiration deadline is extended.

Organizations must define whether accounts are deleted or disabled when no longer needed.

Deleting an account removes the account from the system permanently, whereas disabling an account leaves it in place but marks it as unusable.

Many organizations disable an account after an employee leaves for a period of time prior to deleting the account. This prevents anyone from using the account and allows administrators to reassign files, forward mail, and “clean up” before taking any permanent actions on the account.





Principles of Computer Security, Fifth Edition

## Attribution

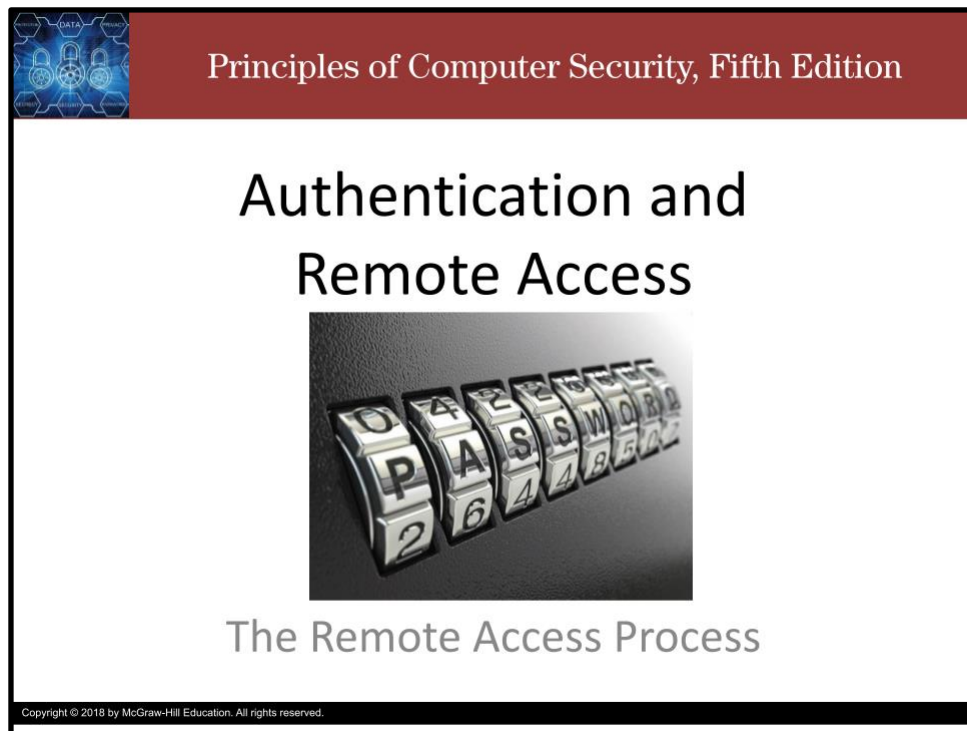
- The course slides are based on slides developed by the textbook authors, W. “Art” Conklin, Ph.D. (University of Houston), and Greg White, Ph.D. (University of Texas at San Antonio).
- These slides remain the intellectual property protected by US Copyright law of the authors and the textbook company.
- There have been some changes to the slide deck.

Copyright © 2018 by McGraw-Hill Education. All rights reserved.

Thank you and take care.


# Authentication and Remote Access: The Remote Access Process

Slide 1



Principles of Computer Security, Fifth Edition


## Authentication and Remote Access



The Remote Access Process

Copyright © 2018 by McGraw-Hill Education. All rights reserved.

Howdy! In this video, we discuss the remote access process.



Principles of Computer Security, Fifth Edition

## The Remote Access Process

- 2 elements of remote access
  - Connection
  - Privileges: Authentication, Authorization, Accounting
- 3 steps to establish privileges
  - Authentication – validate credentials
  - Authorization – grant privileges based on account
  - Accounting – logging access and activity

Copyright © 2018 by McGraw-Hill Education. All rights reserved.

The process of connecting by remote access involves two elements: a temporary network connection and a series of protocols to negotiate privileges and commands.

Once the connection is made, the primary issue is authenticating the identity of the user and establishing proper privileges for that user.


This is accomplished using a combination of protocols and the operating system on the host machine.

The three steps in the establishment of proper privileges are authentication, authorization, and accounting, or simply, AAA.

Authentication is the matching of user-supplied credentials to previously stored credentials on a host machine, and it usually involves an account username and password.

Authorization is the granting of specific permissions based on the privileges held by the account.

Accounting is the collection of billing and other detail records, such as access and activity logs.



Principles of Computer Security, Fifth Edition

## Identification

- **Identification** is the process of ascribing a computer ID to a specific user, computer, network device, or computer process.
  - The identification process is typically performed only once, when a user ID is issued to a particular user.
  - User identification enables authentication and authorization to form the basis for accountability.
  - For accountability purposes, user IDs should not be shared, and for security purposes, they should not be descriptive of job function.

Copyright © 2018 by McGraw-Hill Education. All rights reserved.


Identification is the process of ascribing an identifier to a specific user, computer, network device, or computer process.

The identification process is typically performed only once, when a user ID is issued to a particular user.

This practice enables you to trace activities to individual users or computer processes so that they can be held responsible for their actions. Identification links the logon ID or user ID to credentials that have been submitted previously to either HR or the IT staff. A required characteristic of user IDs is that they must be unique so that they map back to the credentials presented when the account was established.

User identification enables authentication and authorization to form the basis for accountability.

For accountability purposes, user IDs should not be shared, and for security purposes, they should not be descriptive of job function.



Principles of Computer Security, Fifth Edition


## Authentication

- Authentication is the process of binding a specific ID to a specific computer connection.
  - Two items need to be presented to cause this binding to occur—the user ID, and some “secret” to prove that the user is the valid possessor of the credentials.
- Historically, three categories of secrets are used to authenticate the identity of a user:
  - What users know, what users have, and what users are
- Today, an additional category is used: what users do.

Copyright © 2018 by McGraw-Hill Education. All rights reserved.

Authentication is the process of binding a specific ID to a specific computer connection. Two items need to be presented to cause this binding to occur—the user ID, and some “secret” to prove that the user is the valid possessor of the credentials. Historically, three categories of secrets are used to authenticate the identity of a user: What users know, what users have, and what users are. Today, an additional category is used: what users do.

These methods can be used individually or in combination. These controls assume that the identification process has been completed and the identity of the user has been verified. It is the job of authentication mechanisms to ensure that only valid users are admitted. Described another way, authentication is using some mechanism to prove that you are who you claimed to be when the identification process was completed.



Principles of Computer Security, Fifth Edition

## Authentication

- Password is most common authentication method.
- Another method to provide authentication involves the use of something that only valid users should have in their possession.
- The third general method to provide authentication involves something that is unique about you.

Copyright © 2018 by McGraw-Hill Education. All rights reserved.


Passwords are common because they are one of the simplest forms and use user memory as a prime component. Because of their simplicity, passwords have become ubiquitous across a wide range of authentication systems.

For greater security, you can add an element from a separate group, such as a smart card token—something a user has in her possession.

A new method, based on how users perform an action, such as their gait when walking, or typing patterns has emerged as a source of a personal “signature”.

Principles of Computer Security, Fifth Edition

## Basic Authentication



localhost  
User name:   
Password:   
 Remember my password  
OK Cancel

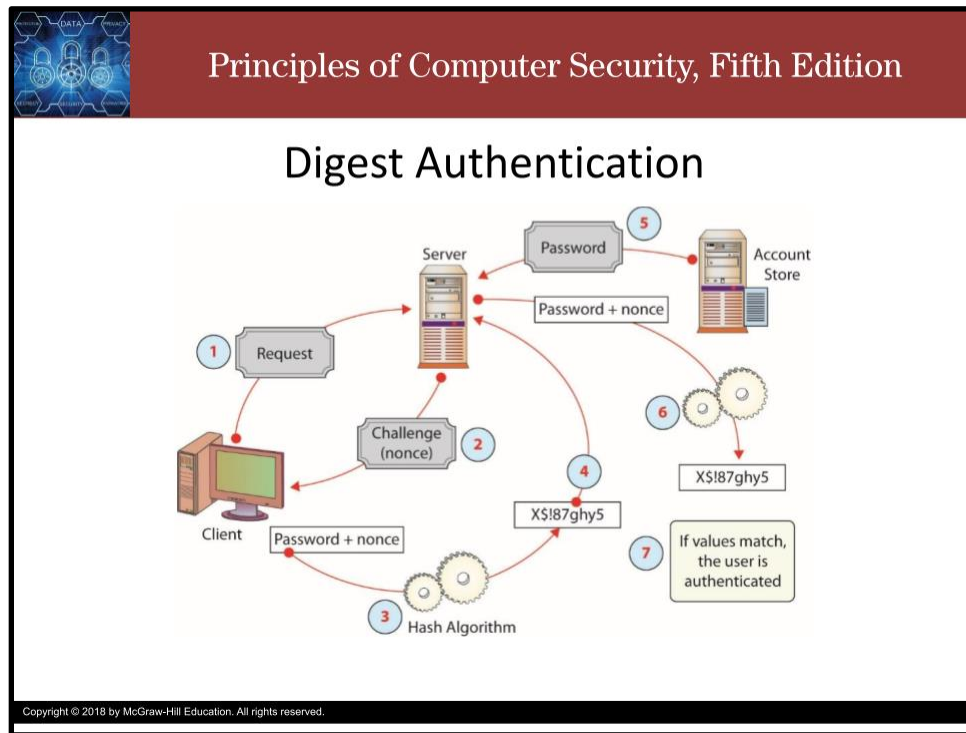
Username and password encoded  
using Base64 encoding and sent to server

```
GET /SomeBasicSite/ HTTP/1.0
Accept: image/gif, image/jpeg, image/png, */*
User-Agent: Mozilla/4.0 (compatible; MSIE 6.0; Windows NT 5.1; .NET CLR 1.1.4322)
Host: SomeBasicSite
If-None-Match: *39d01a8ae1f051a28*
Authorization: Basic YWxpY2U6UGFzc3dvcnQxMjM=
Connection: Keep-Alive
```

→ YWxpY2U6UGFzc3dvcnQxMjM= → alice:Password123

Copyright © 2018 by McGraw-Hill Education. All rights reserved.

Basic authentication is the simplest technique used to manage access control across HTTP. Basic authentication operates by passing information encoded in Base64 form using standard HTTP headers. This is a plaintext method without any pretense of security.



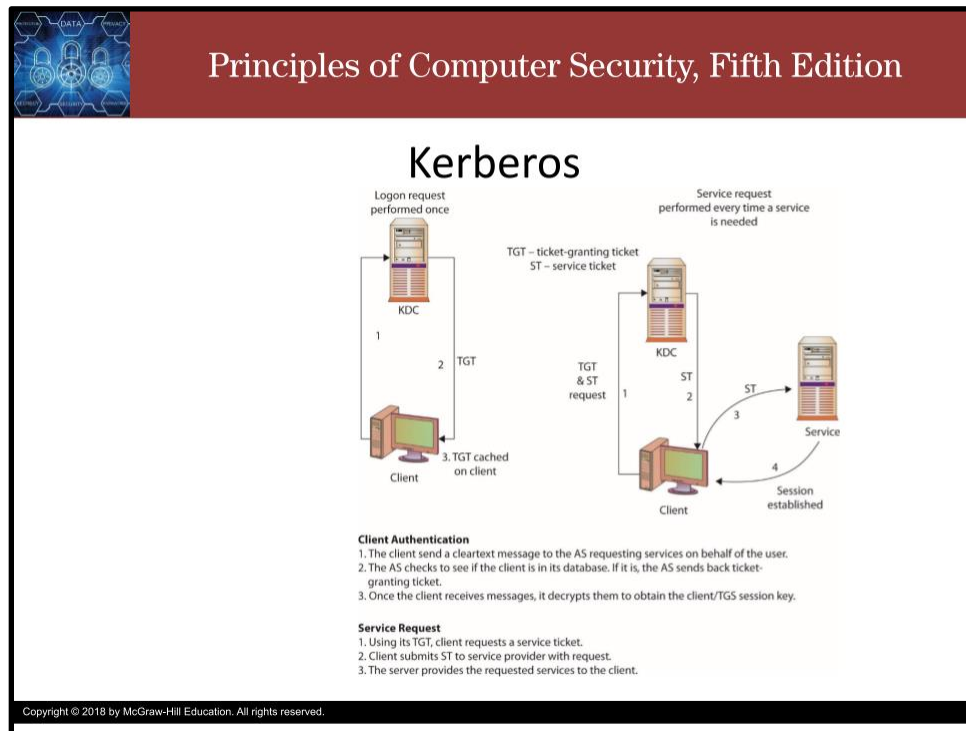
Digest authentication is a method used to negotiate credentials across the Web.

It uses hash functions and a nonce to improve security over basic authentication.

Although it improves security over basic authentication (because passwords are not sent in the clear), it does not provide any significant level of security (because it is subject to man-in-the-middle attacks and potentially replay attacks).

The way digest authentication is shown in this figure: Step 1, The client requests login, step 2, The server responds with a challenge and provides a nonce. Step 3, The client hashes the password and nonce, step 4, the client returns the hashed password to the server, Step 5, The server requests the password from a password store. Step 6, The server hashes the password and nonce, step 7, If both hashes match, login is granted.






Developed as part of MIT's project Athena, Kerberos is a network authentication protocol designed for a client/server environment.

Taking its name from the three-headed dog of Greek mythology, Kerberos is designed to work across the Internet, an inherently insecure environment. Kerberos uses strong encryption so that a client can prove its identity to a server and the server can in turn authenticate itself to the client. A complete Kerberos environment is referred to as a Kerberos realm. The Kerberos server contains user IDs and hashed passwords for all users that will have authorizations to realm services. The Kerberos server also has shared secret keys with every server to which it will grant access tickets.

Kerberos is built around the idea of a trusted third party, termed a key distribution center (KDC), which consists of two logically separate parts: an authentication server (AS) and a ticket-granting server (TGS). The basis for authentication in a Kerberos environment are tickets that serves to prove the identity of users.

Tickets are used in a two-step process with the client. The first ticket is a ticket-granting ticket (TGT) issued by the AS to a requesting client. The client can then present this ticket to the Kerberos server with a request for a ticket to access a specific server. This client-to-server ticket (also called a service ticket) is used to gain access to a server's service in the realm.

Since the entire session can be encrypted, this eliminates the inherently insecure transmission of items such as a password that can be intercepted on the network. Tickets are time-stamped and have a lifetime, so attempting to reuse a ticket will not be successful.



Principles of Computer Security, Fifth Edition

## Mutual Authentication


- Describes a process in which each side of an electronic communication verifies the authenticity of the other.
- This provides a mechanism for each side of a client/server relationship to verify the authenticity of the other to address this issue.
- A common method involves using a secure connection, such as Transport Layer Security (TLS), to the server and a one-time password generator that then authenticates the client.

Copyright © 2018 by McGraw-Hill Education. All rights reserved.

Mutual authentication describes a process in which each side of an electronic communication verifies the authenticity of the other.

This provides a mechanism for each side of a client/server relationship to verify the authenticity of the other to address this issue.

A common method involves using a secure connection, such as TLS, and a one-time password generator that then authenticates the client.



Principles of Computer Security, Fifth Edition

## Certificates

- A method of establishing authenticity of specific objects such as an individual's public key or downloaded software.
- A **digital certificate** is a digital file that is sent as an attachment to a message and is used to verify that the message did indeed come from the entity it claims to have come from.

Copyright © 2018 by McGraw-Hill Education. All rights reserved.


Certificates are a method of establishing authenticity of specific objects such as an individual's public key or downloaded software.

A digital certificate is a digital file that sent as an attachment to a message and is used to verify that the message did indeed come from the entity it claims to have come from.

Principles of Computer Security, Fifth Edition

## Tokens

- Authentication factor that typically takes the form of a physical or logical entity that the user must be in possession of to access their account or certain resources.
- Something-you-have and Something-you-know.
- Several variations on this type of device exist.
- Commonly employed in remote authentication schemes
- Most are physical tokens that display a series of numbers that change every 30 to 90 seconds.



Copyright © 2018 by McGraw-Hill Education. All rights reserved.


A token is an authentication factor that typically takes the form of a physical or logical entity that the user must be in possession of to access their account or certain resources.

Tokens are commonly employed in remote authentication schemes as they provide additional surety of the identity of the user, even users who are somewhere else and cannot be observed.

It functions as both a something-you-have and something-you-know authentication mechanism.

Most tokens are physical tokens that display a series of numbers that changes every 30 to 90 seconds which are used in a challenge/response authentication process.

Several variations on this type of device exist, which all work on the same basic principles.



Principles of Computer Security, Fifth Edition

## Software Tokens

- Provide two-factor authentication but don't require the user to have a separate physical device on hand.
- Some tokens require software clients that store a symmetric key in a secured location on the user's device.
- Other software tokens use public key cryptography.
- Most common form is a software token on the device and the user supplies the rest of the details needed to demonstrate authenticity.


Copyright © 2018 by McGraw-Hill Education. All rights reserved.

Software tokens still provide two-factor authentication but don't require the user to have a separate physical device on hand.

Some tokens require software clients that store a symmetric key (sometimes called a seed record) in a secured location on the user's device (laptop, desktop, tablet, and so on).

Other software tokens use public key cryptography, which often associate a PIN with a specific user's token.

The most common form of software token is for identifying a specific device in addition to a user, in that the software token is on the device and the user supplies the rest of the details needed to demonstrate authenticity.



Principles of Computer Security, Fifth Edition

## OTP/HOTP/TOTP

- **One-Time Password (OTP)** – a password that can only be used once.
- **HMAC-based One-Time Password (HOTP)** uses a counter to create/update the OTP on each request/validation.
  - HMAC stands for Hash-based Message Authentication Code.
- **Time-based One-Time Password (TOTP)** uses a timestamp to create/update the OTP at each timestep.
  - $TOTP(K) = HOTP(K, \text{timestamp})$  <https://www.onelogin.com/learn/otp-totp-hotp>

Copyright © 2018 by McGraw-Hill Education. All rights reserved.


A One-Time Password is a password that can be used only once. Once it is used (or it expires), it can no longer be used.

The algorithms that generate OTPs use two inputs: a secret key and a moving factor. The secret key is established when the account is created on the authentication server.

The secret doesn't change, but the moving factor does. How the moving factor is generated is the main differentiator between HOTP and TOTP.

HOTP stands for HMAC-based One-Time Password. The moving factor in HOTP is a counter that is incremented every time the OTP is validated. The secret key and the counter are combined and hashed to produce the OTP.

TOTP stands for Time-based One-Time Password. The moving factor in TOTP is a timestamp. Rather than updating every time a OTP is used, the TOTP is updated every 30 to 90 seconds. The secret key and the timestamp are combined and hashed on every update to produce the next one time password. Essentially, a TOTP is a HOTP where the counter value is a timestamp.



Principles of Computer Security, Fifth Edition

## Smart Cards


- Increase physical security.
- Something you know + something you have.
- Smart card readers are common kit.

Copyright © 2018 by McGraw-Hill Education. All rights reserved.

Smart cards can increase physical security because they can carry cryptographic tokens that are too long to remember and have too large a space to guess.

Smart cards can find use in a variety of situations where you want to combine something you know (a pin or password) together with something you have (and can't be duplicated, such as a smart card).

Many standard corporate-type laptops come with smart card readers installed, and their use is integrated into the Windows user access system.



Principles of Computer Security, Fifth Edition

## Multifactor Authentication

- **Multifactor Authentication** is the combination of two or more types of authentication
- Also called multiple-factor authentication)
- Five broad categories of authentication can be used:
  - **What you are** (for example, biometrics)
  - **What you have** (for instance, tokens)
  - **What you know** (passwords and other information)
  - Somewhere you are (location)
  - Something you do (physical performance)

Copyright © 2018 by McGraw-Hill Education. All rights reserved.

Multifactor Authentication is the combination of two or more types of authentication. Five broad categories of authentication can be used: Something you know, something you have, something you are, something you do, and somewhere you are. The first three are the main three. Something you know specifically refers to passwords.

The challenge with something you know is that it can be “shared”, sometime even without the user realizing it

Something you have specifically refers to tokens and other items that a user can possess physically.

The challenges with this factor are that you have to have the token with you whenever you wish to be authenticated (so stealing the token is, in effect, a denial of service), and that it relies on interfaces that might not be available for some systems.

Something you are specifically refers to biometrics.

The challenges here are that biometrics tend to be hard to change, so once assigned they become immutable (except by the incessant onward march of time) and that measuring things on a person can be imprecise or inaccurate or both.

Something you do specifically refers to activities.

For example: the movement of the pen and the two-dimensional output of a signature are difficult for others to reproduce.

While useful for authentication, one challenge is capturing the data, similar to biometrics.

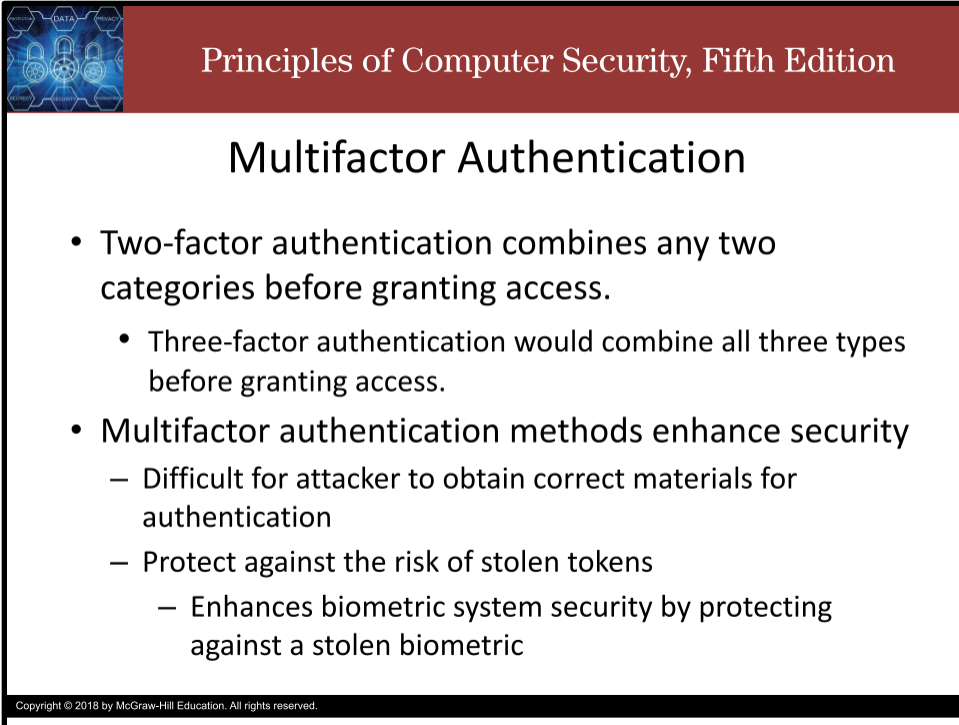


Somewhere you are is a form of something you do: you do be at somewhere. That is, it refers to location.

Various pieces of information can be used for this factor, such as comparing your alleged location (where you claim to be in the authentication request) to your supposed location (where you are expected to be given known patterns and location history) to determine if you are really there, or even should be there.

This factor doesn't work well without accurate location services.

## Slide 16



Principles of Computer Security, Fifth Edition

### Multifactor Authentication

- Two-factor authentication combines any two categories before granting access.
  - Three-factor authentication would combine all three types before granting access.
- Multifactor authentication methods enhance security
  - Difficult for attacker to obtain correct materials for authentication
  - Protect against the risk of stolen tokens
    - Enhances biometric system security by protecting against a stolen biometric

Copyright © 2018 by McGraw-Hill Education. All rights reserved.

Two-factor authentication combines any two categories before granting access, such as something you know, like a password, and something you have, like a TOTP.

Three-factor authentication combines three types before granting access: password and authenticator token and biometrics.

Multifactor authentication methods enhance security because they make it difficult for attacker to obtain correct materials for authentication.

They also protect against the risk of stolen tokens, since an attacker needs all the factors to authenticate. A token on its own is not enough.

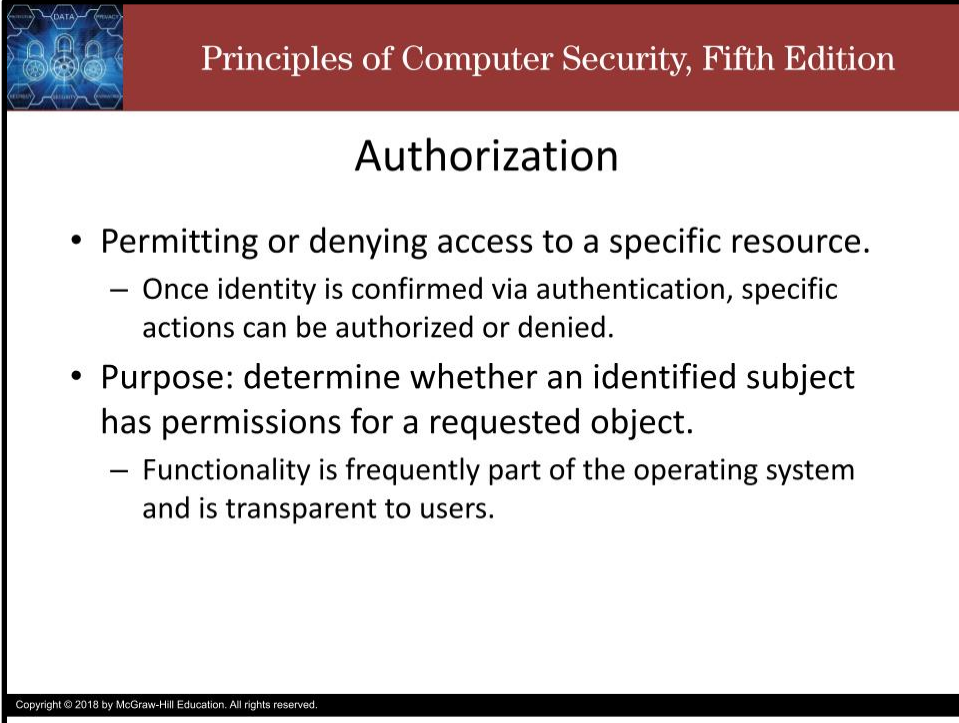
By that same token (ha! Get it?), it enhances biometric system security by protecting against a stolen biometric. Who needs high-tech fingerprint or retina cloning capability when you can just take a finger

or eyeball during an advanced interrogation session? I always tell people to use their pinky fingerprint. A little bit of security by obscurity and a little bit of it would suck to lose an index finger.

Multifactor authentication is one of the best ways to ensure proper authentication and access control.

If you haven't enabled it on most (if not all) of your accounts, I would encourage you to do so. If you are designing a system, I would encourage you to include support for MFA and if you don't outright require it, at least encourage your users to enable it.

## Slide 17



Principles of Computer Security, Fifth Edition

### Authorization

- Permitting or denying access to a specific resource.
  - Once identity is confirmed via authentication, specific actions can be authorized or denied.
- Purpose: determine whether an identified subject has permissions for a requested object.
  - Functionality is frequently part of the operating system and is transparent to users.


Copyright © 2018 by McGraw-Hill Education. All rights reserved.

Authorization is the process of permitting or denying access to a specific resource.

Once identity is confirmed via authentication, specific actions can be authorized or denied.

Many types of authorization schemes are used, but the purpose is the same: to determine whether a given user who has been identified has permissions for a particular object or resource being requested.

This functionality is frequently part of the operating system and is transparent to users.



Principles of Computer Security, Fifth Edition

## Separation of Tasks has Advantages

- Identification | Authentication | Authorization
- Many methods can be used to perform each task.
  - On many systems several methods are concurrently present for each task.
- Separation of these tasks into individual elements allows combinations of implementations to work together.
- Any system or resource that requires authorization can use its own authorization method once authentication has occurred.
  - This makes for efficient and consistent application of security principles.

Copyright © 2018 by McGraw-Hill Education. All rights reserved.


The separation of tasks, from identification to authentication to authorization, has several advantages.

Many methods can be used to perform each task, and on many systems several methods are concurrently present for each task.

Separation of these tasks into individual elements allows combinations of implementations to work together.

Any system or resource, be it hardware (router or workstation) or software (database system), that requires authorization can use its own authorization method once authentication has occurred.

This makes for efficient and consistent application of security principles.



## Principles of Computer Security, Fifth Edition

### Attribution

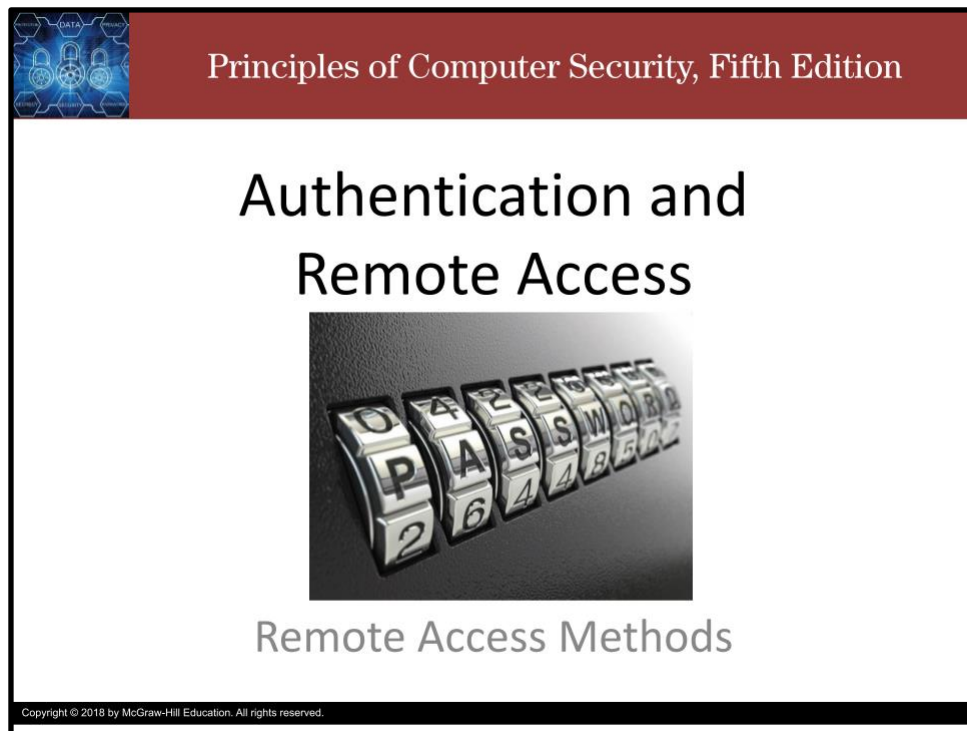
- The course slides are based on slides developed by the textbook authors, W. “Art” Conklin, Ph.D. (University of Houston), and Greg White, Ph.D. (University of Texas at San Antonio).
- These slides remain the intellectual property protected by US Copyright law of the authors and the textbook company.
- There have been some changes to the slide deck.

Copyright © 2018 by McGraw-Hill Education. All rights reserved.

Thank you and take care.


# Authentication and Remote Access: Remote Access Methods

Slide 1



Principles of Computer Security, Fifth Edition


## Authentication and Remote Access



Remote Access Methods

Copyright © 2018 by McGraw-Hill Education. All rights reserved.

Howdy! In this video, we introduce some remote access methods.



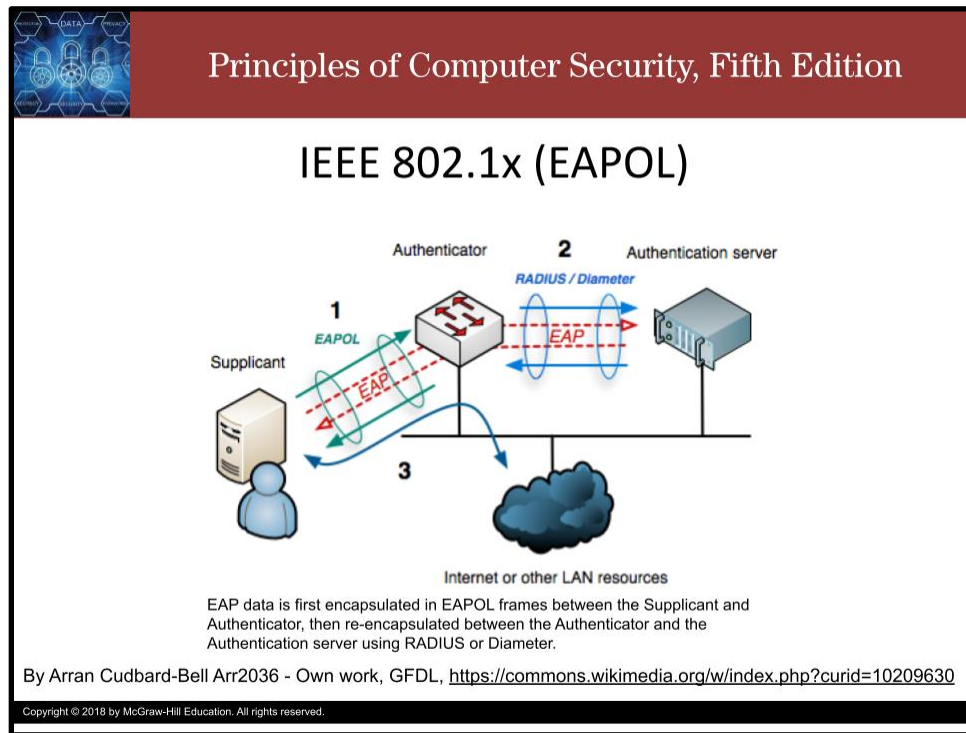
Principles of Computer Security, Fifth Edition

## Remote Access Methods

- When a user requires access to a remote system, the process of remote access is used to determine the appropriate controls.
- This is done through a series of protocols and processes.

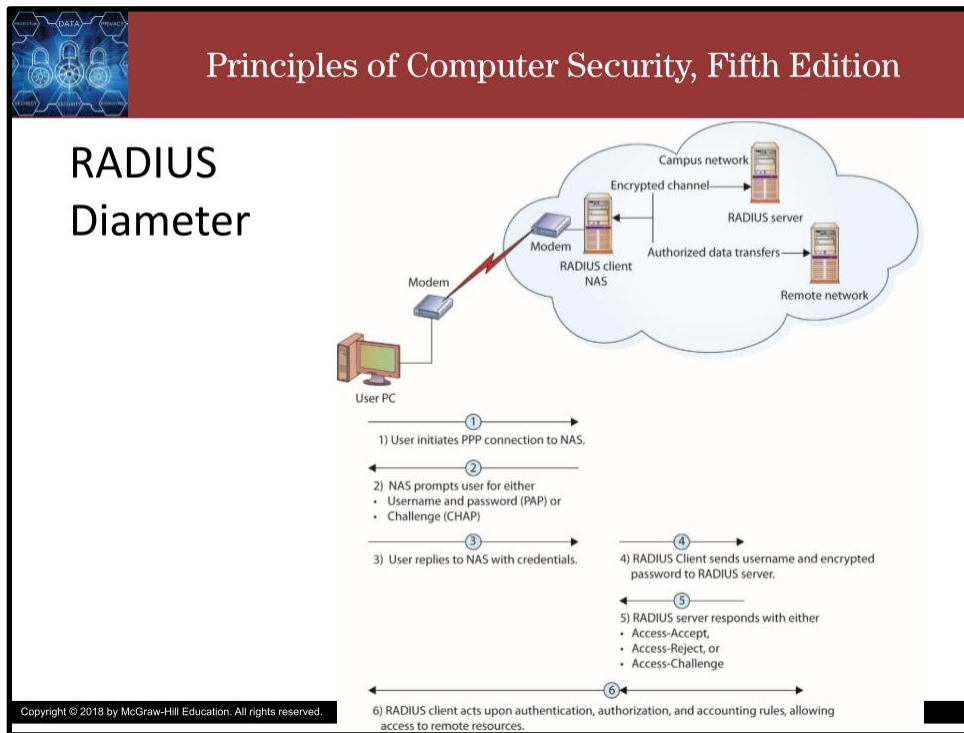
Copyright © 2018 by McGraw-Hill Education. All rights reserved.

When a user requires access to a remote system, the process of remote access is used to determine the appropriate controls. This is done through a series of protocols and processes.



802.1X authentication involves three parties: a supplicant, an authenticator, and an authentication server. The supplicant is a client device (such as a laptop) that wishes to attach to the LAN/WLAN. The term 'supplicant' is also used interchangeably to refer to the software running on the client that provides credentials to the authenticator. The authenticator is a network device that provides a data link between the client and the network and can allow or block network traffic between the two, such as an Ethernet switch or wireless access point; and the authentication server is typically a trusted server that can receive and respond to requests for network access, and can tell the authenticator if the connection is to be allowed, and various settings that should apply to that client's connection or setting. Authentication servers typically run software supporting the RADIUS and EAP protocols. In some cases, the authentication server software may be running on the authenticator hardware.

The authenticator acts like a security guard to a protected network. The supplicant (i.e., client device) is not allowed access through the authenticator to the protected side of the network until the supplicant's identity has been validated and authorized. With 802.1X port-based authentication, the supplicant must initially provide the required credentials to the authenticator - these will have been specified in advance by the network administrator and could include a user name/password or a permitted digital certificate. The authenticator forwards these credentials to the authentication server to decide whether access is to be granted. If the authentication server determines the credentials are valid, it informs the authenticator, which in turn allows the supplicant (client device) to access resources located on the protected side of the network.



The Remote Authentication Dial-In User Service (or Radius) is an AAA protocol. It was designed as a connectionless protocol. As such, UDP is employed as its transport layer protocol. Connection issues are handled by the Radius application.

Radius is a client/server protocol. The client is typically a network access server (or NAS).

Network access servers act as intermediaries, authenticating clients before allowing them to access to a network. The RADIUS server is a process or daemon running on a UNIX or Windows Server machine. Radius encrypts only the users passwords as it travels from the RADIUS client to Radius server. All other information such as the username, authorization, accounting, are transmitted in clear text. Therefore, it is vulnerable to different types of attacks.

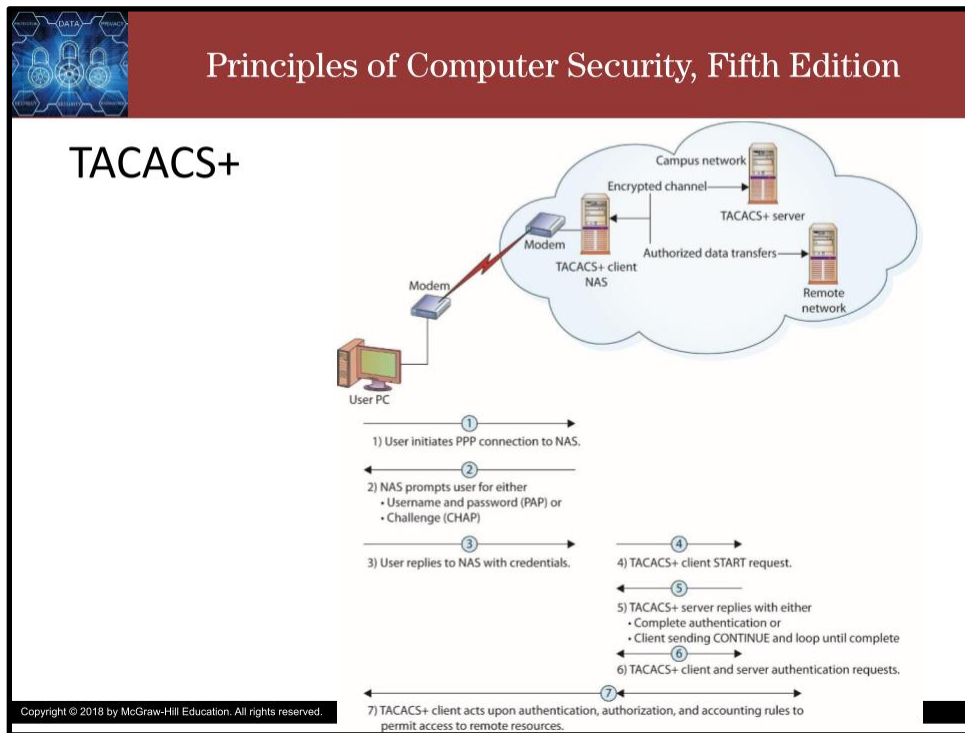
This figure shows the interaction between a user and the RADIUS client and RADIUS server and the steps taken to make a connection. This is how RADIUS handles authentication.

A user login authentication consists of a query (Access-Request) from the RADIUS client and a corresponding response (Access-Accept, Access-Challenge, or Access-Reject) from the RADIUS server. The Access-Challenge response is the initiation of a challenge/response handshake. If the client cannot support challenge/response, then it treats the challenge message as an Access-Reject. The Access-Request message contains the username, obfuscated password, NAS IP address, and port. The message also contains information concerning the type of session the user wants to initiate. Once the RADIUS server receives this information, it searches its database for a match on the username. If a match is found, either a default profile is loaded or an Access-Reject reply is sent to the user. If the entry is found



or the default profile is used, the next phase involves authorization, which is then followed by accounting.

Diameter is the name of an AAA protocol suite that is designed by the IETF to replace the aging RADIUS protocol. Diameter operates like RADIUS in a client/server configuration. It improves upon RADIUS, resolving some discovered weakness. It is a TCP-based service and has more extensive AAA capabilities. It is designed for all types of remote access and it has an improved method of encrypting message exchanges to prevent replay and man-in-the-middle attacks. Taken all together, Diameter, with its enhanced functionality and security, is an improvement on the proven design of the old RADIUS standard.



The Terminal Access Controller Access Control System+ (TACACS+) protocol is the current generation of the TACACS family. Originally TACACS was developed by BBN Planet Corporation for MILNET, an early military network, but it has been enhanced by Cisco, which has expanded its functionality twice. The original BBN TACACS system provided a combination process of authentication and authorization. Cisco extended this to Extended Terminal Access Controller Access Control System (XTACACS), which provided for separate authentication, authorization, and accounting processes. The current generation, TACACS+, has extended attribute control and accounting processes.

One of the fundamental design aspects is the separation of authentication, authorization, and accounting in this protocol. Although there is a straightforward lineage of these protocols from the original TACACS, TACACS+ is a major revision and is not backward-compatible with previous versions of the protocol series.

TACACS+ uses TCP as its transport protocol.


It's a client/server protocol, with the client typically being a NAS and the server being a daemon process on a UNIX, Linux, or Windows server.

This is important to note, for if the user's machine (usually a PC) is not the client (usually a NAS), then communications between PC and NAS are typically not encrypted and are passed in the clear. Communications between a TACACS+ client and TACACS+ server are encrypted using a shared secret that is manually configured into each entity and is not shared over a connection. Hence, communications between a TACACS+ client (typically a NAS) and a TACACS+ server are secure, but the communications between a user (typically a PC) and the TACACS+ client are subject to compromise. Whereas RADIUS only secures (and weakly) the password, TACACS+ secures all the authentication information and therefore does not have the vulnerabilities present in the RADIUS protocol.

The authentication process is illustrated in in this figure, and it begins with a START message from the client to the server. This message may be in response to an initiation from a PC connected to the TACACS+ client. The START message describes the type of authentication being requested (simple plaintext password, PAP, CHAP, and so on). This START message may also contain additional authentication data, such as a username and password. A START message is also sent as a response to a restart request from the server in a REPLY message. A START message always has its sequence number set to 1.

When a TACACS+ server receives a START message, it sends a REPLY message. This REPLY message indicates whether the authentication is complete or needs to be continued. If the process needs to be continued, the REPLY message also specifies what additional information is needed. The response from a client to a REPLY message requesting additional data is a CONTINUE message. This process continues until the server has all the information needed, and the authentication process concludes with a success or failure.

Like RADIUS, TACACS+ also handles authorization and accounting.



Principles of Computer Security, Fifth Edition

## Authentication Protocols

- Numerous authentication protocols have been developed.
  - Some did not enjoy market share.
  - Others have had security issues.
  - Others have been revised and improved in newer versions.

Copyright © 2018 by McGraw-Hill Education. All rights reserved.


Numerous authentication protocols have been developed, used, and discarded in the brief history of computing.

Some did not enjoy market share.

Others have had security issues.

Still others have been revised and improved in newer versions.

It is impractical and unnecessary to cover them all, so only some of the more common ones are covered here.



Principles of Computer Security, Fifth Edition

## L2TP and PPTP


- Layer 2 Tunneling Protocol (L2TP) and Point-to-Point Tunneling Protocol (PPTP) are both OSI Layer 2 tunneling protocols.
  - **Tunneling** is the encapsulation of one packet within another.
    - This allows you to hide the original packet from view.
    - This can be done for both security and practical reasons.

Copyright © 2018 by McGraw-Hill Education. All rights reserved.

Layer 2 Tunneling Protocol (L2TP) and Point-to-Point Tunneling Protocol (PPTP) are both OSI Layer 2 tunneling protocols.

Tunneling is the encapsulation of one packet within another, which allows you to hide the original packet from view or change the nature of the transport.

This can be done for both security and practical reasons.



Principles of Computer Security, Fifth Edition


## Layer 2 Tunneling Protocol (L2TP)

- Internet standard and came from the Layer 2 Forwarding (L2F) protocol, a Cisco initiative designed to address issues with PPTP.
- Designed for use across all kinds of networks
- Can be implemented by both hardware and software
- Designed to work with established AAA services such as RADIUS and TACACS+

Copyright © 2018 by McGraw-Hill Education. All rights reserved.

Layer 2 Tunneling Protocol (L2TP) is an Internet standard.

It was designed for use across all kinds of networks, to be implemented by both hardware and software, and to work with established AAA services such as RADIUS and TACACS+.



Principles of Computer Security, Fifth Edition

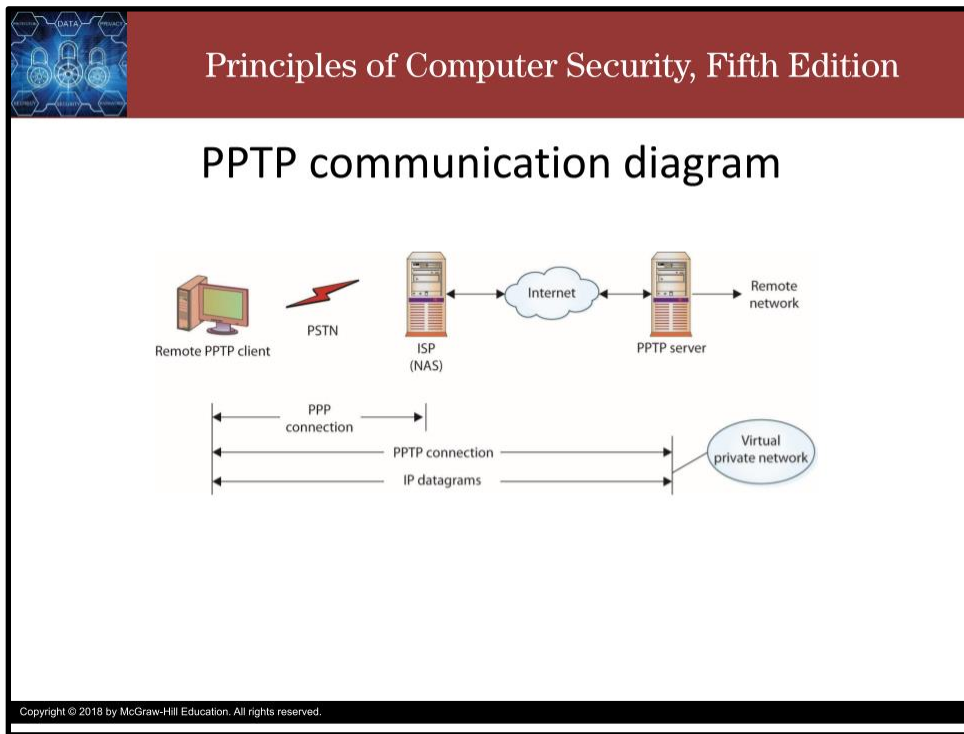
## Point-to-Point Tunneling Protocol (PPTP)

- Network protocol that enables the secure transfer of data from a remote PC to a server by creating a VPN across a TCP/IP network.
- It can also span a public switched telephone network (PSTN) and is thus an economical way of connecting remote dial-in users to a corporate data network.
- For most PPTP implementations, three computers are involved: the PPTP client, the NAS, and a PPTP server.

Copyright © 2018 by McGraw-Hill Education. All rights reserved.

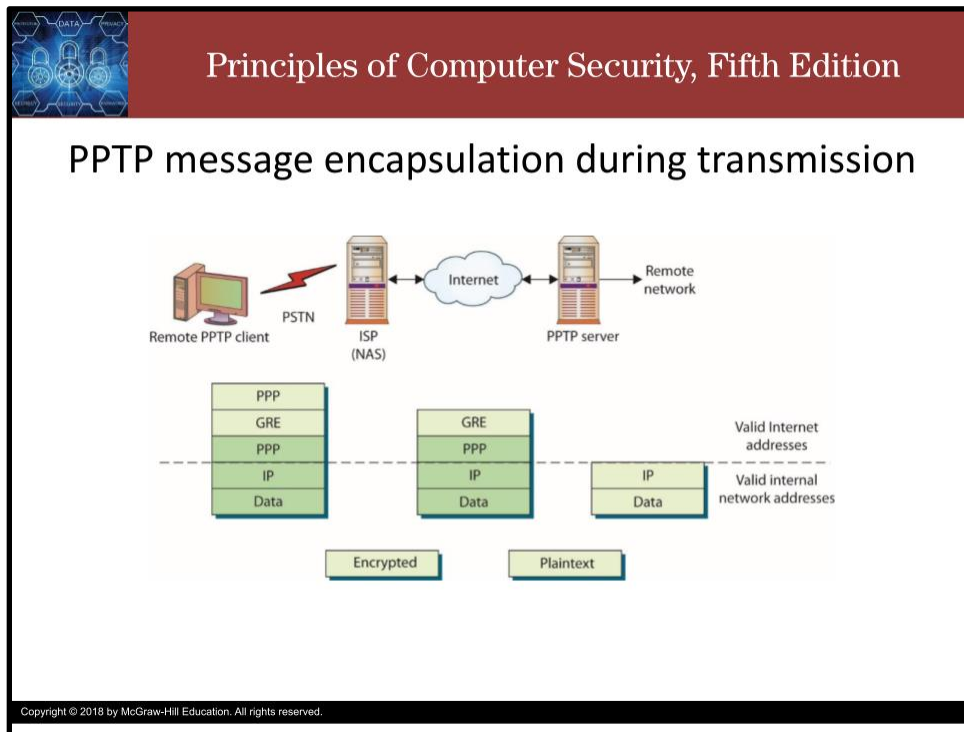
The Point-to-Point Tunneling Protocol (PPTP) is a network protocol that enables the secure transfer of data from a remote PC to a server by creating a VPN across a TCP/IP network.

It can also span a public switched telephone network (PSTN) and is thus an economical way of connecting remote dial-in users to a corporate data network.




For most PPTP implementations, three computers are involved: the PPTP client, the NAS, and a PPTP server, as shown in this figure.





The connection between the remote client and the network is established in stages, as illustrated in this figure.

First the client makes a PPP connection to a NAS, typically an ISP. (In today's world of widely available broadband, if there is already an Internet connection, then there is no need to perform the PPP connection to the ISP.) Once the PPP connection is established, a second connection is made over the PPP connection to the PPTP server. This second connection creates the VPN connection between the remote client and the PPTP server. A typical VPN connection is one in which the user is in a hotel with a wireless Internet connection, connecting to a corporate network. This connection acts as a tunnel for future data transfers.




Principles of Computer Security, Fifth Edition

## Point-to-Point Protocol (PPP)

- Protocol for establishing dial-in connections over serial lines or Integrated Services Digital Network (ISDN) services.
- Several authentication mechanisms: PAP, CHAP, and the Extensible Authentication Protocol (EAP).
- Standardized Internet encapsulation of IP traffic over point-to-point links, such as serial lines.
- The authentication process is performed only when the link is established.

Copyright © 2018 by McGraw-Hill Education. All rights reserved.

The Point-to-Point Protocol (or PPP) is an older, still widely used protocol for establishing dial-in connections over serial lines or Integrated Services Digital Network (ISDN) services. PPP has several authentication mechanisms: PAP, CHAP, and the Extensible Authentication Protocol (EAP), which are protocols used to authenticate the peer device. PPP is a standardized Internet encapsulation of IP traffic over point-to-point links, such as serial lines. The authentication process is performed only when the link is established.




Principles of Computer Security, Fifth Edition

## Extensible Authentication Protocol (EAP)

- A universal authentication framework.
- Frequently used in wireless networks and point-to-point connections
- Can be used for wired authentication
- Most often used in wireless LANs

Copyright © 2018 by McGraw-Hill Education. All rights reserved.

Extensible Authentication Protocol is a universal authentication framework that is frequently used in wireless networks and point-to-point connections. Although EAP is not limited to wireless and can be used for wired authentication, it is most often used in wireless LANs.



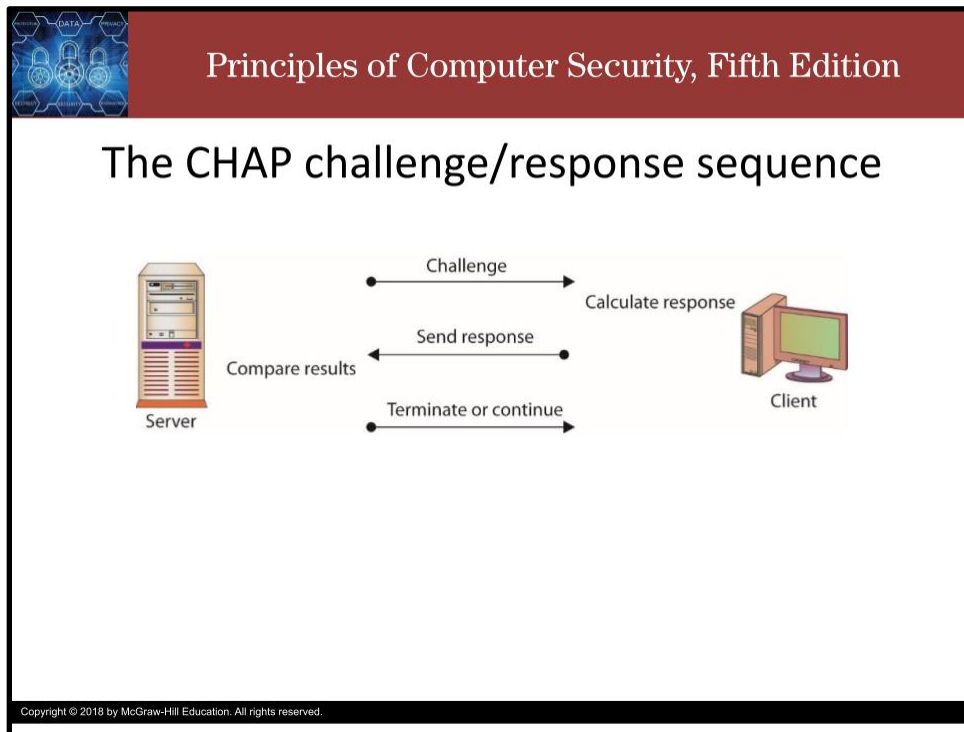
Principles of Computer Security, Fifth Edition

## Challenge-Handshake Authentication Protocol (CHAP)


- Used to provide authentication across a point-to-point link using PPP.
  - Authentication after the link has been established is not mandatory.
  - CHAP is designed to provide authentication periodically through the use of a challenge/response system that is sometimes described as a *three-way handshake*.
  - Microsoft has created two versions of CHAP.
    - Both are broken.

Copyright © 2018 by McGraw-Hill Education. All rights reserved.

Challenge-Handshake Authentication Protocol (or CHAP) is used to provide authentication across a point-to-point link using PPP. Authentication after the link has been established is not mandatory. CHAP is designed to provide authentication periodically through the use of a challenge/response system that is sometimes described as a three-way handshake. Microsoft has created two versions of CHAP. Microsoft has created two versions of CHAP, modified to increase the usability of CHAP across Microsoft's product line. Both versions are broken.



The CHAP *three-way handshake* is illustrated in this figure. The initial challenge (a randomly generated number) is sent to the client. The client uses a one-way hashing function to calculate what the response should be and then sends this back. The server compares the response to what it calculated the response should be. If they match, communication continues. If the two values don't match, then the connection is terminated. This mechanism relies on a shared secret between the two entities so that the correct values can be calculated.



Principles of Computer Security, Fifth Edition

## Telnet


- Telnet is the standard terminal-emulation protocol within the TCP/IP protocol series.
  - Allows users to log in remotely and access resources as if the user had a local terminal connection
  - Offers little security, as usernames, passwords, and all data are passed in cleartext over the TCP/IP connection
  - Makes its connection using TCP port 23
  - Important to control access to Telnet on machines and routers when setting them up
- No authentication.

Copyright © 2018 by McGraw-Hill Education. All rights reserved.

One of the methods to grant remote access to a system is through Telnet. Telnet is the standard terminal-emulation protocol within the TCP/IP protocol series, and it is defined in RFC 854. Telnet allows users to log in remotely and access resources as if the user had a local terminal connection. Telnet is an old protocol and offers little-to-no security. Information, including account names and passwords, is passed in cleartext over the TCP/IP connection.

Telnet makes its connection using TCP port 23. As Telnet is implemented on most products using TCP/IP, it is important to control access to Telnet on machines and routers when setting them up. Failure to control access by using firewalls, access lists, and other security methods, or even by disabling the Telnet daemon, is equivalent to leaving an open door for unauthorized users on a system.

Most implementations of Telnet have no authentication that would ensure communication is carried out between the two desired hosts and not intercepted in the middle.



Principles of Computer Security, Fifth Edition

## Secure Shell (SSH)

- Secure Shell (SSH) is a protocol series designed to facilitate secure network functions across an insecure network.
  - Designed to replace the insecure Telnet application
  - Uses TCP port 22
  - Three major components
    - Transport layer protocol
    - User authentication protocol
    - Connection protocol
  - Very popular in the UNIX environment

Copyright © 2018 by McGraw-Hill Education. All rights reserved.

An SSH connection is an encrypted channel, providing for confidentiality and integrity protection. SSH provides support for securing any network service. SSH has its origins as a replacement for the insecure Telnet application from the UNIX operating system. An original component of UNIX, Telnet allowed users to connect between systems. Although Telnet is still used today, it has some drawbacks. Some enterprising University of California, Berkeley, students subsequently developed the r- commands, such as rlogin, to permit access based on the user and source system, as opposed to passing passwords. This was not perfect either, however, because when a login was required, it was still passed in the clear. This led to the development of the SSH protocol series, designed to eliminate all of the insecurities associated with Telnet, r- commands, and other means of remote access.

SSH opens a secure transport channel between machines by using an SSH daemon on each end. These daemons initiate contact over TCP port 22 and then communicate over higher ports in a secure mode. One of the strengths of SSH is its support for many different encryption protocols. SSH 1.0 started with RSA algorithms, but at the time they were still under patent, and this led to SSH 2.0 with extended support for Triple DES (3DES) and other encryption methods. Today, SSH can be used with a wide range of encryption protocols, including RSA, 3DES, Blowfish, IDEA, CAST128, AES256, and others.

The SSH protocol has facilities to encrypt data automatically, provide authentication, and compress data in transit. It can support strong encryption, cryptographic host authentication, and integrity protection. The authentication services are host-based and not user-based. If user authentication is desired in a system, it must be set up separately at a higher level in the OSI model. The protocol is designed to be flexible and simple, and it is designed specifically to minimize the number of round-trips between systems.

The key exchange, public key, symmetric key, message authentication, and hash algorithms are all negotiated at connection time.

Individual data-packet integrity is assured through the use of a message authentication code that is computed from a shared secret, the contents of the packet, and the packet sequence number.

The SSH protocol consists of three major components:

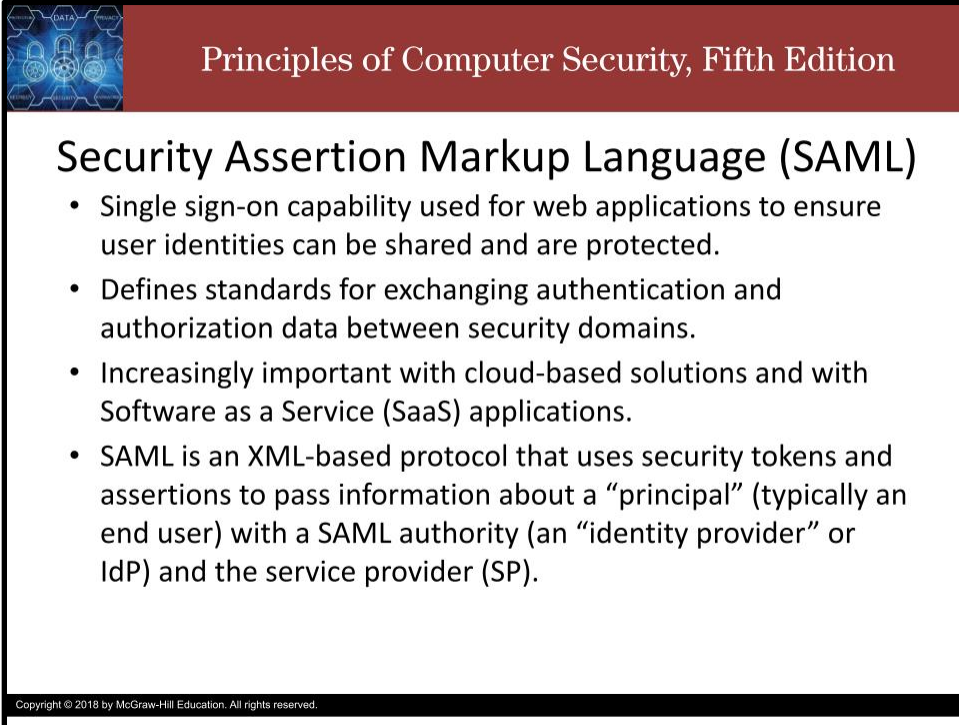
Transport layer protocol provides server authentication, confidentiality, integrity, and compression.

User authentication protocol authenticates the client to the user.

Connection protocol provides multiplexing of the encrypted tunnel into several logical channels.

SSH is very popular in the UNIX environment, and it is actively used as a method of establishing VPNs across public networks. Because all communications between the two machines are encrypted at the application layer by the two SSH daemons, this leads to the ability to build very secure solutions and even solutions that defy the ability of outside services to monitor.

## Slide 18



**Principles of Computer Security, Fifth Edition**

### Security Assertion Markup Language (SAML)

- Single sign-on capability used for web applications to ensure user identities can be shared and are protected.
- Defines standards for exchanging authentication and authorization data between security domains.
- Increasingly important with cloud-based solutions and with Software as a Service (SaaS) applications.
- SAML is an XML-based protocol that uses security tokens and assertions to pass information about a “principal” (typically an end user) with a SAML authority (an “identity provider” or IdP) and the service provider (SP).

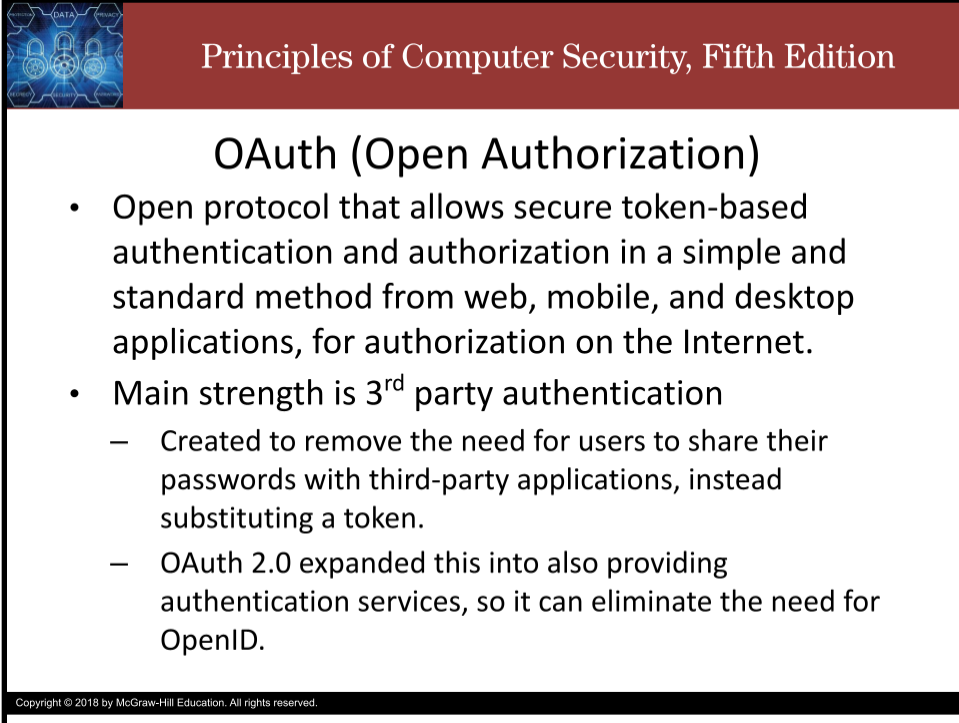
Copyright © 2018 by McGraw-Hill Education. All rights reserved.

Security Assertion Markup Language (or SAML) is an open standard for exchanging authentication and authorization data between parties, in particular, between an identity provider and a service provider. SAML is an XML-based markup language for security assertions. Security assertions are statements that service providers use to make access-control decisions. An important use case that SAML addresses is web-browser single sign-on. SAML defines standards for exchanging authentication and authorization



data between security domains and is increasingly important with cloud-based solutions and with Software as a Service applications.

## Slide 19



**Principles of Computer Security, Fifth Edition**

### OAuth (Open Authorization)

- Open protocol that allows secure token-based authentication and authorization in a simple and standard method from web, mobile, and desktop applications, for authorization on the Internet.
- Main strength is 3<sup>rd</sup> party authentication
  - Created to remove the need for users to share their passwords with third-party applications, instead substituting a token.
  - OAuth 2.0 expanded this into also providing authentication services, so it can eliminate the need for OpenID.

Copyright © 2018 by McGraw-Hill Education. All rights reserved.

OAuth, which stands for open authorization, is an open protocol that allows secure token-based authentication and authorization in a simple and standard method from web, mobile, and desktop applications, for authorization on the Internet.


OAuth 1.0 was developed by a Twitter engineer as part of the Twitter OpenID implementation.

OAuth 2.0 (which is not backward compatible with 1.0) has taken off with support from most major web platforms.

OAuth's main strength is that it can be used by an external partner site to allow access to protected data without having to re-authenticate the user.

It was created to remove the need for users to share their passwords with third-party applications, instead substituting a token.

This was expanded in OAuth 2.0 to also providing authentication services, so it can eliminate the need for OpenID.



Principles of Computer Security, Fifth Edition

## OpenID Connect

- **OpenID Connect** is a simple identity layer on top of the OAuth 2.0 protocol.
- OpenID Connect allows clients of all types to request and receive information about authenticated sessions and end users.
- OpenID is about proving who you are.
  - Authentication
- OpenID is commonly paired with OAuth 2.0

Copyright © 2018 by McGraw-Hill Education. All rights reserved.

OpenID Connect is a simple identity layer on top of the OAuth 2.0 protocol.

It allows clients of all types to request and receive information about authenticated sessions and end users.

OpenID and OAuth are typically used together and have different purposes. OpenID is used for authentication and OAuth is used for authorization.



## Security Token Service

- A **security token** service is responsible for issuing, validating, renewing, and cancelling security tokens.
- Secure tokens solve the problem of authentication across stateless platforms, because user identity must be established with each request.
- Security token service provides the same functionality as OpenID, but unlike OpenID is not patent encumbered
- Use a five-step process for using tokens where the steps are highly scalable and can be widely distributed and even shared.
  1. The user requests access with a username and password
  2. The secure token service validates the user's credentials
  3. The secure token service provides a signed token to the client
  4. The client stores that token and sends it along with every request
  5. The server verifies the token and responds with data

Copyright © 2018 by McGraw-Hill Education. All rights reserved.


A security token service is responsible for issuing, validating, renewing, and cancelling security tokens.

Security tokens solve the problem of authentication across stateless platforms, because user identity must be established with each request.

Security token service provides the same functionality as OpenID, but unlike OpenID is not patent encumbered

The five-step process for using security tokens is shown on the slide.

The steps are highly scalable and can be widely distributed and even shared.



Principles of Computer Security, Fifth Edition

## FTP/FTPS/SFTP

- File Transfer Protocol (FTP) is a plaintext protocol that operates by communicating over TCP between a client and a server.
- **FTPS** is the use of FTP over an TLS-secured channel.
- SSH FTP (**SFTP**) is not FTP over SSH.
- Secure Copy Protocol (**SCP**) is a thing. Use SFTP instead.


Copyright © 2018 by McGraw-Hill Education. All rights reserved.

One of the methods of transferring files between machines is through the use of the File Transfer Protocol (FTP). FTP is a plaintext protocol that operates by communicating over TCP between a client and a server. The client initiates a transfer with an FTP request to the server's TCP port 21. This is the control connection, and this connection remains open over the duration of the file transfer. The actual data transfer occurs on a negotiated data transfer port, typically a high-order port number. FTP was not designed to be a secure method of transferring files. If a secure method is desired, then using FTPS or SFTP is best.

FTPS is the use of FTP over a TLS-secured channel. This can be done either in explicit mode, where an AUTH TLS command is issued, or in implicit mode, where the transfer occurs over TCP port 990 for the control channel and TCP port 989 for the data channel.

SFTP is not FTP over SSH, but rather a completely separate protocol, which is an extension of the SSH protocol.

There is also SCP, which is the secure copy protocol. It is based on SSH, but the developers say it is now out of date and recommend using a more modern protocols like sftp.



Principles of Computer Security, Fifth Edition

## VPNs

- Secure virtual network built on top of a physical network.
- Security lies in the encryption of packet contents between the endpoints
- Typical use of VPN services is a user accessing a corporate data network from a home PC across the Internet.
- The sole purpose of the VPN connection is to provide a private connection between machines.
- Use many different protocols

Copyright © 2018 by McGraw-Hill Education. All rights reserved.

A virtual private network (VPN) is a secure virtual network built on top of a physical network.

The security of a VPN lies in the encryption of packet contents between the endpoints that define the VPN.

The physical network upon which a VPN is built is typically a public network, such as the Internet.

Because the packet contents between VPN endpoints are encrypted, to an outside observer on the public network, the communication is secure, and depending on how the VPN is set up, security can even extend to the two communicating parties' machines.

A typical use of VPN services is a user accessing a corporate data network from a home PC across the Internet. The employee installs VPN software from work on a home PC. This software is already configured to communicate with the corporate network's VPN endpoint; it knows the location, the protocols that will be used, and so on. When the home user wants to connect to the corporate network, they connect to the Internet and then starts the VPN software. The user can then log into the corporate network by using an appropriate authentication and authorization methodology.

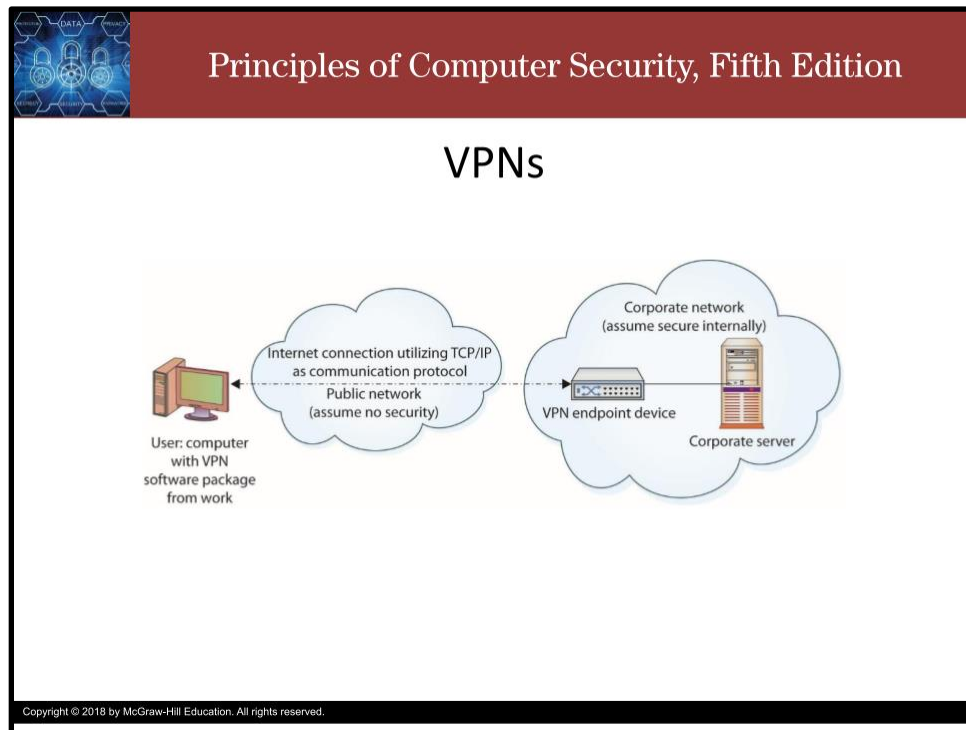
The sole purpose of the VPN connection is to provide a private connection between the machines, which encrypts any data sent between the home user's PC and the corporate network.

Identification, authorization, and all other standard functions are accomplished with the standard mechanisms for the established system.


VPNs can use many different protocols to offer a secure method of communicating between endpoints. Common methods of encryption on VPNs include PPTP, IPsec, SSH, and L2TP. The key is that both

endpoints know the protocol and share a secret. All of this necessary information is established when the VPN is set up. At the time of use, the VPN only acts as a private tunnel between the two points and does not constitute a complete security solution.

## Slide 24



Virtual private networking is not a protocol per se, but rather a method of using protocols to achieve a specific objective—secure communications—as shown in this figure. A user who wants to have a secure communication channel with a server across a public network can set up two intermediary devices, VPN endpoints, to accomplish this task. The user can communicate with their endpoint, and the server can communicate with its endpoint. The two endpoints then communicate across the public network. VPN endpoints can be software solutions, routers, or specific servers set up for specific functionality. This implies that VPN services are set up in advance and are not something negotiated on-the-fly.



Principles of Computer Security, Fifth Edition

## Vulnerabilities of Remote Access Methods


- The primary vulnerability associated with many of these methods of remote access is the passing of critical data in cleartext.
- The strength of the encryption algorithm is also a concern.
- There always exists the possibility that a flaw or defect could open the system to attack.

Copyright © 2018 by McGraw-Hill Education. All rights reserved.

The primary vulnerability associated with many of these methods of remote access is the passing of critical data in cleartext. Plaintext passing of passwords provides no security if the password is sniffed, and sniffers are easy to use on a network. Even plaintext passing of user IDs gives away information that can be correlated and possibly used by an attacker. Plaintext credential passing is one of the fundamental flaws with Telnet and is why SSH was developed. This is also one of the flaws with RADIUS and TACACS+, as they have a segment unprotected. There are methods for overcoming these limitations, although they require discipline and understanding in setting up a system.

The strength of the encryption algorithm is also a concern. Should a specific algorithm or method prove to be vulnerable, services that rely solely on it are also vulnerable. To get around this dependency, many of the protocols allow numerous encryption methods, so that should one prove vulnerable, a shift to another restores security.

As with any software implementation, there always exists the possibility that a flaw or defect could open the system to attack. Programming errors have been corrected in most software packages to close holes that made systems vulnerable, and remote access functionality is no exception. Critical flaws have been found in almost every product. The important issue is not the presence of software design and implementation errors, for as software continues to become more complex, this is an unavoidable issue. The true key is vendor responsiveness to fixing the errors once they are discovered.



## Principles of Computer Security, Fifth Edition

### Attribution

- The course slides are based on slides developed by the textbook authors, W. "Art" Conklin, Ph.D. (University of Houston), and Greg White, Ph.D. (University of Texas at San Antonio).
- These slides remain the intellectual property protected by US Copyright law of the authors and the textbook company.
- There have been some changes to the slide deck.

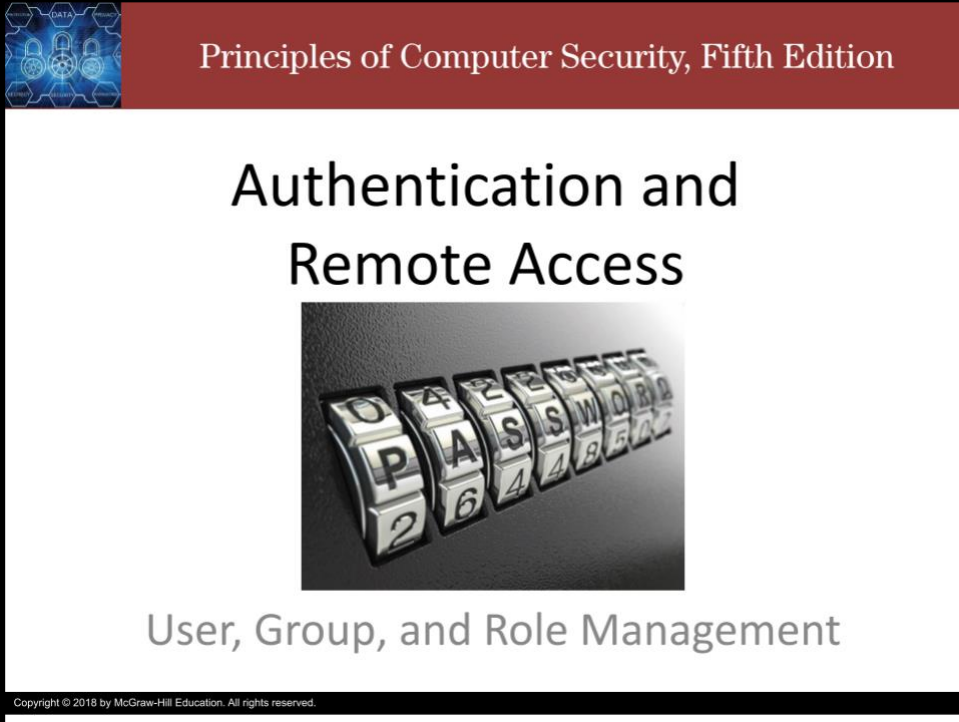
Copyright © 2018 by McGraw-Hill Education. All rights reserved.

Thank you and take care.




# Authentication and Remote Access: User, Group and Role Management

Slide 1



Principles of Computer Security, Fifth Edition

## Authentication and Remote Access




User, Group, and Role Management

Copyright © 2018 by McGraw-Hill Education. All rights reserved.

Howdy! In this video, we discuss user, group, and role management.

## Slide 2



Principles of Computer Security, Fifth Edition

### Introduction

- **Privileges** mean you have the ability to “do something” on a computer.
- **Privilege management** is the process of restricting a user’s ability to interact with the computer system.
- **Remote access** enables users outside a network to have network access and privileges as if they were inside the network.
- **Authentication** is the process of establishing a user’s identity to enable the granting of permissions.

Copyright © 2018 by McGraw-Hill Education. All rights reserved.


**Privileges** mean you have the ability to “do something” on a computer.

Essentially, everything a user can do to or with a computer system falls into the realm of **privilege management**. **Privilege management** occurs at many different points within an operating system or even within applications running on a particular operating system.

Remote access is another key issue for multiuser systems in today’s world of connected computers. Isolated computers, not connected to networks or the Internet, are rare these days. **Remote access** enables users outside a network to have network access and privileges as if they were inside the network.

**Authentication** is the process of establishing a user’s identity to enable the granting of permissions. To establish network connections, a variety of methods are used, the choice of which depends on network type, the hardware and software employed, and security requirements.

## Slide 3



Principles of Computer Security, Fifth Edition


### User, Group, and Role Management

- To effectively manage privileges, a mechanism for separating people into distinct entities (**users**) is required.
- It is convenient and efficient to be able to lump users together when granting many different people (**groups**) access to a resource at the same time.
- It is useful to be able to grant or restrict access based on a person's job or function within the organization (**roles**).

Copyright © 2018 by McGraw-Hill Education. All rights reserved.

To effectively manage privileges, a mechanism for separating people into distinct entities is required so you can control access on an individual level. At the same time, it is convenient and efficient to be able to lump users together when granting many different people access to a resource at the same time. At other times, it is useful to be able to grant or restrict access based on a person's job or function within the organization. While you can manage privileges on the basis of individuals alone, managing user, group, and role assignments together is even more convenient, efficient, and useful.

## Slide 4



### Principles of Computer Security, Fifth Edition

## User

- Single individual who access a computer system.
- Every **user** has a unique **username**
  - Must be easy to use
- **Permissions** control what the user is allowed to do with objects on the system.
  - Default = None
- Special users with special permissions
  - Administrator / root – all permissions
  - Nobody – no permissions
  - mail, news, backup, www-data

Copyright © 2018 by McGraw-Hill Education. All rights reserved.

The term **user** generally applies to any person accessing a computer system.

In privilege management, a user is a single individual, such as Alice, or Bob.

The **user** is generally the lowest level addressed by privilege management and the most common area for addressing access, rights, and capabilities.

When accessing a computer system, each user is associated with a username for identification purposes.

A **username** is a unique alphanumeric identifier the user will use to identify themselves when accessing the system.

Some systems allow users to create their own accounts, such as web applications. But, usually, accounts must be created for users by a privilege user, like a system administrator.

The system policy may allow users to request a username, or it may require that all usernames follow a specific format.

When developing a scheme for selecting **usernames**, you should keep in mind that usernames must be unique to each user, but they must also be fairly easy for the user to remember.

Once the account is created, specific permissions can be granted to the user.

**Permissions** control what the user is allowed to do with objects on the system.

Note that the default permission level is typically null – no permissions. This follows the principle of fail safe defaults. If an attacker manages to get a new user account created for themselves, it will have no permissions and will therefore be less useful to the attacker than an account with a higher level of permissions.

Not every account on a system is necessarily tied to a person in a one-to-one manner.

Some accounts are used for special functions and may be used by more than one person when the need arises.

These “special” user accounts typically have much more access and control.

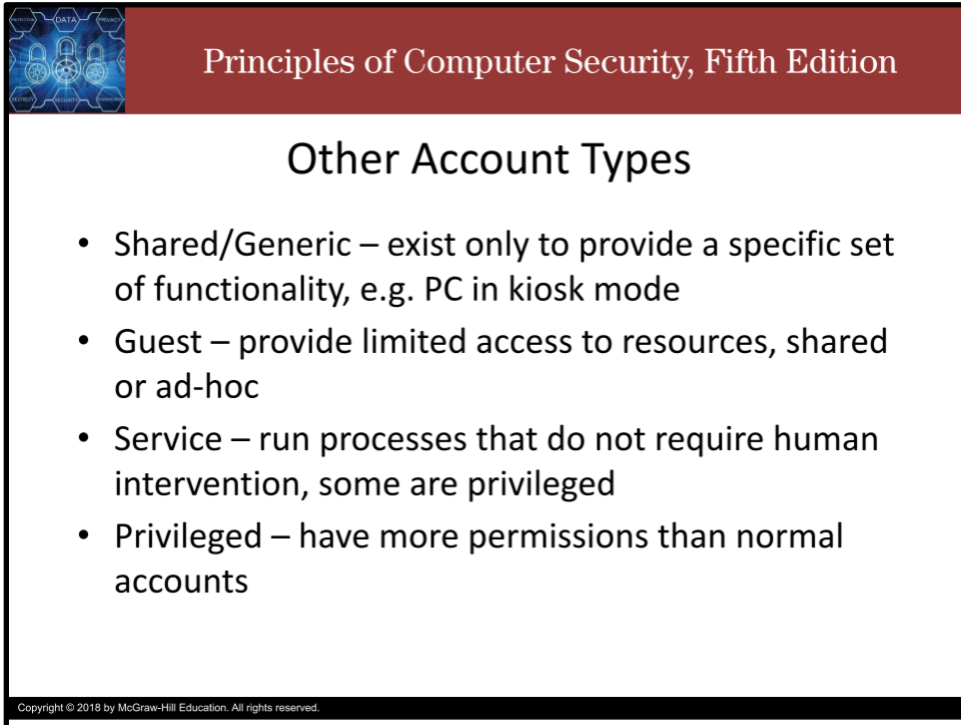
One extreme example of such an account is the superuser account (called administrator windows and root on unix), which has all the permissions. Anything that can be done on the system, the superuser can do. As such, the superuser account is a juicy target for attackers and must be protected with a very strong password.

The other extreme is the account with no permissions, which has the username “nobody” on unix.

Other examples are user accounts for managing email, backup, internet, and so on.

This follows the principles of least privilege and separation of duties. Some accounts are created with a single purpose and given only the permissions required to fulfill that purpose.

## Slide 5



The slide features a dark red header with the text "Principles of Computer Security, Fifth Edition" in white. To the left of the header is a small graphic of a blue hexagonal grid with various icons. The main content area is white with a black border. The title "Other Account Types" is centered in black. Below the title is a bulleted list of four account types. At the bottom left of the slide, there is a small copyright notice: "Copyright © 2018 by McGraw-Hill Education. All rights reserved."

Principles of Computer Security, Fifth Edition

### Other Account Types

- Shared/Generic – exist only to provide a specific set of functionality, e.g. PC in kiosk mode
- Guest – provide limited access to resources, shared or ad-hoc
- Service – run processes that do not require human intervention, some are privileged
- Privileged – have more permissions than normal accounts

Copyright © 2018 by McGraw-Hill Education. All rights reserved.

**Shared accounts** go against the specific treatise that accounts exist so that user activity can be tracked.

Shared accounts are sometimes called generic accounts and exist only to provide a specific set of functionality.

For example: a PC running in kiosk mode, with the browser limited to specific sites as an information display.

Being able to associate the activity to a specific individual is not particularly useful.

**Guest accounts** are frequently used on corporate networks to provide visitors' access to the Internet and common corporate resources, like projectors and printers.

Guest accounts are restricted in their network capability to a defined set of machines with a defined set of access.

Guest accounts may be implemented as shared accounts, or as temporary user accounts which are associated with a specific individual for a specific period of time.

**Service accounts** are used to run processes that do not require human intervention to start, stop, or administer.

Some operating systems may not even allow them to log into the system since there is no person who is supposed to use them, which limits the attack vectors that can be applied to these accounts.

You can also apply time restrictions for service accounts, such as a process that runs batch jobs only at night.

Service accounts that run in an elevated privilege mode should receive extra monitoring and scrutiny.

**Privileged accounts** are typically admin-level accounts and represent risk in that they are powerful.

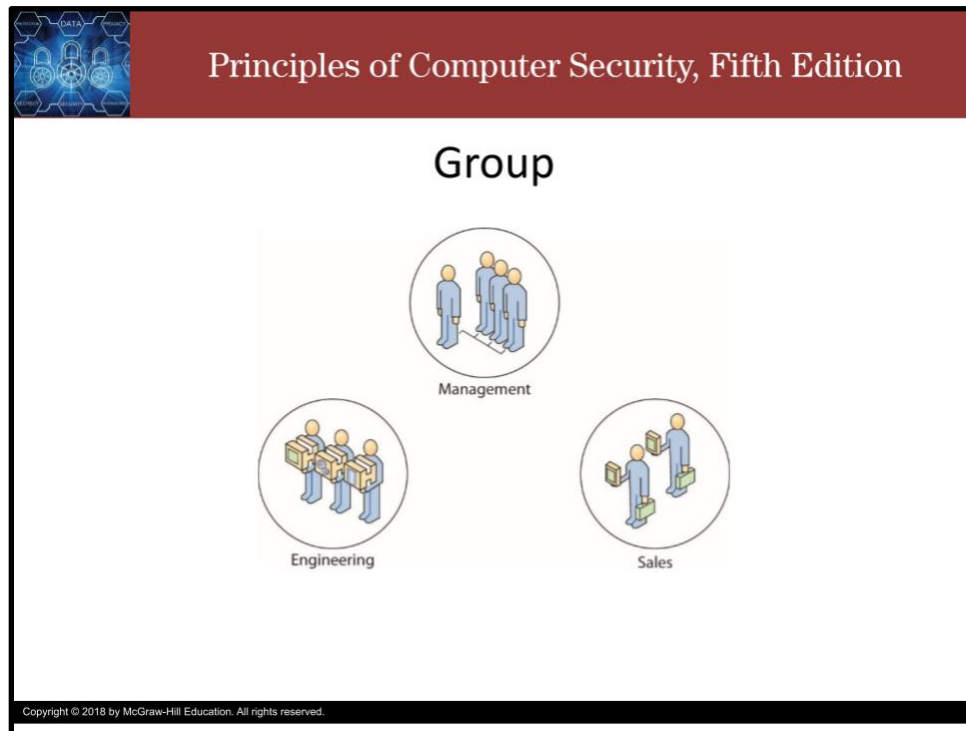
Actually, all accounts are privileged. The question is how much. The nobody account is least privileged, it can do nothing, not even login. The root account is the highest privileged account of all. It is all-powerful. Every other account has privilege somewhere in the middle.

But, when we talk of privileged accounts, or elevated privilege, what we mean are those that have greater than normal user access permissions.

Privileged accounts require regular real-time monitoring, if at all possible, and should always be monitored when operating remotely.


There may be reasons why system administrators are acting via a remote session, but when they are, the purposes should be known and approved.

## Slide 6



Under privilege management, a **group** is a collection of users with some common criteria, such as a need for access to a particular dataset or group of applications. Adding a user to a group will automatically allow that user to access the resources to which the group has access privileges. That is, a user “inherits” the permissions of the group as soon as they are placed in that group. Groups make the tasks of assigning and managing permissions easier, since permissions can be assigned to the group in one operation and all members of the group will inherit that permission. A common approach to grouping users is based on job function.

## Slide 7



Principles of Computer Security, Fifth Edition

### Role

- A **role** is usually synonymous with a job or set of functions.
- Security admins need to accomplish specific functions
  - All security admins need the same permissions.
  - For simplicity and efficiency, permissions are assigned to the role “security admin”, and individuals whose job role includes security admin are assigned the role and inherit the permissions.
- Role-based access control makes it easier to manage permissions

Copyright © 2018 by McGraw-Hill Education. All rights reserved.

A **role** is usually synonymous with a job or set of functions.

For example, the **role** of security admin in Microsoft SQL Server may be applied to someone who is responsible for creating and managing logins, reading error logs, and auditing the application. This role can belong to many people, all of whom need the same permissions.

Security admins need to accomplish specific functions, such as creating and modifying accounts, accessing error logs, and starting and stopping services.


Anyone serving in the role of security admin needs the same rights and privileges as every other security admin.

So, for simplicity and efficiency, rights and privileges can be assigned to the role of security admin, and anyone assigned to fulfill that role automatically has the correct rights and privileges to perform the required tasks.

Also, when a user changes roles, their access permissions will automatically change appropriately once they are removed from their old role and added to their new role, or, rather their old role is removed from their account and their new role is added, depending on the presentation of the role-based access control interface.



## Slide 8



### Principles of Computer Security, Fifth Edition

## Attribution


- The course slides are based on slides developed by the textbook authors, W. "Art" Conklin, Ph.D. (University of Houston), and Greg White, Ph.D. (University of Texas at San Antonio).
- These slides remain the intellectual property protected by US Copyright law of the authors and the textbook company.
- There have been some changes to the slide deck.

Copyright © 2018 by McGraw-Hill Education. All rights reserved.

Thank you and take care.


# Authentication and Remote Access: Domain Password Policy

Slide 1



Principles of Computer Security, Fifth Edition

## Authentication and Remote Access




Domain Password Policy

Copyright © 2018 by McGraw-Hill Education. All rights reserved.

Howdy! In this video, we discuss domain password policy.

## Slide 2



Principles of Computer Security, Fifth Edition

### Domain Password Policy

(Windows-specific)

- **Domain password policy** – password policy for a specific domain.
  - Usually falls under a Group Policy Object (GPO)
- **Domain controller** – computer that responds to security authentication requests.
- **Domain** – logical group of computers that share a central directory database.
  - DB contains user accounts and security info for all resources in the domain

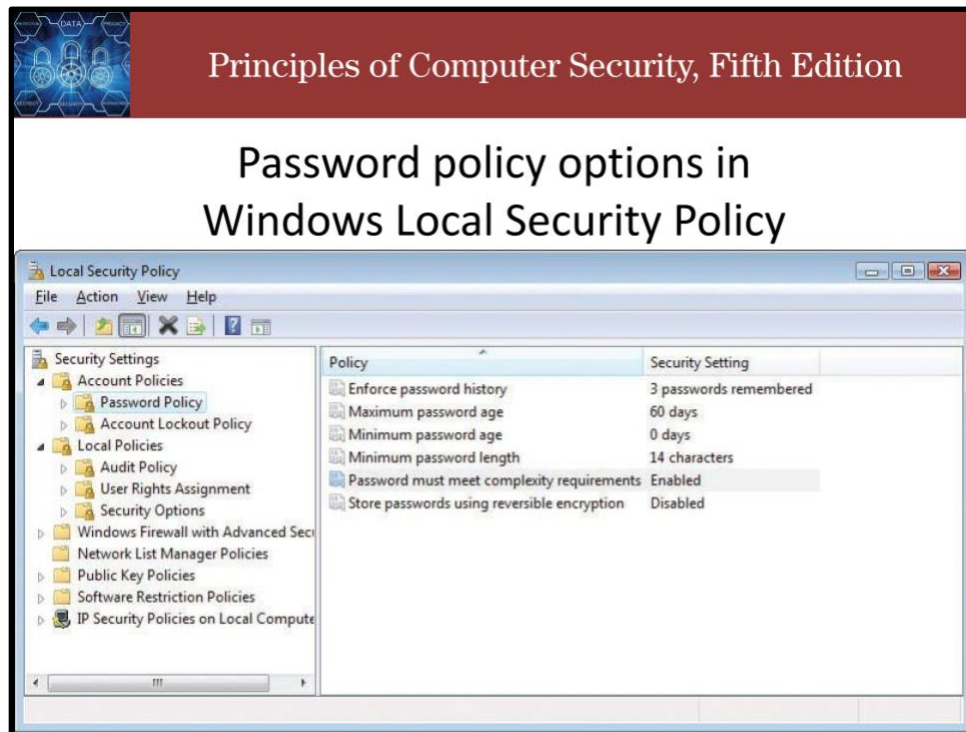
Copyright © 2018 by McGraw-Hill Education. All rights reserved.

A domain password policy is a password policy for a specific domain. The domain controller is a computer that responds to security authentication requests, such as logging into a computer. The domain password policy usually falls under a group policy object (GPO) and has several elements, which we will see on the next slide. Domains are logical groups of computers that share a central directory database, known in Windows as the Active Directory database.

The database contains information about the user accounts and security information for all resources identified within the domain. Each user within the domain is assigned his or her own unique account (that is, a domain is not a single account shared by multiple users), which is then assigned access to specific resources within the domain.

In operating systems that provide domain capabilities, the password policy is set in the root container for the domain and applies to all users within that domain. Setting a password policy for a domain is similar to setting other password policies in that the same critical elements need to be considered (password length, complexity, life, and so on). If a change to one of these elements is desired for a group of users, a new domain needs to be created because the domain is considered a security boundary. In a Windows operating system that employs Active Directory, the domain password policy can be set in the Active Directory Users and Computers menu in the Administrative Tools section of the Control Panel.

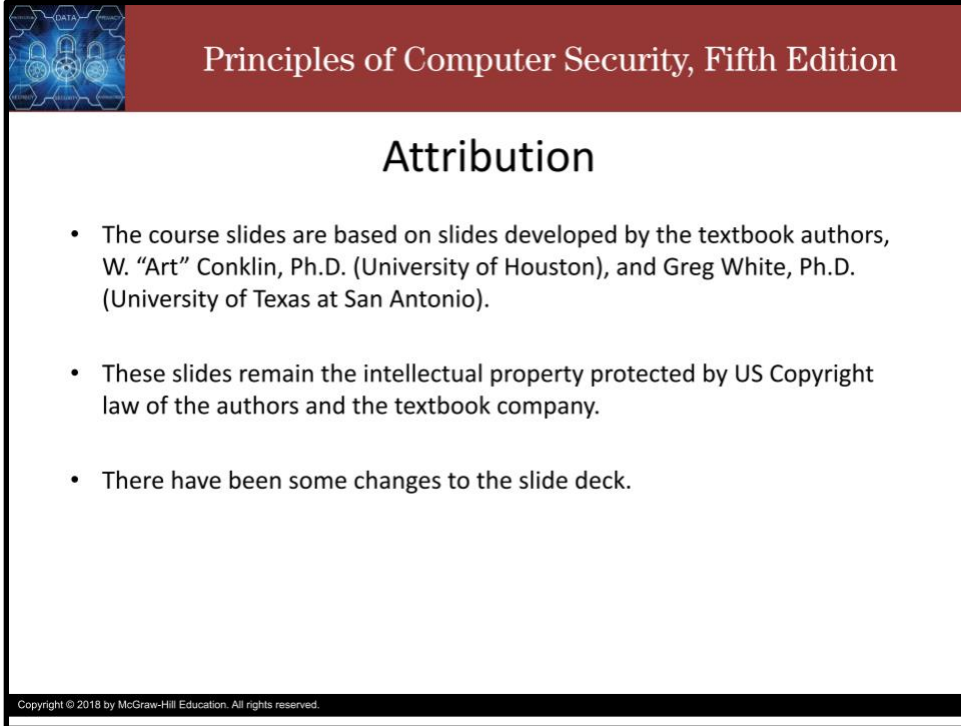
## Slide 3



This is the Local Security Policy tool in Windows. In it, we can see the elements of the group policy object for passwords. The “Enforce password history” policy tells the system how many passwords to remember and does not allow a user to reuse an old password. The “Maximum password age” policy specifies the maximum number of days a password may be used before it must be changed. The “Minimum password age” policy specifies the minimum number of days a password must be used before it can be changed again. The “Minimum password length” policy specifies the minimum number of characters that must be used in a password. The “Password must meet complexity requirements” policy specifies that the password must meet the minimum length requirement and have characters from at least three of the following four groups: English uppercase characters (A through Z), English lowercase characters (a through z), numerals (0 through 9), and non-alphabetic characters (such as !, \$, #, %).

The “Store passwords using reversible encryption” policy determines whether the system will store the password using encryption, or hashing. Encrypted passwords can be decrypted and are actually not much more secure than storing the password in plaintext. This policy should only be enabled when applications use protocols that require the user’s password for authentication (such as the Challenge-Handshake Authentication Protocol, or CHAP). Best practice with passwords is to never store them in any reversible format. That means using many iterations of a strong cryptographic hash function and large random salt values.

## Slide 4



The slide features a dark red header with a blue and white graphic of a network and a padlock on the left. The title "Principles of Computer Security, Fifth Edition" is centered in the header. The main content area is white with a black border, containing the title "Attribution" and a bulleted list. A small copyright notice is at the bottom left of the slide.

Principles of Computer Security, Fifth Edition

### Attribution

- The course slides are based on slides developed by the textbook authors, W. “Art” Conklin, Ph.D. (University of Houston), and Greg White, Ph.D. (University of Texas at San Antonio).
- These slides remain the intellectual property protected by US Copyright law of the authors and the textbook company.
- There have been some changes to the slide deck.

Copyright © 2018 by McGraw-Hill Education. All rights reserved.

Thank you and take care.

# Authentication and Remote Access: Single Sign-On

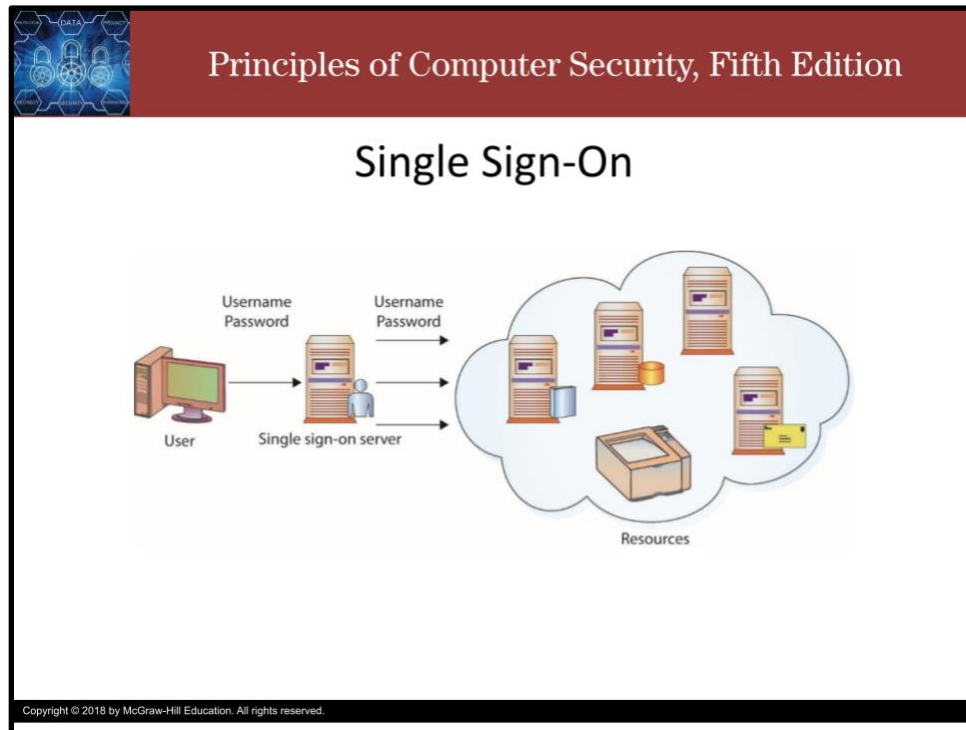
Slide 1



The slide features a dark red header with the text "Principles of Computer Security, Fifth Edition" in white. Below the header, the main title "Authentication and Remote Access" is centered in a large, black, sans-serif font. Underneath the title is a photograph of a mechanical cipher device with several rotors, showing the letters "O P A S S W O R D" and numbers "4 2 6 4 8 5 4 2" on the rotors. Below the photograph, the text "Single Sign-On" is centered in a gray, sans-serif font. At the bottom left of the slide, there is a small copyright notice: "Copyright © 2018 by McGraw-Hill Education. All rights reserved."

Howdy! In this video, we briefly introduce single sign-on.

## Slide 2




**Single sign-on (SSO)** is a form of authentication that involves the transferring of credentials between systems. SSO allows a user to transfer their credentials, so that logging into one system acts to log them into all of them. Once the user has entered a user ID and password, the single sign-on system passes these credentials transparently to other systems so that repeated logons are not required. Put simply, you supply the right username and password once and you have access to all the applications and data you need, without having to log in multiple times (and possibly having to keep track of several passwords). From a user standpoint, **SSO** means you need to remember only one username and one password.

From an administration standpoint, SSO can be easier to manage and maintain. From a security standpoint, SSO can be even more secure, as users who need to remember only one password are less likely to choose something too simple or something so complex they need to write it down. In reality, SSO is usually a little more difficult to implement than vendors would lead you to believe. To be effective and useful, all your applications need to be able to access and use the authentication provided by the SSO process.

The more diverse your network, the less likely this is to be the case. If your network, like most, contains different operating systems, custom applications, and a diverse user base, SSO may not even be a viable option. As a case in point, TAMU has used SSO for many years but only recently has the implementation been improved to cover more systems. Maybe only a year ago, signing into Howdy and eCampus required two separate authentications.

## Slide 3



Principles of Computer Security, Fifth Edition

### Attribution

- The course slides are based on slides developed by the textbook authors, W. "Art" Conklin, Ph.D. (University of Houston), and Greg White, Ph.D. (University of Texas at San Antonio).
- These slides remain the intellectual property protected by US Copyright law of the authors and the textbook company.
- There have been some changes to the slide deck.

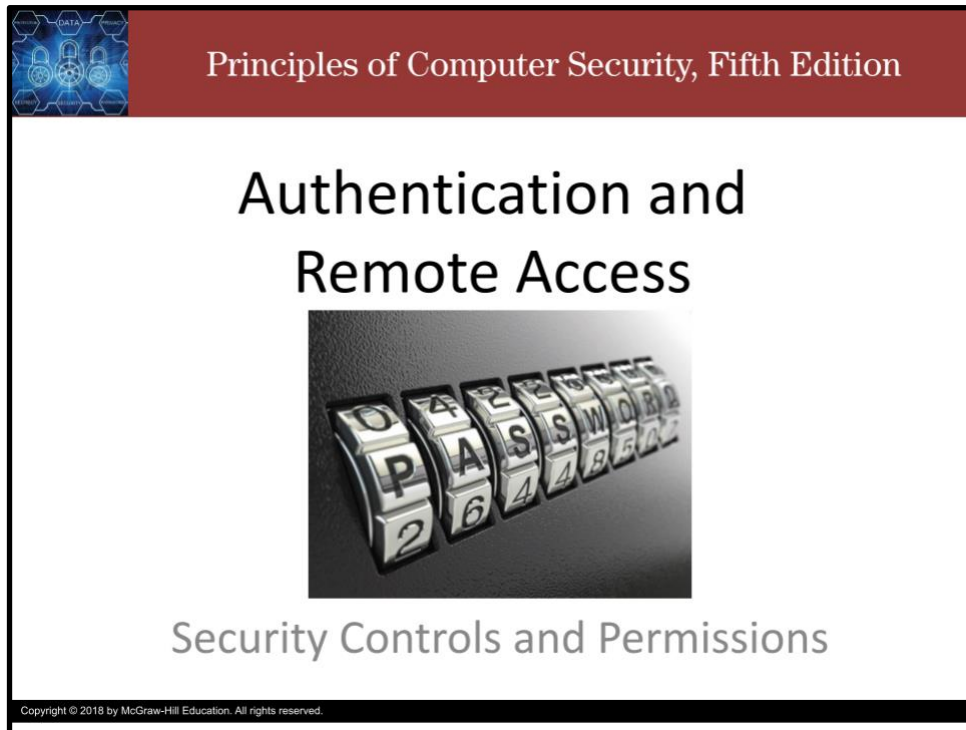
Copyright © 2018 by McGraw-Hill Education. All rights reserved.

Thank you and take care.



# Authentication and Remote Access: Security Controls and Permissions


Slide 1



The image shows the cover of a book titled "Principles of Computer Security, Fifth Edition". The cover has a dark red header with the title in white. Below the header, the main title "Authentication and Remote Access" is centered in a large, black, sans-serif font. Underneath the main title is a graphic of a password lock with several dials showing the characters "O", "4", "E", "2", "6", "W", "O", "R", "D". Below the lock graphic, the subtitle "Security Controls and Permissions" is written in a smaller, grey, sans-serif font. In the bottom left corner of the cover, there is a small copyright notice: "Copyright © 2018 by McGraw-Hill Education. All rights reserved."

Howdy! In this video, we discuss security controls and permissions.

## Slide 2



Principles of Computer Security, Fifth Edition

### Security controls and permissions


- OS uses permissions and rights to control access.
  - Permission: what actions on what objects
  - Rights: what actions on the system itself
- Example: Windows
  - Uses permissions and rights to control access
    - Permissions: read, write, execute, etc.
    - Rights: log in, remote access, logging, etc.
  - Access and use of peripherals can be controlled using permissions.
  - Remember: Least Privilege.

Copyright © 2018 by McGraw-Hill Education. All rights reserved.

Most operating systems use the concepts of permissions and rights to control and safeguard access to resources. Permissions control what a user is allowed to do with objects on the system. Rights define the actions the user can perform on the system itself. The Windows operating system provides an example.

Windows uses permissions and rights to control access to files, folders, and information resources, as well as what actions the user is allowed to perform on the system. Permissions are things like reading, writing, and executing files. Rights are things like logging on, remote access, and viewing and modifying logs. Files are not the only things that can be controlled with permissions. Even access and use of peripherals such as printers can be controlled using permissions. A very important concept to consider when assigning rights and privileges is the principle of least privilege, which requires that users be given the absolute minimum of rights and privileges required to perform their authorized duties.

## Slide 3



Principles of Computer Security, Fifth Edition

### Access Control

- **Access control list (ACL)**
  - For each object, list the subjects that have access and specify their permissions on the object
  - Fast to list all subjects with permissions on an object
  - Slow to list all objects on which a subject has permissions
  - Not very big
- **Access control matrix**
  - For every (subject, object) pair, specify the permissions that the subject has on the object.
  - Slow to list all subjects with permissions on an object
  - Slow to list all objects on which a subject has permissions
  - Very big

Copyright © 2018 by McGraw-Hill Education. All rights reserved.

The term **Access control list (ACL)** is used in more than one manner in the field of computer security. When discussing routers and firewalls, an ACL is a set of rules used to control traffic flow into or out of an interface or network.

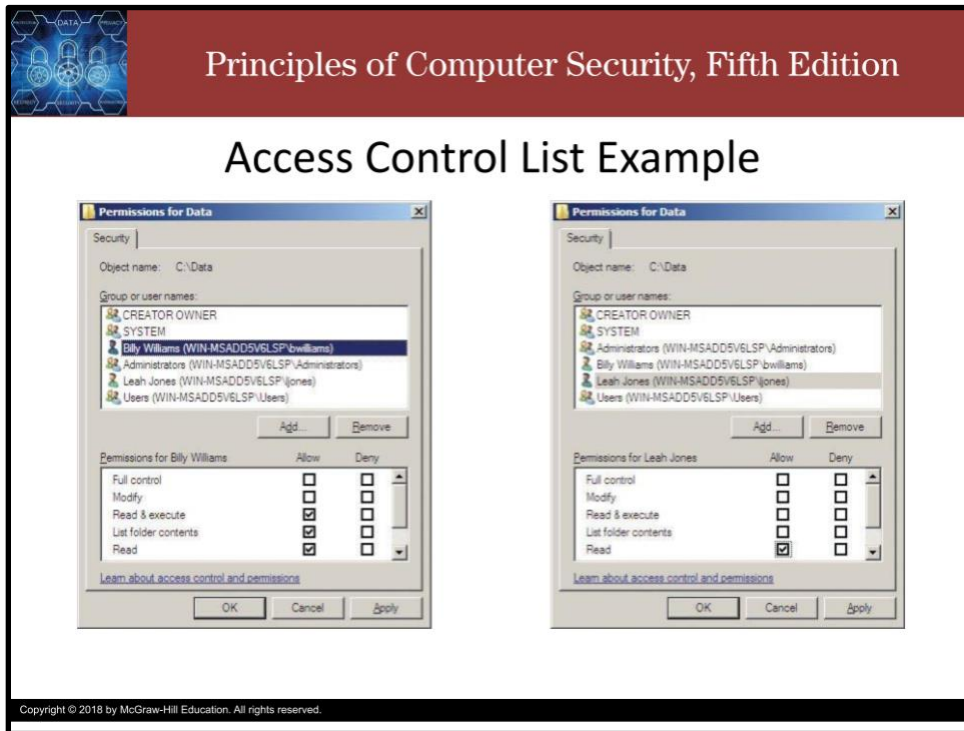
When discussing system resources, an ACL lists permissions attached to an object. Every object has a list of subjects and their permission on the object.

An alternative to an access control list is an **Access control matrix**.

The matrix is a very simple approach that consists of a giant table with subjects down one side and object across the other.

At every cell of the table, corresponding to a particular subject and a particular object, are the permissions that the subject has on the object.

Access control matrices are seldom used in computer systems because they are extremely costly in terms of storage space and processing.



This figure shows the access control list for the Data folder. On the left, the user identified as Billy Williams has Read & Execute, List Folder Contents, and Read permissions, meaning that this user can open the folder and see what's in the folder.

On the right are the permissions for a user identified as Leah Jones, who has only Read permissions on the same folder.

## Slide 5

Principles of Computer Security, Fifth Edition

### Access Control Matrix Example

	Process 1	Process 2	File 1	File 2	Printer
Process 1	Read, write, execute		Read, write	Read	Write
Process 2	Execute	Read, write, execute	Read, write	Read, write	Write

Copyright © 2018 by McGraw-Hill Education. All rights reserved.


This is a very small access control matrix for a system that has 2 subjects and 5 objects (2 of which are the subjects themselves, since all subjects are also objects).

Both subjects have all permissions over themselves. Process 2 has execute permission over process 1.

Process 1 can only read file 2, while process 2 can also write.

Imagine if the system had hundreds of subjects and thousands of objects. The table would be massive.

And, considering that most subjects would have no permissions on most objects, much of the table would be empty, which means a lot of wasted space.



Principles of Computer Security, Fifth Edition

## Mandatory Access Control (MAC)

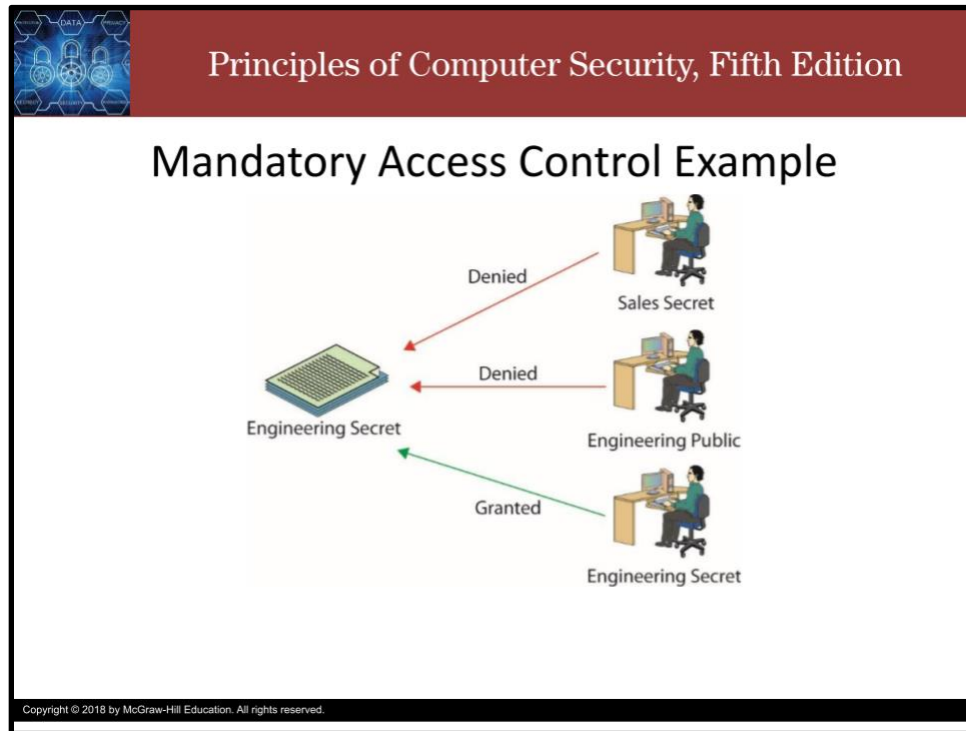
- **Mandatory access control (MAC)** is the process of controlling access to information based on the sensitivity of that information and whether or not the user is operating at the appropriate sensitivity level and has the authority to access that information.
  - Information and resources labeled with a sensitivity level
  - Users assigned a clearance level
  - Access control and sensitivity labels required in a MAC system

Copyright © 2018 by McGraw-Hill Education. All rights reserved.

**Mandatory access control (MAC)** restricts access based on the sensitivity of the information and whether or not the user has the authority to access that information.


Under a **MAC** system, each piece of information and every system resource is labeled with its sensitivity level. Users are assigned a clearance level that sets the upper boundary of the information and devices that they are allowed to access. The access control and sensitivity labels are required – one might even say “mandatory” -- in a MAC system. Labels are defined and then assigned to users and resources. Users must then operate within their assigned sensitivity and clearance levels—they don’t have the option to modify their own sensitivity levels or the levels of the information resources they create. Due to the complexity involved, MAC is typically run only on systems where security is a top priority. The Bell-LaPadula security model supports mandatory access control.

## Slide 7



This figure illustrates MAC in operation. The information resource on the left has been labeled “Engineering Secret,” meaning only users in the Engineering group operating at the Secret sensitivity level or above can access that resource. The top user is operating at the Secret level but is not a member of Engineering and is denied access to the resource. The middle user is a member of Engineering but is operating at a Public sensitivity level and is therefore denied access to the resource. The bottom user is a member of Engineering, is operating at a Secret sensitivity level, and is allowed to access the information resource.

## Slide 8



Principles of Computer Security, Fifth Edition

### Discretionary Access Control

- **Discretionary access control (DAC)** is the process of using file permissions and optional ACLs to restrict access to information based on a user's identity or group membership.
  - Most common access control system and is commonly used in both UNIX and Windows operating systems.
  - Under the DAC model, the file's owner can change the file's permissions any time they want.

Copyright © 2018 by McGraw-Hill Education. All rights reserved.

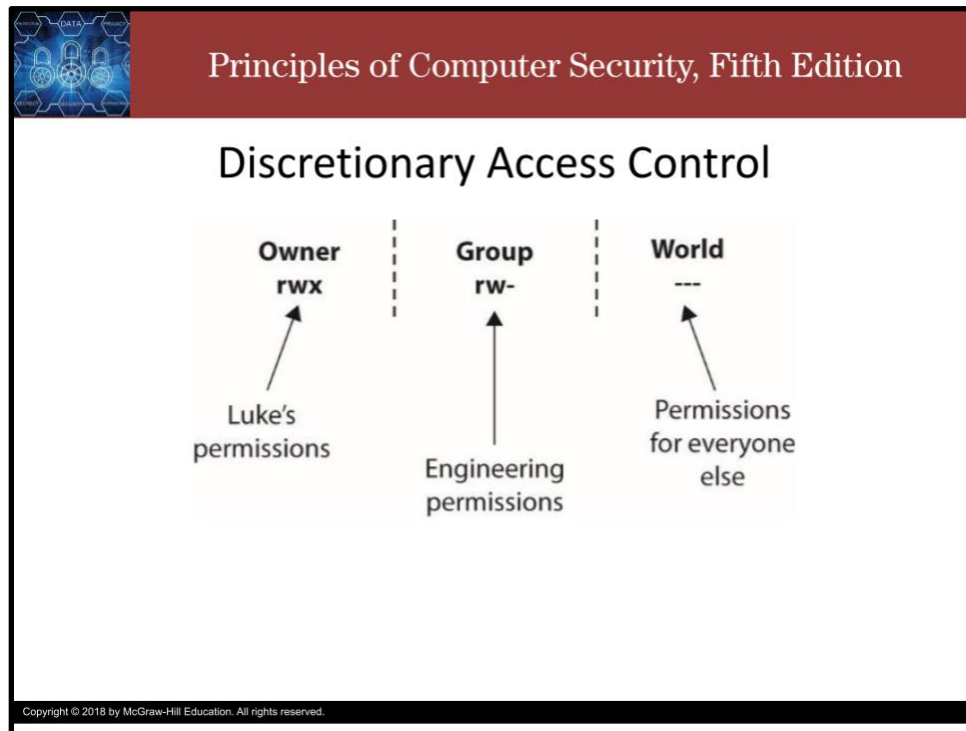
**Discretionary access control (DAC)** uses file permissions and optional ACLs to restrict access to information based on a user's identity or group membership.

DAC is a very commonly used access control system and is used in UNIX, Windows and Apple operating systems.

Under DAC, each user may, at their discretion, modify the permissions on any object of which they are the owner.



## Slide 9




In UNIX, a file's permissions are usually displayed as a series of nine characters, with the first three characters representing the owner's permissions, the second three characters representing the group permissions, and the last three characters representing the permissions for everyone else, or for the world. This concept is illustrated in this figure.

The basic permissions are Read, Write, and eXecute. If the letter is present, then the subject has the permission. If not, the subject does not have the permission.

Suppose a file is owned by Luke with group permissions for Engineering (because Luke is part of the Engineering group), and the permissions on that file are as shown in this figure. This would mean that: Luke can read, write, and execute the file (specified by the rwx in the owner section).

Members of the Engineering group can read and write the file but not execute it (specified by the rw- in the group section).

The world has no access to the file and cannot read, write, or execute it (specified by the --- in the world section).



Principles of Computer Security, Fifth Edition

## Role-Based Access Control (RBAC)

- **Role-based access control (RBAC)** is the process of managing access and privileges based on the user's assigned roles.
- RBAC is the access control model that most closely resembles an organization's structure.
- Under RBAC, you must first determine the activities that must be performed and the resources that must be accessed.
  - When a role is assigned to a specific user, the user gets all the rights and privileges assigned to that role.


Copyright © 2018 by McGraw-Hill Education. All rights reserved.

**Role-based access control (RBAC)** is an approach to managing access and privileges based on the user's assigned roles.

In this scheme, instead of each user being assigned specific access permissions for the objects associated with the computer system or network, that user is assigned a set of roles. The roles are in turn assigned the access permissions necessary to perform the tasks associated with the role. Users will thus be granted permissions to objects in terms of the specific duties they must perform—not just because of a security classification associated with individual objects.

Most organizations with many employees probably use RBAC because it most closely matches their organizational structure and can be used to implement MAC or DAC.

Role-based access control works best when most users have only one or a few roles. When many users each have many roles, the risk of compromise is very high.



Principles of Computer Security, Fifth Edition

## Rule-Based Access Control

- In **rule-based access control**, access is either allowed or denied based on a set of predefined rules.
- Each object has an associated ACL (much like DAC), and when a particular user or group attempts to access the object, the appropriate rule is applied.
- A good example for rule-based access control is permitted logon hours.
  - Many operating systems give administrators the ability to control the hours during which users can log in.

Copyright © 2018 by McGraw-Hill Education. All rights reserved.


**Rule-based access control** is yet another method of managing access and privileges (and unfortunately shares the same acronym as role-based access control; if you see “RBAC”, think role-based first).

In this method, access is either allowed or denied based on a set of predefined rules. Each object has an associated ACL (much like DAC), and when a particular user or group attempts to access the object, the appropriate rule is applied.

A good example for rule-based access control is permitted logon hours.

Many operating systems give administrators the ability to control the hours during which users can log in. For example, a bank may allow its employees to log in only between the hours of 8 A.M. and 6 P.M. Monday through Saturday. If a user attempts to log in outside of these hours, 3 A.M. on Sunday for example, then the rule will simply reject the login attempt.

## Slide 12



Principles of Computer Security, Fifth Edition

### Attribution


- The course slides are based on slides developed by the textbook authors, W. “Art” Conklin, Ph.D. (University of Houston), and Greg White, Ph.D. (University of Texas at San Antonio).
- These slides remain the intellectual property protected by US Copyright law of the authors and the textbook company.
- There have been some changes to the slide deck.

Copyright © 2018 by McGraw-Hill Education. All rights reserved.

Thank you and take care.


# Authentication and Remote Access: Preventing Data Loss or Theft

Slide 1



Principles of Computer Security, Fifth Edition

## Authentication and Remote Access



Preventing Data Loss or Theft

Copyright © 2018 by McGraw-Hill Education. All rights reserved.

Howdy! In this video, we briefly introduce data loss prevention.


## Slide 2



The slide features a dark red header with the text "Principles of Computer Security, Fifth Edition" in white. Below the header, the title "Preventing Data Loss or Theft" is centered in black. The main content area contains a stylized illustration: a tan building with a black data icon on its side, a black silhouette of a person running away from the building, and a black prohibition sign (a circle with a diagonal slash) to the right. An arrow points from the running person towards the prohibition sign. In the bottom left corner of the slide, there is a small copyright notice: "Copyright © 2018 by McGraw-Hill Education. All rights reserved."

Today's hackers are after intellectual property, business plans, competitive intelligence, personal information, credit card numbers, client records, or any other information that can be sold, traded, or manipulated for profit or advantage. This has created a whole industry of technical solutions labeled as data loss prevention (DLP) solutions. The best DLP solution is a combination of security elements, some to secure data at rest, such as encryption and access control, and some in the form of real-time monitoring, such as intrusion detection systems and traffic analysis.

## Slide 3



Principles of Computer Security, Fifth Edition

### Attribution

- The course slides are based on slides developed by the textbook authors, W. "Art" Conklin, Ph.D. (University of Houston), and Greg White, Ph.D. (University of Texas at San Antonio).
- These slides remain the intellectual property protected by US Copyright law of the authors and the textbook company.
- There have been some changes to the slide deck.

Copyright © 2018 by McGraw-Hill Education. All rights reserved.

Thank you and take care.