



European IT Certification Curriculum Self-Learning Preparatory Materials

EITC/IS/ACSS
Advanced Computer Systems Security



This document constitutes European IT Certification curriculum self-learning preparatory material for the EITC/IS/ACSS Advanced Computer Systems Security programme.

This self-learning preparatory material covers requirements of the corresponding EITC certification programme examination. It is intended to facilitate certification programme's participant learning and preparation towards the EITC/IS/ACSS Advanced Computer Systems Security programme examination. The knowledge contained within the material is sufficient to pass the corresponding EITC certification examination in regard to relevant curriculum parts. The document specifies the knowledge and skills that participants of the EITC/IS/ACSS Advanced Computer Systems Security certification programme should have in order to attain the corresponding EITC certificate.

Disclaimer

This document has been automatically generated and published based on the most recent updates of the EITC/IS/ACSS Advanced Computer Systems Security certification programme curriculum as published on its relevant webpage, accessible at:

<https://eitca.org/certification/eitc-is-acss-advanced-computer-systems-security/>

As such, despite every effort to make it complete and corresponding with the current EITC curriculum it may contain inaccuracies and incomplete sections, subject to ongoing updates and corrections directly on the EITC webpage. No warranty is given by EITCI as a publisher in regard to completeness of the information contained within the document and neither shall EITCI be responsible or liable for any errors, omissions, inaccuracies, losses or damages whatsoever arising by virtue of such information or any instructions or advice contained within this publication. Changes in the document may be made by EITCI at its own discretion and at any time without notice, to maintain relevance of the self-learning material with the most current EITC curriculum. The self-learning preparatory material is provided by EITCI free of charge and does not constitute the paid certification service, the costs of which cover examination, certification and verification procedures, as well as related infrastructures.

TABLE OF CONTENTS

Mobile security	4
Mobile device security	4
Mobile app security	12
Security analysis	21
Symbolic execution	21
Network security	30
Web security model	30
Network security	40
Secure channels	41
Certificates	42
Implementing practical information security	43
Information security in real life	43
Messaging	44
Messaging security	44
Security of storage	45
Untrusted storage servers	45
Timing attacks	46
CPU timing attacks	46

EITC/IS/ACSS ADVANCED COMPUTER SYSTEMS SECURITY DIDACTIC MATERIALS**LESSON: MOBILE SECURITY****TOPIC: MOBILE DEVICE SECURITY**

Mobile Device Security

Mobile device security is a crucial aspect of cybersecurity, and in this material, we will explore the topic in detail. Over the course of two subsequent materials, we will delve into both the platform and application level security of mobile devices. Today, we will focus on the lower levels of the stack and then we will discuss application level security, specifically in the context of Android and Apple.

Mobile device security has gained significant importance due to the fact that it is generally considered to be better than desktop security. This is partly because mobile devices came later, allowing developers to learn from past security issues and improve their security measures. By studying mobile device security, we can gain valuable insights and learn from their advancements.

One interesting aspect to consider is the concept of enclave support. Enclave support involves running sensitive computations separately from the operating system, thereby minimizing the need to trust the operating system entirely. This concept is also prevalent in the design of mobile device security, where a more hardened implementation is often preferred over a software-based approach.

To begin, let's consider the threat model associated with mobile device security. The core issue we need to address is the scenario where an attacker steals the phone and attempts to access the data stored on it. This threat model assumes that the phone is locked when stolen and that it is password protected. If the phone is not locked or lacks password protection, the attacker can easily unlock the phone and access the data.

To mitigate this threat, Apple, for instance, aims to ensure that even if the attacker has physical access to the stolen phone, they cannot read the user's files. This objective requires significant engineering and design efforts on Apple's part. While the threat model may seem straightforward, the actual implementation requires meticulous planning and execution.

Why does Apple invest so much effort into mobile device security? There are several plausible reasons. Firstly, Apple has customers who rely on their devices to store sensitive information, such as medical or financial data. It is crucial for Apple to provide these customers with robust security measures to protect their confidentiality. Additionally, there are ordinary customers who want assurance that their passwords and personal information are secure, especially when accessing sensitive services like online banking.

Furthermore, Apple likely sees a competitive advantage in offering better security measures for their devices. By addressing this specific threat model effectively, they can position themselves as a trusted provider of secure mobile devices.

Mobile device security is a critical aspect of cybersecurity. By studying the advancements made in this field, we can learn valuable lessons and apply them to other areas of security. Apple's extensive efforts in mobile device security highlight the importance of protecting user data and meeting the needs of customers who require strong security measures.

Mobile device security is a crucial aspect of cybersecurity, especially considering the increasing reliance on smartphones for personal and professional purposes. In this context, it is essential to understand the potential attacks that can compromise the security of mobile devices.

One of the primary concerns is an exhaustive password search. If an attacker gains physical access to the phone, they can attempt various password combinations to gain unauthorized access. Most pin codes are relatively short, making it easier for attackers to guess the correct combination. To mitigate this risk, it is necessary to implement measures that limit the number of password guesses, ensuring that attackers cannot easily break into the device.

Another significant concern is the impersonation of fingerprints or face IDs. Mobile devices use biometric authentication methods to enhance security. However, attackers may attempt to impersonate these biometrics

to gain unauthorized access. To address this issue, robust security measures must be in place to detect and prevent such impersonation attempts.

The removal of flash chips from mobile devices poses another security risk. Flash chips store sensitive data, and attackers may attempt to extract this data by removing the chips and inserting them into another device. To counter this threat, it is crucial to implement measures that protect the integrity and confidentiality of data stored in flash chips.

Bugs in the kernel and driver software of mobile operating systems also pose significant security concerns. The kernel is a critical component of the operating system, and any vulnerabilities in it can be exploited by attackers. Similarly, bugs in driver software can provide attackers with unauthorized control over the device. It is essential to regularly update and patch the operating system and drivers to mitigate these risks.

Physical access to a mobile device opens up the possibility of booting an attacker's operating system on the device. This grants the attacker unrestricted control over the device and its functionalities. To prevent this, secure boot mechanisms should be implemented to ensure that only authorized and verified software runs on the device.

In addition to secure boot, it is crucial to protect against the downgrade or rollback of the operating system. If an attacker can revert the device to a previous version of the operating system with known vulnerabilities, they can exploit these vulnerabilities to gain unauthorized access. Measures should be in place to prevent such downgrades and ensure that devices are always running the latest secure version of the operating system.

It is worth noting that mobile device security has significantly improved over the years. Measures discussed in this material have made it considerably harder for attackers to steal data from mobile devices. The effectiveness of these measures is evident from the fact that even law enforcement agencies, such as the FBI, have faced difficulties in accessing data from locked devices. This demonstrates the success of the security measures implemented by mobile device manufacturers.

Mobile device security is a critical aspect of cybersecurity. Understanding the potential attacks and implementing appropriate security measures is crucial to protect sensitive data and ensure the integrity of mobile devices.

Mobile device security is a crucial aspect of cybersecurity, especially considering the sensitive information that is stored and transmitted through these devices. In order to ensure the security of mobile devices, advanced computer systems employ various mechanisms and components.

One of the key components in mobile device security is the secure enclave chip. This chip is responsible for running sensitive computations and protecting against potential attacks. It is separate from the main CPU and is connected to it through a dedicated processor. The main CPU, on the other hand, runs the operating system and applications.

The secure enclave chip has its own read-only memory (ROM) which stores the necessary information for it to start running. This memory is inaccessible from the rest of the phone, ensuring its security. Additionally, the secure enclave chip contains a unique identifier (UID) which is a 256-bit AES key. This key is fused into the hardware during manufacturing and cannot be read or written by anyone, including the secure processor itself. However, data can be encrypted or decrypted using this key.

To protect user data, the secure enclave chip manages a set of keys that are used for encryption. Whenever the operating system wants to write or access user files, the data is immediately encrypted with these keys. The encryption and decryption processes are handled by a DMA (Direct Memory Access) engine that sits between the main memory and the flash chips containing user data. This ensures that data flowing between memory and flash storage is always encrypted.

The encrypted data stored in the flash chips provides an additional layer of security, as unauthorized access to the data is prevented. Only the secure enclave chip has access to the keys required for decryption, guaranteeing that the data remains secure even if the operating system is compromised.

Mobile device security relies on the presence of a secure enclave chip that handles sensitive computations and

manages encryption keys. This chip is separate from the main CPU and ensures the encryption of data stored in flash memory. By employing these mechanisms, mobile devices can provide a higher level of security for user data.

When it comes to mobile device security, one of the key aspects to consider is secure boot. Secure boot ensures that when a mobile device is turned on, the secure enclave processor runs Apple code that was specifically written for it, rather than code supplied by an attacker. This is crucial in preventing unauthorized access to the device. Another important issue is the division of responsibilities between the operating system and the enclave processor. The enclave processor is typically responsible for user authentication, such as PIN entry, to ensure that the operating system does not have access to sensitive information like the PIN. This requires a secure path from the user to the enclave processor to protect against potential attacks.

Data encryption is also a critical component of mobile device security. The encryption process involves setting up a key by the enclave processor and encrypting the data on the file system using an AES engine. However, the encryption process becomes more complex when the device is in lock mode, as certain background applications may still require access to specific files. For example, downloading email attachments in the background should still be possible. This adds another layer of complexity to the data encryption story.

To understand the secure boot process, it is important to note that booting a mobile device involves multiple stages. The first stage is the boot ROM, which is read-only and contains code burned into it during manufacturing. The boot ROM initiates the boot process and hands over control to subsequent bootloaders. The reason for having multiple stages is to allow for updates to the operating system without modifying the boot ROM. If all the boot code were in the boot ROM, updating the operating system would be impossible.

Mobile device security requires a robust secure boot process, clear division of responsibilities between the operating system and the enclave processor, and careful consideration of data encryption, especially in lock mode. These measures help protect against unauthorized access and ensure the confidentiality and integrity of user data.

In the context of mobile device security, it is crucial to ensure that the code running on the device is trustworthy. One of the key components in this process is the boot ROM, which contains important information, including Apple's public key. This public key is accessible to anyone and is used to verify the authenticity of the code during the boot process.

However, it is important to note that even with these security measures in place, there are still potential attack surfaces. One such surface is the theft of Apple's private key. If an attacker gains access to this key, they can potentially upgrade any phone with malicious code, compromising the device's security.

It is important to understand that no security plan is completely foolproof. There will always be trade-offs between security, performance, and usability. In this case, the protection of Apple's private key is of utmost importance. While the exact details of how it is stored and protected are not provided, it is likely that it is not connected to the internet and is safeguarded by a robust protection scheme.

During the boot process, the boot code verifies the signature of the iBoot code using Apple's private key. If the signature checks out, there is confidence that the code is supplied by Apple. This process is repeated for every piece of code that runs after the boot process, ensuring that each code segment is signed with Apple's private key.

After the boot process, the iBoot code initializes various sensors and other components before booting the enclave kernel. The enclave kernel follows the same process of verifying the code's signature using Apple's private key.

While these measures provide a high level of confidence in the authenticity of the code running on the device, it is important to be aware that no security system is perfect. There will always be potential vulnerabilities that can be exploited by determined attackers.

Mobile device security is a crucial aspect of cybersecurity, as mobile devices have become an integral part of our daily lives. In this didactic material, we will explore the differences between mobile device security and desktop security, and why mobile security is considered to be more advanced.

When it comes to mobile device security, one of the key components is the concept of enclaves. Enclaves are secure areas within the device's hardware where sensitive data and processes are stored and executed. These enclaves have their own dedicated processors and run code that is produced by the device manufacturer, such as Apple. This ensures that the code running on the device is authentic and trustworthy.

The process of booting up a mobile device involves several steps. First, the enclaves kernel is checked to ensure its integrity. Once the kernel is verified, the enclaves processor starts running the enclaves code. At this point, the device is operational, and sensors and other components are activated. However, the main CPU, which runs the operating system (iOS) and applications, is not yet running.

To fully boot up the device, there are additional steps that need to be taken. This includes running UEFI (Unified Extensible Firmware Interface) code, which is a standard for booting up devices. The UEFI code is responsible for booting up iOS and launching the applications. It is important to note that the enclaves continue to run in parallel with the main CPU, ensuring the security of the device.

Mobile device security has several advantages over desktop security. Firstly, mobile devices have the advantage of starting from scratch in terms of design. This means that the manufacturers were able to incorporate the lessons learned from past security incidents in desktop environments. As a result, mobile devices were able to implement stronger security measures right from the beginning.

Secondly, desktop environments are more open and diverse compared to the standardized nature of mobile devices. Desktops can run different operating systems and are produced by various manufacturers. This diversity makes it more challenging for attackers to exploit vulnerabilities consistently across different systems. On the other hand, mobile devices are more homogeneous, making it easier to implement consistent security measures.

Another factor that contributes to the advanced security of mobile devices is the portability factor. Mobile devices, such as iPhones, are more likely to be stolen or misplaced compared to desktop computers. This means that attackers have easier access to mobile devices, increasing the need for robust security measures.

In terms of technical mechanisms, mobile device security often relies on cryptographic techniques. Before running any code, a cryptographic hash is computed over the code to generate a short byte or bit string. This string is then signed by the device manufacturer. When the code is executed, the signature is validated, and the hash is compared to the computed hash of the stored code. If the signatures do not match, the device enters recovery mode, allowing the user to perform a factory reset or restore the device using iTunes.

Mobile device security has evolved to be more advanced compared to desktop security. This is due to several factors, including the clean slate design of mobile devices, the standardized nature of the hardware and software, and the increased focus on portability. By incorporating enclaves, cryptographic techniques, and strict validation processes, mobile devices have become more secure and trustworthy.

After successfully recovering a mobile device and ensuring that the correct software is installed, there is still a potential security issue to address - the downgrade attack. In this type of attack, an attacker gains access to an older version of the software that is signed with the device manufacturer's private key. When the device boots up, it runs this older version of the software, which may have known vulnerabilities that can be exploited by the attacker.

To mitigate this risk, mobile device manufacturers, such as Apple, have implemented a protection mechanism called the downgrade protection attack plan. This plan utilizes a unique identifier called an EC ID, which is stored in the read-only memory of the device. The EC ID is a 64-bit number that uniquely identifies the device.

Here is an overview of how the downgrade protection attack plan works:

1. When a user decides to upgrade their mobile device, they send a message to the device manufacturer's central server (e.g., Apple). This message includes the EC ID, a nonce (a random number used only once), and measurements (cryptographic hashes) of the version of the software they want to install or upgrade to.
2. The device manufacturer verifies the EC ID and checks the current software version running on the device.

Based on this information, they determine the latest version available for that specific EC ID. If the requested software version is older than the latest version, it is not provided.

3. The device manufacturer then sends a response message to the device, which includes the measurements, EC ID, and nonce - all signed with the manufacturer's private key.
4. Upon receiving the response message, the device checks the signature, verifies the EC ID matches its own, and checks the nonce to ensure it is not a replayed message from an older version.
5. If all the checks pass, the device proceeds with the upgrade.

This scheme enhances security by preventing the installation of older software versions that may contain known vulnerabilities. It also prevents the installation of software intended for a different device (due to the EC ID check) and protects against replay attacks (due to the nonce check).

By implementing the downgrade protection attack plan, mobile device manufacturers can ensure that only authorized and up-to-date software versions are installed on their devices, reducing the risk of exploitation through known vulnerabilities.

Mobile device security is a critical aspect of cybersecurity, particularly when it comes to advanced computer systems security. One important consideration in mobile security is the management of software versions on mobile devices. Storing the version number in flash memory and checking against it can be a potential solution. However, attackers may attempt to change the version number, which poses a challenge. While the version number is typically stored in memory, it is difficult to modify without proper access.

Another important component of mobile device security is the secure enclave. The secure enclave is responsible for user authentication and limiting password guessing. It serves as a central function in the mobile device's security system. Trusting the main CPU with user authentication can be risky, as any exploit or bug in the operating system could compromise user authentication and potentially give attackers access to sensitive data. Therefore, user authentication is performed on the secure enclave separately from the main CPU.

To ensure the security of the communication between the sensor (e.g., fingerprint sensor) and the secure enclave, a shared cryptographic key is used. This key is established during the manufacturing process of the mobile device. The sensor and the enclave processor both have access to this shared key, which is used to establish a secure communication channel. This prevents potential attacks where a compromised operating system could record sensitive data, such as fingerprints, and replay them later to gain unauthorized access to the device.

Mobile device security involves managing software versions, utilizing a secure enclave for user authentication, and protecting the communication between sensors and the secure enclave. These measures help safeguard mobile devices against potential threats and ensure the security of user data.

Mobile device security is a critical aspect of cybersecurity, especially considering the sensitive data that these devices often store. In this material, we will discuss the advanced security measures implemented in mobile devices, such as iPhones, to protect user data.

These devices have a specialized component called the enclave processor, which plays a crucial role in securing the communication and data stored on the device. The enclave processor acts as a computing engine and encrypts and signs the sensor data using a unique key called KS. It then sends this encrypted data to the main CPU and memory system for further processing.

The enclave processor is responsible for verifying the integrity and confidentiality of the data it receives. It has the necessary key (KS) to decrypt the data and ensure that the sensor data remains secure. This means that the operating system does not have access to the actual fingerprint or other sensitive information. Instead, it only sees the encrypted version of the data.

To protect against replay attacks, the sensor establishes a secure connection, such as an SSL or SSH connection, with the enclave processor. This secure channel ensures the confidentiality, integrity, and freshness of the data being transmitted. Although the specific cryptographic protocol used is not described in the material,

it is essential for maintaining a secure communication channel.

The enclave processor also enforces a guessing limit to prevent unauthorized access to the device. If an attacker attempts to guess the password or fingerprint multiple times and fails, the device will wipe all data. This limits the attacker's budget for guessing and makes it extremely difficult to gain access to the device.

Some users may question the effectiveness of biometric authentication methods like Face ID or fingerprint ID, which typically have a relatively low probability of false positives. However, these methods are still considered secure due to the additional protection provided by the password or passcode. In the case of iPhones, users are required to enter the passcode after a limited number of unsuccessful attempts with biometric authentication.

The primary reason for incorporating biometric authentication methods is to encourage users to lock their devices when not in use. Apple believes that the password or passcode is the strongest form of authentication, and the biometric methods serve as a convenient way to lock the device. By making it easier for users to lock their devices, Apple aims to ensure that sensitive data remains protected when the device is not in use.

Mobile device security, particularly the security of sensitive data, is a crucial aspect of cybersecurity. Advanced security measures, such as the enclave processor and biometric authentication methods, are implemented to protect user data. These measures ensure the confidentiality, integrity, and freshness of the data, making it extremely difficult for attackers to gain unauthorized access to the device.

In the realm of mobile device security, one crucial aspect is the protection of user data through encryption and user authentication. When a phone enters recovery mode, it undergoes a hardware reset. However, the specific details of what occurs during recovery mode are not fully known. In the event that a user forgets their password or PIN, there is no way to recover access to the device other than performing a factory reset and going through the recovery mode operation.

To prevent unauthorized access, the iPhone employs a security feature called the enclave processor, which enforces limits on password guessing. However, it is possible to bypass these limits by exploiting the enclave processor. User authentication is closely tied to encryption. When a user locks their phone, the key used to encrypt and decrypt their data is no longer stored on the device. This ensures that even if an attacker gains access to a locked phone, the user key has been erased and the encrypted data remains secure.

To address the challenge of reestablishing the user key upon unlocking the phone, a scheme is employed. The user key is made dependent on the password or PIN code. The key generation process involves encrypting the Unique Identifier (UID) of the device along with the password several times, resulting in the final key. This ensures that only a user who possesses the correct password can generate the corresponding key. Additionally, the key is also dependent on the UID, making it impossible to read the data by transferring the flash chips to a different phone.

This scheme also serves to prevent unauthorized access to data. By encrypting the UID and password multiple times, it becomes computationally intensive to guess the password or construct the key. This mechanism adds an extra layer of security to protect against brute-force attacks.

However, there is a challenge when it comes to background applications. The iPhone allows applications to run even when the phone is locked. This introduces the need for data encryption in the background. The specifics of how this is achieved are not discussed in the provided material.

Mobile device security involves various measures to protect user data. Recovery mode allows for hardware resets, but the exact details are unknown. Forgetting the password or PIN results in the need for a factory reset. User authentication and encryption are tightly integrated, with the user key being dependent on the password and UID. This prevents unauthorized access and data decryption. The scheme also makes it computationally intensive to guess the password or construct the key. However, there are challenges when it comes to running background applications and data encryption.

Mobile device security is a critical aspect of cybersecurity, especially considering the increasing use of smartphones and the sensitive information they store. One challenge in mobile security is how to protect user data while still allowing background applications to function effectively. This is particularly important when it comes to handling email attachments or incoming phone calls.

To address this issue, a common practice is to implement a technique called key wrapping. Key wrapping involves encrypting one key with another key, allowing for secure delegation of access. In the context of mobile device security, this means that sensitive user keys are not directly shared with background applications. Instead, a wrapped key is provided, which can be decrypted by the application to gain access to the necessary data.

In the case of file encryption, a top-level key, referred to as the key FS, is used for the file system. This key is stored on the flash chip and is responsible for encrypting the file system metadata, which includes information about file locations and other relevant details. During the boot process, the key FS is loaded into the DMA engines or the operating system, allowing for read and write operations on the metadata.

To ensure the security of the key FS, it is encrypted using a unique identifier (UID) associated with the device. This encryption makes it difficult to extract the flash chip and use it in another device, as the encrypted key cannot be decrypted without the corresponding UID. Additionally, this encryption provides a cost-effective method of erasing the content of the disk. By discarding the key FS, the data on the file system becomes inaccessible.

In addition to the key FS, each file has its own metadata, which is also encrypted using the key FS. This metadata includes the data protection level, which determines the level of security applied to the file. The operating system does not have direct access to the key FS but can request the enclave processor to retrieve the metadata and set up the necessary protections.

Within the metadata, there are specific keys known as the K data protection file keys. These keys are generated based on the data protection level and are used to encrypt and decrypt the actual file contents. By using different keys for different data protection levels, a layered approach to security is achieved, allowing for granular control over file access.

Mobile device security employs key wrapping techniques to protect user data while allowing background applications to function. The key FS is used to encrypt file system metadata, and each file has its own metadata encrypted using the key FS. Data protection levels determine the keys used to encrypt and decrypt file contents, providing a layered approach to security.

In the realm of mobile device security, one crucial aspect to consider is the protection of files and data stored on the device. Mobile operating systems employ various techniques to ensure the security of these files, with one such technique being the use of Key Derivation Functions (KDFs) and Key File Systems (KFS).

To understand how this works, let's delve into the concept of KFS encryption. The KFS itself is encrypted using a KDF, which means that in order for the operating system to access and read a file, it needs access to the KDF. This intermediate step allows for the implementation of different key delegation schemes, as the KFS can be computed in different ways depending on the level of protection required for the file.

There are several protection levels that can be applied to files in mobile device security. The first level is called "complete," where the KDF is derived from the user's pin code. In this case, when the phone is locked, the pin code is erased, and the files cannot be read or written by background applications. This is the most secure level of protection.

The next level is "complete unless open," which allows background applications to write or append to files but not read them. The enclave processor constructs a special KDF for these applications, which sits in the memory of the processor while the phone is locked. This level of protection is commonly used for applications like email, where downloading messages and writing them into the file system is necessary.

The third level is "until first authentication," where the KDF is constructed the first time the user types their password after booting the phone. Even if the phone is locked, files protected using this level of protection can still be read and written by the operating system. However, it is important to note that this level of protection is intended for third-party applications and should not be used for highly sensitive data.

The fourth level of protection is essentially no protection other than encryption with the base file system key. This level is typically used for background applications. However, it is worth mentioning that there are

challenges in implementing this level of protection, which require key chaining to ensure its effectiveness.

Mobile device security employs various levels of file protection, with encryption being a fundamental aspect. Key Derivation Functions (KDFs) and Key File Systems (KFS) play a crucial role in securing files and data on mobile devices. Understanding these concepts is essential for designing robust security measures for mobile devices.

Mobile device security is a critical aspect of cybersecurity, as these devices often contain sensitive data and are vulnerable to various attacks. One such attack is the physical tampering of the device, which is expensive and challenging to execute. The attacker would need to replace the device's read-only memory (ROM) with a new ROM, while ensuring that the tampering remains undetected. This process is made difficult by the tamper-resistant design of modern devices.

However, the most commonly exploited attack surface for iPhones is through software vulnerabilities. Attackers target bugs in iOS or the driver software running on the device's main CPU. Gaining control over iOS allows attackers to access sensitive data stored in the device's memory. This is because the filesystem runs on the main CPU, and when an iOS app requests access to encrypted data, the enclaves processor sets up the appropriate decryption key. When the device is unlocked, an attacker with control over iOS can read the user's unencrypted data from the main memory.

It is important to note that iOS only protects user data when the device is locked. When the device is unlocked, iOS needs access to the data to run applications, making it a profitable target for attackers. However, attacking iPhones in this manner is challenging due to the security measures implemented by Apple.

Zero-day exploits, which are previously unknown vulnerabilities, can be used to exploit iOS and gain control over the device. These exploits are highly valuable and can cost up to a million dollars on the market. Companies specializing in manufacturing zero-day exploits sell them to interested buyers. The price of these exploits has slightly decreased in recent months due to the discovery of multiple bugs.

Another potential attack vector is targeting the booting process of the device. If an attacker can manipulate the booting process and insert their own code, they can gain control over the device. By booting their own software, the attacker can execute malicious actions on the device.

Mobile device security, particularly for iPhones, is a complex and challenging field. Apple has implemented various security measures to protect user data, making it difficult for attackers to exploit vulnerabilities. However, software bugs and zero-day exploits can still be used to gain control over iOS and access user data. Understanding these attack vectors is crucial for developing effective countermeasures and ensuring the security of mobile devices.

EITC/IS/ACSS ADVANCED COMPUTER SYSTEMS SECURITY DIDACTIC MATERIALS**LESSON: MOBILE SECURITY****TOPIC: MOBILE APP SECURITY**

In this material we will consider mobile phone security, focusing specifically on the security of mobile applications. To provide some context, we previously covered securing the device itself against theft and potential compromises to the operating system. For this material, we will assume that the operating system is secure and the device has not been stolen.

When it comes to securing applications on mobile devices, there are several security properties to consider. One important aspect is the protection of data stored on the device. For example, you may have personal contacts or photos stored on your device, and it is crucial to determine which applications should have access to this data.

In addition to data protection, mobile devices also have various sensors and peripherals, such as cameras and location information. Applications may need to access these sensors and peripherals, raising questions about how to ensure secure access.

Furthermore, we will explore situations where different applications need to interact with each other. It is important to establish a security framework that governs these interactions and protects against potential vulnerabilities.

To delve deeper into the topic, we have chosen to discuss Android and iOS platforms. The reason behind this choice is not to make a judgment on which one is better, but rather because the iOS paper provides a more detailed explanation of the lower-level security measures, while the Android paper focuses on the app-level isolation.

Before we explore how Android isolates applications from each other, it is helpful to understand the alternative models for running applications. In the world of desktop applications, the security model is different. Operating systems like Windows or Linux do provide some isolation, but it is primarily based on the user level rather than individual apps. Permissions are set based on user access, and applications run with full user privileges. This means that if you run an app on your desktop, it has access to all your files, certificates, passwords, and other processes running on your computer. There is no isolation between apps in the traditional desktop sense.

While this lack of isolation may seem concerning, there are some advantages to this model. It allows for easy sharing of files between applications and provides some level of isolation between multiple users in a multi-user computer system.

In contrast to desktop apps, we now have mobile apps like Android and iOS. These platforms operate under a different security model. We will explore the specifics of these models in the upcoming materials.

Mobile app security is a crucial aspect of cybersecurity, especially in the context of advanced computer systems security. Before the emergence of Android, mobile devices ran on embedded Windows systems. However, these systems lacked strong security measures. Another model that was considered was using web applications for mobile phones. This would have eliminated the need for a new app model and instead used JavaScript-based web apps. Web apps offer app isolation, meaning that each app operates independently and cannot access data from other apps. This model seemed reasonable, as it provided security through isolation.

However, there were certain challenges with using web apps. One of the challenges was the requirement for an online server and the need to communicate with it. But even during the development of Android, there were already browsers that supported offline web apps. This challenge could have been mitigated by saving HTML and JavaScript files locally. Another challenge was the lack of a traditional file system in web apps. While data could be stored in web apps, it was not the same as having a dedicated file system. Despite these challenges, the lack of sharing capabilities was a major issue with web apps. It was difficult to share data between different web apps, such as transferring photos from one app to another. While some popular apps had integrated sharing capabilities, there was no general-purpose solution for sharing data between web apps.

Android aimed to address these challenges and provide a solution for mobile app security and sharing. Android

apps are similar to Linux processes, with each app being a binary executable file that runs as a process on the Linux kernel. Each app has its own directory within the Android file system, where it can store and access its files and data. Additionally, Android apps have a manifest file that contains important information about the app, such as its permissions and components. This manifest file plays a crucial role in the security and functionality of Android apps.

The emergence of Android revolutionized mobile app security by introducing a new app model that addressed the limitations of web apps. Android apps operate as individual processes on the Linux kernel, with each app having its own directory for file storage. The manifest file provides essential information about the app, ensuring its security and functionality. By overcoming the challenges of web apps, Android has become a robust platform for secure and efficient mobile app development.

An Android application consists of various components, including code, a manifest file, and a private key signature from the developer. These components are bundled together and signed to create an Android application package (APK) file. When you interact with an Android app, you are essentially running this APK file.

Android has a strict isolation plan to ensure security. Each application runs in its own isolated environment, with its own data directory and process. Android configures the system in a way that these applications cannot tamper with each other. This isolation is achieved by assigning a unique user identifier (UID) to each process and setting permissions on the directories of each application.

While these applications are isolated, there is still a need for sharing data between them. Android achieves this through a mechanism called intents. Intents are messages sent between applications to request or share data. To send an intent, the application sends it to the kernel, which then passes it to a component called the reference monitor. The reference monitor is responsible for enforcing the security policy and determining where the intent should go. It may modify the intent before forwarding it to the recipient application.

The reference monitor plays a crucial role in app-to-app sharing and is in charge of enforcing security policies. It reads messages from the kernel, analyzes them, and may modify them before re-injecting them into the kernel for delivery.

Android applications are designed to run in isolated environments to ensure security. Intents are used to facilitate sharing between applications, with the reference monitor overseeing the process and enforcing security policies.

In mobile app security, one important aspect is the handling of intent messages. An app may need to send a message to another app that is not currently running. This is where the reference monitor comes in. Its job is to start the recipient app as a running Linux process so that it can receive and process the message. Android provides infrastructure for apps to handle these messages, and this process is typically written in Java.

Android developers use a Java library to interact with the reference monitor and dispatch intents. This library is automatically started for the app when the reference monitor determines that a message is being received. It takes care of the details of handling incoming messages. Within an application, Android uses components to specify where messages should go. These components are named by strings and are part of the Java process. However, some apps may want to move beyond Java and send the intent to another location. Android allows apps to spawn different processes and forward messages, as long as they are running under the same user ID and have the same level of security.

In Android, there are several well-known component types. Activities are components that represent screens in the app. You can send an intent to an activity to trigger it to appear on the screen. Other component types include database components and service components for remote procedure call (RPC) servers.

While Android does not restrict the use of languages other than Java, the standard bindings and infrastructure are written in Java. When an app is started, the reference monitor launches a Java Virtual Machine (VM) with some shim code. From there, the app can choose to fork off and execute a regular Linux elf binary. The Java bindings are also useful for interacting with intent messages. Android does not directly expose the kernel API for dealing with intents, so developers are encouraged to use the Java library for convenience and compatibility across different Android versions and devices.

The choice of using Java in Android development was driven by the need for portability. When Android was first developed, the future processors for mobile phones were uncertain. By using Java, the developers ensured that apps could easily run on different hardware platforms without requiring extensive recompilation. This portability aspect was a significant benefit at the time.

Mobile app security involves understanding how intent messages are handled in Android. The reference monitor starts recipient apps as running Linux processes, and the Java library provides the necessary infrastructure for apps to handle these messages. Component types, such as activities, help disambiguate where messages should go within an app. While Android allows flexibility in using different languages, Java is widely used due to its convenience, compatibility, and portability.

Mobile app security is a crucial aspect of cybersecurity in today's digital landscape. In this didactic material, we will explore the concept of mobile app security, specifically focusing on the security of intents in Android applications.

Intents in Android apps serve as a means of communication between different components within an application or between different applications. To understand the security implications of intents, it is important to delve into their internal structure.

An intent can be thought of as having three main fields. The first field is the component name, which specifies the destination of the intent. In Android, each app has a unique name, and components within an app can be referenced using a Java-style naming convention. For example, "com.google.android.dialer" could refer to the dialing application on an Android phone.

The second field is the action field, which is a string that defines the desired action to be performed by the recipient of the intent. Actions are completely defined by the apps themselves, allowing each app to specify its own set of actions. Android provides a predefined set of actions that cover common functionalities such as making a phone call, viewing a document, or looking at a map.

The third field is the data field, which is similar to a URI or URL. It provides additional information or payload for the intent. For example, if the desired action is to make a phone call, the data field could contain the phone number to be dialed.

One interesting aspect of Android intents is that the component name is optional. If the component name is not specified, the intent is called an implicit intent. In such cases, the responsibility of determining the appropriate recipient falls on the reference monitor. The reference monitor examines the action and data fields of the intent and consults its table of installed applications to identify the suitable recipient. If there is a single matching application, the intent is sent directly to that application. However, if multiple applications can handle the intent, the user is prompted to choose the desired application.

This flexibility of implicit intents allows for the seamless integration of multiple applications, even if they are unaware of each other's existence. For example, if a user wants to view a PDF document and has multiple PDF viewers installed, the reference monitor will prompt the user to select the preferred viewer for that particular action and data type.

By understanding the structure and behavior of intents, developers can design their applications to ensure secure communication and prevent unauthorized access or misuse of sensitive data.

Mobile app security, particularly regarding intents in Android applications, is a critical aspect of cybersecurity. By comprehending the internal structure of intents and the role of the reference monitor in handling them, developers can enhance the security of their applications and ensure seamless interoperability between different apps.

In the context of mobile security, it is important to understand the differences between iOS and Android in terms of configurability and flexibility. While iOS has some level of configurability, Android takes it to a different level by offering many more configurable options. For example, on Android, users can change their home screen, open the phone dialer, or choose a different web browser, which are not possible on iOS. Android allows for a standard set of actions for various functions, such as opening a map. On the other hand, iOS is more opinionated and restricts certain customization options.

This difference in configurability can be attributed to Google's desire to make Android applicable in a wider range of settings. Android aims to be flexible and adaptable for various use cases, even beyond smartphones. For instance, Android can be used in devices like refrigerators or vending machines, where the vendor may want to replace the home screen with a customized interface.

The difficulty in modifying Android's functionality lies in the fact that the kernel is designed to prevent unauthorized actions. The Android framework configures the kernel to give each app a separate User ID (UID) and access to its own files, while restricting access to kernel-neutral data structures and other apps. As long as the kernel is set up correctly and free of bugs, it is difficult to break its security measures. The kernel only allows the sending of intents, which are monitored and controlled by the reference monitor, preventing unauthorized access.

Permissions play a crucial role in Android's security model. Permissions are controlled through the use of labels, which are strings that define the privileges and protections associated with a component. For example, the label "dial perm" represents the permission required for applications or components involved in making phone calls. Each application has a list of privileges, which are positive actions or capabilities granted to the application. The reference monitor is aware of these privileges.

On the recipient side, components within applications have specific roles and permissions. For instance, a "dial" component within the phone application may require the "dial perm" permission. The reference monitor verifies the permissions of the sender and recipient before allowing communication between them.

Android offers more configurability and flexibility compared to iOS. Its security model relies on the kernel's isolation capabilities, permissions, and the reference monitor's control over intents. Understanding these concepts is crucial in ensuring the security of mobile applications.

In the context of mobile app security, one important aspect to consider is the protection of sensitive functionalities, such as making phone calls, within an application. To achieve this, Android provides a mechanism called permissions. Permissions are labels associated with components in an Android application, and they determine which privileges a caller must have in order to interact with that component.

For example, let's consider a phone app that has a component responsible for receiving intents to make a phone call. To protect this functionality, the developer can associate a permission label with this component, such as "dial_perm". This label acts as a restriction, allowing only callers with the "dial_perm" privilege to send an intent to this particular component. Other components within the phone app may have different permission labels, but this specific component is protected by the "dial_perm" permission.

Now, you might wonder where these permission labels come from. In Android, even the permissions themselves are defined by applications. This means that application developers have the flexibility to define their own set of permissions, allowing them to replace or redefine system components as needed. However, Android does provide a number of standard permissions that come pre-installed on your phone, such as "dial_perm" for the built-in phone dialer and "camera" for the camera app.

The information about permissions is stored in the application manifest. The manifest is an important part of an Android application and contains various details about the application, including permissions. There are three main aspects related to permissions in the manifest:

1. Defining new permissions: An application can define new permission labels to indicate the resources or functionalities it provides. For example, the phone app may define a permission label indicating which apps are allowed to make phone calls. This helps other apps and users understand and agree on the app's intended usage.
2. Requesting privileges: An application can request a list of permission labels that it needs to function properly. This list is included in the manifest and is presented to the user during the installation process. The user then has the choice to approve or reject these requested permissions. This user consent is an important aspect of Android's security model.
3. Inter-app communication: When one app wants to interact with another app's protected functionality, it needs

to request the appropriate permission. For example, if a third-party app wants to make phone calls using the phone app, it needs to request the "dial_perm" privilege.

Permissions play a crucial role in mobile app security by allowing developers to control and protect sensitive functionalities within their applications. Android provides a flexible system where permissions can be defined, requested, and granted based on user consent.

Mobile app security is a crucial aspect of ensuring the overall security of mobile devices. In order to understand mobile app security, it is important to first understand the concept of app privileges and how they are managed within the Android operating system.

Internally, a phone app has various privileges that allow it to interact with the device's hardware and software components. However, these privileges do not automatically transfer to other apps that are installed on the device. Instead, other apps can only send an intent to the phone app, which then exercises its own privileges internally. This means that the new app does not directly access the phone app's privileges or data.

To facilitate this communication between apps, Android uses a messaging system. Each app has its own files in the file system, and one app can send a message to another app to request access to certain data. However, direct access to another app's files is not allowed. Instead, the requesting app must send a message to the target app, which will then read the data and send back a response, if necessary. This messaging system ensures that all communication between apps is mediated by a reference monitor, which helps prevent unauthorized access to sensitive data.

The Android manifest file plays a crucial role in defining the privileges and protection labels for each component of an app. By default, if a component is not specified in the manifest, it is considered private and cannot be accessed by other apps, regardless of their permissions. However, if an app wants to allow other apps to send messages to a specific component, it can specify the protection label for that component in the manifest. This allows other apps with the required permissions to send messages to the protected component.

In order for apps to communicate effectively, they need to agree on certain conventions. For example, they need to agree on the action string that represents a specific functionality, such as making a phone call. Android provides a list of standard action strings to ensure consistency among apps. By using these standard action strings, apps can communicate with each other without any issues. Additionally, apps need to request the necessary privileges in their manifest files to perform certain actions. For example, an app that wants to make phone calls needs to request the "dial" privilege.

It is important to note that the app initiating the message may not know which specific phone app is installed on the device. This is why it is crucial for apps to agree on the action string and the required privileges. By doing so, any app that supports the agreed-upon action can handle the message and perform the necessary actions.

Mobile app security in the context of Android involves managing app privileges and communication between apps through a messaging system. The Android manifest file plays a crucial role in defining app privileges and protection labels for each component. By adhering to standard action strings and requesting the necessary privileges, apps can securely communicate with each other.

Mobile app security is an important aspect of cybersecurity, especially in the context of advanced computer systems security. In the Android framework, apps are defined using a manifest file which specifies the permissions and privileges that the app requires. This manifest file is crucial because it determines what actions the app can perform on the device.

One interesting feature of Android is that any app can request permissions to access various device features such as the camera, location, and SMS messages. This flexibility allows for customization and innovation but also raises concerns about potential misuse of these permissions. Unlike iOS, which has a more restrictive approach, Android gives users more control over their device by allowing them to grant or deny these permissions.

Initially, Android had a plan to inform users about the list of privileges an app required during the installation process. However, this approach was deemed insufficient, and a new plan was devised. In the current approach, the app still declares its permissions in the manifest file, but when the app attempts to use a privilege for the

first time, a dialog box pops up on the user's device, asking for their consent. This mechanism allows users to make informed decisions about granting or denying specific privileges to apps.

It is worth noting that any app can request any privilege, with one exception. The exception is related to implicit intents, which are defined in the manifest file to handle certain types of actions. For example, an alarm clock app could declare that it can handle any intent, which could potentially lead to unintended consequences. When multiple apps can handle the same intent, the user is prompted to choose which app to use. This can be problematic if a malicious app is handling intents that it shouldn't.

The question of whether Android's approach to app permissions is good or bad is still open for debate. On one hand, it provides users with more control over their devices and allows for customization. On the other hand, it raises concerns about potential misuse and the possibility of granting permissions to malicious apps.

To address these concerns, Android offers a mode similar to iOS, where users can choose to only install apps from Google's Play Store. This ensures that apps have been carefully reviewed and approved by Google developers. This approach is aimed at providing a more secure environment for users who prefer a stricter control over app permissions.

Mobile app security is a critical aspect of cybersecurity. Android's approach to app permissions allows for flexibility and customization but also raises concerns about potential misuse. Users have the power to grant or deny permissions to apps, but it is important to be cautious and make informed decisions. Android offers options for stricter control over app permissions, such as installing apps only from Google's Play Store.

Google's data center is an important aspect of mobile app security. It allows apps to run in virtual machines and access the network. Google plays a significant role in catching potential security threats. However, the mindset of focusing on securing app components rather than the messages being sent can be questioned. For example, when opening a document, it is the document itself that needs to be protected, not necessarily the component. There are extensions in the model that attempt to address this issue, although they may not be as effective as the current plan.

One potential reason why the Google Play Store app on Android devices requests certain privileges is for sound or keyboard input. However, it is uncertain if this is the exact reason. This raises the question of whether time of use confirmation would be a better approach than upfront permission approval. With time of use confirmation, users have the opportunity to deny permissions if they are not relevant or understood. This may result in app crashes, but at least the user has the ability to say no.

In Android, there is a labeling system for permissions. When defining a new label, a string is chosen along with a description and a type. The description is what appears when installing an app, informing the user about the permission being requested. The registration of labels is first-come, first-served, which can cause issues if important labels are not registered first. Additionally, subsequent apps cannot redefine existing labels, which can create problems if a malicious app defines a permission that conflicts with a legitimate app, such as Facebook.

Android categorizes permissions into three types: normal, dangerous, and signature. Normal permissions are not considered crucial and do not prompt the user for consent. Examples include changing the wallpaper or ringtone. Dangerous permissions require explicit consent from the user and are considered more critical. Signature permissions are special cases where only apps from the same developer can request them.

Mobile app security involves considering the role of Google's data center, the mindset of component security versus message security, the use of permissions and their types, and the potential benefits of time of use confirmation.

Mobile app security is an important aspect of cybersecurity. In this context, permissions play a crucial role in determining what actions an app can perform on a user's device. While some permissions may seem unnecessary or intrusive, they serve a purpose in terms of defense-in-depth and preventing potential issues.

One type of permission is the dangerous permission, which includes actions such as sending SMS messages or making phone calls. These permissions prompt the user for approval either during the installation process or when the action is requested. This ensures that the user is aware of and consents to these potentially risky

actions.

Another type of permission is the signature permission. These permissions are unique to apps from the same developer and are restricted to only those apps. They allow for secure communication between different apps from the same developer while preventing unauthorized access. For example, if a developer has multiple apps, they may want them to share certain information internally, but not allow any random app on the user's device to access that information.

The process of installing an app involves the Google Playstore app downloading the app's APK file and then sending an intent to the reference monitor, which is responsible for managing the installation process. The reference monitor acts as a gatekeeper, prompting the user for approval before installing the app. This ensures that the user has the final say in whether or not to allow the app to be installed.

App signatures serve an important purpose in ensuring the authenticity and integrity of the app. When an app is installed for the first time on an Android device, the user does not know who the app should be signed by. App signatures provide a way for the user to verify the source and legitimacy of the app. This is especially important in preventing the installation of malicious or unauthorized apps.

Mobile app security relies on permissions and app signatures to protect user devices and data. Dangerous permissions prompt the user for approval for potentially risky actions, while signature permissions ensure secure communication between apps from the same developer. App signatures provide a means of verifying the authenticity and integrity of an app before installation.

When it comes to mobile app security, there are several important aspects to consider. One of these is the use of signatures. Signatures serve two purposes: signature permissions and updates.

Signature permissions are crucial because they determine the trustworthiness of the app. When you install an app for the first time, you trust that the public key used for the signature is valid. However, for subsequent updates, it is essential that they are signed by the same public key as the initial install. This ensures that the updates are legitimate and not tampered with.

In practice, there are two common approaches to ensuring the trustworthiness of app installations. The first is to only install apps from trusted sources like Google's App Store. The second is to be cautious about downloading apps from websites, ensuring that the website is secure (HTTPS) and trusted.

Another important consideration is the handling of dependencies. In the Android world, there is no built-in support for managing complicated dependency structures. Apps cannot rely on the framework to install dependencies automatically. Instead, developers must inform users about any required dependencies and ask them to install them separately.

Android employs a mandatory access control (MAC) system to enforce security policies on top of existing application code. The MAC system separates the application code from the policy enforcement, which is done by the reference monitor. This approach allows for easier analysis of the security policies and provides protection against certain types of vulnerabilities even if the application code is buggy.

One advantage of the MAC system is that it allows for the analysis of the security policies on a device. Users can determine which apps have access to certain resources, such as contacts or PDF files.

In a MAC system, default policies can be set to be secure, ensuring that no component can communicate with any other component unless explicitly allowed. This adds an extra layer of protection against unauthorized access.

Mobile app security, particularly in the context of Android, involves considering signature permissions, handling dependencies, and implementing a mandatory access control system. By following best practices such as installing apps from trusted sources and being cautious when downloading from websites, users can mitigate potential security risks. The MAC system provides an additional layer of protection by separating the application code from the security policies and enforcing them to prevent unauthorized access.

In the context of mobile app security, it is important to understand the concept of access control policies and

how they are enforced in different operating systems. In this didactic material, we will discuss the differences between mandatory access control (MAC) and discretionary access control (DAC) in relation to mobile app security.

In MAC, there are manifests that specify the access control policies, and a reference monitor enforces these policies on the app's code. This means that the policies are set externally and cannot be changed by the app itself. On the other hand, DAC, which is commonly found in Unix-like operating systems, relies on permissions set on files and apps. Apps can set their own permissions, which can lead to potential security vulnerabilities if the app code is buggy or malicious.

Android, as a mobile operating system, takes a different approach. It aims to help app developers do the right thing by default. Android uses a form of MAC, where the access control policies are defined in manifests and enforced by a reference monitor. This ensures that the defaults are secure and reduces the risk of apps unintentionally introducing vulnerabilities.

However, there are some exceptions to Android's MAC model. One such exception is related to remote procedure calls (RPC) between services. In Android, an RPC component can accept messages from other apps. While the manifest allows setting a label for protecting the RPC component, it is often necessary to have finer-grained control over who can access the component. Android addresses this by allowing the RPC component to check the caller and make decisions based on the caller's identity. This exception deviates from the strict MAC model but provides more flexibility in managing access to RPC services.

Another exception in Android's permission model is related to broadcast intents. Android uses broadcast intents to send messages between apps. While MAC policies are typically enforced for access control, Android allows apps to send broadcast intents without explicitly checking the access control policies. This means that anyone can send a broadcast intent, which can introduce potential security risks. However, Android internally implements a more complex plan to handle these broadcast intents securely.

These exceptions highlight the trade-offs that mobile operating systems like Android make to balance security and flexibility. While MAC policies work well at a coarse-grained scale, finer-grained policy decisions often require cooperation from the app's code.

Android's approach to mobile app security involves the use of mandatory access control (MAC) policies enforced by a reference monitor. This helps ensure secure defaults and reduces the risk of apps introducing vulnerabilities. However, there are exceptions to this model, such as the handling of RPC between services and broadcast intents, which provide more flexibility but require additional security considerations.

In the context of mobile security, one of the challenges is handling intents, which are messages used to communicate between different components of an Android application. The problem arises when the intent itself contains sensitive information that needs to be protected. For example, when a text message is received on an Android phone, an SMS received intent is sent to all apps that subscribe to this intent. The payload of this intent contains the actual text message. The question then becomes, who should be able to receive these intents?

To address this issue, Android introduced a way to annotate the intent itself with a label. This label is similar to a permission and is attached to the intent when it is sent. The label specifies which apps have the privilege to receive the intent. For example, there is a permission called "receive SMS" that can be used as a label. If an app has this permission in its privileged list, it is allowed to receive the SMS intent.

This approach allows for protecting the contents of the intent rather than just the component that receives it. However, it may seem slightly clunky compared to other security mechanisms in Android. The reason for this is that Android aims to give users the ability to make choices about which apps can communicate with each other. Hardcoding knowledge of different apps into the code would make the system more rigid and less flexible.

In addition to protecting intents, Android also introduced enterprise policies that allow for multiple user profiles on a phone. These profiles can have different access control policies, similar to a mandatory access control system. For example, an app running in a company user profile may have different privileges compared to the same app running in a personal user profile.

Mobile app security involves protecting the contents of intents and implementing access control policies to ensure that only authorized apps can receive sensitive information. Android provides mechanisms such as labeling intents and enterprise policies to address these security challenges.

In the realm of mobile app security, it is essential to have a comprehensive plan in place to ensure the protection of sensitive data and prevent unauthorized access. One approach to achieving this is through the implementation of mandatory access control (MAC). MAC works effectively at a coarse granularity, allowing for broad policy decisions to be imposed on existing apps. However, when it comes to fine-grained situations, such as fine-grained access control or other complex scenarios, it becomes more challenging to enforce policies on an app from the outside.

To address these challenges, a sophisticated plan for application-level access controls is necessary. This plan should include a secure kernel and secure boot, as discussed in our previous session. Additionally, it should incorporate mechanisms for interposing app interactions and facilitating seamless integration of new apps while considering user decisions. Despite the potential problems that may arise, this approach has proven to be successful in practice, striking a balance between flexibility and security for users.

A robust mobile app security plan requires a combination of mandatory access control, a secure kernel, secure boot, and a well-thought-out strategy for managing app interactions. By implementing these measures, organizations can enhance the security of their mobile applications and protect sensitive data from unauthorized access.

EITC/IS/ACSS ADVANCED COMPUTER SYSTEMS SECURITY DIDACTIC MATERIALS**LESSON: SECURITY ANALYSIS****TOPIC: SYMBOLIC EXECUTION**

Good afternoon, everyone. Today's topic is finding bugs in computer code and the concept of privilege separation. In previous materials, we discussed containing code within boxes to prevent any potential bugs from causing exploits. However, an alternative approach is to eliminate the boxes altogether or use a complementary approach. This material will focus on more advanced methods of bug finding, particularly symbolic execution.

Symbolic execution is an interesting approach to finding bugs, specifically deep bugs that cannot be easily detected through ordinary testing or fuzzing. The basic idea behind symbolic execution is to explore all the branches and paths in a program using symbolic representations instead of concrete values. Initially, this approach seems impractical for real-world programs due to the potential explosion of paths. However, the material will describe various techniques that make symbolic execution feasible for large-scale programs.

This material is important because symbolic execution will be utilized in Lab 3 to find bugs in a code called Zoo Bar. Additionally, it is worth noting that from a security perspective, a bug can be considered equivalent to an exploit, although this is not always the case. From a defensive standpoint, it is crucial to eliminate bugs, as they may potentially lead to exploits. However, it should be noted that finding an exploit is often challenging even if a bug has been identified. This difficulty was evident in Lab 1, where discovering a buffer overrun bug was not straightforward, and exploiting it proved to be equally complex.

To address this challenge, companies, governments, and researchers actively search for exploits in software. Many companies have dedicated teams focused on identifying exploitable bugs, and a significant research community examines both open-source and closed-source software to find bugs and potential exploits. The general sentiment is to learn from these bugs and exploits. As a result, security researchers are encouraged to report bugs and document them. This information is then collected and categorized to facilitate learning and improve security practices.

For instance, MITRE operates the Common Vulnerabilities and Exposures (CVE) website, where individuals can report bugs that have the potential to be exploited. These bugs are carefully recorded and assigned unique identification numbers. The website allows users to search for bugs by specifying keywords, such as "Linux," resulting in a comprehensive list of CVEs related to Linux, along with detailed descriptions and, in some cases, explanations of the exploits.

This material introduces the concept of symbolic execution as an advanced bug-finding technique. It highlights the significance of eliminating bugs as a defensive measure and the challenges associated with finding and exploiting bugs. Additionally, it emphasizes the importance of reporting and documenting bugs to facilitate learning and improve security practices.

Symbolic execution is an approach used in advanced computer systems security for security analysis. It involves analyzing the code of a software system to identify and understand potential vulnerabilities and security bugs. The goal is to find and fix these bugs to improve the overall security of the software.

One way to identify security bugs is through the use of Common Vulnerabilities and Exposures (CVEs). These are publicly known vulnerabilities that have been categorized based on their severity. For example, a high or critical security bug indicates a significant security impact, while a medium or low bug may have a lesser impact. The security research community is open about sharing information on bugs and exploits to encourage learning and software improvement.

When a security researcher finds a bug or exploit, they have a responsibility to disclose it to the software supplier. There is typically a protocol for disclosure, which includes giving the supplier a certain amount of time to fix the issue before making it public. This encourages vendors to take action and fix the bug before it becomes widely known.

While the ideal goal is to eliminate all bugs from code, this is not always possible. Verification is a research area that aims to eliminate bugs through the specification and proof of software implementation. However, this approach is still in development and not widely used for everyday software systems.

Testing is the most common way to find bugs in software. By writing test cases, developers can identify and eliminate bugs that have unintended consequences. However, test cases may not cover all possible scenarios, and some bugs may go undetected.

Another approach used in the security community to find bugs is fuzzing. Fuzzing involves generating unusual inputs to the software to see if they trigger any bugs or vulnerabilities. This technique helps identify bugs that may not be found through traditional testing methods.

Symbolic execution, along with other approaches like verification, testing, and fuzzing, plays a crucial role in improving the security of computer systems. By identifying and fixing bugs, developers can create more secure software that protects users' data and privacy.

Symbolic execution is a technique used in cybersecurity to analyze and identify bugs in computer systems. It involves systematically generating random inputs to drive programs into unusual scenarios, testing them to see if any issues arise. This method is particularly effective at finding known bugs that may not be easily triggered through traditional testing methods.

One challenge with symbolic execution is that it may not uncover deeper bugs that are buried within the program. These bugs require specific values for each branch of the program to be triggered. To address this, a technique called "fuzzing" can be used in conjunction with symbolic execution. Fuzzing involves guessing the correct values for each branch to uncover these deep bugs.

The approach advocated in the paper is to combine symbolic execution with fuzzing to enhance the effectiveness of bug detection. The setup involves a program written in the C programming language, with potential bugs. The attacker has control over the input to the program, similar to a web application where the attacker can manipulate the URL parameters. The goal of symbolic execution is to determine if there is a path between the input that could trigger a bug.

The paper focuses on three types of bugs: crashes, out-of-bounds memory accesses, and application-specific bugs. Crashes are bugs that cause the program to crash, indicating that the bug has been triggered. Out-of-bounds memory accesses can lead to buffer overruns. Application-specific bugs are those that the programmer is specifically concerned about and wants to ensure certain conditions are met. These bugs are typically checked using assert statements in the code.

To use symbolic execution, the programmer needs to either define crash conditions or insert assert statements in the program. The symbolic execution itself does not know what constitutes a bug, it only explores all possible execution paths. The decision on whether a bug has occurred is based on either a crash or the triggering of an application-specific assert statement.

In order to perform symbolic execution, the application must mark certain input variables as symbolic. These variables are the ones that will be solved for during execution.

The goal of symbolic execution is to discover interesting bugs that other tools may not find. It is not meant to prove the absence of bugs, but rather to identify bugs that are difficult to uncover. The success of the symbolic execution approach is measured by its ability to find bugs that other tools cannot detect.

Symbolic execution combined with fuzzing is a powerful technique for analyzing computer systems and identifying bugs. By systematically exploring different execution paths and manipulating input variables, it can uncover bugs that may not be easily triggered through traditional testing methods.

Symbolic execution is an advanced technique used in cybersecurity to analyze computer systems security. It involves computing symbolic values instead of concrete values for variables in a program. This allows us to explore different possible paths and conditions within the program.

The first concept in symbolic execution is the computation of symbolic values. Instead of assigning specific values to variables, we consider variables with unknown values but with constraints. For example, we can mark all inputs as symbolic and try to find concrete inputs that trigger specific bugs.

The second idea is to create path conditions. At every if statement in the program, there is a condition being tested. This condition may involve symbolic values from the input. The idea is that testing the if statement implicitly raises a constraint on the input value. A path condition is a sequence of statements and conditions that the code went through to reach a particular point in the execution of the program. Path conditions typically involve symbolic values that depend on the input.

Finally, we use a solver to determine if a branch is possible. For every encountered if statement, we compute the path condition for both the true and false cases. These path conditions form a set of constraints that we pass to the solver. The solver will either find a concrete instance that satisfies the condition, indicating that the branch can be taken, or it will determine that there is no solution, meaning that the branch could never have been reached in a real execution. This optimization helps avoid path explosion, where the number of possible paths exponentially increases.

To illustrate these concepts, let's consider a trivial example. We have a program with seven statements that reads variables X and Y. The program tests if X is big enough, and if true, it sets Y as the result. Then, it checks if X is smaller than Y, and if true, it increments X. Finally, if the sum of X and Y is equal to 7, it raises an error.

To analyze this program using symbolic execution, we mark X and Y as symbolic inputs. The question we want to answer is, for what values of X and Y might the error be triggered? Even though this program is simple, it still requires some thought to determine the specific values of X and Y that would trigger the error.

Symbolic execution is a powerful technique in cybersecurity that allows us to analyze computer systems security. By computing symbolic values, creating path conditions, and using a solver, we can explore different paths and conditions in a program to identify potential security vulnerabilities.

Symbolic execution is a technique used in cybersecurity to analyze the security of computer systems. It provides a principled way to explore all possible paths of a program and identify actions that may trigger vulnerabilities or security breaches.

The process of symbolic execution involves two main steps: a compile-time step and a runtime step. During the compile-time step, the program is instrumented by the compiler to replace non-symbolic values with calls to a runtime system. This allows the runtime system to handle the execution of symbolic expressions during runtime.

At runtime, the symbolic execution process becomes more complex. The running program, or application process, encounters various branches and if-statements. When a branch is reached, the runtime system takes control and checks if there are concrete values that satisfy the conditions of the branch. To do this, it sends constraints to a separate process called a solver. The solver, such as STP (as mentioned in the paper), is responsible for determining if a solution exists for the given constraints.

The solver receives two sets of constraints: one for the true case and one for the false case. The runtime system wants to explore both paths and sends the constraints to the solver to determine if there is a solution for each case. If a solution is found, it means there are concrete values that satisfy the constraints. If no solution is found, it means there are no concrete values that satisfy the constraints.

In some cases, the solver may not be able to compute a solution within a certain time limit. In such situations, the solver returns a "don't know" response. The runtime system treats this response differently depending on the specific case. In many cases, if a solution cannot be found within a specified time limit, the runtime system decides not to explore that path further.

It is important to note that if a path is not explored further, there is a possibility of missing bugs or vulnerabilities that may exist along that path. This is because there could be a solution to the constraints if the software was run for a longer period of time or if more computational resources were allocated to the solver.

Symbolic execution is a powerful technique for analyzing computer systems' security. It allows for the exploration of all possible paths and helps identify potential vulnerabilities. However, it is not without limitations, as the solver may not always be able to compute a solution within a reasonable time frame.

Symbolic execution is a technique used in cybersecurity to analyze the security of computer systems. It involves

executing a program with symbolic inputs instead of concrete values, allowing for the exploration of multiple paths and the identification of potential vulnerabilities.

During symbolic execution, the program's execution path is determined by the values of symbolic variables. These variables represent unknown values that can take on any possible value. As the program executes, symbolic execution generates constraints based on the operations performed on these symbolic variables.

One approach to symbolic execution is called "depth-first search." In this strategy, the program explores a particular path until it reaches the end, and then backtracks to explore other paths. However, this strategy can lead to a problem known as "loop exploration." If the program contains a loop with a symbolic loop bound, the exploration can get stuck in an infinite loop, resulting in limited coverage of other paths.

To address this issue, another strategy called "breadth-first search" can be used. In this strategy, the program creates two processes, one for the true branch and one for the false branch, and explores both paths simultaneously. This allows for a more comprehensive exploration of different paths in the program.

However, neither depth-first search nor breadth-first search alone can guarantee complete coverage of all possible paths. To overcome this limitation, a heuristic-based approach is often employed. The scheduler, a component of symbolic execution tools, decides which process to run next based on certain criteria. For example, the scheduler may prioritize processes that have not been explored extensively, ensuring a broader coverage of paths.

Let's consider an example to understand how symbolic execution works in practice. Suppose we have two variables, X and Y, and an empty path condition. Initially, X is assigned the symbolic value "alpha" and Y is assigned the symbolic value "beta." As the program executes, the path condition is updated based on the constraints imposed by the operations performed on X and Y.

In our example, after the first statement, X remains "alpha" and Y remains "beta," and the path condition is still empty. Moving to the second statement, we explore the true branch. The state after executing the second statement with the true branch will have X as "alpha," Y as "beta," and the path condition will be "alpha > beta." This process creates a new process to explore the next statement.

Symbolic execution continues in this manner, exploring different paths and updating the path condition based on the constraints generated. The goal is to identify interesting bugs or vulnerabilities that may not be found by other tools. The scheduler plays a crucial role in deciding which process to run next, ensuring a balanced exploration of different paths.

Symbolic execution is a powerful technique in cybersecurity for analyzing the security of computer systems. It allows for the exploration of multiple paths and the identification of potential vulnerabilities. By employing different strategies such as depth-first search, breadth-first search, and heuristic-based scheduling, symbolic execution tools can provide valuable insights into the security of advanced computer systems.

In the context of advanced computer systems security and security analysis, one technique that is commonly used is symbolic execution. Symbolic execution is a method that allows us to analyze the behavior of a program by executing it with symbolic values instead of concrete ones. This approach can help us identify potential vulnerabilities or errors in the code.

During symbolic execution, the program is executed with symbolic values instead of concrete ones. These symbolic values represent variables in the program and can take on any value that satisfies certain constraints. As the program executes, these constraints are updated based on the operations performed on the symbolic values.

In the given material, the speaker discusses the state of a table after executing certain statements in the program. They mention the use of symbolic values for variables, such as x and y, and how the path condition is updated based on the execution of different branches of the code.

The speaker also highlights the optimization aspect of symbolic execution. In certain cases, where the path condition indicates that a branch is not feasible, the program can skip exploring that branch, leading to improved efficiency.

The material emphasizes the importance of understanding the if statement in symbolic execution. By understanding how the if statement affects the path condition and the execution of different branches, one can gain a better understanding of the overall process of symbolic execution.

The speaker concludes by mentioning that the discussed concepts will be relevant in future sessions, indicating the significance of understanding symbolic execution for advanced computer systems security.

To summarize, symbolic execution is a powerful technique used in advanced computer systems security for security analysis. It involves executing a program with symbolic values instead of concrete ones, allowing for the identification of potential vulnerabilities or errors. By understanding the impact of the if statement on the path condition and the execution of different branches, one can effectively analyze the behavior of a program.

Symbolic execution is a technique used in cybersecurity to analyze and identify vulnerabilities in computer systems. It involves executing a program symbolically, meaning that instead of using actual values, symbolic values are used to represent inputs. By doing this, it becomes possible to explore different execution paths and identify potential security issues.

The process of symbolic execution involves running a program with symbolic inputs and observing the program's behavior. The goal is to identify inputs that can trigger specific paths in the program that may lead to security vulnerabilities. This is done by analyzing the program's code and identifying conditions that depend on user inputs.

To illustrate the concept of symbolic execution, let's consider an example. Suppose we have a program that takes two inputs, let's say the values 2 and 4. By running the program symbolically, we can determine that these inputs will cause the program to crash. This information is valuable because it allows us to identify potential vulnerabilities and fix them before they can be exploited by attackers.

Symbolic execution is not limited to simple programs. It can also be applied to complex software systems. For example, the Berkeley Packet Filter (BPF) is a piece of software that runs inside the Linux kernel and is used by programs like tcpdump. BPF filters are written in a programming language and are used to filter packets coming from the network. Symbolic execution can be used to analyze these filters and identify potential vulnerabilities.

In a research paper, a group of researchers demonstrated the effectiveness of symbolic execution by using it to generate filters that caused the Linux kernel to crash. The filters they generated had a large offset value that, when combined with a small length value, caused an integer overflow. This overflow led to the program reading or writing to random memory locations, which can be a serious security issue.

The researchers' findings highlight the power of symbolic execution in identifying vulnerabilities in real-world software systems. By using this technique, they were able to uncover a bug in the BPF interpreter, a well-studied and widely used piece of software.

Symbolic execution is a valuable technique in cybersecurity for analyzing computer systems and identifying potential security vulnerabilities. By running programs symbolically and exploring different execution paths, it becomes possible to uncover bugs and weaknesses that could be exploited by attackers.

Symbolic execution is a technique used in cybersecurity to analyze the security of advanced computer systems. It involves the use of a solver to solve symbolic expressions or formulas. The solver takes these symbolic expressions and attempts to solve them, providing answers to the equations.

The solver plays a crucial role in symbolic execution as it is responsible for solving the equations. In the case of the discussed approach, the solver is used to solve the equations derived from the symbolic expressions. The paper emphasizes the importance of optimizing the solver to ensure quick and efficient solutions to the equations.

While the paper goes into detail about the optimizations and tricks used to improve the solver's performance, for the purposes of this material, we will treat the solver as a black box. However, it is important to note that the solver used in the paper is a custom-built solver specifically designed for the problem at hand. In our class, we will be using a solver called C3, which is similar in spirit and widely used for various problems.

To provide a better understanding of the solver, let's consider a simple example using C3. In a Python shell, we can import C3 and create symbolic values, such as x and y . We can then use C3 to solve constraints involving these symbolic values. For example, we can send the constraint " $x + 7 = y$ " to C3, and it will provide us with the solution, indicating that for $y = 0$ or $x = 7$, the constraint holds true.

It is important to note that the solver is capable of handling more complex constraints beyond simple integer equations. In fact, the solver represents symbolic values as bit arrays rather than integers. For example, a symbolic value like x is represented as a bit array, with each bit representing a byte or a bit of the value. This allows for more flexibility in representing and solving equations.

The solver is a crucial component in symbolic execution, responsible for solving the equations derived from symbolic expressions. It is optimized to provide quick solutions and can handle complex constraints involving bit arrays. By using a solver like C3, we can effectively analyze the security of advanced computer systems.

In the field of advanced computer systems security, one important aspect is security analysis. One technique used in security analysis is symbolic execution. Symbolic execution involves analyzing the behavior of a program by using symbolic values instead of concrete values. This allows us to reason about the program's behavior in a more abstract and general way.

In the context of cybersecurity, symbolic execution can be used to analyze the security of computer systems. One particular use case is in solving boolean expressions or formulas. When we have a boolean expression, we can use a solver to find instances for the variables in the expression that make the expression true.

A solver is a tool that can reason about boolean expressions and find solutions to them. When working with symbolic execution, the solver needs to have internal knowledge of how to perform operations like addition, subtraction, and multiplication. It needs to understand the concept of carrying over in addition and the complexities of multiplication. The solver then converts the symbolic expression into a boolean constraint, which is solved by a set solver.

To illustrate this concept, let's consider the example of a circuit. In a circuit, we have inputs and outputs. The inputs are boolean values, and the circuit performs boolean operations on these inputs to produce an output. Solving a boolean expression is like finding inputs that make the circuit output true. This can be formulated as a boolean formula that the solver can solve.

However, not all boolean expressions are easy to solve. Some expressions, like multiplication circuits, are more complex than addition circuits. Solving them requires more computational resources and time. For example, multiplying three variables together may be harder to solve than adding them together.

In some cases, solving certain boolean expressions may be practically impossible. For instance, if we have a boolean expression that represents the SHA-256 hash of a value, it would be extremely difficult to find an input value that produces a specific hash. This is because the SHA-256 function is designed to be secure and difficult to invert.

In general, solving boolean expressions is an NP-complete problem, which means it is computationally difficult. While solvers can provide solutions, they often rely on optimizations and heuristics to improve efficiency. It's important to note that these techniques may not be applicable to all types of boolean expressions encountered in real-world programs.

Symbolic execution and the use of solvers are valuable tools in the field of computer systems security. They allow us to reason about the behavior and security of computer systems by analyzing boolean expressions and finding solutions to them. However, solving boolean expressions can be computationally challenging, and some expressions may be practically impossible to solve.

Symbolic execution is a technique used in computer systems security to analyze the behavior of a program by executing it symbolically, instead of using concrete inputs. This allows for the identification of potential vulnerabilities and security flaws. One important aspect of symbolic execution is the use of constraint solvers to reason about the program's execution paths.

Constraint solvers are algorithms that can determine the values of variables in a program based on a set of constraints or conditions. In the context of symbolic execution, these solvers are used to explore different execution paths and identify possible inputs that lead to certain behaviors or outcomes. However, solving these constraints can be challenging and time-consuming, especially when dealing with complex expressions or data types.

The paper discussed in this material focuses on the use of symbolic execution for analyzing security in advanced computer systems. It emphasizes the importance of optimizations to improve the efficiency of the constraint solver and obtain faster answers. One of the key topics covered in the paper is the analysis of integers, excluding bit arrays and strings. The paper also briefly mentions the topic of race conditions, which can pose challenges in symbolic execution.

When dealing with arrays, the paper explains that the solver's performance depends on the type of indexing used. If the array is indexed by a concrete value, the solver can provide a quick answer. However, if the array is indexed by a symbolic value, the solver faces difficulties due to the need to test multiple possible concrete indexes. This situation leads to a large number of conditions that need to be evaluated, resulting in a more complex and time-consuming analysis.

The paper also highlights the challenges posed by symbolic pointers. Since a pointer can potentially point to any memory location, reasoning about its behavior becomes harder. The authors discuss the importance of addressing this challenge and spend significant time discussing race conditions and pointers in the context of symbolic execution.

In some cases, when information about a pointer is limited, the authors propose a technique where a concrete value is chosen for the pointer, effectively converting it from a symbolic value to an incomplete value. While this approach may help in certain scenarios, it also introduces the risk of missing potential errors or vulnerabilities that could be identified with a fully symbolic analysis.

Symbolic execution is a powerful technique for analyzing the security of advanced computer systems. The paper discussed in this material focuses on the challenges and optimizations associated with symbolic execution, with a particular emphasis on integers, arrays, race conditions, and pointers. By understanding these concepts, researchers and practitioners can gain insights into the complexities of security analysis and improve the effectiveness of their approaches.

Symbolic execution is a powerful technique used in advanced computer systems security for security analysis. It involves analyzing a program by executing it symbolically, rather than with concrete values. This allows for the exploration of different program paths and the identification of potential vulnerabilities or bugs.

One limitation of symbolic execution is that it may miss some possible values that could be true for a pointer, resulting in incomplete analysis. However, it is still a valuable tool for finding bugs and vulnerabilities in software.

When it comes to floating point numbers, symbolic execution does not model them at all. This is because computing the addition of two floating point numbers is complex and non-associative. Additionally, the behavior of floating point numbers can vary depending on their size. However, this is not a major issue for system software, as they typically do not use floating point numbers.

The execution tree in symbolic execution grows as the program branches. Every time a branch is encountered, two paths are created: one for the true statement and one for the false statement. This can result in an exponential growth of the execution tree. Symbolically executing large programs, such as the Linux kernel, is feasible but can be time-consuming due to the large number of if statements present.

Symbolic execution tools, like the one discussed in the material, can be reasonably quick for certain cases. However, there are situations where symbolic execution may take a long time or even fail to terminate. This is especially true when there are many branches to explore. Optimizations are employed to control the execution time and prevent infinite exploration.

In symbolic execution, the symbol 'S' represents a symbolic value, while 'C' represents a concrete value. Symbolic values are like variables that can take on different values during execution, while concrete values are

fixed and known. Symbolic execution allows for the analysis of both symbolic and concrete values.

It is important to note that symbolic execution does not find all bugs. There are scenarios in which it may miss certain vulnerabilities or bugs. However, it is still a valuable tool for bug finding and security analysis.

Symbolic execution is a powerful technique in the field of cybersecurity and advanced computer systems security. It allows for the analysis of program paths and the identification of potential vulnerabilities or bugs. While it has its limitations, it remains a valuable tool for security analysis.

Symbolic execution is an advanced technique used in cybersecurity for analyzing computer systems' security. It involves executing a program symbolically rather than with concrete values. By doing so, symbolic execution can explore all possible paths and identify potential vulnerabilities that may be hidden deep within the code.

During symbolic execution, the program's constraints are sent to a solver, which attempts to find concrete values that satisfy those constraints. However, if the solver times out or fails to find any concrete values within a specified timeframe, it assumes that the constraints cannot be solved and ignores that branch of the program. This approach may overlook potential bugs or vulnerabilities that the solver did not find.

Symbolic execution is particularly useful for finding bugs that may be difficult to trigger using traditional methods like fuzzing or random inputs. It can explore multiple branches of a program and identify errors that may be buried deep within the code. This deep analysis can uncover bugs that would otherwise be challenging to detect.

In the context of computer security, it is crucial to assume that a bug in the code may indicate the presence of an exploit. Constructing exploits can be complex and may involve multiple bugs and careful manipulation of the code. Therefore, bug finding is a critical aspect of computer security research, and significant time and effort are devoted to this area.

Symbolic execution helps in constructing possible inputs that could trigger bugs hidden deep within the code. It assists in identifying vulnerabilities that may not be easily found through testing or fuzzing. However, symbolic execution is not foolproof and has its limitations. There is a risk of path explosion, where the number of possible paths becomes too large to handle effectively. To mitigate this risk, symbolic execution engines selectively mark variables as symbolic inputs, focusing on those that are relevant to the analysis.

In lab three, you will implement a symbolic execution engine. However, there is a challenge when dealing with libraries or database calls. In the previous materials, it was assumed that the entire program was available for symbolic execution. Inlining the library into the program is one way to handle this, but it can be costly and impractical for larger programs. To address this issue, a technique called "combi-colic execution" will be used. Combi-colic execution combines symbolic and concrete execution. When encountering a library or database function that cannot handle symbolic values, a concrete value is chosen, and the function is invoked with that value. This allows the program to interact with real databases and explore potential bugs.

Symbolic execution is a powerful technique for analyzing computer systems' security. It helps in identifying bugs and vulnerabilities that may be hidden deep within the code. However, it has its limitations, such as path explosion. The research community invests significant time and effort in bug finding to improve security. In lab three, you will implement a symbolic execution engine and explore the challenges of dealing with libraries and database calls.

Symbolic execution is a technique used in cybersecurity to analyze computer systems' security. It involves automatically finding bugs in code that can potentially be exploited by attackers. While there are tools available for symbolic execution, such as XE and Klee, they are mainly research prototypes and not widely used in practice. However, some companies have developed their own tools based on symbolic execution for internal use.

One of the key considerations in cybersecurity is whether to make security analysis tools and techniques public or keep them closed. This debate revolves around the idea of transparency versus confidentiality. Making tools and techniques public can help the good guys in identifying and fixing vulnerabilities, but it also means that attackers can access and potentially exploit them. The hope is that by making these tools and techniques available to the good guys, they can stay ahead of the bad guys.

In the context of symbolic execution, assuming that any answer returned by a database is a symbolic value may not make sense in all cases. This is because attackers have control over the URLs but not necessarily over the database itself. Therefore, blindly assuming all answers from the database are symbolic values may not be accurate.

Symbolic execution is a valuable technique in cybersecurity for identifying and fixing vulnerabilities in computer systems. While there are tools available, they are still in the research stage and not widely adopted in practice. The debate between public and closed security analysis tools continues, with arguments for both transparency and confidentiality. It is important to carefully consider the assumptions made in symbolic execution, as blindly assuming symbolic values from a database may not always be accurate.

EITC/IS/ACSS ADVANCED COMPUTER SYSTEMS SECURITY DIDACTIC MATERIALS**LESSON: NETWORK SECURITY****TOPIC: WEB SECURITY MODEL**

Today's lecture will focus on web security, specifically the security of web browsers and the challenges they face when interacting with various websites over the internet. The main issue is that a user's web browser interacts with multiple websites, including potentially malicious ones, and there is a lot of code and content from different sources running within the browser. This includes HTML, JavaScript, CSS, flash, and Java, among others.

On the positive side, web browsers have mechanisms in place to isolate the different content and prevent websites from accessing sensitive information on the user's computer. However, there are still important aspects that need to be protected, such as preventing compromised websites from compromising other websites or accessing sensitive data.

The complexity of web security arises from the evolution of the web itself. Initially, the web was designed to support text documents and images, and security was not a major concern. However, as the web evolved, more sensitive applications emerged, such as dynamic applications and the use of cookies and JavaScript. Unfortunately, many of the security mechanisms in place today were not initially designed with these advancements in mind.

The reading material for this topic can be quite overwhelming, as it highlights the challenges and complexities of web security. The lack of sufficient security mechanisms in the early design stages of the web has resulted in a situation where securing web browsers and protecting user data is a difficult task.

In the upcoming lectures, we will delve deeper into web security, discussing topics such as how the web deals with network problems, the functioning of certificates, and issues related to TLS (Transport Layer Security). These lectures will provide a comprehensive understanding of the various aspects of web security and the measures taken to mitigate the risks associated with browsing the internet.

The field of cybersecurity focuses on protecting computer systems, networks, and data from unauthorized access or damage. One area of cybersecurity is network security, which aims to secure the communication and data transmission between different devices and systems. Another area is web security, which specifically focuses on securing web applications and websites.

One challenge in cybersecurity is retrofitting security measures after a system has been designed and implemented. This can be difficult because adding security to an existing system that was not initially designed with security in mind can be a complex task. This difficulty is highlighted in the case study discussed in this material.

The evolution of web security can be seen as a progression from desktop security to mobile security. Desktop security, inherited from older operating systems, was not as sophisticated as the web and lacked many of the security mechanisms needed for today's problems. On the other hand, mobile security, developed after the web, learned from the mistakes made in web security and implemented a more comprehensive security plan.

There are several reasons why web security is complicated. One reason is the path that the web took to reach its current state. The need for compatibility with existing applications and the reluctance to break them poses constraints on introducing new security features. Additionally, the lack of a central authority overseeing web standards has led to decentralized decision-making by browser vendors, resulting in inconsistencies in security practices.

Another constraint in web security is the prevalence of sharing and interaction between different components of web applications. This complexity makes it challenging to provide strong security guarantees while maintaining functionality and interactivity.

In the context of this material, the focus is on the threat model within the browser box. The discussion centers on the interactions and vulnerabilities that can occur within the browser itself, rather than the network connections to external servers. The assumption is made that the network is secure, and the focus is on how

attackers can exploit vulnerabilities within the browser.

The threat model includes an attacker who owns a domain, such as "attacker.com," and a victim who is using a web browser on their computer. The assumption is that the victim visits the attacker's domain, which can lead to potential security breaches.

Web security is a complex field that requires retrofitting security measures, learning from past mistakes, addressing compatibility issues, and managing the complexities of sharing and interaction. Understanding the threat model within the browser box is crucial for developing effective security measures.

In the context of web security, network security plays a crucial role in protecting users from potential threats. One aspect of network security is the web security model, which aims to disentangle websites within a browser and ensure that they operate independently. This model is known as the same origin policy.

The same origin policy is designed to prevent websites from accessing or manipulating data from other websites. It achieves this by logically separating websites into different units, similar to processes in UNIX. Each unit is referred to as an origin and is identified by its URL.

An origin is determined by extracting the protocol scheme (HTTP or HTTPS), the domain name, and the port number from the URL. For example, the origin of "http://foo.com/bar/add-onx.html" would be "http://foo.com". This allows the browser to differentiate between different origins and enforce restrictions on their interactions.

The same origin policy is essential for ensuring web security. By isolating websites within their respective origins, the risk of unauthorized access or manipulation of data is significantly reduced. However, it is important to note that the same origin policy is not foolproof, and websites must still implement additional security measures to protect against potential vulnerabilities.

It is worth mentioning that the same origin policy assumes that the browser is trustworthy and has implemented the policy correctly. Additionally, it assumes that the browser is free from bugs or vulnerabilities that could compromise its security. While browsers are complex software and may have bugs, addressing these issues is beyond the scope of this lecture.

In this lecture, we have focused on the threat model of web security, specifically the same origin policy. We have discussed the importance of isolating websites within their respective origins to ensure security. However, it is important to note that network attacks and other related topics will be discussed in future lectures.

Understanding web security and implementing secure web applications requires a comprehensive understanding of the same origin policy and other relevant concepts. While this lecture provides a high-level overview, further reading and study are necessary to grasp the intricacies of web security.

The same-origin policy is a fundamental concept in web security that aims to protect users from malicious activities on the internet. It is designed to prevent scripts from one origin (website) from accessing or manipulating resources from another origin without explicit permission.

In the context of the same-origin policy, an origin is defined by a combination of the URL and the port number. For example, the origin "https://www.example.com:443" represents the website running on the domain "example.com" and listening on port 443. The path component of the URL does not matter for the purposes of determining the origin.

Under the same-origin policy, scripts running in a web browser are assigned an origin, representing the website they are running on behalf of. These scripts can interact with objects or resources, such as other pages or URLs, which are also assigned an origin. The rule is simple: a script is allowed to access a resource only if their origins match.

To illustrate this, let's consider an example. Suppose we have two websites, "gmail.com" and "attacker.com", and a web browser running both of them. The browser creates separate origins for each website: "https://gmail.com:443" and "https://attacker.com:443". This separation ensures that the attacker's scripts cannot tamper with the Gmail window.

In the same-origin policy, the concept of a window is equivalent to a process in Unix. Each website runs in its own window, and these windows are visually represented as tabs in modern browsers. Scripts running in one window can manipulate its own contents but are prohibited from accessing or modifying the contents of other windows.

For example, a script running on "a.com" can manipulate its own page but cannot access or modify the contents of the Gmail page. This prevents unauthorized access and protects user privacy.

The same-origin policy extends beyond the browser environment. In JavaScript, for instance, there is an API called XMLHttpRequest that allows developers to make network requests. When a script running in a browser requests a URL, the browser examines the origin associated with that URL. If the origin of the script matches the origin of the requested URL, the request is allowed. Otherwise, it is rejected.

This broader application of the same-origin policy ensures that requests made by scripts running on a website can only be made to URLs that belong to the same origin. This helps to prevent cross-origin attacks and unauthorized data access.

While the same-origin policy provides a strong security mechanism, it also allows for exceptions that enable sharing and collaboration between websites. These exceptions, however, also introduce challenges in securing web applications, as they can be exploited by attackers.

The same-origin policy is a crucial security measure in web development that restricts scripts from different origins from accessing or manipulating each other's resources. By enforcing this policy, web browsers can protect users from unauthorized data access and malicious activities.

In the context of web security, the same-origin policy plays a crucial role in ensuring the isolation of websites and protecting user data. However, there are certain exceptions and complexities associated with this policy that need to be understood.

The designers of the same-origin policy had to consider various factors while formulating this security measure. Although it may seem ad hoc, the policy aims to control exceptions and prevent unauthorized access between websites. For example, a link created by a dot-com website can potentially lead a user to Gmail, highlighting the challenge of isolating websites completely.

To better comprehend the same-origin policy, it is important to delve into its underlying principles. The policy relies on the mindset of its designers, who believed that this approach would be effective in safeguarding user data. While it may not be foolproof, understanding the rationale behind its implementation is crucial.

In addition to browser-side exceptions, there are also server-side exceptions that can be set up to allow requests from specific sources. This means that a server can authorize requests from a specific domain, such as a.com, based on its expectations and preferences. This introduces a level of complexity to the policy, making it a multifaceted system.

Moving beyond the policy itself, it is essential to explore the network aspect of web security. The browser, which is under the control of the same-origin policy designers, connects to various websites. However, these connections are not enforced by any mechanism. Both the browser and the websites it connects to have an interest in identifying each other. The browser wants to ensure that it is communicating with the correct website, while the websites want to verify the identity of the browser.

To address this, different mechanisms are employed for server identification and client identification. The browser uses a network encryption protocol known as TLS (Transport Layer Security) along with certificates to authenticate the server. This combination, often referred to as HTTPS, ensures that the browser is indeed communicating with the intended server. However, the process of client identification is more intriguing.

When a server receives a request from a browser, it needs to determine the identity of the client. In this scenario, cookies have emerged as the prevailing mechanism. Cookies are sent along with requests and provide the server with information about the user issuing the request. Stealing a cookie can lead to user impersonation, making it a potential security vulnerability.

Understanding how cookies work is crucial in comprehending their role in maintaining the separation of websites within a browser. Logically, the browser maintains a table that contains various cookies. Each cookie consists of a domain, a key, and a value. For example, a cookie for the domain mit.edu might have a key-value pair where the key represents the user's Athena username. Storing multiple cookies for different keys allows for differentiation between various libraries or services.

While cookies serve as an important mechanism for authenticating users to the server, they also pose security risks. If a cookie is stolen, an attacker can impersonate the user. This vulnerability has been a source of security problems in the web.

The same-origin policy is a critical component of web security, aiming to isolate websites and protect user data. Despite its ad hoc nature and the existence of exceptions, understanding the mindset of its designers helps to comprehend its purpose. Additionally, the network aspect of web security involves mechanisms such as TLS and certificates for server identification, while cookies play a significant role in client identification. Recognizing the vulnerabilities associated with cookies is essential for ensuring robust web security.

Cookies play a crucial role in web security models, allowing different applications to share information and authenticate users. When a browser interacts with servers over the network, these servers can give cookies to the browser. For example, when logging into a website like Gmail, the web server will provide a cookie, which acts as a short-term password for the user's account. This session identifier allows users to access their account without having to type in their password every time they click on a link.

The way cookies work is that the server sends a cookie with a session ID to the browser, which stores it in a cookie table. Whenever the browser sends a request to a particular server, it retrieves the matching cookies from the table and sends them back to the server. This allows the server to identify the user and grant access to their account.

However, there are some security concerns with cookies. The same origin policy, which restricts access to resources from different origins, applies to cookies as well. The domain of the cookie should align with the origin, but there are cases where this alignment is not perfect. For example, the protocol (HTTP or HTTPS), port number, and path may not match exactly. This mismatch can lead to security vulnerabilities.

Another issue is that all matching cookies are sent on every request to a web server, regardless of the origin that triggered the request. This means that if a user clicks on a link from one website to another, all their cookies will be sent to the second website. This sharing of cookies between sites can lead to potential attacks and security problems.

Furthermore, there is a question of who can set a cookie. Browsers have not always been consistent in enforcing rules regarding cookie setting. This inconsistency can be exploited by attackers to give a user's browser a session that corresponds to the attacker, leading to potential security breaches.

Cookies are an essential component of web security models, allowing for user authentication and information sharing between applications. However, there are security concerns related to the alignment of cookies with the same origin policy, the sharing of cookies between sites, and the control over who can set cookies. These issues can lead to potential vulnerabilities and attacks.

Network Security and Web Security Model

In the realm of cybersecurity, one aspect that requires attention is network security. When using various websites, it is essential to understand the potential risks associated with logging in. Attackers may attempt to gain unauthorized access to your Google or Gmail account, which can lead to various consequences. For instance, if an attacker signs you into their Google account, they can collect information about your searches and browsing habits.

This type of attack can also occur on other websites where you may not notice any user-specific changes. The server on these sites can collect and store data in the attacker's account instead of your own. Consequently, the attacker can gain insights into your browsing activities over time. While this may seem like a subtle attack, it is a concern for many websites.

To better understand the risks, let's consider a hypothetical scenario. Suppose you visit a.com and wonder if it can set a cookie for the domain gmail.com. This scenario is plausible since a.com is a legitimate domain, and gmail.com is a suffix. However, allowing a.com to set a cookie for the suffix "com" would enable attackers to force users into their Google accounts. Therefore, this scenario poses a significant security risk.

To prevent such attacks, browsers implement rules regarding cookie settings. One such rule is that a single dot in a domain is considered bad. However, two dots are acceptable. For example, a.com can set a cookie for "co.uk" since it has two dots. But allowing a.com to set a cookie for "com" would be too risky.

The rules regarding cookie settings can be complex, as they involve a list known as the public suffix list. This list contains all the entities that can register domain names under a specific suffix. For instance, "co.uk" is on the list, while "google.com" is not. The public suffix list is extensive and includes various entities such as school districts and counties. It is a constantly evolving list, as anyone can submit entries.

The public suffix list is managed by a website called publicsuffix.org, which functions similarly to GitHub. Domain owners can submit pull requests to add their domains to the list, provided they can prove ownership. Browser vendors regularly update their browsers with the latest version of the public suffix list before shipping them to users.

Network security and web security models play a crucial role in protecting users from potential attacks. Understanding the risks associated with logging into websites and the complexities of cookie settings is essential for maintaining a secure online presence.

Domains play a crucial role in ensuring security in computer systems, particularly in the context of network and web security. When users host content on a domain, it is important for organizations not to host their own content on the same domain. This separation of content is achieved through the use of different origins or domains.

For example, in Gmail, when a user downloads an attachment, the attachment URL does not come from google.com or gmail.com. Instead, Gmail provides a link to download the attachment from a different origin called Google user content com. This separation of origins ensures that the attachments are hosted separately from the main Gmail domain, with different security mechanisms in place to prevent any security breaches. Similar practices can be observed in other companies like Facebook, where the main domain (e.g., facebook.com) is used for storing cookies, while a different domain (e.g., FB CDN net) is used for distributing user-uploaded photos.

The concept of privilege separation is important in web security. It involves having a different process, user ID (UID), or container for untrusted content. In the context of the web, this separation is achieved through the use of different origins or domains. Simply having a different hostname is not sufficient due to the cookie problem. To ensure a separate user ID, organizations need to register a separate com domain.

However, it is worth noting that some domains may not be completely isolated from each other. For example, websites under the same domain (e.g., dot MIT ODU) may share cookies, even though they may have different subdomains. This can lead to potential security risks, as one website could potentially convince a user's browser to log them into another website's account.

Certificates are used to enhance security to some extent. However, in certain cases, intermediate entities between certificates and the actual website may be involved. For example, MIT uses certificates for logging into a system called touchstone or Shibboleth, which then sets a cookie for the desired website. This reliance on cookies from an intermediary entity introduces potential vulnerabilities.

Exceptions to the same origin policy exist in order to facilitate sharing on the web. These exceptions can be exploited by adversaries, leading to various attacks on websites. One such exception is the ability to create links. When a website creates a link to another origin, such as a href="http://gmail.com", and a user clicks on that link, the target origin (in this case, gmail.com) is loaded in the browser window or tab along with all the cookies associated with that origin. This can potentially allow an attacker to access the target origin with the user's cookies, raising concerns about potential control over the target origin.

Domains serve as security entities in computer systems, particularly in network and web security. They provide

a means to separate content and ensure privilege separation. However, there are exceptions to the same origin policy that can be exploited by adversaries. It is important for organizations to be aware of these exceptions and take appropriate measures to mitigate potential security risks.

In the context of advanced computer systems security, network security plays a crucial role in ensuring the protection of data and information transmitted over networks. One aspect of network security is web security, which involves the protection of web applications and websites from various threats and vulnerabilities.

One important concept in web security is the web security model. The web security model is designed to define the rules and policies that govern the behavior of web browsers when interacting with web content from different origins. An origin is defined by the combination of a protocol, domain, and port number.

The web security model includes a concept called the same-origin policy. The same-origin policy states that web browsers should only allow scripts and resources from the same origin to access each other's content. This policy helps prevent malicious websites from accessing sensitive information from other websites.

However, there are some exceptions to the same-origin policy. One exception is the ability to load images from different origins. This exception exists because, in the early days of the web, it was not practical to copy and host images on every website. Instead, websites could refer to images hosted on other sites, reducing the need for duplication. Despite this exception, loading images from different origins can still pose security risks, as the browser will send all cookies associated with the origin to fetch the image.

Another exception to the same-origin policy is the ability to load content from links. When a user clicks on a link, the browser will load the content from the linked URL. This content may have side effects, such as deleting emails or modifying data. To mitigate this risk, the browser changes the origin to the linked URL's origin, preventing the content from accessing resources from the original origin.

It is worth noting that although a website can load images or display content from other origins, the browser enforces fine-grained security policies. For example, a website can display an image from another origin but cannot access the pixels of that image. This prevents websites from leaking sensitive information from images.

However, it is important to be aware that over time, the web security model has evolved, and new techniques and technologies have emerged that may introduce new vulnerabilities. For example, CSS and layout tricks can allow websites to determine the dimensions of images loaded from other origins, even though they cannot access the image content itself. This can potentially be exploited in certain scenarios, such as determining the size of sensitive images based on how they are rendered on a webpage.

The web security model, including the same-origin policy, is an important component of network security. It helps protect web applications and websites from unauthorized access and data leakage. However, it is essential to stay updated with the evolving web security landscape to address new vulnerabilities and ensure the ongoing protection of web content.

Network Security: Cross-Site Request Forgery (CSRF)

In the field of web security, one common attack that poses a significant threat is called Cross-Site Request Forgery (CSRF). This attack takes advantage of a browser's behavior of sending cookies to every request, regardless of the origin of the request. To better understand this attack, let's consider an example scenario.

Imagine you are logged into your bank's website and you visit a different website, let's call it dot-com. This dot-com website contains an image tag with a source pointing to the bank's website, along with some instructions to transfer money to a recipient, who happens to be the attacker. The browser, upon encountering this image tag, will send a request to the bank's website, including all the cookies associated with your session. From the bank's perspective, this request appears legitimate, as it contains the necessary cookies to identify you as a user.

This attack is possible because the browser does not differentiate between requests initiated from different origins. As a result, the bank's website is unaware that the request originated from the dot-com website and not from the user directly. This lack of distinction allows attackers to exploit the trust between the browser and the target website, potentially leading to unauthorized actions being performed on behalf of the user.

It's important to note that this attack can be executed using both GET and POST requests. While the example above demonstrates a GET request, a similar approach can be taken with POST requests. To initiate a POST request, the attacker can create a form on the dot-com website with the necessary input fields and use JavaScript to synthetically submit the form, simulating a user's action.

The implications of CSRF attacks can be severe, as they can lead to unauthorized actions being performed on the user's behalf. The server receiving the request has no way of distinguishing whether the user intentionally initiated the action or if it was a result of a CSRF attack. This poses a challenge in terms of implementing effective countermeasures.

Various solutions have been proposed to mitigate CSRF attacks. These include the use of additional headers to indicate the origin of the request, or implementing rules to prevent the sending of cookies in cross-origin requests. These solutions have been explored in extensions to web security standards, which aim to address the vulnerabilities associated with CSRF attacks.

Cross-Site Request Forgery (CSRF) is a prevalent attack in web security that exploits the trust between a browser and a target website. By tricking the browser into sending requests with the user's cookies, attackers can perform unauthorized actions on the user's behalf. Implementing effective countermeasures is crucial to mitigate the risks associated with CSRF attacks.

Modern browsers have implemented a security feature called same-site cookies. These cookies are not sent by default when other origins try to use or send a request. In older browsers, cookies were sent along with matching requests, but modern browsers have an extra column in the cookie table that specifies whether it is a same-site cookie. Same-site cookies are only sent if the request is being made by the same domain, and they are not sent in cases where the request is coming from a different page.

However, implementing same-site cookies can cause certain functionalities to break. For example, the Facebook like button may not work because it requires clicking from a different page. Similarly, links may not work as expected, and users may have to log in again when clicking on a bookmarked page. These issues arise because many websites were not designed to handle the restrictions imposed by same-site cookies. The challenge lies in balancing security measures with maintaining backwards compatibility.

To defend against cross-site request forgery (CSRF) attacks, a common approach is to use a secret token. This token is stored on the server and is not accessible to the attacker's domain. In addition to the usual cookies and passwords, the server stores a randomly generated token for each user. When a request is made to transfer money, the request must include this token, which should match the token stored for the user in the server's table. This ensures that the request is legitimate and not forged.

To implement this defense mechanism, the server sends an HTML page to the client with a form to transfer money. The form includes a hidden field with the token value specific to the user. When the form is submitted, the server checks if the token in the request matches the token stored for the user. If the tokens match, the transfer is allowed. However, if an attacker tries to perform the same attack, their token guess will not be valid, as they do not have access to the secret token value.

Same-site cookies and secret tokens are two security measures used to enhance network and web security. Same-site cookies prevent cookies from being sent to different origins by default, reducing the risk of unauthorized access. Secret tokens add an additional layer of security by requiring a valid token in requests to prevent CSRF attacks. By implementing these measures, websites can better protect user data and prevent malicious activities.

In the context of advanced computer systems security, network security plays a crucial role in protecting sensitive information and preventing unauthorized access to networks. One aspect of network security is web security, which involves safeguarding web applications and preventing attacks such as cross-site request forgery (CSRF).

CSRF is a type of attack where an attacker tricks a user's browser into making an unintended request to a vulnerable website. This can lead to unauthorized actions being performed on behalf of the user without their knowledge or consent. To mitigate this risk, a web security model has been developed to prevent CSRF attacks.

The web security model relies on the use of tokens, also known as anti-CSRF tokens or synchronizer tokens. These tokens are unique for each user and are only displayed when the user is logged in. They serve as a secret value that is required to successfully submit a form or perform an action on the website. Attackers cannot obtain these tokens unless they are logged in as the user, making it difficult for them to exploit CSRF vulnerabilities.

When a user requests a page from a website, the server includes the token value in the response. This value is then included in subsequent requests made by the user, ensuring that the request is legitimate and not initiated by an attacker. If an attacker tries to initiate a request from a different origin, the browser will block it, preventing the CSRF attack from being successful.

Web application frameworks, such as Node.js, Django, or React, often handle the implementation of CSRF protection automatically. These frameworks take care of generating and validating tokens, reducing the burden on web application developers. However, developers must still be aware of potential vulnerabilities and understand the underlying mechanisms to build secure web applications.

Another approach that has been considered is storing the token in a separate cookie, along with the session ID cookie. This would eliminate the need for the server to maintain a separate table for tokens. However, this approach has its own drawbacks. Attackers may find ways to inject their own chosen token into a victim's browser, allowing them to bypass the CSRF protection.

Web security models, such as the use of anti-CSRF tokens, are essential for protecting web applications from CSRF attacks. These tokens ensure that requests are legitimate and prevent unauthorized actions on behalf of users. Web application developers should be familiar with the mechanisms behind CSRF protection and the potential vulnerabilities associated with it.

The same-origin policy is a fundamental security feature in web browsers that restricts how web pages or scripts loaded from one origin can interact with resources from a different origin. However, there are exceptions to this policy that can be exploited by attackers. One such exception is related to Cross-Site Request Forgery (CSRF) attacks.

In a CSRF attack, an attacker tricks a user's browser into making a request to a vulnerable website on which the user is authenticated. To prevent CSRF attacks, websites often include a CSRF defense mechanism that generates a unique token for each user session. This token is then included in every request made by the user, and the server verifies its authenticity before processing the request. However, if the implementation of the CSRF defense is flawed, hackers can bypass it and carry out CSRF attacks.

Another exception to the same-origin policy involves JavaScript code itself. JavaScript code can be included in a web page using the `<script>` tag. The question arises: what happens if a web page from one origin includes JavaScript code from a different origin? In this case, the browser allows the code to be loaded and executed, even though it is from a different origin. The code is then executed in the same origin as the web page itself.

This exception can be compared to using a shared library or calling code from a different source in a process. The JavaScript code, regardless of its origin, runs in the same context as the web page, with all the same privileges. If the code contains a bug, vulnerability, or malicious code, it can potentially take over the entire web page and its functionalities.

In the case of images, we were concerned about whether a web page could access the contents of an image from a different origin. However, when it comes to JavaScript code, it is a bit different. Although the web page cannot directly access the literal ASCII representation of the code, it can infer a lot about the code's structure and functionality. JavaScript has reflection features that allow it to inspect functions, conditions, and other aspects of the code.

To protect against potential attacks, it is important that any sensitive information or secrets do not resemble JavaScript code. If a secret accidentally resembles JavaScript code, an attacker could load it and examine the entire code of the web page, potentially exposing sensitive information. This is particularly concerning because secrets are often loaded with cookies, making them accessible to attackers.

A recommended practice to mitigate this risk is to include an infinite loop at the beginning of the response that makes the JavaScript interpreter continuously loop. This prevents the accidental execution of any code that might resemble JavaScript. For example, if a web page exposes data that can be fetched by providing cookies without any CSRF protection, the response should start with an infinite loop statement, followed by the actual data.

Understanding the exceptions to the same-origin policy in web security is crucial for developing secure web applications. By being aware of how JavaScript code can be loaded from different origins and the potential risks associated with it, developers can implement appropriate security measures to protect against attacks.

Cross-Site Scripting (XSS) is a type of attack that exploits vulnerabilities in web applications. In this attack, an attacker uses a user's browser to send requests to a target website, such as Facebook. The goal is to trick the website into sending malicious code to other users.

To understand how XSS works, let's look at an example. Suppose the attacker wants to target Facebook. They don't need to own a website themselves; instead, they can use their own browser to send requests to Facebook. The attacker relies on the fact that Facebook allows users to share data, such as posts.

The attacker might post an article with a script source, something like `<script src="malicious-url"></script>`. If Facebook doesn't properly handle this input, it might send the script source to other users when they request the list of postings. From the victim's browser's perspective, it receives the script source tag, which can run JavaScript code from the attacker's chosen URL.

This means that the attacker can control what JavaScript code runs in the victim's browser when they visit facebook.com. Depending on the sensitivity of the victim's account, this can lead to unauthorized actions or data breaches.

To prevent XSS attacks, web applications need to implement proper security measures. One approach is to sanitize user input and restrict the use of certain tags, such as angle brackets. However, this can limit the formatting options available to users.

Another approach is to filter tags, allowing only specific ones like bold or italics, but blocking dangerous ones like script tags. However, this can be challenging because there are many ways to embed JavaScript code in HTML, and it's difficult to predict which ones browsers will execute.

Another defense mechanism is the use of HTTP-only cookies. These cookies are only sent from the browser to the server and are not accessible to JavaScript running in the browser. While this seems like a good idea to prevent XSS attacks, it is not a foolproof solution.

XSS attacks highlight the importance of designing web security models that can effectively isolate and protect against such vulnerabilities. It is crucial for web developers to implement proper input validation, output encoding, and secure coding practices to mitigate the risk of XSS attacks.

Network security is a critical aspect of cybersecurity, especially when it comes to web security. In the context of web security, one of the vulnerabilities that attackers exploit is Cross-Site Request Forgery (CSRF). This attack involves tricking a user's browser into performing unwanted actions on a targeted website without their knowledge or consent.

One way attackers can execute CSRF attacks is by stealing the victim's cookie, which contains authentication information, and sending it to the attacker. With the victim's cookie, the attacker gains unauthorized access to the victim's account. To mitigate this risk, developers initially attempted to prevent JavaScript from accessing cookies, assuming it would provide sufficient protection. However, this approach proved to be only a partial defense.

Another solution that emerged later is the HTTP Content Security Policy (CSP). CSP is an opt-in mechanism that allows web applications or web pages to set a special header in the victim's browser. This header informs the browser that the website does not expect to receive certain types of content, such as scripts, from specific origins. By configuring the CSP policy, developers can restrict the loading of JavaScript code to only trusted sources, enhancing browser security against various attacks.

CSP offers a range of settings that developers can opt into to improve browser security. For example, they can specify that only JavaScript code served from particular origins should be executed, preventing the loading of JavaScript libraries from untrusted sources. Many websites have adopted CSP as a security measure, although its implementation can be challenging for complex sites like Facebook. Consequently, CSP is not enabled by default.

Web security has evolved over time to address vulnerabilities like Cross-Site Request Forgery. While attempts to fix the web's security issues have had mixed success, the introduction of the HTTP Content Security Policy (CSP) has significantly improved browser security. By implementing CSP, web developers can specify which origins are allowed to serve JavaScript code, reducing the risk of attacks.

EITC/IS/ACSS ADVANCED COMPUTER SYSTEMS SECURITY DIDACTIC MATERIALS

LESSON: NETWORK SECURITY

TOPIC: NETWORK SECURITY

This part of the material is currently undergoing an update and will be republished shortly.

EITC/IS/ACSS ADVANCED COMPUTER SYSTEMS SECURITY DIDACTIC MATERIALS

LESSON: NETWORK SECURITY

TOPIC: SECURE CHANNELS

This part of the material is currently undergoing an update and will be republished shortly.

EITC/IS/ACSS ADVANCED COMPUTER SYSTEMS SECURITY DIDACTIC MATERIALS

LESSON: NETWORK SECURITY

TOPIC: CERTIFICATES

This part of the material is currently undergoing an update and will be republished shortly.

EITC/IS/ACSS ADVANCED COMPUTER SYSTEMS SECURITY DIDACTIC MATERIALS

LESSON: IMPLEMENTING PRACTICAL INFORMATION SECURITY

TOPIC: INFORMATION SECURITY IN REAL LIFE

This part of the material is currently undergoing an update and will be republished shortly.

EITC/IS/ACSS ADVANCED COMPUTER SYSTEMS SECURITY DIDACTIC MATERIALS

LESSON: MESSAGING

TOPIC: MESSAGING SECURITY

This part of the material is currently undergoing an update and will be republished shortly.

EITC/IS/ACSS ADVANCED COMPUTER SYSTEMS SECURITY DIDACTIC MATERIALS

LESSON: SECURITY OF STORAGE

TOPIC: UNTRUSTED STORAGE SERVERS

This part of the material is currently undergoing an update and will be republished shortly.

EITC/IS/ACSS ADVANCED COMPUTER SYSTEMS SECURITY DIDACTIC MATERIALS

LESSON: TIMING ATTACKS

TOPIC: CPU TIMING ATTACKS

This part of the material is currently undergoing an update and will be republished shortly.