



# **European IT Certification Curriculum Self-Learning Preparatory Materials**

EITC/IS/WASF  
Web Applications Security Fundamentals



This document constitutes European IT Certification curriculum self-learning preparatory material for the EITC/IS/WASF Web Applications Security Fundamentals 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/WASF Web Applications Security Fundamentals 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/WASF Web Applications Security Fundamentals 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/WASF Web Applications Security Fundamentals certification programme curriculum as published on its relevant webpage, accessible at:

<https://eitca.org/certification/eitc-is-wasf-web-applications-security-fundamentals/>

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

<b>Introduction</b>	<b>4</b>
Introduction to web security, HTML and JavaScript review	4
<b>Web protocols</b>	<b>12</b>
DNS, HTTP, cookies, sessions	12
<b>Session attacks</b>	<b>22</b>
Cookie and session attacks	22
<b>Same Origin Policy</b>	<b>32</b>
Cross-Site Request Forgery	32
Exceptions to the Same Origin Policy	41
<b>Cross-site scripting</b>	<b>51</b>
Cross-Site Scripting (XSS)	51
Cross-Site Scripting defenses	61
<b>Web fingerprinting</b>	<b>71</b>
Fingerprinting and privacy on the web	71
<b>DoS, phishing and side channels</b>	<b>81</b>
Denial-of-service, phishing and side channels	81
<b>Injection attacks</b>	<b>92</b>
Code injection	92
<b>TLS attacks</b>	<b>103</b>
Transport layer security	103
<b>HTTPS in the real world</b>	<b>116</b>
HTTPS in the real world	116
<b>Authentication</b>	<b>125</b>
Introduction to authentication	125
WebAuthn	138
<b>Managing web security</b>	<b>151</b>
Managing security concerns in Node.js project	151
<b>Server security</b>	<b>161</b>
Server security: safe coding practices	161
Local HTTP server security	171
<b>DNS attacks</b>	<b>182</b>
DNS rebinding attacks	182
<b>Browser attacks</b>	<b>192</b>
Browser architecture, writing secure code	192
<b>Practical web applications security</b>	<b>205</b>
Securing web applications with modern platform features	205

**EITC/IS/WASF WEB APPLICATIONS SECURITY FUNDAMENTALS DIDACTIC MATERIALS****LESSON: INTRODUCTION****TOPIC: INTRODUCTION TO WEB SECURITY, HTML AND JAVASCRIPT REVIEW**

Welcome to the introduction to web security. In this class, we will cover the fundamentals of web application security, including an overview of HTML and JavaScript. The goal of this course is to provide you with a solid understanding of web security principles and techniques.

Throughout the course, we will have a series of assignments, mostly programming assignments, to help you apply the concepts you learn. Additionally, there may be a written assignment with questions to test your knowledge. Assignment 0 will be released tonight, so make sure to check it out and get a feel for what to expect.

We will also have guest lectures, which will make up about a third of the class. These lectures will be delivered by industry professionals who will share cutting-edge web security topics with you. Some of the confirmed speakers include representatives from the Brave web browser and the Google Chrome team. We are still finalizing the lineup, but rest assured, we have some exciting speakers in store for you.

For questions and discussions, we will be using Piazza, an online platform. If you have any feedback or suggestions for the class, there is a form on the website where you can submit them anonymously.

Now, let's talk about assignment 0. The purpose of this assignment is to assess your HTML and JavaScript knowledge, which will be essential for the rest of the course. Whether this is your first security class or not, assignment 0 will help you gauge your skills and identify any areas that may need improvement.

Assignment 0 consists of three workshops: HTML, JavaScript, and Node.js. These workshops are interactive and will require you to solve exercises using a terminal interface. The HTML and JavaScript workshops should be relatively easy for most of you, especially if you have recently taken a web development course like CS 142. However, if you are a bit rusty, don't worry. Just work through the exercises, and it should not be too challenging.

The Node.js workshop, on the other hand, may be new to many of you. We will cover various aspects of Node.js, and the exercises will guide you through the process. The workshop interface is user-friendly and provides instant feedback on your solutions.

If you have any questions or need clarification, feel free to ask. We are here to help you throughout this course, and we hope you find the material engaging and informative.

Web security is an important aspect of cybersecurity, particularly when it comes to web applications. In this course, we will explore the fundamentals of web security, with a focus on the introduction to web security, HTML, and JavaScript review.

The course was originally offered as CS 241 secure web programming in 2011. It was taught by Dan Binet and John Mitchell. The class provided valuable insights into web security and was well-received by students. The material covered in the course was designed to make web security approachable and hands-on. It emphasized the fact that anyone can learn and explore web security concepts.

One of the key takeaways from the course is the idea of having a security mindset. This mindset involves considering security implications while writing code and examining various aspects of web applications. By adopting this mindset, individuals can better understand potential vulnerabilities and actively search for them.

During the course, students were encouraged to experiment and explore different aspects of web security. For example, one student discovered a vulnerability in the local storage API. The API allowed storing a string in the user's browser, with a supposed limit of five megabytes. However, the student found a way to bypass this limit and fill up the user's browser, potentially causing their computer to run out of space or crash. This highlighted the importance of thorough testing and considering potential security risks in API design.

Another interesting topic covered in the course was "cliff jacking" attack. In this attack, users were led to

believe they were playing a game by clicking on a moving button. However, an invisible window positioned under the user's mouse captured their clicks without their knowledge. This demonstrated the potential for deceptive practices and the need for user awareness.

Web security is an ever-evolving field, and mistakes can happen. By learning about these mistakes and vulnerabilities, individuals can contribute to a safer web environment. This course aims to reintroduce the valuable content covered in the original CS 241 class, with a focus on web security fundamentals.

#### Web Applications Security Fundamentals - Introduction

In this class, we will be discussing the fundamentals of web security, with a focus on web applications. We will begin by introducing the concept of web security and reviewing the basics of HTML and JavaScript.

One important topic we will cover is clickjacking attacks. These attacks involve tricking users into clicking on something they did not intend to click on, leading to potential security risks. We will learn how to defend against such attacks.

Having the mindset of an attacker is crucial when it comes to writing secure code. By thinking like an attacker, we can identify potential vulnerabilities in our code and address them before they can be exploited. Likewise, understanding the techniques and thought processes of attackers is essential for designing secure systems.

On the other hand, the defender mindset is equally important. To effectively defend against attacks, we must understand how they are carried out and what techniques attackers employ. It is impossible to defend something if we do not know how it can be attacked.

When it comes to the difficulty of attacking versus defending, it is generally harder to defend. Attackers only need to find one vulnerability, while defenders must protect against all possible attack vectors. Additionally, it is challenging to ensure code security, as there may be hidden vulnerabilities or bugs that are not immediately apparent.

In the real world, it is impossible to guarantee that code is completely secure. However, we can adopt techniques and practices that increase the security of our systems. Throughout this course, we will explore various strategies for living in a world where perfect code security is unattainable.

As an incentive for students, I have introduced an extra credit policy. If you discover a security vulnerability in any system during the quarter, you will receive extra credit. The amount of credit will vary depending on the severity and impact of the vulnerability. This policy is intended to encourage exploration and creativity in the field of cybersecurity.

It is important to note that responsible disclosure is essential when reporting vulnerabilities. I recommend familiarizing yourself with the Stanford bug bounty program website, which outlines the requirements for reporting bugs in Stanford websites. Remember to exercise caution and not engage in any activities that could cause harm or legal issues.

This extra credit opportunity is entirely optional, and it is possible to excel in the course without participating. However, I encourage you to consider it as it can enhance your learning experience and contribute to a more engaging class environment.

In the upcoming sessions, we will delve deeper into the technical aspects of web security. Today's session served as a general overview, laying the foundation for our future discussions.

Web security is a crucial topic in the field of computer security. It is important to understand the reasons why web security is challenging. One reason is the presence of buggy code in web applications. Many programs contain bugs that are not necessarily security-related but can still impact the functionality of the program. Even if a program appears to be functioning correctly, it can still have security vulnerabilities.

Another factor that contributes to the difficulty of web security is social engineering. Social engineering involves manipulating individuals to gain access to sensitive information. For example, an attacker may impersonate someone and request passwords or other confidential data. Social engineering has proven to be surprisingly

effective and can be used to exploit vulnerabilities.

The motivation for exploiting vulnerabilities is often financial gain. There is a marketplace for buying and selling vulnerabilities, where individuals can profit from stolen information or by taking advantage of security flaws. Some individuals make a living by actively searching for vulnerabilities in web applications.

It is important to differentiate between vulnerabilities and exploits. A vulnerability refers to a flaw or weakness in a system that can be exploited. An exploit, on the other hand, is a technique or code that takes advantage of a vulnerability to gain unauthorized access or control over a system. Finding vulnerabilities is just the first step, as there are individuals who specialize in weaponizing these vulnerabilities and others who exploit them for malicious purposes.

Once a machine is compromised, there are various actions that can be taken. These include taking the machine offline, launching attacks on competitors, stealing sensitive information such as credit card numbers, creating physical cards with stolen information, or even using the compromised machine to carry out further attacks. Each of these actions represents a different marketplace within the realm of cybercrime.

Web security is challenging due to the presence of buggy code, the effectiveness of social engineering, and the financial incentives for exploiting vulnerabilities. Understanding these factors is crucial in developing effective strategies to protect web applications from potential threats.

Web security is a crucial aspect of cybersecurity, especially when it comes to web applications. In this material, we will provide an introduction to web security, as well as review important concepts related to HTML and JavaScript.

One common technique used by attackers is to utilize someone else's IP address to send spam emails. By doing so, the spam is more likely to bypass filters and appear legitimate. This highlights the importance of ensuring the security of web servers, as they can be exploited for malicious purposes.

Denial-of-service (DoS) attacks are another prevalent threat. Attackers overwhelm a website with an excessive amount of traffic, causing it to go offline. They then demand a ransom in exchange for restoring the site's functionality. This type of attack can result in financial losses for the targeted organization.

Web attacks can target either client machines or web servers. When targeting client machines, attackers exploit vulnerabilities in web browsers used by users visiting a site. On the other hand, attacking web servers allows hackers to infect numerous machines by compromising a single server. This is particularly effective when popular sites are targeted, as a significant percentage of visitors may be using outdated and vulnerable browsers.

Data theft is a serious concern in web security. Attackers aim to steal sensitive information, such as social security numbers, payment card details, health insurance data, and driver's license information. These attacks can result in millions of individuals having their personal information compromised.

Ransomware attacks have become increasingly prevalent in recent years. Attackers gain unauthorized access to computers and encrypt the victim's files. They then demand a ransom, typically in cryptocurrency, to provide the decryption key and restore access to the files. Cryptocurrencies have facilitated the monetization of these attacks, making them more appealing to cybercriminals.

Political attacks are also on the rise. These attacks target political organizations and institutions, aiming to gain unauthorized access to sensitive information or disrupt operations. The DNC hack serves as an example of such attacks.

Moving on to web security fundamentals, the same-origin policy is a crucial concept to understand. It is a model that governs web security and ensures that websites are isolated from each other. For instance, when browsing Wikipedia and accessing online banking in the same browser, the same-origin policy prevents Wikipedia from accessing the banking information. This policy, although seemingly common sense, required careful engineering to enforce.

Web security is a critical aspect of cybersecurity, particularly when it comes to web applications. Understanding

the various threats, such as spam, DoS attacks, data theft, ransomware, and political attacks, is essential for effectively protecting web servers and client machines. Additionally, grasping the concept of the same-origin policy is crucial in ensuring the isolation and security of websites.

Web security is a crucial aspect of ensuring the safety and integrity of web applications. In this didactic material, we will explore the fundamentals of web security, with a focus on the introduction to web security, as well as a review of HTML and JavaScript.

One important consideration in web security is how different sites interact with each other. In a browser context, we often encounter mashups, where content from one source, such as an ad network, is displayed on another site. It is essential to ensure that the code running in these contexts cannot manipulate or interfere with the content outside of its designated area.

Another perspective to consider is web security from the server angle. When writing a web server, the goal is to prevent attacks and unauthorized access. This involves defining a set of rules that determine what actions clients are allowed to perform. However, the challenge lies in the fact that the server does not have control over the client's code, which runs in the browser. Clients can modify their code and even generate HTTP requests that a normal browser would not send. Therefore, server-side decisions must account for potential client-side modifications.

To illustrate this, let's consider an example where a user sends a Perl request to a server. Even if the server does not provide a button or an interface to trigger that specific request, a user can still send it directly. This highlights the importance of considering potential client-side actions when designing server-side security measures.

On the client-side, we must also address the security of the code running in the user's browser. It is crucial to ensure that a user cannot be attacked by malicious code embedded in URLs or links. For instance, a technique called cross-site scripting can trick a web application into executing unauthorized code. This can have severe consequences, such as leaking sensitive information or compromising the user's browsing experience.

To emphasize the impact of cross-site scripting, let's examine an example involving a McDonald's website. In this case, an attacker found a vulnerability that allowed them to extract users' passwords simply by visiting the site. This demonstrates how easily a user can be tricked into visiting a malicious URL, even without direct interaction with the attacker. It is essential to prevent any formulation of URLs that could cause the user's browser to execute unintended code.

Web security encompasses various aspects, including the interaction between different sites, server-side security, and client-side code security. By understanding these fundamentals, we can develop robust web applications that protect against common vulnerabilities and ensure the safety of users' data and browsing experience.

Web applications are vulnerable to various security threats, and it is crucial to understand these threats in order to protect against them. One common vulnerability is the ability for attackers to inject malicious code into a website's database, which then gets executed when the page is rendered. This allows attackers to exploit the privileges of logged-in users, such as creating new admin accounts. By targeting WordPress sites, for example, attackers can gain access to a large number of installations. It is important to note that even if users have strong passwords and secure configurations, they can still fall victim to social engineering attacks. Therefore, it is essential to design systems that are secure even in the presence of social engineering. Additionally, web users should be aware of the extensive tracking and data profiling that occurs on the web. Despite being technically literate, individuals can still be vulnerable to these practices. An example of a dubious attack is typo squatting, where an attacker registers a name similar to a popular package or website. By relying on users' typos, attackers can trick them into installing their malicious code. In one case, an undergraduate student registered package names similar to popular ones and monitored how many people mistakenly installed his package. Although this attack did not cause any harm, it raised concerns about the security of the open-source ecosystem. It is important for users to regularly audit the code they run on their computers to ensure its integrity. By being vigilant and taking necessary precautions, individuals can protect themselves against web application security threats.

Web security is a challenging problem due to several factors. One of the main difficulties arises from the fact



that web applications often rely on code from various sources, making it virtually impossible to audit all of it. This poses a risk as malicious code can be introduced, either intentionally or accidentally, compromising the security of the application. Additionally, the constantly changing nature of the code further complicates the auditing process.

To address this issue, JavaScript package managers like NPM have implemented measures to prevent potential attacks. For instance, when registering a package, NPM now checks for names that are similar to popular packages and denies registration to avoid confusion and potential security vulnerabilities. However, even with such measures in place, there are still ways for attackers to bypass them.

Protecting users from trackers is another aspect of web security. Chrome extensions, for example, can be installed to monitor the domains a browser connects to. By highlighting the domains visited directly by the user and those that receive information from the visited site, users can gain insight into the tracking practices prevalent on the web.

Web security is challenging due to the technical decisions made during the design of the web. The web was initially created in a different era, primarily for academic use, leading to unforeseen consequences as it evolved into the complex platform it is today. This is similar to the challenges faced in networking, where the internet was initially designed for academic use, and the consequences are still being dealt with.

Fundamentally, web security is difficult because browsers aim to execute code from various sources securely. The goal is to allow users to run code from untrusted individuals without any negative consequences. This ambitious objective, combined with the need for different sites to interact with each other in the same user context, adds to the complexity of web security. Additionally, the desire for high-performance web applications introduces low-level features and APIs that need to be secured, further complicating the task.

Web security is a complex and challenging problem due to the need to run untrusted code securely, the interaction between different sites, the constantly changing nature of code, and the desire for high-performance web applications. Efforts have been made to mitigate these challenges, but the evolving nature of the web and the inherent design decisions continue to pose difficulties.

The evolution of the web has led to its purpose becoming more complex. Unlike other platforms such as iOS, Android, Mac, and Windows, the web has strict backwards compatibility requirements. This means that changes cannot be made that would break old websites. Browser vendors are reluctant to make changes that could potentially lead to websites breaking and losing users. Additionally, there are countless old websites that cannot be fixed or broken due to their historical significance. While this is one of the great aspects of the web, it also makes ensuring good security challenging.

Developers of modern web applications must navigate a tangle of technologies that have been developed over time and pieced together haphazardly. Each component of the web application stack, from HTTP requests to browser-side scripts, comes with important yet subtle security consequences. It is crucial for developers to understand these consequences in order to keep users safe.

The model for the web has changed significantly since its creation. For example, the use of secure connections was originally thought to be necessary only for transmitting sensitive information like credit card numbers. However, this has changed due to factors such as browser plugins that can intercept cookies on local networks. The revelation of government surveillance, as seen in the Snowden revelations, has also played a major role in the shift towards using HTTPS. Websites are now being pushed by browsers to adopt HTTPS, and browsers even display warnings for sites that are not secure.

The browser has an incredibly challenging task of allowing various functionalities while ensuring security. It must allow known malicious actors to download code, spawn processes, open sockets, and connect to servers, all while protecting the user and their data. The browser also faces competing pressures regarding its role. Some believe it should be a simple document viewer with fewer features and less JavaScript, while others want it to be a full-fledged operating system that can handle any task safely.

The web's evolution and strict backwards compatibility requirements present challenges for web security. Developers must navigate a complex landscape of technologies and understand the security consequences of each component. The shift towards using HTTPS has been driven by factors such as government surveillance



and browser warnings. The browser itself faces the difficult task of balancing functionality and security. Understanding these challenges is essential for developers to ensure the safety of web applications and users.

Web security is an important aspect of cybersecurity, especially when it comes to web applications. As web applications become more complex and feature-rich, they also become more vulnerable to attacks. The more interactions and APIs there are, the higher the chances of potential bugs and security vulnerabilities.

However, despite these challenges, the web is a robust platform that has evolved over time to incorporate various security measures. It is the most attacked platform, yet it has proven to be resilient. As Ilya Gregorek, an engineer at Google, said, "It's all too easy to criticize lament and create paranoid scenarios about the unsound security foundations of the web, but the truth is all of that criticism is true and yet the web is proven to be an incredibly robust platform."

In this course, we will learn how to build secure web systems by discussing secure architecture, patterns, and concepts. We will also explore ways to ensure that even if one component fails, the overall system remains secure. This involves having multiple backups and contingency plans in place.

The course will cover various topics, starting with the browser security model and the same-origin policy, which is a crucial security concept for web security. We will then delve into ways to attack the client (user's browser) and the server. Secure authentication will also be a key focus, with a guest lecture from GitHub's security team on their authentication methods, including biometrics.

Additionally, we will explore secure connections through TLS (HTTP) and learn from a guest lecture by Emily and Chris from Google Chrome's security team on how they ensure TLS security in Chrome. Lastly, we will cover how to write secure code.

The course will conclude with a final exam, which will account for 25% of the total grade. There will be no midterm. The course can be counted as a CS elective, and it will be included on the program sheet.

Before delving into the course material, we will begin with a quick review of web technologies, specifically HTML. HTML is used to format and structure web documents. We use tags to define headings, paragraphs, lists, and other elements. For example, the "ul" tag represents an unordered list, while the "li" tag represents a list item.

By understanding the fundamentals of web security, including the browser security model, authentication, and secure connections, we can build robust and secure web applications.

URLs are an essential part of web applications and understanding their structure is crucial for web security. A URL consists of several components, including the scheme, host name, port number, path, query, and fragment.

The scheme refers to the protocol being used, such as HTTP or HTTPS. The host name represents the domain name of the server being connected to. The port number specifies the communication channel on the server. The path, which used to be a literal path to a file on the server's file system, now often represents the name of the request being made. The query is used for dynamic server endpoints and includes parameters that configure the type of request. The query starts with a question mark and uses key-value pairs separated by ampersands. If an ampersand needs to be included in the key or value, it needs to be escaped.

The fragment component is less common and refers to a specific part of the document that is labeled with a name. It allows jumping directly to that part of the page. A new specification is being considered that allows for more flexibility in linking to specific parts of a page without the need for predefined anchors.

In HTML, there are different ways to specify URLs. The full URL includes the complete address. A relative URL is based on the current page and appends the specified path to the current URL. If a trailing slash is present, it indicates that the URL is within a specific folder. Without the trailing slash, the URL is treated as a replacement for the current folder. An absolute URL starts with a slash and represents the root of the website.

Understanding URL structure is essential for web security as it helps distinguish between user input and commands. This knowledge is crucial in preventing attacks that exploit vulnerabilities in web applications.

Web security is an important aspect of cybersecurity. In this class, we will focus on the fundamentals of web application security, starting with an introduction to web security, as well as a review of HTML and JavaScript.

When it comes to web security, understanding the structure and vulnerabilities of web applications is crucial. Web applications are built using HTML, which stands for Hypertext Markup Language. HTML consists of various tags that define the structure and content of a web page. Some important tags we will discuss in this class include:

- Link: The link tag is used to include an external CSS file, which is responsible for styling the web page. It allows us to separate the design from the content.
- Style: The style tag is used for inline CSS, where we can directly write CSS code within the HTML document.
- Script: The script tag is used to include JavaScript code in the web page. JavaScript is a programming language that allows us to add interactivity and dynamic behavior to web applications. It can be either referenced from an external source or written directly within the HTML document.

Understanding these tags is essential because they are often the target of attacks. By exploiting vulnerabilities in HTML, CSS, and JavaScript, attackers can gain unauthorized access to sensitive information or manipulate the behavior of web applications.

It is worth noting that the decisions made when these tags were introduced still have consequences today. The web cannot break backward compatibility with old websites, which means we have to work around these decisions to ensure the security of our web applications.

JavaScript, in particular, is a powerful and flexible language. It allows developers to quickly prototype and experiment with ideas. However, this flexibility can also lead to potential security risks. For example, JavaScript allows redefining the value of "undefined," which can cause unexpected behavior in code.

Despite its initial flaws, JavaScript has evolved over time and now offers many features that developers desire. However, the flexibility of the language can sometimes create more complexity and potential vulnerabilities. It is important to be aware of these trade-offs when developing secure web applications.

Understanding web security, HTML, and JavaScript is crucial for building secure web applications. By familiarizing ourselves with the structure and vulnerabilities of web applications, as well as the features and potential risks of HTML and JavaScript, we can better protect our applications from attacks.

#### Web Security Fundamentals: Introduction to Web Security, HTML, and JavaScript Review

Web security is a crucial aspect of cybersecurity, particularly when it comes to protecting web applications. In this material, we will explore the basics of web security, as well as review important concepts related to HTML and JavaScript.

Web security is essential because web applications are vulnerable to various threats, such as unauthorized access, data breaches, and malicious attacks. Understanding the fundamentals of web security is crucial for developers and system administrators to ensure the safety and integrity of web applications.

HTML (Hypertext Markup Language) is the standard language used for creating web pages. It provides the structure and layout of the content on a web page. HTML tags are used to define the elements of a web page, such as headings, paragraphs, images, and links. It is important to note that HTML itself does not have any security features. However, developers must be aware of potential security vulnerabilities that can arise from improper use of HTML tags, such as cross-site scripting (XSS) attacks.

JavaScript, on the other hand, is a programming language that allows developers to add interactivity and functionality to web pages. It is executed on the client-side, meaning it runs directly in the user's browser. JavaScript can be used to validate user input, manipulate HTML elements, and interact with web servers. However, it is important to use JavaScript securely to prevent security risks, such as cross-site scripting (XSS) and cross-site request forgery (CSRF) attacks.

Node.js is a runtime environment that allows developers to run JavaScript on the server-side. It was created to extend the capabilities of JavaScript beyond the browser. Node.js provides additional functionality, such as file

system access, HTTP requests, and networking capabilities. This enables developers to build scalable and efficient web applications. Understanding the role of Node.js in web development is important for ensuring the security and performance of web applications.

It is crucial to differentiate between JavaScript as a language and the APIs (Application Programming Interfaces) provided by browsers or Node.js. JavaScript itself is a powerful and flexible language, but the additional APIs provided by browsers and Node.js introduce additional functionality and capabilities. The Document Object Model (DOM) API, for example, allows developers to manipulate the structure and content of web pages. However, it is important to note that not all APIs are part of the JavaScript language itself. Some APIs, such as those related to the DOM, are provided by the browser, while others are specific to Node.js.

Web browsers have evolved over time, and many old and outdated APIs have been removed. However, some legacy APIs still exist, which can pose security risks. For example, the `window.open` API allows pop-ups to be displayed, which can be exploited by malicious actors. It is important for developers to be aware of these APIs and their potential security implications.

Understanding the fundamentals of web security, as well as the basics of HTML, JavaScript, and Node.js, is crucial for building secure and robust web applications. By following best practices and staying updated on the latest security threats and vulnerabilities, developers can ensure the safety and integrity of their web applications.

Web security is a crucial aspect of cybersecurity, especially when it comes to protecting web applications. In this material, we will provide an introduction to web security, along with a review of HTML and JavaScript.

Web security involves implementing measures to protect web applications from various threats, such as unauthorized access, data breaches, and malicious attacks. It is essential to ensure the confidentiality, integrity, and availability of web application resources.

HTML (Hypertext Markup Language) is the standard markup language used for creating web pages. It provides the structure and content of a webpage. Understanding HTML is essential for web developers and security professionals to identify potential vulnerabilities and implement appropriate security controls.

JavaScript is a programming language commonly used for adding interactivity and dynamic features to web pages. It runs on the client-side, allowing web applications to respond to user actions. However, JavaScript can also introduce security risks if not properly implemented and secured.

One important aspect of web security is ensuring that user input is properly validated and sanitized to prevent malicious code injection, such as cross-site scripting (XSS) attacks. XSS attacks occur when an attacker injects malicious scripts into a web application, which are then executed by unsuspecting users' browsers.

Another common vulnerability is cross-site request forgery (CSRF), where an attacker tricks a user into performing unintended actions on a web application. This can be mitigated by implementing CSRF tokens and validating requests to ensure they originate from trusted sources.

Web security also involves protecting sensitive data transmitted between the client and the server. This can be achieved through secure communication protocols like HTTPS, which encrypts data to prevent eavesdropping and tampering.

In addition to these fundamental concepts, it is important to understand other web security topics, such as session management, cookies, and the same-origin policy. These topics will be covered in more detail in future materials.

Web security is a critical aspect of cybersecurity, particularly in safeguarding web applications from various threats. Understanding HTML and JavaScript is essential for identifying vulnerabilities and implementing appropriate security controls. By following best practices and staying updated on the latest security techniques, we can help ensure the safety and integrity of web applications.

**EITC/IS/WASF WEB APPLICATIONS SECURITY FUNDAMENTALS DIDACTIC MATERIALS****LESSON: WEB PROTOCOLS****TOPIC: DNS, HTTP, COOKIES, SESSIONS**

When you type a URL into your browser and press Enter, several steps occur to establish a connection with the server. The first step is a DNS query. DNS (Domain Name System) is a system that translates user-friendly domain names into IP addresses. IP addresses are used to address all nodes connected to the internet, while domains are the human-readable strings we type into our browsers. The browser needs to perform this translation to determine which server to contact.

The DNS query process involves the browser making a request to a DNS server. The DNS server receives the query and responds with the IP address for the requested domain. This process is crucial because the browser cannot establish a connection without knowing the IP address.

At a high level, the browser sends a query to the DNS server, asking for the IP address of a specific domain (e.g., stanford.edu). The DNS server, also known as a recursive resolver, takes on the responsibility of looking up the IP address for the domain. It performs the necessary work to find the answer and sends it back to the client. This process may involve multiple steps, which is why it is called a recursive resolver.

The recursive resolver starts by sending the query to the root nameservers. These nameservers are hard-coded into every DNS resolver and are operated by different organizations. The recursive resolver then delegates authority to another nameserver responsible for a subset of domains within the top-level domain (TLD). In the case of stanford.edu, the recursive resolver asks the TLD nameserver for the IP address. The TLD nameserver, unable to handle a listing of all domains, responds by directing the recursive resolver to the authoritative nameserver for stanford.edu. The authoritative nameserver finally provides the IP address for Stanford University.

Once the IP address is obtained, the next step is to establish an HTTP connection with the server. This connection allows the browser to retrieve and display the requested web page.

When you type a URL into your browser, the browser initiates a DNS query to translate the domain name into an IP address. The DNS server, acting as a recursive resolver, performs the necessary steps to find the IP address and returns it to the browser. With the IP address, the browser can establish an HTTP connection with the server to retrieve the web page.

When it comes to web application security, understanding web protocols is crucial. In this material, we will focus on three important protocols: DNS, HTTP, and sessions.

DNS, or Domain Name System, is responsible for translating domain names into IP addresses. When we enter a website URL in our browser, DNS resolves that domain name to the corresponding IP address. This IP address tells us which server to connect to. However, DNS queries are sent in plain text, making them vulnerable to attacks. Any device on the internet between us and the recursive resolver can intercept and modify these queries. This allows an attacker to redirect our connection to a different server, potentially leading to malicious activities.

One common attack in this architecture is called a DNS spoofing attack. In this attack, the attacker intercepts the DNS query and responds with a different IP address. As a result, we unknowingly connect to the wrong server. To defend against this attack, we can use Transport Layer Security (TLS). TLS provides encryption and ensures the integrity and authenticity of the response. When we see HTTPS in our browser, it means that TLS is being used to secure the connection.

Another way attackers can exploit web protocols is by manipulating HTTP requests and responses. By changing the DNS records of a target domain, an attacker can redirect visitors to their own web server. This is known as phishing. In a phishing attack, the attacker can collect users' passwords or engage in other malicious activities. Additionally, attackers can insert JavaScript code into web pages to mine cryptocurrencies using the visitors' computing resources. This can lead to financial gain for the attacker at the expense of the user's electricity consumption.

To carry out these attacks, attackers can target various points in the DNS system. They can infect users' systems with malware that changes the local DNS settings, redirecting domains to different destinations. They can also compromise the DNS resolver itself, which is typically provided by the internet service provider. If the resolver gets hacked, it can start providing incorrect answers to DNS queries.

Understanding web protocols is essential for web application security. DNS, HTTP, and sessions play a significant role in how web applications function and how they can be exploited. By implementing proper security measures like TLS and being aware of potential attack vectors, we can protect ourselves and our systems from cyber threats.

#### Web Protocols - DNS, HTTP, Cookies, Sessions

Web protocols play a crucial role in ensuring the security of web applications. In this lesson, we will focus on three important web protocols: DNS, HTTP, and cookies.

DNS, or Domain Name System, is responsible for translating human-readable domain names, such as stanford.edu, into IP addresses that computers can understand. However, DNS can be vulnerable to attacks. For example, hackers can compromise DNS name servers, causing them to provide incorrect IP addresses. This can lead to users unknowingly connecting to malicious servers.

Another vulnerability lies in the DNS name service provided by certain companies. If an account with such a company is compromised, attackers can manipulate DNS records and redirect users to malicious servers.

HTTP, or Hypertext Transfer Protocol, is the foundation of communication on the web. It defines how messages are formatted and transmitted between clients (browsers) and servers. However, HTTP alone does not provide strong security measures.

One way to enhance security is by using HTTPS, which stands for Hypertext Transfer Protocol Secure. HTTPS adds an additional layer of encryption to HTTP, making it more difficult for attackers to intercept and manipulate data. HTTPS relies on digital certificates to verify the authenticity of servers. Certificates are issued by trusted authorities, such as Let's Encrypt, after verifying domain ownership.

However, even HTTPS is not foolproof. Attackers can exploit vulnerabilities in the certificate issuance process. For example, if a malicious server is able to trick the certificate authority into issuing a certificate for a domain it does not control, it can establish trust with the client's browser.

Cookies and sessions are mechanisms used to maintain stateful interactions between clients and servers. Cookies are small pieces of data stored on the client's browser, while sessions are maintained on the server. They are often used to store user authentication information.

However, cookies can also be exploited by attackers. For example, if an attacker gains access to a user's cookie, they can impersonate the user and gain unauthorized access to their account.

Web protocols such as DNS, HTTP, cookies, and sessions are fundamental to web application security. Understanding their vulnerabilities and implementing appropriate security measures is crucial to protect users' data and ensure the integrity of web applications.

DNS, HTTP, cookies, and sessions are fundamental components of web protocols that play a crucial role in web applications security. In this didactic material, we will discuss the basics of these protocols and their implications for cybersecurity.

DNS, or Domain Name System, is responsible for translating domain names into IP addresses. When you type a domain name in your browser, it sends a DNS query to a recursive resolver, which then resolves the domain name to the corresponding IP address. However, DNS queries are sent in plain text, making them vulnerable to interception and monitoring by internet service providers (ISPs) or other malicious actors.

One issue related to DNS is DNS hijacking, where ISPs intercept invalid domain name queries and redirect users to a different page instead of showing an error. This practice is intended to provide a more useful experience for users, but it can also be exploited by attackers. For example, in the case of a web comic called "questionable

content," the ISP injected itself into the browsing experience, causing a buggy implementation. This kind of behavior is unexpected from an ISP and can be considered a security risk.

HTTP, or Hypertext Transfer Protocol, is the protocol used for communication between web servers and clients. It is also sent in plain text, which means that sensitive information, such as login credentials, can be intercepted by attackers. To mitigate this risk, websites can use HTTPS, which adds a layer of encryption to the communication. HTTPS relies on digital certificates to verify the authenticity of the server, ensuring that the client is communicating with the intended website.

Cookies and sessions are mechanisms used to maintain state and track user interactions on websites. Cookies are small pieces of data stored on the client's browser, while sessions are server-side data that store information about the user's session. These mechanisms help websites remember user preferences and enable personalized experiences. However, cookies can also be used for tracking and profiling users, raising privacy concerns.

To address privacy issues related to DNS and web browsing, users can set their DNS settings to a trusted company that has a privacy policy in place. One example is the CloudFlare DNS service, which claims not to sell users' DNS queries to third parties. Additionally, there is ongoing development of DNS over HTTP, which aims to encrypt DNS queries to prevent interception and monitoring.

Understanding the fundamentals of web protocols, such as DNS, HTTP, cookies, and sessions, is essential for ensuring web applications' security. By being aware of the potential risks and implementing appropriate measures, users and website owners can protect sensitive information and maintain privacy online.

#### Web Protocols: DNS, HTTP, Cookies, Sessions

Web protocols are essential for the functioning of web applications and ensuring secure communication between clients and servers. Three key protocols that play a crucial role in web application security are DNS (Domain Name System), HTTP (Hypertext Transfer Protocol), and sessions.

DNS, or Domain Name System, is responsible for translating human-readable domain names into IP addresses. When a user enters a URL into their browser, the browser sends a DNS query to a DNS server to obtain the IP address associated with the domain name. This process allows the browser to establish a connection with the correct server. It is important to note that DNS queries are typically unencrypted, which can pose security risks. However, some browsers now support DNS over HTTP, which encrypts DNS queries for improved security.

HTTP, or Hypertext Transfer Protocol, is the protocol used for communication between clients (such as browsers) and servers. When a user requests a web page or resource, the browser sends an HTTP request to the server, which then responds with the requested content. HTTP requests consist of several components, including the method (e.g., GET, POST), the path, and the protocol version. HTTP is a text-based protocol, making it easy to read and understand. The response from the server includes a status code, indicating the success or failure of the request.

Sessions play a crucial role in web application security by allowing servers to maintain stateful connections with clients. When a user logs into a web application, a session is created, and a unique session identifier is generated. This identifier is typically stored in a cookie, a small piece of data sent from the server to the client and included in subsequent requests. The server can then use the session identifier to authenticate and authorize the client for subsequent requests, ensuring secure access to protected resources.

Understanding these web protocols is essential for web developers and security professionals to build and maintain secure web applications. By implementing secure DNS practices, such as DNS over HTTP, and understanding the intricacies of HTTP requests and responses, developers can ensure the confidentiality and integrity of data transmitted between clients and servers. Additionally, proper management of sessions and cookies is crucial for preventing unauthorized access to sensitive information.

DNS, HTTP, and sessions are fundamental web protocols that play a vital role in web application security. By understanding and implementing best practices for these protocols, developers can enhance the security of their web applications and protect user data.



HTTP is a stateless protocol used for communication between web servers and clients. It allows clients to send requests to servers and receive responses. HTTP is independent of the underlying transport protocol, meaning it can be used over different communication channels, although TCP is commonly used for its reliability.

One important characteristic of HTTP is that it is stateless. This means that each request made by a client is independent and has no relation to previous or future requests. The protocol itself does not maintain any state. However, web applications can implement state on top of HTTP using techniques such as cookies and sessions.

When interacting with web services, it is common to encounter stateful behavior. For example, when logging into a website and clicking on links, the server remembers the user's login state, allowing them to view different pages while remaining logged in. This is achieved through the use of cookies, which are small pieces of data stored on the client's side and sent with each request to the server.

HTTP status codes are used to indicate the outcome of a request. There are different categories of status codes, with the most common being in the 200 range, indicating a successful request. Status codes in the 300 range indicate redirection, where the server directs the client to another destination. Status codes in the 400 range indicate client errors, while status codes in the 500 range indicate server errors.

One specific status code worth mentioning is 206, which is used in response to a range request. A range request allows the client to request a specific subset of a resource, such as a portion of a video or audio file. The server can then respond with the requested portion, allowing the client to start playing the media before it is fully downloaded. This is useful for seeking to specific parts of media files without having to download the entire file sequentially.

In addition to 206, there are other important status codes to be aware of. For example, 301 and 302 are used for redirection. The main difference between them is that 301 indicates a permanent redirect, while 302 indicates a temporary redirect. When a client receives a 301 redirect, it will automatically redirect to the new URL without checking the original URL again. However, with a 302 redirect, the client will check the original URL again because the redirect may only be temporary.

Understanding the basics of HTTP, its stateless nature, and the use of status codes is crucial for web application security. By implementing proper protocols and handling status codes correctly, developers can ensure the secure and efficient communication between web servers and clients.

#### Web Protocols - DNS, HTTP, Cookies, Sessions

Web protocols are essential in the functioning and security of web applications. In this didactic material, we will explore three important web protocols: DNS, HTTP, and cookies. Additionally, we will discuss the concept of sessions and their role in web application security.

DNS (Domain Name System) is responsible for translating domain names into IP addresses. When a user enters a URL in their browser, the browser sends a DNS request to a DNS server to obtain the IP address associated with that domain name. This IP address is then used to establish a connection with the web server hosting the requested website. DNS plays a crucial role in the proper functioning of web applications by enabling users to access websites using human-readable domain names.

HTTP (Hypertext Transfer Protocol) is the protocol used for communication between web browsers and web servers. It defines the format of the messages exchanged between the client (browser) and the server. HTTP requests are made by the client to retrieve resources such as HTML pages, images, or scripts from the server. The server responds to these requests with HTTP status codes, indicating the success or failure of the request. Some commonly encountered HTTP status codes include:

- 301: This code is used for permanent redirects. It instructs the client to redirect to a different URL permanently. This is useful when implementing functionality on a website that requires redirecting users to different pages randomly.
- 304: This code is used when the browser has cached a resource and wants to check if it has been modified since the last request. If the resource has not been modified, the server responds with a 304 status code, indicating that the client already has the latest version of the resource.



HTTP requests can also include headers that provide additional information to the server. For example, the client can include a header indicating the version of a resource it has, allowing the server to determine if the resource has changed since that version.

Cookies are small pieces of data stored on the client-side by the web server. They are used to maintain state and track user interactions with the website. When a user visits a website, the server can set a cookie on the client's browser. This cookie is then sent back to the server with subsequent requests, allowing the server to identify and personalize the user's experience. Cookies can store information such as user preferences, shopping cart contents, or authentication tokens.

Sessions are a higher-level concept built on top of cookies. A session represents a logical connection between a client and a server. When a user logs in to a website, a session is created, and a unique session identifier is stored in a cookie on the client's browser. This session identifier is used to associate subsequent requests from the client with the corresponding session on the server. Sessions are crucial for implementing secure web applications as they allow servers to maintain user authentication and authorization information across multiple requests.

DNS, HTTP, cookies, and sessions are fundamental web protocols and concepts that play crucial roles in the functioning and security of web applications. Understanding these protocols and their interactions is essential for developers and security professionals working in the field of cybersecurity.

#### Web Protocols: DNS, HTTP, Cookies, Sessions

Web protocols are essential for the functioning of web applications and ensuring the security of user data. In this section, we will discuss three fundamental web protocols: DNS, HTTP, and cookies.

DNS (Domain Name System) is responsible for translating domain names into IP addresses. When a user types a domain name in their browser, the DNS server resolves the domain name to the corresponding IP address, allowing the browser to establish a connection with the server hosting the website.

HTTP (Hypertext Transfer Protocol) is the protocol used for communication between a client (browser) and a server. It defines the format and rules for exchanging data over the internet. HTTP requests are made by the client to retrieve resources from the server, while HTTP responses are sent by the server to provide the requested resources.

HTTP headers play a crucial role in HTTP communication. They contain additional information that is not part of the response itself. The headers consist of key-value pairs, allowing the client and server to exchange relevant information. Two commonly used headers are the "Host" and "User-Agent" headers.

The "Host" header specifies the name of the server the client wants to communicate with. Most servers require this header for a successful response. Without it, the server may not respond at all.

The "User-Agent" header identifies the client's browser and operating system. This information helps the server determine the client's identity and can be useful for various purposes, such as detecting search engine crawlers or blocking malicious requests.

Another useful header is the "Referer" (misspelled as "Refer" in HTTP) header. It indicates the webpage from which the current request originated. This information is valuable for tracking user behavior and analyzing referral traffic.

Cookies are small pieces of data sent by the server and stored on the client's machine. They are used to maintain session state and store user-specific information. When a server sends a "Set-Cookie" header with a value, the client saves it and includes it in future requests to the same server. This allows the server to identify the client and provide personalized services.

Cookies can be used for various purposes, such as remembering user preferences, tracking user sessions, and implementing authentication mechanisms. However, it's important to note that cookies can also pose security risks if not handled properly.

Understanding web protocols like DNS, HTTP, and cookies is crucial for web application security. These protocols enable efficient communication between clients and servers, facilitate personalized user experiences, and ensure the integrity of user data.

#### Web Protocols - DNS, HTTP, Cookies, Sessions

Web protocols play a crucial role in the security of web applications. In this didactic material, we will discuss three important web protocols: DNS, HTTP, and cookies. These protocols are essential for the proper functioning and security of web applications.

DNS (Domain Name System) is responsible for translating domain names into IP addresses. It acts as a phonebook of the internet, allowing users to access websites using human-readable domain names. When a user enters a domain name in their browser, DNS is used to resolve the domain name to the corresponding IP address. This process is essential for establishing a connection between the user's device and the web server hosting the website.

HTTP (Hypertext Transfer Protocol) is the foundation of data communication on the web. It defines how messages are formatted and transmitted between clients (browsers) and servers. HTTP operates on a request-response model, where clients send requests to servers, and servers respond with the requested data. HTTP requests contain important information, such as headers and cookies, which are crucial for web application security.

Cookies are small pieces of data stored on the user's device by websites. They are used to track user sessions, store user preferences, and enhance user experience. When a user logs into a web application, the server generates a session token and sends it to the user's browser as a cookie. This token is then included in subsequent requests to identify the user and maintain their session. Cookies can also be used for other purposes, such as tracking user behavior and personalizing content.

HTTP headers play a significant role in web application security. They provide additional information to the server and the client, enabling them to communicate effectively and securely. For example, the "Accept-Language" header allows the client to specify the preferred language for the response. This header can be used to customize the content based on the user's language preference.

Response headers, sent by the server, provide valuable information to the client. Headers like "Last-Modified" and "Expires" help in caching and managing the freshness of data. Caching can improve performance by storing a copy of the response on the client's device, reducing the need to fetch the same data repeatedly. However, caching can also lead to issues, such as serving outdated content to different users. The "Vary" header helps mitigate this problem by instructing the client to consider specific headers when deciding whether to use a cached response.

The "Set-Cookie" header allows the server to set a cookie on the client's device. This header is crucial for managing user sessions and storing user-specific data. By setting cookies, web applications can provide personalized experiences and maintain user authentication across multiple requests.

The "Location" header is used to redirect clients to a different URL. When a server sends a "300 Multiple Choices" response, it includes the "Location" header with the URL where the client should redirect. This mechanism is commonly used for handling URL changes and indicating new locations for resources.

Understanding and properly implementing these web protocols is essential for web application security. DNS ensures the correct resolution of domain names, HTTP enables effective communication between clients and servers, and cookies and headers play vital roles in session management and data exchange.

#### Web Protocols - DNS, HTTP, Cookies, Sessions

Web protocols are essential for the functioning of web applications and ensuring secure communication between clients and servers. In this section, we will discuss three fundamental web protocols: DNS, HTTP, and TLS.

DNS (Domain Name System) is used to translate domain names (like example.com) into IP addresses. It allows clients to locate the correct server to establish a connection. DNS is the first step in the web protocol stack.

Once the IP address is obtained through DNS, the client establishes a connection using TCP (Transmission Control Protocol). TCP ensures reliable and ordered delivery of messages between the client and the server.

TLS (Transport Layer Security) provides encryption for secure communication between the client and the server. While optional, TLS is highly recommended to protect sensitive data transmitted over the network.

HTTP (Hypertext Transfer Protocol) is the protocol used for transferring data over the established TCP connection. It is commonly used to transfer website content such as HTML, CSS, and JavaScript. However, it can be used to transfer any type of data. HTTP is the final step in the web protocol stack.

To demonstrate the process of making an HTTP client from scratch, we can use a package like Node.js. By creating a TCP socket and connecting to the server's IP address (or domain name), we can send an HTTP request and receive a response.

The request consists of a string that includes the HTTP method (e.g., GET), the requested path (e.g., /), and the host header. The host header is crucial for the server to understand the client's request.

Once the request is sent, the server responds with data. In our example, we can simply print the response to the standard output. The response is received through the socket, which acts as a stream for data from the server.

By understanding the web protocol stack and the process of making an HTTP client, we gain insight into the underlying mechanisms of web applications. This knowledge is essential for web developers and cybersecurity professionals to ensure secure and efficient communication between clients and servers.

#### Web Protocols - DNS, HTTP, Cookies, Sessions

Web protocols play a crucial role in the security of web applications. In this section, we will discuss three important web protocols: DNS, HTTP, and cookies. Understanding these protocols is essential for ensuring the security of web applications.

DNS (Domain Name System) is responsible for translating domain names into IP addresses. When you enter a domain name in your web browser, DNS resolves the domain name to the corresponding IP address. This is necessary because computers communicate using IP addresses, not domain names. DNS allows us to use user-friendly domain names instead of remembering complex IP addresses.

HTTP (Hypertext Transfer Protocol) is the foundation of data communication on the web. It defines how messages are formatted and transmitted between web servers and clients. HTTP is a stateless protocol, meaning that each request is independent of previous requests. HTTP requests consist of a request line, headers, and an optional body. The response from the server includes a response line, headers, and a body. HTTP provides the basic structure for web communication.

Cookies are small pieces of data stored on the client's computer by the web server. They are used to maintain state and track user activity. When a user visits a website, the server can send a cookie to the client, which is then stored and sent back with subsequent requests. Cookies can contain information such as session IDs, user preferences, and shopping cart contents. However, it is important to note that cookies can also be used for malicious purposes, such as tracking user behavior or stealing sensitive information.

Sessions are a way to maintain state between multiple requests from the same client. When a user logs into a website, a session is created on the server. The server generates a unique session ID, which is then stored as a cookie on the client's computer. The session ID is used to identify the user in subsequent requests. Sessions allow websites to provide personalized experiences and secure access to restricted areas.

Understanding these web protocols is crucial for web application security. By understanding how DNS works, we can ensure that our web applications are connecting to the correct servers. HTTP provides the foundation for secure communication between clients and servers. Cookies and sessions help maintain state and provide

personalized experiences for users.

DNS, HTTP, cookies, and sessions are fundamental web protocols that play a crucial role in web application security. By understanding these protocols, developers can build secure and reliable web applications.

When browsing the internet, web browsers make multiple requests to retrieve resources from websites. This process involves several steps to ensure the proper functioning and rendering of web pages. One of the fundamental protocols used in this process is the Domain Name System (DNS).

The first step when a user types a URL and presses Enter is the DNS lookup. This involves translating the domain name into an IP address. The browser then opens a socket to the IP address on port 80, the default port for HTTP communication. It sends a request to the server and awaits a response.

Upon receiving the response, the browser does not immediately display the entire web page. Instead, it parses the HTML content to determine its structure and constructs the Document Object Model (DOM). The DOM represents the nodes and elements that make up the web page.

Once the initial DOM is constructed, the browser begins rendering the page. However, the initial DOM may be missing some resources, such as images or CSS files. To fetch these resources, the browser sends additional requests to the server. This process continues until all the required resources are obtained.

During the rendering process, the browser waits for responses from the server for each request it made. As the responses come back, the browser incorporates the resources into the page, ensuring proper rendering.

The process of retrieving and rendering web pages involves DNS lookup, sending requests for resources, parsing HTML to construct the DOM, rendering the page, and fetching additional resources as needed. This process ensures that web pages are displayed correctly and efficiently.

While this overview focuses on the web protocols DNS and HTTP, it is important to note the role of cookies in web applications security. Cookies are small pieces of data stored on the client-side by the server. They are used for various purposes, such as maintaining user sessions and storing user preferences.

When a server wants to set a cookie, it includes the "Set-Cookie" header in the response. This header contains a key-value pair that represents the cookie's data. The browser then stores this cookie and includes it in subsequent requests to the same server. This allows the server to recognize and remember the client's information.

Cookies can be used for authentication, allowing users to log in to websites and maintain their session. However, cookies can also be vulnerable to attacks. It is essential to implement proper security measures to protect against cookie-related threats.

Understanding web protocols such as DNS and HTTP, as well as the role of cookies in web applications, is crucial for ensuring secure and efficient web browsing. By comprehending the processes involved in retrieving and rendering web pages, users can better understand the underlying mechanisms of the internet.

To ensure the security of web applications, it is crucial to understand the fundamentals of web protocols such as DNS, HTTP, cookies, and sessions. In this didactic material, we will explore these concepts in detail.

DNS (Domain Name System) is a protocol used to translate human-readable domain names into IP addresses. It acts as a directory for the internet, allowing users to access websites using domain names instead of IP addresses. When a user enters a domain name in their web browser, the browser sends a DNS query to a DNS server, which responds with the corresponding IP address. This IP address is then used to establish a connection with the web server hosting the website.

HTTP (Hypertext Transfer Protocol) is the foundation of data communication on the web. It is a request-response protocol, where clients (web browsers) send HTTP requests to servers, and servers respond with HTTP responses. HTTP requests consist of a method (GET, POST, PUT, DELETE, etc.), a URL, headers, and an optional body. HTTP responses contain a status code, headers, and a response body. This protocol enables the transfer of various types of data, including HTML, images, videos, and more.

Cookies are small pieces of data stored on the client-side by web browsers. They are used to store information about the user or their interactions with a website. When a user visits a website, the server can set cookies in the HTTP response headers. These cookies are then sent back to the server in subsequent requests, allowing the server to identify and personalize the user's experience. However, cookies can also pose security risks if not properly implemented or secured.

Sessions are a way to maintain stateful interactions between a web server and a client. When a user logs in to a website, a session is created on the server-side, and a unique session identifier (usually stored in a cookie) is sent to the client. This identifier is used to associate subsequent requests from the client with the corresponding session on the server. Sessions are commonly used to store user authentication information and other session-specific data.

To illustrate the implementation of these concepts, let's consider an example of a bank login page. The login form on the webpage consists of two input fields: one for the username and another for the password. The form is designed to submit its data to a login endpoint, which is a URL responsible for authenticating the user.

On the server-side, we can use a framework like Express (built on top of Node.js) to handle HTTP requests. By instantiating an instance of Express and specifying the port to listen on, we can create an HTTP server. We can then define routes using Express, which are responsible for handling specific URLs and HTTP methods.

In our example, we would define a route for the login endpoint using the POST method. When a user submits the login form, a POST request is sent to this endpoint. The server, upon receiving the request, can access the request object (req) to retrieve information such as the URL, headers, and request body. In this case, we would extract the username and password from the request body.

To provide the user with the login page, we can use the file system module (FS) in Node.js to read the index.html file and send it as a response. This can be achieved using the createReadStream function, which allows us to stream the file's content to the response object, ensuring efficient data transmission.

By understanding the fundamentals of web protocols like DNS and HTTP, and concepts like cookies and sessions, developers can build secure web applications that protect user data and ensure a smooth user experience.

In web application security, understanding web protocols is crucial. Three important web protocols are DNS, HTTP, and cookies. DNS (Domain Name System) is responsible for translating domain names into IP addresses. HTTP (Hypertext Transfer Protocol) is the protocol used for communication between web servers and clients. Cookies are small pieces of data stored on the client's browser that allow the server to remember information about the user.

To implement a login endpoint in a web application, we need to handle the form data submitted by the user. One way to make it easier to read and parse the form data is by using a module called body-parser, which is part of Express. By using body-parser, we can access the submitted data as an object with key-value pairs representing the form fields.

To demonstrate this, we can create a user database using an object called "users". Each user object will have a username and password. When a user tries to log in, we can retrieve the username and password from the request body and compare them with the values in the user database. If the password matches, we can set a cookie to indicate that the user is logged in.

To set the cookie, we can use the cookie function provided by Express. We need to provide a cookie name and value. In this case, we can use "username" as the cookie name and set it to the logged-in user's username. This cookie will be sent in the response header with the "Set-Cookie" field.

Additionally, we can respond to the user with a success or failure message. If the login is successful, we can inform the user with a "success" message. Otherwise, we can inform them with a "fail" message.

By implementing these steps, we can handle user login in a web application and use cookies to maintain the user's session.

## Web Protocols: DNS, HTTP, Cookies, Sessions

Web protocols play a crucial role in the security of web applications. In this material, we will discuss three important web protocols: DNS, HTTP, and cookies. We will also touch upon the concept of sessions.

DNS (Domain Name System) is responsible for translating domain names into IP addresses. It acts as a phonebook for the internet, allowing users to access websites by typing in easy-to-remember domain names instead of complex IP addresses. However, DNS can also be exploited by attackers to redirect users to malicious websites. It is important to ensure DNS security to protect users from such attacks.

HTTP (Hypertext Transfer Protocol) is the foundation of data communication on the web. It enables the transfer of various types of data between clients (web browsers) and servers. However, HTTP is an unencrypted protocol, making it vulnerable to eavesdropping and tampering. To address this issue, HTTPS (HTTP Secure) was introduced, which encrypts the data exchanged between clients and servers using SSL/TLS protocols. It is essential to use HTTPS to protect sensitive information transmitted over the web.

Cookies are small pieces of data stored on a user's computer by websites they visit. They are used to remember user preferences, track user activity, and maintain session information. However, cookies can be manipulated by attackers to impersonate users or gain unauthorized access to their accounts. Therefore, it is crucial to implement proper security measures when handling cookies.

Sessions are a way to maintain stateful communication between clients and servers. They allow servers to keep track of user information and provide personalized experiences. However, session hijacking is a common attack where an attacker steals a user's session ID and impersonates them. To prevent session hijacking, session IDs should be securely generated, transmitted over HTTPS, and regularly rotated.

Understanding and implementing secure web protocols is essential for protecting web applications from various attacks. DNS security, HTTPS adoption, proper handling of cookies, and secure session management are key factors in ensuring the security of web applications.



**EITC/IS/WASF WEB APPLICATIONS SECURITY FUNDAMENTALS DIDACTIC MATERIALS****LESSON: SESSION ATTACKS****TOPIC: COOKIE AND SESSION ATTACKS**

A session is a way for a server to keep track of user data across multiple requests. In order to implement sessions, servers use cookies. Cookies are set by the server and included in every request made by the browser. This allows the server to correlate requests and maintain state.

There are several use cases for sessions. One example is the login feature on websites. When a user logs in, the server sets a cookie to remember their authentication status. This allows the user to navigate the site while staying logged in. Another use case is a shopping cart. Even if a user is not logged in, the server can still maintain a session for the items in the cart. Lastly, cookies can be used for tracking purposes. Ad servers can set a unique cookie for each user, allowing them to track the websites the user visits.

It's important to note that cookies are only sent to the same site that set them. This prevents cookies from being sent to every site the user visits. However, there are some complexities and rules surrounding this, which will be discussed further.

In terms of the technical implementation, the server sets a cookie with a key-value pair, such as "username=user123". This cookie is then sent back by the browser with every subsequent request. The server can use this cookie to identify the user and maintain the session.

It's worth mentioning that the example given in the material is insecure and should not be used in practice. Sending sensitive information like usernames and passwords in plain text is a security risk. Proper security measures should be implemented to protect user data.

Sessions and cookies play a crucial role in web application security. They allow servers to maintain state and provide personalized experiences for users. However, it's important to implement them securely to ensure the protection of user data.

**Web Applications Security Fundamentals - Session attacks - Cookie and session attacks**

In web applications, session attacks are a common type of security vulnerability that attackers exploit to gain unauthorized access or manipulate user sessions. One such attack is the cookie and session attack.

Session attacks target the session management system of a web application. Session management is crucial for maintaining user authentication and authorization during a user's interaction with the application. It allows the server to recognize and differentiate between different users.

One method of session management is through the use of cookies. Cookies are small pieces of data that the server sends to the client's browser, which then stores them. The client's browser includes the cookies in subsequent requests to the server, allowing the server to identify and authenticate the user.

In a cookie and session attack, the attacker attempts to exploit vulnerabilities in the session management system by tampering with or stealing cookies. By gaining unauthorized access to a user's session, the attacker can impersonate the user, perform actions on their behalf, or gain access to sensitive information.

To understand why this attack is possible, we need to consider the concept of ambient authority. Ambient authority refers to the access control based on a global and persistent property of the requester. In the case of web applications, the cookie header attached to each request serves as the property that grants access.

Imagine an alternative scenario where every action on a website requires the user to provide their username and password with each request. This would be cumbersome and inefficient. Cookies provide a more seamless and user-friendly experience, as the server can recognize the user based on the attached cookie.

While cookies are the most common method of session management, there are three other less common methods: IP address-based authentication, built-in HTTP authentication, and token-based authentication.



IP address-based authentication involves allowing or denying access based on the source IP address of the HTTP request. This method is used by institutions like Stanford libraries to restrict access to specific networks.

Built-in HTTP authentication is an older and less user-friendly method that presents a username and password prompt at the top of the browser. However, cookies have become the preferred method due to their flexibility and better user experience.

Token-based authentication is a more modern approach that involves generating and validating tokens instead of using cookies. However, it is not widely used.

It is important to note that while IP address-based authentication may seem secure, it is still susceptible to spoofing attacks. Changing one's IP address or using a VPN can bypass this type of authentication.

In the previous demonstration, we saw an example of a simple bank website. The website used cookies for session management. If a user had a valid cookie, they were considered authenticated. Otherwise, they were redirected to a login form.

The login form accepted a username and password and compared them to the values stored in a password database. If the credentials matched, a cookie with the username was issued, and the user was redirected back to the home page.

Session attacks, such as cookie and session attacks, highlight the importance of secure session management in web applications. Developers must implement robust security measures to protect against unauthorized access and ensure the integrity of user sessions.

#### Web Applications Security Fundamentals - Session Attacks: Cookie and Session Attacks

In web application security, session attacks are a common concern. One type of session attack is a cookie attack, where an attacker manipulates cookies to gain unauthorized access to a user's session. In this didactic material, we will focus on the fundamentals of cookie and session attacks.

A cookie is a small piece of data that is stored on a user's computer by a web browser. It is used to store information about the user's session, such as login credentials or preferences. Cookies are sent by the browser to the server with each request, allowing the server to identify the user and maintain their session.

In the given material, we see an example of a simple web application that displays a user's balance. The application uses cookies to store the user's username. However, the implementation is insecure, as anyone can manipulate the cookie to impersonate another user.

To demonstrate this vulnerability, the material shows how an attacker can change the username in the cookie to gain access to another user's session. By simply modifying the cookie value to "Bob," the attacker can view Bob's balance.

To mitigate this vulnerability, we need to implement proper session management techniques. One way to do this is by using session tokens instead of storing sensitive information directly in cookies. A session token is a unique identifier generated by the server and associated with the user's session. It is stored in the cookie, but it does not contain any sensitive information.

When the server receives a request with a session token, it can use this token to retrieve the user's session data from a secure storage location, such as a database. By separating the sensitive information from the cookie, we can prevent attackers from tampering with the session data.

Additionally, it is essential to properly handle session expiration and logout functionality. In the material, the logout route is added to clear the cookie and redirect the user to the home page. This ensures that the user's session is effectively terminated.

It is worth noting that clearing a cookie is not a built-in functionality. Instead, we set an expiration date in the past to instruct the browser to delete the cookie. This method is not ideal, but it is a common practice due to the limitations of cookies.

To summarize, session attacks, such as cookie attacks, pose a significant risk to web application security. By implementing proper session management techniques, such as using session tokens and handling logout functionality correctly, we can mitigate these vulnerabilities and protect user sessions.

In web applications, session attacks can pose a serious threat to the security of user data. One type of session attack is known as a cookie and session attack. In this type of attack, an attacker tries to manipulate or modify the values stored in cookies, which are small pieces of data that are sent from a website and stored on the user's browser.

To mitigate the risk of cookie and session attacks, a technique called cookie signing can be used. Cookie signing involves using cryptographic primitives to ensure that the value stored in a cookie cannot be modified by an attacker. This technique relies on three algorithms: G (generator), S (signer), and V (verifier).

When a user logs in to a web application, the server generates a public key and a secret key using the G algorithm. The server then validates the user's credentials and if they are correct, it generates a tag using the S algorithm. The value being signed is typically a unique identifier for the user, such as their username. The server then sends both the value and the tag back to the user's browser as cookies.

When the user visits other pages on the website, the browser sends both the value and the tag back to the server. The server can then use the V algorithm to verify if the tag matches the value that was originally signed. If the verification is successful, it means that the value has not been modified by an attacker.

There are two important properties of a signature scheme that are desired: correctness and security. Correctness means that the verification function should always return true for a valid tag and input. Security means that it should be computationally hard for an attacker to create a valid tag for an arbitrary input.

To ensure the security of the signature scheme, it is important not to reuse keys for different purposes. Additionally, it is recommended to use different algorithms for different purposes to avoid potential vulnerabilities.

Cookie signing is a technique used to protect against cookie and session attacks in web applications. It involves using cryptographic primitives to sign the value stored in a cookie, ensuring that it cannot be modified by an attacker. By verifying the signature on subsequent requests, the server can ensure the integrity of the data stored in the cookies.

In web applications, session attacks can pose a serious threat to the security and integrity of user data. One type of session attack is a cookie and session attack, where an attacker tries to manipulate or tamper with the session cookies used by the application.

To prevent such attacks, it is important to ensure the authenticity and integrity of the session cookies. One way to achieve this is by using encryption. By encrypting the session cookies, we can make sure that the data stored in them cannot be easily understood or modified by an attacker.

In the context of web applications, the client (browser) usually has access to the session cookies. This allows the browser to display personalized information to the user, such as their username. However, it is crucial to ensure that the client cannot tamper with the session cookies or modify the data stored in them.

One approach to achieve this is by using a signing system. When the server sends a cookie to the client, it includes a signature that is generated using a secret key. This signature ensures that the cookie has not been tampered with during transmission. When the client sends the cookie back to the server, the server can verify the signature to ensure its authenticity.

To implement this approach, the server needs to generate a secret key and securely store it. This secret key is used to generate the signature for each cookie. The server also needs to include the "signed" option when setting the cookie, which tells the framework to handle the signing process automatically.

When the server receives a cookie from the client, it can use the "signed cookies" property instead of the regular "cookies" property to access the cookie's data. This property only includes cookies that have a valid

signature, ensuring their integrity. By checking if the username is present in the "signed cookies" object, the server can determine if the cookie has been tampered with.

By implementing these measures, web applications can protect against cookie and session attacks, ensuring the security and integrity of user data.

Web applications security involves protecting web applications and their users from various types of attacks. One common type of attack is known as session attacks, which can include cookie and session attacks. In this didactic material, we will explore the fundamentals of session attacks and how they can be mitigated.

Session attacks aim to exploit vulnerabilities in the session management process of web applications. By gaining unauthorized access to a user's session, attackers can impersonate the user and perform malicious actions. Cookie and session attacks are two specific types of session attacks that we will focus on.

In a web application, sessions are typically managed using cookies. Cookies are small pieces of data stored on the user's browser that contain information about the session. They are used to maintain the user's state and enable seamless interaction with the application.

During a cookie attack, an attacker tries to manipulate the values stored in the user's cookies to gain unauthorized access. This can involve modifying the cookie values to impersonate another user or tampering with the integrity of the data stored in the cookies.

Session attacks can also target the session ID, which is a unique identifier assigned to each user's session. By stealing or guessing a valid session ID, an attacker can bypass authentication mechanisms and gain unauthorized access.

To protect against cookie and session attacks, various security measures can be implemented. One common approach is to use cryptographic signatures to ensure the integrity and authenticity of the cookies. By signing the cookie values with a secret key, the server can verify the integrity of the data and detect any tampering attempts.

However, this approach has its limitations. If an attacker manages to steal the user's cookie, they can use it to impersonate the user indefinitely, even after the user has logged out or changed their password. This is because the signature remains valid as long as the secret key is not changed.

To address this issue, an alternative approach can be used. Instead of relying on cryptographic signatures, a session ID can be generated for each user and stored in a database. When a user logs in, they are assigned a unique session ID, which is then used to identify their session. This session ID is not tied to any specific user data and can be easily invalidated by removing it from the database upon logout.

By using this approach, the server can effectively destroy the session by deleting the corresponding session ID from the database. This ensures that even if an attacker possesses a valid session ID, it becomes useless once it has been invalidated.

Session attacks, including cookie and session attacks, pose a significant threat to web application security. By understanding the vulnerabilities associated with session management and implementing appropriate security measures, such as cryptographic signatures or session ID invalidation, web applications can better protect their users' sessions and prevent unauthorized access.

Session attacks, specifically cookie and session attacks, are important topics in web applications security. In this material, we will discuss the fundamentals of session attacks and how they can be mitigated.

Session attacks occur when an attacker gains unauthorized access to a user's session, allowing them to impersonate the user and perform malicious actions. One common type of session attack is cookie and session attacks.

In a typical web application, a user's session is identified by a unique session ID, often stored in a cookie. The session ID is used to associate the user's actions with their session data on the server. However, if an attacker can guess or obtain a valid session ID, they can impersonate the user and perform actions on their behalf.

To mitigate session attacks, it is important to ensure that session IDs are not easily guessable or obtainable. One way to achieve this is by using a large space of possible session IDs, making it difficult for attackers to find a valid session ID through guessing. By increasing the size of the space, the probability of guessing a valid session ID becomes extremely low.

Another important aspect of session security is the ability to invalidate or destroy session data when a user logs out. By removing the corresponding entry from the server's database, the session ID becomes invalid, preventing any further unauthorized access. This is different from older systems where session data was signed, making it difficult to invalidate a session once it was stolen.

However, it is worth noting that if an attacker has persistent access to a user's computer, such as through malware, they can still maintain access to the session even after the user logs out. This is a challenging problem to address and falls beyond the scope of this discussion.

One advantage of the current design is the ability to remotely manage sessions. For example, popular platforms like Google and Facebook allow users to view and delete active sessions from their accounts. This feature is made possible by the design we are discussing, where session data is stored in a database and can be easily accessed and modified.

To implement session security, we can use a sessions object, which acts as a mapping between session IDs and user names. When a user logs in, instead of assigning a user name cookie, we generate a unique session ID and assign it as a session ID cookie. This session ID is stored in the sessions object, associating it with the user's name.

When a user visits the website and provides their session ID, we can use it to look up their corresponding user name in the sessions object. This allows us to authenticate the user and provide them with the appropriate access and privileges.

Session attacks, specifically cookie and session attacks, pose a significant threat to web applications. By implementing proper session security measures, such as using a large space of session IDs and invalidating sessions upon logout, we can greatly reduce the risk of unauthorized access and impersonation. Additionally, the ability to remotely manage sessions provides users with greater control over their account security.

A session attack is a type of cybersecurity attack that targets the session management mechanism in web applications. In this attack, an attacker tries to gain unauthorized access to a user's session by exploiting vulnerabilities in the session management process.

One common type of session attack is the cookie and session attack. In this attack, the attacker tries to steal or manipulate the session ID stored in the user's cookie to impersonate the user and gain unauthorized access to their account.

To understand how this attack works, let's consider an example. Suppose we have a web application that uses session IDs to authenticate users. When a user logs in, the server generates a unique session ID for that user and stores it in a cookie on the user's browser. The server also associates the session ID with the user's account in a database.

In a cookie and session attack, the attacker tries to steal the session ID from the user's cookie. They can do this by exploiting vulnerabilities in the web application, such as cross-site scripting (XSS) or cross-site request forgery (CSRF). Once they have the session ID, they can use it to impersonate the user and gain unauthorized access to their account.

To protect against cookie and session attacks, web developers should implement proper security measures. One common approach is to use secure session management techniques, such as:

1. Generating a strong and random session ID: The session ID should be difficult to guess or predict. It should be a long, random string of characters that is unique for each user.
2. Encrypting the session ID: The session ID should be encrypted before storing it in the user's cookie. This

prevents attackers from easily extracting the session ID from the cookie.

3. Validating the session ID: The server should validate the session ID received from the user's cookie. It should check if the session ID exists in the database and if it is associated with a valid user account.

4. Implementing session expiration: Sessions should have a limited lifespan and should expire after a certain period of inactivity. This prevents attackers from using stolen session IDs for an extended period of time.

5. Using secure cookies: Developers should use secure cookies that are only transmitted over HTTPS connections. This prevents attackers from intercepting the cookie and stealing the session ID.

By implementing these security measures, web developers can protect their web applications from cookie and session attacks, ensuring the confidentiality and integrity of user sessions.

In web applications security, session attacks are a common concern. One type of session attack is a cookie and session attack. In this attack, the attacker tries to manipulate or exploit the session cookies used by the web application to maintain user sessions.

To understand cookie and session attacks, let's first discuss how sessions work in web applications. When a user accesses a web application, a unique session ID is assigned to them. This session ID is typically stored in a cookie on the user's browser. The server uses this session ID to identify and track the user's session.

In the transcript, the speaker mentions the issue of multiple users in the same household appearing as the same person due to the shared IP address. This can be a problem because web applications may mistakenly associate the actions of different users with a single session.

To address this issue, the speaker suggests using Transport Layer Security (TLS), which is a protocol that ensures secure communication between the client (browser) and the server. By using TLS, the browser automatically handles the creation and management of session cookies, reducing the risk of session attacks.

The speaker also mentions the concept of signing cookies. Signing a cookie involves adding a digital signature to the cookie value to ensure its integrity and prevent tampering. This can be useful in scenarios where the server wants to provide a promotion or discount to a user, but wants to prevent the user from modifying the cookie to exploit the offer.

However, the speaker clarifies that signing cookies is not relevant to the current discussion on session ID selection. Instead, the speaker suggests using a better method to generate session IDs. They propose using the "random bytes" function from the crypto module in Node.js to generate a random string of bytes. This random string can then be converted into a string format suitable for use as a session ID.

The speaker emphasizes the importance of using a secure cryptographic key randomizer to ensure the uniqueness and unpredictability of session IDs. They mention that collisions (duplicate session IDs) are highly unlikely to occur, but it is possible to add additional checks for extra safety.

The speaker then demonstrates the implementation of the suggested changes, including importing the necessary crypto module and generating session IDs using the "random bytes" function. They also mention the use of base64 encoding to make the session IDs easier to read.

Finally, the speaker suggests testing the new session ID generation by logging in as different users and observing the differences in the session IDs. They confirm that the new method provides unguessable and distinct session IDs, which is crucial for preventing session attacks.

To mitigate cookie and session attacks, it is important to use secure methods for generating and managing session IDs. By implementing proper session ID selection techniques, web applications can enhance their security and protect user sessions from exploitation.

Web Applications Security Fundamentals - Session attacks - Cookie and session attacks

Web applications often use cookies to store and retrieve information about a user's session. However, the

implementation of cookies can be vulnerable to attacks if not properly secured. In this lesson, we will explore the fundamentals of cookie and session attacks in web applications.

Cookies were first implemented in 1994 by a developer working on the Netscape browser. At that time, there was no formal specification for how cookies should work, and it was more like a brainstormed idea than a well-defined standard. Over time, attempts were made to establish a common understanding and behavior for cookies across different browsers, but these attempts were often overly ambitious and did not accurately reflect the existing behavior of popular browsers.

In 2011, a group of people finally wrote an RFC (Request for Comments) that browsers started to follow, providing a more standardized description of how cookies should work. However, due to the ad hoc development of the cookie design, there are still inconsistencies and security concerns associated with their usage.

Cookies have various attributes that can be added to them. One such attribute is the "expires" attribute, which allows the developer to specify an expiration date for the cookie. If no expiration date is set, the cookie will last as long as the user's browser session is active. However, in practice, some browsers have added a feature called session restoration, which can cause cookies without an expiration date to persist even after the browser is closed and reopened.

To ensure that a cookie is cleared when the user closes their browser, it is recommended to set an explicit expiration date. Relying on the browser's session restoration feature may not guarantee the desired behavior.

It is important to note that cookies have a different security model compared to other aspects of web applications, such as the same origin policy. This difference in security models can lead to unique vulnerabilities and potential security risks.

Understanding the fundamentals of cookie and session attacks is crucial for web application security. Properly implementing and securing cookies, including setting appropriate expiration dates, can help mitigate the risk of session-based attacks.

#### Web Applications Security Fundamentals - Session attacks - Cookie and session attacks

In web application security, session attacks are a common concern. One type of session attack involves manipulating cookies. Cookies are small pieces of data that are stored on the user's browser and are used to maintain session information.

One attribute of cookies is called "path". The path attribute allows you to set a cookie that will only be sent to the server when the user is on a specific path of the website. For example, if a user is on a blog hosted at example.com/blog, you might want to set a cookie that is only sent when the user is on pages that start with "/blog". This helps to limit the scope of the cookie to specific parts of the website.

Another attribute is called "domain". The domain attribute allows you to scope the cookie to a broader domain than the one that returned the cookie. For example, if a user logs in to a website at stanford.edu and receives a session ID cookie, the website might want to allow subdomains of stanford.edu to also access the cookie. This can be achieved by setting the domain attribute to "stanford.edu". Without this attribute, the cookie would only be available to the specific subdomain that set it.

To set these attributes, you can simply append them to the cookie header. If you need to set multiple cookies, you will need to include multiple "Set-Cookie" headers.

In terms of expiration, it used to be common practice to set cookies to last for a very long time, such as until the year 2038. However, this can raise privacy concerns. It is now recommended to set the expiration date to a shorter period, such as a month or two in the future. This allows you to reset the cookie whenever the user revisits the website, without the need for long-term storage.

To delete a cookie, you can simply set the expiration date to a time in the past. This effectively removes the cookie from the user's browser.



Accessing cookies in JavaScript can be a bit tricky. The API for accessing cookies is not very intuitive. To set a cookie in JavaScript, you would use the syntax `document.cookie = name=value`. However, if you want to add a new cookie without overwriting the old one, you need to be careful. Adding a new cookie does not overwrite the old one, but rather appends it to the existing list of cookies. Therefore, when retrieving the value of `document.cookie`, you will get a semicolon-delimited string of all the cookies.

It is important to be aware of these aspects of cookie and session attacks to ensure the security of web applications.

Web applications are vulnerable to various attacks, including session hijacking, which occurs when an attacker steals a user's cookies and impersonates them. One way this can happen is if the web application is not using HTTPS, as the attacker can intercept the unencrypted traffic and obtain the cookies. They can then use these cookies to send requests to the server, fooling it into thinking that they are the legitimate user.

To illustrate this attack, imagine a client sending a request to the server with a session ID. In a secure scenario, the server would respond with the requested page containing private information. However, in the case of a session hijacking attack, an attacker intercepts the request, sends the same request to the server, and receives the private web page intended for the user. This effectively allows the attacker to be logged in as the user.

In 2010, an interesting incident occurred where a Firefox extension was created to put network cards into monitor mode. This mode allowed the extension to capture unencrypted traffic flowing through the network. By connecting to a public Wi-Fi network, the extension could intercept requests made by other users, extract their session IDs, and present them in a user-friendly interface. This was done to raise awareness about the importance of using HTTPS for secure communication.

To mitigate session hijacking attacks, two solutions are commonly employed. The first is to use HTTPS throughout the entire web application, ensuring that all communication is encrypted. This prevents attackers from intercepting and manipulating traffic. The second solution involves setting the "secure" flag when setting a cookie. This flag ensures that the cookie is only sent over an encrypted connection. If a user visits a non-HTTPS page, the cookie will not be sent, reducing the risk of session hijacking.

Web applications are vulnerable to session hijacking attacks, where an attacker steals a user's cookies and impersonates them. This can occur if the application does not use HTTPS, allowing the attacker to intercept unencrypted traffic. To mitigate this risk, it is crucial to use HTTPS throughout the application and set the "secure" flag for cookies.

#### Web Applications Security Fundamentals - Session attacks - Cookie and session attacks

In web application security, session attacks are a common concern. One type of session attack involves stealing cookies, which can be used to gain unauthorized access to a user's session. This can happen when an attacker is able to run their code within a website and extract the cookie information.

To steal the cookie, the attacker can create a HTTP GET request to a specific URL. One way to do this is by embedding the request in an image source. When the browser loads the page, it will send a request to the attacker's server, allowing them to obtain the cookie. This is a serious security risk, as the attacker can then use the stolen cookie to impersonate the user and gain access to their session.

To defend against this type of attack, web developers can set the "HTTP Only" flag for cookies. This flag prevents the cookie from being accessed by JavaScript code, making it more difficult for attackers to steal the cookie. By adding the "HTTP Only" flag, the cookie will be included in the headers of requests, but JavaScript code will not be able to access it.

Another aspect to consider is the "path" attribute of cookies. The "path" attribute allows developers to specify that a cookie should only be sent to a particular URL or path. However, it is important to note that the "path" attribute does not protect the cookie from unauthorized reading on related URLs. This means that if an attacker can run their code on a related URL, they may still be able to access the cookie.

Session attacks involving cookies are a significant security risk in web applications. It is crucial for developers to implement measures such as using the "HTTP Only" flag and carefully considering the "path" attribute to protect



against these attacks.

Web applications are vulnerable to various types of attacks, including session attacks. In this session, we will specifically focus on cookie and session attacks.

When a user logs into a web application, a session is created to track their interactions with the application. This session is often identified by a unique session ID, which is stored in a cookie on the user's browser. The server uses this session ID to authenticate and authorize the user for subsequent requests.

However, attackers can exploit vulnerabilities in the session management process to gain unauthorized access to a user's session. One such vulnerability is known as a cookie and session attack.

In a cookie and session attack, the attacker tries to manipulate the session cookie to gain access to another user's session. This can be done by stealing the session ID or by tricking the browser into sending the session cookie to a different path or domain.

To illustrate this attack, let's consider an example. Suppose we have a web application hosted on a URL, and the session cookie is set to be sent only to a specific path, such as "/cs160". If a user tries to access a different path, such as "/cs253", the browser will not send the session cookie.

However, attackers can bypass this restriction by creating an iframe, which is a small window to another page, and adding it to their own page. By setting the source of the iframe to the URL of the target web application, the attacker can effectively load the target application within their own page.

Once the iframe is added, the attacker can access the content document of the iframe, which represents the document of the target application. By accessing the document's cookie property, the attacker can read the session cookie of the target application.

This attack works due to a security mechanism called the same origin policy. The browser allows access to the content document of an iframe only if the domain, port, and protocol of the iframe's source URL match those of the parent page. In our example, since both the parent page and the target application are hosted on the same domain and use the same protocol, the same origin policy allows the attacker to access the target application's cookies.

To mitigate cookie and session attacks, web developers should implement proper session management techniques. This includes securely generating and storing session IDs, using secure cookies with appropriate attributes such as the "Secure" and "HttpOnly" flags, and validating session IDs on each request to ensure they are associated with a valid session.

Additionally, developers should be aware of the same origin policy and ensure that sensitive operations or data are not accessible from iframes or other cross-origin contexts.

Cookie and session attacks pose a significant threat to the security of web applications. By exploiting vulnerabilities in session management, attackers can gain unauthorized access to user sessions. It is crucial for web developers to implement robust session management techniques and adhere to security best practices to protect against these attacks.

Web applications security is a crucial aspect of cybersecurity. In this material, we will discuss session attacks, specifically focusing on cookie and session attacks.

Session attacks exploit vulnerabilities in web applications to gain unauthorized access to user sessions. One common type of session attack is the cookie attack. Cookies are small pieces of data stored on the user's browser to track their session information. However, if an attacker can manipulate or steal a user's cookie, they can impersonate the user and gain unauthorized access.

During the material, it was mentioned that manipulating the page is not always possible due to certain security measures. However, this does not guarantee protection against cookie attacks. Attackers can still run their own JavaScript directly on the page, bypassing some security measures.

Furthermore, the speaker highlighted the potential security issues related to subdomains. Subdomains are treated as different host names, which means they have separate cookies and sessions. This can be exploited by attackers to target specific subdomains and gain unauthorized access.

To mitigate session attacks, web applications should implement strong security measures. These may include encryption of cookies, using secure HTTPS connections, and implementing proper session management techniques. It is also recommended to regularly update and patch web applications to address any known vulnerabilities.

Session attacks, particularly cookie attacks, pose a significant threat to web application security. Web developers and security professionals must be aware of these vulnerabilities and take appropriate measures to protect user sessions and sensitive information.

**EITC/IS/WASF WEB APPLICATIONS SECURITY FUNDAMENTALS DIDACTIC MATERIALS****LESSON: SAME ORIGIN POLICY****TOPIC: CROSS-SITE REQUEST FORGERY**

The same origin policy is an important aspect of web application security. It is designed to prevent malicious websites from accessing sensitive data or performing unauthorized actions on behalf of a user. In this context, we will explore the same origin policy in relation to cookies and cross-site request forgery (CSRF) attacks.

Cookies are small pieces of data that websites store on a user's browser. They are commonly used to maintain user sessions and store user preferences. However, if not properly secured, cookies can be vulnerable to attacks. The same origin policy plays a crucial role in mitigating these risks.

The same origin policy states that a website can only access resources (such as cookies) from the same origin (domain, protocol, and port) as the website itself. This means that a website can only access cookies that were set by itself or by another website with the same origin. This policy prevents malicious websites from accessing cookies set by other websites, thereby protecting sensitive user information.

However, there are certain scenarios where the same origin policy can be bypassed. One such scenario is when a website includes a frame or iframe from another domain. In this case, the parent website can access the DOM (Document Object Model) of the framed website, and thus access its cookies. This can be exploited by attackers to perform CSRF attacks.

Cross-site request forgery (CSRF) is an attack where a malicious website tricks a user's browser into making a request to another website on which the user is authenticated. If the target website relies solely on cookies for authentication, the request will be treated as legitimate by the server, leading to unauthorized actions on behalf of the user.

To prevent CSRF attacks, web developers should implement additional security measures such as CSRF tokens. A CSRF token is a unique value that is generated by the server and included in each form or request. When the server receives a request, it checks if the CSRF token matches the one expected. If not, the request is considered unauthorized and rejected.

In addition to CSRF tokens, there are other best practices for securing cookies. For example, cookies should be set with the "secure" flag, which ensures that they are only sent over secure HTTPS connections. Cookies should also be set with the "HttpOnly" flag, which prevents them from being accessed by JavaScript code running in the browser. This helps protect against attacks such as cross-site scripting (XSS).

It is important to note that relying solely on the "path" attribute of cookies for security is not recommended. The same origin policy does not consider the path attribute when determining if a website can access a cookie. Therefore, it is best practice to explicitly set the path to the root of the website to ensure cookies are visible only where intended.

The same origin policy is a fundamental security principle in web applications. It restricts websites from accessing cookies and resources from other domains, protecting user data from unauthorized access. However, developers must be aware of potential bypasses, such as CSRF attacks, and implement additional security measures like CSRF tokens. Properly securing cookies with the "secure" and "HttpOnly" flags is also essential to mitigate risks.

**Web Applications Security Fundamentals - Same Origin Policy - Cross-Site Request Forgery**

In web application security, one important concept to understand is the Same Origin Policy. This policy is a security mechanism implemented by web browsers to prevent web pages from making requests to a different origin or domain. The purpose of this policy is to protect users from malicious websites that may try to access sensitive information or perform unauthorized actions on their behalf.

However, there is a vulnerability known as Cross-Site Request Forgery (CSRF) that can bypass the Same Origin Policy. CSRF attacks occur when an attacker tricks a user's browser into making a request to a target website, using the user's credentials and session information. This can lead to unauthorized actions being performed on

the user's behalf.

Let's consider an example to better understand how CSRF attacks work. Imagine a scenario where a user is logged into their online banking account. They visit a malicious website that embeds an image from a specific URL. Unbeknownst to the user, this URL contains a request to transfer \$1,000 from their account to the attacker's account.

Even though the request is sent as an image, the user's browser will still include any cookies associated with the target website, in this case, the online banking website. The server, upon receiving the request, sees the valid session ID and assumes that the user has initiated the transfer. Consequently, the server processes the request and transfers the funds to the attacker's account.

This attack is possible because the browser automatically attaches the necessary cookies to any requests made to websites that have associated cookies. The attacker takes advantage of this behavior to forge a request on behalf of the user, tricking the server into thinking it is a legitimate action.

To summarize, CSRF attacks exploit the trust between a user's browser and a target website by tricking the browser into making unauthorized requests. This can lead to actions being performed without the user's consent or knowledge, potentially resulting in financial loss or other security breaches.

It is important for web developers and security professionals to be aware of CSRF vulnerabilities and implement appropriate defenses. One common defense is to use the HTTP POST method instead of GET for sensitive operations, as POST requests are not automatically triggered by browsers. Additionally, implementing anti-CSRF tokens can help validate the authenticity of requests and prevent unauthorized actions.

By understanding the risks associated with CSRF attacks and implementing proper security measures, web applications can better protect users' data and prevent unauthorized actions.

Cross-Site Request Forgery (CSRF) is a type of attack that exploits the trust a website has in a user's browser. In this attack, an attacker tricks a user into unknowingly making a request to a website they are authenticated on, without the user's consent or knowledge. This can be done through various means, such as sending a targeted email to the user or luring them to a malicious website.

The Same Origin Policy (SOP) is a fundamental security concept in web applications that prevents malicious websites from accessing data from other websites. It ensures that a web page can only access resources (such as cookies, local storage, or DOM) from the same origin (protocol, domain, and port) as the page itself. This policy is enforced by web browsers and helps protect users from various types of attacks, including CSRF.

In a CSRF attack, the attacker crafts a malicious request that will be automatically sent by the user's browser when they visit a different website. This request can be designed to perform actions on the target website, such as creating a new admin user or executing arbitrary code on the server. If the user is logged in as an admin, the attacker can exploit this to gain unauthorized access and perform malicious activities.

To successfully execute a CSRF attack, the attacker needs to know the path of the request and assume that the user is logged in. If the user is not logged in, the attack will fail. However, this can be circumvented in scenarios where the attacker has obtained information about the user's membership in a particular website, such as through phishing attacks or other means of gathering information.

It is important to note that in a CSRF attack, the attacker does not need to read the response from the server to determine the success of the attack. The damage is done once the request is sent and received by the server. The response may not be relevant to the attacker's goals.

To demonstrate the concept of CSRF, let's consider an example. Suppose we have a bank website with a feature that allows users to transfer money to other users. We will add this functionality to the website and show how it can be exploited through a CSRF attack.

First, we modify the code of the bank website to include a form for transferring money. The form has two fields: the amount to transfer and the recipient's username. When the form is submitted, it sends a POST request to the "/transfer" endpoint.

Next, we implement the `"/transfer"` route on the server. When a user posts to this endpoint, the server checks if they have an active session. If not, the request is rejected. If the user is authenticated, the server processes the transfer based on the amount and recipient specified in the request.

An attacker can craft a malicious webpage that includes a hidden form, pre-filled with the necessary information to perform a transfer. When a user visits this webpage while authenticated on the bank website, their browser automatically sends the request to the `"/transfer"` endpoint, effectively transferring money to the attacker's desired account.

This example demonstrates how a CSRF attack can be used to exploit the trust a website has in a user's browser and perform unauthorized actions on behalf of the user. It highlights the importance of implementing countermeasures, such as CSRF tokens, to mitigate this type of attack.

CSRF is a serious security vulnerability that can allow attackers to perform unauthorized actions on a user's behalf. Understanding the Same Origin Policy and implementing proper countermeasures are crucial in ensuring the security of web applications.

The Same Origin Policy is a fundamental security concept in web applications that restricts how a document or script loaded from one origin can interact with a resource from another origin. An origin is defined by the combination of the protocol, domain, and port number. The Same Origin Policy ensures that resources from different origins cannot access or manipulate each other's data without explicit permission.

One common vulnerability that can occur due to a violation of the Same Origin Policy is Cross-Site Request Forgery (CSRF). CSRF attacks exploit the trust that a web application has in a user's browser. In a CSRF attack, an attacker tricks a victim into performing an unwanted action on a web application without their knowledge or consent.

In this example, we have a simple transfer system where a user can transfer money from their account to another user's account. The system uses a form to submit the transfer request to the server. However, there is a vulnerability that allows an attacker to forge a request and perform unauthorized transfers.

To demonstrate this vulnerability, the attacker sets up a separate server on port 9999 and creates a page that looks innocent, like a cat picture gallery. However, hidden in the page is a form that mimics the transfer form of the target web application. The attacker pre-fills the form with the desired transfer amount and target user.

When a user visits the attacker's page and interacts with the innocent-looking content, the hidden form is automatically submitted in the background without the user's knowledge. The form is submitted to the target web application's transfer endpoint, causing the server to process the forged transfer request.

To mitigate this vulnerability, web developers should implement measures such as CSRF tokens. A CSRF token is a unique value generated by the server and included in the form. When the form is submitted, the server checks if the CSRF token matches the expected value, ensuring that the request is legitimate.

By implementing CSRF tokens, web applications can protect against unauthorized actions performed by attackers using CSRF attacks. It is essential for developers to be aware of the Same Origin Policy and its implications to ensure the security of web applications.

The Same Origin Policy is a fundamental security concept in web applications that aims to protect users from malicious attacks. One specific threat that the Same Origin Policy addresses is Cross-Site Request Forgery (CSRF). CSRF occurs when an attacker tricks a user's browser into making an unintended request to a target website on which the user is authenticated.

To understand how CSRF works, let's consider an example. Suppose a user is logged into their online banking account and visits an attacker-controlled website. The attacker's website contains a hidden form that, when submitted, initiates a transfer of funds from the user's account to the attacker's account. The attacker entices the user to unknowingly submit the form by using various techniques, such as social engineering or misleading design.

The Same Origin Policy prevents this type of attack by enforcing that web browsers only allow requests to be made to the same origin as the currently loaded web page. An origin is defined by the combination of the protocol (e.g., HTTP, HTTPS), domain, and port number. In our example, the user's online banking website has a different origin than the attacker's website, so the user's browser will block the request to transfer funds due to the Same Origin Policy.

However, attackers have found ways to bypass the Same Origin Policy in certain scenarios. One technique involves the use of HTML frames or iframes. By embedding an invisible frame within their website, the attacker can load the target website's login page and submit the hidden form programmatically. This technique makes the attack less noticeable and increases the chances of success.

To mitigate this type of attack, web developers can implement additional security measures. One approach is to include a token, known as a CSRF token, in forms that perform sensitive actions, such as fund transfers or account modifications. The CSRF token is a unique value generated by the server and included in the form. When the form is submitted, the server verifies that the CSRF token matches the expected value, ensuring that the request originated from the same website.

Another approach is to use the SameSite attribute for cookies. The SameSite attribute allows developers to specify whether a cookie should be sent with requests originating from other websites. By setting the SameSite attribute to "Strict" or "Lax" for session cookies, developers can prevent the automatic attachment of cookies to cross-origin requests, thereby reducing the risk of CSRF attacks.

The Same Origin Policy is a crucial security mechanism in web applications that prevents Cross-Site Request Forgery attacks. Developers must be aware of the potential vulnerabilities and implement additional security measures, such as CSRF tokens and SameSite attributes for cookies, to enhance the protection of their web applications.

The Same Origin Policy is a fundamental security concept in web applications that aims to protect users from malicious attacks. It ensures that web pages from different origins (e.g., different domains) cannot access each other's data without explicit permission. One important aspect of the Same Origin Policy is the handling of cookies.

By default, web browsers will send cookies along with every request if they match the request. However, there are three settings that can be used to control this behavior. The default setting is "none," which means that cookies will always be sent along with the request if they match. Another setting is "lax," which applies to sub resource requests. These are requests for resources like images or forms that are embedded within a page. With the "lax" setting, cookies will not be attached to sub resource requests originating from another site, but they will be attached to top-level navigation requests.

For those who want to take an extra step in ensuring security, there is the "strict" mode. In this mode, cookies are never sent if the user was sent to the site from another site. This means that even innocent requests, like clicking a link to a bank's homepage from a blog post, will not have cookies attached if they originated from another site. While this setting is very aggressive, it provides an additional layer of protection.

It is important to note that the "lax" setting covers most cases and is actually a good compromise between security and usability. Despite its name, it is a better option than the default "none" setting. However, it can still break certain use cases. For example, if a site uses an iframe to load a comment box from a social media platform like Facebook, enabling the "lax" setting would prevent the cookies from being attached to the request. This would result in the user being logged out of the comment box and having to manually enter their credentials to use it.

Similarly, advertisers who rely on tracking user behavior may not want to enable the "lax" setting as it would prevent cookies from being attached to ad requests. This would hinder their ability to gather user information for targeted advertising.

It is worth mentioning that the landscape around the Same Origin Policy and its related settings is evolving. There is a proposal by Google to make the "same-site=lax" setting the default for all browsers. This change is driven by the increasing concerns about security threats and user privacy. The proposal suggests that cookies should be treated as "same-site=lax" by default, and those explicitly asserting "same-site=none" should also be



honored.

The Same Origin Policy plays a crucial role in web application security. Understanding the different settings that control cookie behavior is essential for developers and website administrators. While the default behavior is to send cookies with every request, the "lax" setting provides a good compromise between security and usability. The "strict" mode offers the highest level of protection but may break certain use cases. It is important to stay informed about the evolving landscape of web security to ensure the best practices are followed.

The Same Origin Policy is a fundamental security concept in web applications that aims to prevent malicious websites from accessing sensitive data from other websites. It states that a web browser should only allow scripts and resources from the same origin to interact with each other.

The Same Origin Policy works by comparing the origin of the current web page (the scheme, host, and port) with the origin of the requested resource. If the origins match, the browser allows the interaction. If the origins do not match, the browser blocks the interaction.

However, there is a vulnerability called Cross-Site Request Forgery (CSRF) that can bypass the Same Origin Policy. CSRF occurs when an attacker tricks a victim into performing an action on a website without their knowledge or consent. This can lead to unauthorized actions being executed on the victim's behalf.

To mitigate CSRF attacks, web developers can implement certain security measures. One approach is to use a technique called "synchronizer token pattern." In this pattern, a unique token is generated for each user session and included in the web form. When the form is submitted, the server checks if the token matches the one associated with the user session. If they match, the request is considered legitimate.

Another approach is to add additional security attributes to cookies. For example, the "Secure" attribute ensures that the cookie is only sent over a secure HTTPS connection, preventing passive eavesdropping. The "HttpOnly" attribute restricts access to the cookie from client-side scripts, reducing the risk of cookie theft through cross-site scripting (XSS) attacks. The "SameSite" attribute specifies whether the cookie should be sent in cross-site requests, helping to prevent CSRF attacks.

Web developers can set the expiration time for cookies to a reasonable duration, such as 30 days. By resetting the expiration time with each user visit, the cookie remains valid for the specified duration. This helps balance convenience for users with the need to minimize the risk of unauthorized access.

When implementing these security measures, it is important to ensure that all attributes of the cookie are consistent, including the name, expiration, and security attributes. In some cases, the use of separate cookies with the same name and attributes may be required to properly clear the cookie.

Understanding and implementing the Same Origin Policy and taking steps to mitigate CSRF attacks are essential for ensuring the security of web applications. By following best practices and utilizing security attributes in cookies, developers can enhance the protection of user data and prevent unauthorized actions.

The Same Origin Policy is a fundamental security concept in web applications that aims to protect users from malicious attacks. It states that two web pages from different sources should not have access to each other's data or functionality, unless explicitly allowed.

One of the main aspects of the Same Origin Policy is the restriction on cookies. Cookies are small pieces of data that websites store on a user's browser. They are used to maintain user sessions and store user preferences. However, cookies are subject to the Same Origin Policy, meaning that they can only be accessed by web pages from the same domain.

Another important aspect of the Same Origin Policy is the use of security headers by web browsers. These headers provide information about the origin of a request and how it was triggered. This allows browsers to make more intelligent decisions regarding the source of a request and whether it should be allowed.

The Same Origin Policy also addresses various scenarios and determines what should be allowed and what should be denied. For example, it allows websites to link to each other, as this is a fundamental aspect of the web. However, embedding one website within another can be risky, as it can lead to potential security



vulnerabilities. Therefore, the policy restricts the ability to modify the contents of an embedded website.

Submitting forms is generally allowed under the Same Origin Policy, although it can be a potential security risk. Embedding images from other sites is also allowed, as it is a common practice to avoid duplicating resources. However, caution should be exercised to ensure that the embedded resources are trusted and secure.

The Same Origin Policy is a crucial security measure for web applications. It helps prevent cross-site request forgery (CSRF) attacks, where an attacker tricks a user into performing actions on a different website without their consent. By enforcing the Same Origin Policy, web browsers ensure that websites can only access data and functionality from the same origin, providing a safer browsing experience for users.

The Same Origin Policy is a vital security concept in web applications. It restricts the access of web pages from different sources to each other's data and functionality, ensuring user privacy and protection against malicious attacks.

The Same Origin Policy is a fundamental security concept in web applications. Its purpose is to ensure that web pages from different sources cannot interfere with each other. This policy plays a crucial role in maintaining the security and integrity of web applications.

At a high level, the web can be thought of as an operating system, and each origin can be seen as a separate process within that operating system. Just as an operating system keeps processes isolated from each other, the browser acts as the operating system and enforces the Same Origin Policy to prevent interference between origins.

The basic idea behind the Same Origin Policy is that two separate JavaScript contexts, one from one origin and another from a different origin, should not be able to access each other. The determining factor for whether two contexts are considered the same origin is the tuple of the protocol, hostname, and port associated with each context.

If the protocols, hostnames, and ports of two contexts are the same, they are considered the same origin and are allowed to access each other. On the other hand, if these properties differ, the contexts are considered different origins and are not allowed to interfere with each other.

It is worth noting that the implementation of the Same Origin Policy is not overly complicated. In fact, it is relatively easy to understand and explain. One can even write a function to determine whether two URLs are considered the same origin or not.

However, the complexity arises when determining when to enforce different security checks between different origins and how to define the boundaries of a document. These considerations are crucial in deciding how much interaction should be allowed between origins that do not cooperate with each other.

To illustrate the concept of the Same Origin Policy and its implications, a demonstration was conducted. In the demonstration, an iframe was created and set to the URL of a different origin. It was observed that although different origins can embed each other, they cannot access each other's documents. This means that while embedding is allowed, modifying or accessing the embedded document is not permitted.

The Same Origin Policy is a vital security measure in web applications. It ensures that different origins cannot interfere with each other, thereby protecting the integrity and security of web pages.

When it comes to web application security, one important concept to understand is the Same Origin Policy (SOP) and its implications on cross-site request forgery (CSRF) attacks. The SOP is a fundamental security model of the web that ensures that web pages from different origins cannot access each other's data without explicit permission. An origin is defined by the combination of the protocol, domain, and port number of a web page.

The SOP prevents malicious websites from making unauthorized requests on behalf of a user to another website. This is important because it protects users from attackers who might try to exploit their trust in a legitimate website to perform malicious actions. CSRF attacks, in particular, rely on the fact that browsers automatically include cookies associated with a target website in requests sent to that website.

To illustrate the SOP and CSRF, let's consider an example. Suppose we have a web page hosted on the domain "253site.com" that includes an iframe pointing to a different domain, "dampena.com". The iframe displays content from "dampena.com" within the page hosted on "253site.com". Now, the question is, should the code running on "253site.com" be able to modify the location of the iframe and access data from "dampena.com"?

Some participants in a discussion believe that it should be allowed, while others think it shouldn't. The concern raised is that if the code on "253site.com" could modify the iframe's location, it might be possible for an attacker to trick users into interacting with a different website without their knowledge. This is a valid concern because users might not notice the change in the URL and unknowingly perform actions on an unintended website.

To test whether it is possible to modify the iframe's location, the instructor hits enter and observes that it actually works. Even though the code running on "253site.com" is not allowed to modify the iframe directly, it can still navigate the entire frame. This behavior is interesting and worth noting.

Next, the instructor introduces another scenario. Suppose the code on "253site.com" wants to fetch data from a different domain, "access.stanford.edu", using the Fetch API. The question is, should the code be able to read the response from "access.stanford.edu" and do something with it?

Some participants argue that it should be allowed because they believe the code is not navigating the window but merely trying to retrieve the response as a string to analyze its content. However, when the instructor runs the code, a message appears stating that the request has been blocked by the browser's CORS policy. This policy prevents the code from attaching cookies associated with "access.stanford.edu" to the request, ensuring that sensitive information is not exposed to unauthorized websites.

The instructor highlights the importance of this policy by mentioning the example of making a request to Gmail. If an arbitrary website could make a request to Gmail with the user's cookies attached, it could potentially access personal information and compromise the user's privacy and security. The instructor emphasizes that while it is possible for a server to make a request to Gmail, it is not the same as allowing any website to do so.

The instructor assures the participants that the SOP is implemented consistently across browsers, with a few past implementation bugs. The SOP has been around since 1995, making it a stable and widely accepted security model for web applications.

The Same Origin Policy is a fundamental security model of the web that restricts web pages from different origins from accessing each other's data without explicit permission. This policy helps protect users from cross-site request forgery attacks, where malicious websites attempt to perform unauthorized actions on behalf of users. By preventing unauthorized access to sensitive information, the SOP plays a crucial role in maintaining the security and privacy of web applications.

The Same Origin Policy is a fundamental security concept in web applications that ensures the protection of user data and prevents unauthorized access. It restricts interactions between different web origins, such as domains, protocols, and ports, to prevent malicious activities like Cross-Site Request Forgery (CSRF).

Under the Same Origin Policy, web browsers allow scripts from the same origin to access each other's resources, such as cookies, DOM (Document Object Model) elements, and local storage. However, scripts from different origins cannot access each other's resources due to security concerns.

Although the Same Origin Policy is a crucial security measure, there are exceptions for backward compatibility reasons. For example, forms can be submitted across origins, as we saw earlier. This exception allows data transfer between sites, but it can also introduce vulnerabilities if not properly secured.

To enforce the Same Origin Policy, web browsers implement security mechanisms like the Cross-Origin Resource Sharing (CORS) API. CORS allows servers to specify which origins are allowed to access their resources, providing a more fine-grained control over cross-origin interactions.

In some cases, the Same Origin Policy may be too restrictive or too permissive for specific requirements. For instance, a site may want to cooperate with another site and share resources, or it may need to keep different sites owned by different organizations separate. In such cases, there are ways to relax the Same Origin Policy.

One approach is to use an old API called "document.domain." This API allows two cooperating sites to set their domain to a common value, effectively bypassing the Same Origin Policy. However, this approach has significant drawbacks and is generally considered a bad idea due to security vulnerabilities and potential abuse.

Another approach to relaxing the Same Origin Policy is through the implementation of Cross-Origin Resource Sharing (CORS) headers on the server-side. By configuring the appropriate CORS headers, a server can specify which origins are allowed to access its resources, enabling controlled cross-origin interactions while maintaining security.

It is important to note that relaxing the Same Origin Policy should be done cautiously, considering the potential security implications. Defaulting to a strict Same Origin Policy is recommended unless there are specific requirements for cooperation between different sites.

The Same Origin Policy is a critical security measure in web applications that restricts interactions between different origins. It ensures the protection of user data and prevents unauthorized access. While there are exceptions for backward compatibility reasons, it is essential to enforce the Same Origin Policy and only relax it when necessary, using secure methods like CORS.

The Same Origin Policy is a fundamental security measure in web applications that aims to protect users from malicious attacks, such as Cross-Site Request Forgery (CSRF). This policy restricts how a web page or script can interact with resources from another origin, which includes different domains, protocols, and ports.

One of the key aspects of the Same Origin Policy is that both the requesting page and the target resource must agree to participate in the interaction. In other words, they must opt-in to allow cross-origin communication. This means that simply having the same domain name is not enough to bypass the Same Origin Policy.

To opt-in, the requesting page must set its document domain to match the domain of the target resource. This ensures that both pages are aware of and agree to the cross-origin interaction. It is important to note that the protocol and port must also match for the comparison to be valid.

To initiate the check, the requesting page must have a reference to the target resource. This can be achieved through an iframe or a newly opened window. When the requesting page attempts to access the target resource, the Same Origin Policy is enforced. If the document domain has been properly set and the protocol and port match, the interaction is allowed to proceed.

It is worth mentioning that cookies play a role in cross-origin communication. If the target resource has set a cookie that is accessible to the requesting page, it can further facilitate the interaction.

While the Same Origin Policy is an effective security measure, there are some potential workarounds. One such workaround involves using the fragment identifier in the URL. By creating an iframe and manipulating the fragment identifier, the requesting page can navigate within the target resource without causing a page reload. However, this approach is not recommended as it can introduce security vulnerabilities.

The Same Origin Policy is a crucial security measure in web applications that prevents unauthorized cross-origin communication. It requires both the requesting page and the target resource to opt-in by matching their document domains, protocols, and ports. While there are some workarounds, they should be avoided due to potential security risks.

The Same Origin Policy is a security measure implemented in web browsers to prevent malicious websites from accessing sensitive data or performing unauthorized actions on behalf of a user. It restricts how web pages can interact with each other based on their origin, which is determined by the combination of protocol, domain, and port.

One of the ways to bypass the Same Origin Policy is by using a communication channel between a parent and a child iframe. This technique, although outdated and not recommended, was used in the past when there were no other alternatives available. The parent can communicate with the child iframe by modifying the fragment identifier of the URL and having the child read and display the updated value.

To implement this communication channel, two files are required: "child" and "parent". In the parent file, a script is included to interact with the child iframe. The script retrieves the value of the fragment identifier from the window location and decodes it for better rendering. It then updates the content of a div element with the decoded value. This process is repeated every tenth of a second to check for changes in the fragment identifier. The parent file also embeds the child iframe using a different origin to simulate separate sites.

To enable communication from the parent to the child, an input field is added in the parent file. The script listens for input events on the input field and updates the iframe source by appending the encoded value of the input to the URL's fragment identifier. This process allows the parent to send messages to the child iframe without navigating to a new page.

To observe this communication in action, two servers need to be started on ports 4000 and 4001. Accessing the parent file on port 4000 will embed the child iframe on port 4001. Typing in the input field of the parent file will update the content of the child iframe with a slight delay. The source of the child iframe can be inspected to see the updated URL with the encoded value.

Although this technique demonstrates cross-origin communication, it is not recommended for production use due to security concerns and the availability of more secure alternatives like the PostMessage API. The PostMessage API allows secure communication between different origins by sending strings or objects between two windows or iframes. It handles complex objects and even handles cycles within the objects.

The Same Origin Policy is a vital security measure in web applications that restricts interactions between different origins. While the parent-child iframe communication technique provides a historical perspective on bypassing the Same Origin Policy, it is not recommended for use in modern web development. The PostMessage API offers a more secure and reliable method for cross-origin communication.

The Same Origin Policy is a fundamental security concept in web applications that restricts how documents or scripts from one origin can interact with resources from another origin. This policy is enforced by web browsers to prevent malicious attacks, such as Cross-Site Request Forgery (CSRF), where an attacker tricks a user's browser into performing unwanted actions on a different website.

One way to bypass the Same Origin Policy and enable communication between different origins is by using the postMessage API. This API allows web pages from different origins to securely exchange messages. By sending messages using postMessage, web pages can cooperate and share useful information without violating the Same Origin Policy.

To use the postMessage API, a few changes need to be made to the code. Instead of copying large amounts of data between pages, transferable objects can be used. Transferable objects allow an object to be given to another site, instantly transferring ownership. This enables zero-copy transfers of data, improving performance.

To implement the postMessage API, the following steps can be followed:

1. Get a handle on the window of the iframe.
2. Call the postMessage function with the message to be sent and the origin of the site to communicate with.
3. Remove unnecessary code that pulls data every hundred milliseconds.
4. Register the iframe's interest in receiving message events using the onmessage event handler.
5. Update the content of the div with the data received in the message.
6. Add a check to ensure that the message came from the expected origin to prevent potential security issues.

By following these steps, web pages can securely communicate and exchange information across different origins, bypassing the Same Origin Policy. This allows for improved collaboration and functionality between web applications.

The Same Origin Policy is a crucial security measure in web applications. However, by using the postMessage API and transferable objects, developers can enable secure communication and overcome the restrictions imposed by the Same Origin Policy.

**EITC/IS/WASF WEB APPLICATIONS SECURITY FUNDAMENTALS DIDACTIC MATERIALS****LESSON: SAME ORIGIN POLICY****TOPIC: EXCEPTIONS TO THE SAME ORIGIN POLICY**

## Web Applications Security Fundamentals - Same Origin Policy - Exceptions to the Same Origin Policy

In web application security, one important concept is the Same Origin Policy (SOP). The SOP is a security measure implemented by web browsers to prevent malicious websites from accessing sensitive data or performing unauthorized actions on behalf of the user. It states that web pages can only interact with resources from the same origin, which is defined as the combination of the protocol, domain, and port.

However, there are certain exceptions to the SOP that allow controlled communication between different origins. One such exception is the use of the `postMessage` API. This API enables two websites to cooperate and send messages to each other. To establish communication, one site must have a reference to the other site, which can be achieved by embedding the site in an `iframe`.

Here's an example of how the `postMessage` API can be used: Suppose we have a website called "Access" that wants to display the name of a logged-in user. The user information is stored on "login.stanford.edu", which is on a different origin. To enable communication between the two sites, "Access" listens for a message containing the username. It then embeds an `iframe` to "login.stanford.edu" and expects the `iframe` to post a message to its parent window (which is "Access"). The message will contain the name of the logged-in user.

However, this approach is insecure because it can be exploited by attackers. If an attacker's site, such as "attacker.stanford.edu", embeds "login.stanford.edu", it can also post a message to the parent window and obtain the name of the logged-in user. This can lead to privacy concerns and potential tracking of user activity.

To mitigate these security risks, it is important to validate the destination of the messages being sent. By checking the origin of the message, websites can ensure that they are only sending messages to trusted sources. This validation can be done by using the browser's built-in functionality, which allows specifying the expected origin of the message.

The Same Origin Policy is a fundamental security measure in web applications that restricts communication between different origins. However, there are exceptions to this policy, such as the `postMessage` API, which allows controlled communication between trusted websites. It is crucial to validate the destination of messages to prevent unauthorized access to sensitive information.

The Same Origin Policy is a fundamental security concept in web applications that restricts the communication between different origins. An origin is defined by the combination of the protocol, hostname, and port number. The Same Origin Policy ensures that scripts from one origin cannot access or manipulate resources from a different origin, thus preventing unauthorized access and protecting user data.

However, there are certain exceptions to the Same Origin Policy that allow controlled communication between different origins. One such exception is the `postMessage` API, which enables secure messaging between windows or frames of different origins. The `postMessage` API allows scripts from one origin to send messages to scripts from a different origin, as long as both origins explicitly agree to communicate.

To ensure the integrity and security of the messaging, the source of the messages needs to be validated. The event object, provided by the `postMessage` API, includes a property called `origin`, which specifies the origin from which the message originated. By checking this `origin` against the expected origin, the receiving script can validate the source of the message and ensure that it is coming from a trusted origin.

It is important to note that the `postMessage` API design makes it easier to forget to validate the source of the messages. In contrast to the previous API, which required passing a parameter for validation, the `postMessage` API does not have such a requirement. This can lead to potential security vulnerabilities if the validation step is overlooked.

In scenarios where an untrusted site embeds a trusted site, there is a risk that the untrusted site can gain access to sensitive information. For example, if an untrusted site embeds a login page from a trusted site, it

may trick the user into entering their credentials, which can then be captured by the untrusted site. In this case, the trusted site is unable to prevent the leakage of sensitive information, as it has already provided the information to the untrusted site.

Conversely, there is also a risk when a trusted site embeds an untrusted site. If the trusted site accepts messages from any source without proper validation, an attacker can exploit this vulnerability by sending malicious messages to the trusted site. This can lead to unauthorized access or manipulation of data on the trusted site.

To mitigate these risks, it is crucial to implement proper validation of the message source using the origin property of the event object. By validating the origin, the receiving script can ensure that messages are only accepted from trusted sources, preventing unauthorized access and manipulation.

The Same Origin Policy is a vital security measure in web applications that restricts communication between different origins. The `postMessage` API provides an exception to this policy, allowing secure messaging between different origins. However, it is essential to validate the source of messages using the origin property to ensure the integrity and security of the communication.

The Same Origin Policy is a fundamental security concept in web applications that restricts how different origins (domains, protocols, and ports) can interact with each other. It ensures that resources (such as cookies, local storage, and DOM) from one origin are not accessible or manipulated by another origin without explicit permission.

Exceptions to the Same Origin Policy exist to enable certain functionalities while still maintaining security. One exception is the ability to embed images from other origins in a web page. For example, if we are on `example.com` and we embed a CSS file from `otherone.com`, this is allowed. This is commonly used when embedding CSS files for Google fonts.

Another exception is the ability to embed scripts from other origins. For instance, if a script is loaded from `other3.com` but executed in the context of the current page, it is allowed. This means that the script runs as if it were pasted directly into the page. However, it is important to note that this does not grant access to private user information. It simply allows the execution of code within the current page's context.

These exceptions were not explicitly designed but were added by browsers over time. They allow for different origins to interact with each other, even without the permission of the other origin. Examples of these exceptions include pages embedding images or submitting forms to other origins. While these exceptions may seem to violate the same-origin policy, they were grandfathered in and allowed to function due to their existing usage.

It is crucial to understand that requests made to other origins carry the ambient authority of the cookies attached to them. This means that when an image is embedded from `target.com` on `attacker.com`, the request will include the cookies associated with `target.com`. This can be exploited in attacks, such as phishing, where an attacker can include an image URL that reveals the avatar of the currently logged-in user on a social network. This can make the attacker's page appear more convincing.

The Same Origin Policy is a vital security measure in web applications. Exceptions exist to enable certain functionalities, such as embedding images and scripts from other origins. However, it is important to be aware of the potential security risks associated with these exceptions, particularly when it comes to the handling of cookies and user information.

The Same Origin Policy is a fundamental security concept in web applications that restricts how web pages can interact with each other. It ensures that a web page can only access resources (such as cookies, data, or scripts) from the same origin (domain, protocol, and port) as the page itself. This policy helps prevent malicious websites from accessing sensitive information or executing unauthorized actions on behalf of the user.

However, there are certain exceptions to the Same Origin Policy that allow cross-origin interactions under specific circumstances. One such exception is the use of same-site cookies. When a website attaches the "same-site" attribute to a cookie, it means that the cookie will only be attached when the request is made from the same site. If a request is made from a different site, the cookie will not be attached, thus preventing



unauthorized access.

There are two modes of same-site cookies: "lax" and "strict". The "lax" mode allows cookies to be attached when the request is made from a different site, but only if the request is a top-level navigation or a safe HTTP method (such as GET). The "strict" mode, on the other hand, completely restricts the attachment of cookies when the request is made from a different site.

Another exception to the Same Origin Policy is the use of the "referrer" header. This header is sent by the browser to the server, indicating the page that made the request. It is also used when clicking on a link to indicate the previous page. Some servers may consider the referrer header to determine if the request is coming from a trusted source. However, relying solely on the referrer header can be problematic, especially when dealing with cached images. If the image is already cached by the browser, the server may not perform the necessary checks and allow unauthorized access.

To mitigate these security risks, it is important for websites to implement proper security measures. One way to protect against cross-origin attacks is to use same-site cookies and specify the appropriate mode (lax or strict). Additionally, websites can prevent themselves from being embedded in iframes by setting the appropriate response headers.

By understanding the same-origin policy and its exceptions, web developers can ensure the security and integrity of their web applications, protecting user data and preventing unauthorized access.

The Same Origin Policy is a fundamental security concept in web applications that restricts the interaction between different origins or domains. It ensures that scripts and resources from one origin cannot access or manipulate data from another origin, thereby preventing malicious attacks such as cross-site scripting (XSS) and cross-site request forgery (CSRF).

However, there are certain exceptions to the Same Origin Policy that allow limited communication between different origins. One such exception is the use of same-site cookies. Same-site cookies are cookies that are only sent by the browser to the same site that originated them. This means that even if a resource is loaded from a different origin, the same-site cookie will still be sent along with the request. This allows the server to authenticate and authorize the request based on the cookie, ensuring secure communication between different origins.

Another scenario where the Same Origin Policy can be bypassed is when web applications host popular JavaScript libraries, such as jQuery, on a central Content Delivery Network (CDN) server. Since these libraries are widely used and likely already present in the user's cache, the idea is to take advantage of the cached version to improve performance. Even though the referring site may be different, the cached script can still be utilized. This approach relies on the fact that the script is loaded from the same origin as the referring site, even though it is hosted on a separate CDN server.

However, there are challenges in implementing these exceptions. One challenge is that the "Referer" header, which is used to indicate the referring site, can cause issues with caching. If the browser treats different referring sites as separate cache entries, it can break the caching mechanism. Another challenge is that sites can opt out of sending the "Referer" header entirely, which defeats the purpose of the referring site scheme. This feature was added for privacy reasons, but it can also hinder the effectiveness of the Same Origin Policy exceptions.

For example, in the case of Google Docs, where URLs are used as secret keys to access documents, it would be a privacy concern if the referring site leaked this information. To address this issue, Google Docs does not send the "Referer" header for outgoing links from the document, ensuring that the secret key remains confidential.

The Same Origin Policy is a crucial security measure in web applications that restricts communication between different origins. However, there are exceptions to this policy, such as the use of same-site cookies and the caching of popular JavaScript libraries from a central CDN server. These exceptions aim to improve performance and user experience while maintaining security. However, challenges such as caching issues and privacy concerns need to be carefully addressed when implementing these exceptions.

The Same Origin Policy is an important security feature in web applications that restricts how different web

pages can interact with each other. It ensures that a web page can only access data from the same origin, which includes the same protocol, domain, and port number. This policy helps prevent malicious websites from accessing sensitive information or performing unauthorized actions on behalf of the user.

However, there are some exceptions to the Same Origin Policy that allow certain interactions between web pages from different origins. One such exception is the ability for an attacker to set cookies for a different origin. This means that an attacker who has control over one website can manipulate the cookies of another website, even if they have different origins. This can lead to serious security vulnerabilities, as the attacker can potentially read or modify sensitive information.

Another exception to the Same Origin Policy is related to DNS hijacking. In this scenario, an internet service provider (ISP) intercepts a request for an invalid domain name and returns a hijacked page instead. This hijacked page may contain malicious code or ads. If the page includes JavaScript, it can read and manipulate cookies from other origins, leading to potential security breaches.

It's important to note that the rules around cookies are different from the Same Origin Policy. Specifically, more specific domain names can manipulate cookies from less specific domain names. This is because the rules for cookies were established before the Same Origin Policy was introduced, and they have not been updated to align with it.

To mitigate these vulnerabilities, the use of HTTPS can help protect against DNS hijacking. By mandating HTTPS for the entire website, it becomes more difficult for attackers to intercept and manipulate requests for invalid domain names.

In the next lecture, we will explore ways to manipulate the Same Origin Policy to make it more specific or less specific for different use cases. We will also discuss an attack called cross-site script inclusion, which exploits weaknesses in the Same Origin Policy.

Before we move on, let's address a question regarding a peculiar observation in Gmail. Upon inspecting the requests and responses made by Google, a set of strange characters was noticed at the beginning of the response. This raised curiosity about its purpose. We will delve into this further to understand why such characters are included and their potential role in defending against attacks.

The Same Origin Policy is a crucial security measure in web applications that restricts interactions between different web pages. Exceptions to this policy, such as cookie manipulation and DNS hijacking, can introduce vulnerabilities. Understanding these exceptions and implementing appropriate security measures, like HTTPS, can help protect against potential attacks.

Web applications often embed images and scripts from other sites, but there are limitations to what can be done with these embedded materials. The same origin policy is a fundamental security measure that restricts how web pages can interact with each other. It ensures that a web page can only access data from the same origin (domain, protocol, and port) as itself, preventing malicious websites from accessing sensitive information from other sites.

However, there are exceptions to the same origin policy that allow certain interactions between different origins. One such exception is the ability for sites to submit forms to each other. This means that one site can send data to another site through a form submission. Another exception is the ability for sites to embed images and scripts from other sites. This allows a site to display images or execute scripts hosted on a different origin.

But what if a site wants to prevent these interactions altogether? Is there a way to firewall a site off from others, so that no one can link to it, submit forms to it, or embed its content? Unfortunately, preventing other sites from linking to your site is not possible. Linking is a fundamental aspect of the web, and anyone can create a link to your site. While search engines like Google may consider the reputation of sites that link to you, it is not within your control to prevent others from linking to your site.

However, there are ways to handle incoming requests and control how they are processed. For example, when someone clicks a link to your site, you can return an error page or redirect them elsewhere. This can be useful if you want to handle specific situations, such as when a competitor is linking to your site and you want to display different content to them. By examining the request headers, specifically the 'Referer' header, you can

determine where the request is coming from and respond accordingly.

It is important to note that the 'Referer' header can be manipulated by the sender. There are different values that can be set for the 'referrer policy', which determines how the 'Referer' header is handled. The default policy is to send the full URL of the page that had the link, but there are other options available. For example, you can choose to never send the 'Referer' header or only send the origin (domain, protocol, and port) without revealing the specific page. This is particularly important when linking from a secure site (HTTPS) to an insecure site (HTTP), as you don't want to expose the specific page information to everyone on the same network.

While it is not possible to prevent other sites from linking to your site, there are ways to handle incoming requests and control how they are processed. By examining the 'Referer' header and setting the appropriate 'referrer policy', you can determine how to respond to different requests and protect your site's integrity.

The Same Origin Policy is a fundamental security concept in web applications that restricts how different origins can interact with each other. An origin is defined by the combination of protocol, domain, and port. The Same Origin Policy ensures that resources from one origin cannot access or manipulate resources from another origin, unless explicitly allowed.

However, there are some exceptions to the Same Origin Policy. One exception is the use of the "Referer" header. When a user clicks on a link, the browser sends a request to the destination URL. The site that contains the link can specify how the "Referer" header should be set. This allows the destination site to know that the user came from the site with the link, but it does not reveal the specific page or path.

For example, Google Docs sets the "Referer" header to indicate that a user came from their site. When a user clicks on a link within Google Docs, the destination site will know that the user came from Google Docs, but it will not have access to the specific document or page that the user was on. This helps protect the privacy of users' browsing activities.

There are also other values that can be set for the "Referer" header, depending on the specific requirements of the site. For instance, a site may choose to send the full URL when the user is on the same origin, but send nothing to other origins. These values allow site owners to make trade-offs between privacy and tracking user behavior.

It is important to note that relying solely on the "Referer" header for security purposes is not recommended. The "Referer" header can be easily manipulated or disabled by attackers. Therefore, it is crucial to implement additional security measures to prevent embedding of your site or protect against other types of attacks, such as clickjacking.

Clickjacking is a type of attack where an attacker embeds a legitimate site within a malicious site and tricks users into performing unintended actions. By manipulating the layout and positioning of elements, the attacker can make users unknowingly interact with the embedded site.

To prevent clickjacking attacks, web developers can implement measures such as the X-Frame-Options header or the Content Security Policy (CSP). The X-Frame-Options header allows site owners to specify whether their site can be embedded within an iframe. The CSP provides a more flexible and powerful way to define the policies for content loading and execution within a web page.

The Same Origin Policy is a crucial security concept in web applications that restricts interactions between different origins. The "Referer" header is one exception to this policy, allowing sites to know where users came from without revealing specific details. However, relying solely on the "Referer" header for security purposes is not recommended, and additional measures should be implemented to prevent embedding and other types of attacks.

The Same Origin Policy is a fundamental security concept in web applications. It restricts how a document or script loaded from one origin (e.g., a website) can interact with resources from another origin. This policy is designed to prevent malicious websites from accessing sensitive data or performing unauthorized actions on behalf of the user.

Under the Same Origin Policy, two web pages are considered to have the same origin if they have the same

protocol (e.g., HTTP or HTTPS), domain, and port number. By default, web browsers enforce this policy strictly, meaning that scripts running in one origin cannot access resources from another origin.

However, there are certain exceptions to the Same Origin Policy. One notable exception is the ability to embed content from other origins using iframes. An iframe is an HTML element that allows one webpage to embed another webpage within itself. This can be useful for displaying content from external sources, such as a map or a video.

Attackers can exploit this exception to perform clickjacking attacks. Clickjacking involves overlaying an invisible or translucent frame on top of a visible element, such as a button. When the user clicks on the visible element, they are actually clicking on a hidden element from a different origin. This allows attackers to trick users into performing unintended actions, such as making a purchase on a different website without their knowledge.

To make the hidden frame invisible, developers can set its opacity to zero using CSS properties. This makes the frame transparent but still functional. If the frame was fully visible, the user would only see a portion of the embedded website and would not notice the hidden button.

It's important to note that if the target website has implemented the Same Site Cookies feature, it can prevent clickjacking attacks. Same Site Cookies ensure that requests from embedded frames are treated as sub-resource requests and not as requests from the top-level origin. This prevents the attacker's website from accessing the user's session cookies and performing unauthorized actions on the target website.

The Same Origin Policy is a crucial security measure in web applications. It prevents malicious websites from accessing or manipulating sensitive data. However, exceptions to this policy, such as iframes, can be exploited by attackers to perform clickjacking attacks. Developers should be aware of these vulnerabilities and implement additional security measures, such as Same Site Cookies, to protect against such attacks.

The Same Origin Policy is a fundamental security concept in web applications that restricts how resources on a web page can interact with resources from other origins. It is designed to prevent malicious websites from accessing sensitive information or performing unauthorized actions on behalf of the user.

Under the Same Origin Policy, web pages can only interact with resources (such as scripts, stylesheets, or data) that originate from the same domain, protocol, and port number. This means that a web page from one origin cannot access or modify resources from a different origin.

Exceptions to the Same Origin Policy exist to allow legitimate cross-origin interactions. One such exception is the ability to embed content from other origins using iframes. However, this exception can be exploited by attackers to perform clickjacking attacks, where they trick users into interacting with hidden or disguised elements on a web page.

To prevent clickjacking attacks, a technique called frame busting was commonly used in the past. Frame busting involves the framed website detecting if it is being displayed within a frame and taking action accordingly. However, research has shown that most frame busting techniques are ineffective and can be easily bypassed.

A more effective approach is to use the X-Frame-Options HTTP header. By including this header in the server's response, web developers can control whether their website can be framed by other origins. The header can have different values to specify the desired behavior. For example, "DENY" indicates that the website should not be framed at all, while "SAMEORIGIN" allows framing only by pages from the same origin.

When a browser receives a response with the X-Frame-Options header, it checks the value and enforces the specified framing policy. If the website requests not to be framed, the browser will simply display an empty frame. This prevents clickjacking attacks by denying the attacker's ability to embed the target website within their malicious page.

It is important to note that the order of the requests is crucial in this process. The attacker's server sends a request to the attacker's site, which then includes an iframe to the target site. The browser first loads the attacker's site, and then attempts to load the target site within the iframe. If the target site has specified a framing policy that does not allow framing by other origins, the browser will deny the request and display an

empty frame.

The Same Origin Policy is a critical security mechanism in web applications that restricts cross-origin interactions. Exceptions to this policy, such as iframes, can be exploited for clickjacking attacks. To prevent such attacks, the X-Frame-Options header can be used to control the framing behavior of a website. By specifying the appropriate value in the header, web developers can protect their websites from being framed by malicious origins.

The Same Origin Policy is a fundamental security concept in web applications that helps protect against unauthorized access to sensitive information. It ensures that web pages from different origins cannot interact with each other unless they have the same origin. An origin consists of the combination of the protocol, domain, and port number.

To enforce the Same Origin Policy, servers include a response header called "X-Frame-Options" in their HTTP responses. This header specifies whether the server allows its content to be framed by other origins. The server typically includes the "X-Frame-Options" header with a value of "SAMEORIGIN", which means that the content can only be framed by pages from the same origin.

However, there are exceptions to the Same Origin Policy. For example, if a server includes the "X-Frame-Options" header with a value of "ALLOW-FROM", it can specify specific origins that are allowed to frame its content. This provides more flexibility but should be used with caution to prevent potential security risks.

Browsers play a crucial role in enforcing the Same Origin Policy. When a browser receives a response from a server, it checks the "X-Frame-Options" header to determine whether the content can be framed. If the browser detects that the content should not be framed, it prevents the content from being displayed within a frame on another origin.

It is important to note that the Same Origin Policy primarily protects against attacks from legitimate users using non-hacked browsers. The goal is to prevent malicious websites from framing legitimate websites and tricking users into performing unintended actions, such as sending invalid requests or disclosing sensitive information.

However, it is essential to understand that the Same Origin Policy cannot fully protect against all types of attacks. It is crucial to implement additional security measures on the server-side to ensure the safety of web applications. Servers should never trust client-side data and should validate and sanitize all input to prevent potential vulnerabilities.

When considering the effectiveness of the Same Origin Policy, it is essential to consider the market share of browsers that support this feature. Older browsers may not support the "X-Frame-Options" header, making them vulnerable to attacks. In such cases, alternative frame-busting techniques may need to be employed.

The Same Origin Policy is a critical security measure in web applications that restricts interactions between different origins. It helps protect against unauthorized access and malicious actions. By understanding its limitations and implementing additional security measures, developers can enhance the overall security posture of their web applications.

The Same Origin Policy (SOP) is a fundamental security concept in web applications that restricts how documents or scripts loaded from one origin (domain, protocol, and port) can interact with resources from another origin. The SOP is enforced by web browsers to prevent malicious websites from accessing sensitive data or executing unauthorized actions on behalf of the user.

Exceptions to the Same Origin Policy can occur in certain scenarios, allowing limited interaction between different origins. One common exception is when two origins have the same domain, but different subdomains. In this case, the SOP considers them as separate origins, unless they explicitly opt-in to share resources using Cross-Origin Resource Sharing (CORS) headers.

Another exception is when an origin embeds itself within its own site. This can be exploited by attackers to perform clickjacking attacks. For example, an attacker can embed a copy of a trusted website, such as MySpace, within an ad on another website. When a user clicks on the ad, they unknowingly interact with the embedded MySpace, allowing the attacker to perform malicious actions on their behalf.



To mitigate this risk, web browsers now enforce a stricter interpretation of the SOP, ensuring that all frames within a chain of embedded origins must have the same origin. If any frame within the chain has a different origin, the entire frame is blocked to prevent clickjacking attacks.

The browser is responsible for loading and managing the frames within a web page, allowing it to easily check the origins of each frame it loads. This check is performed when a framed resource requests a particular policy, and the browser verifies if the requested policy matches the origin of the frame. This process allows the browser to enforce the SOP and protect users from malicious actions.

Preventing form submissions from unauthorized sources is another important aspect of web application security. While same-site cookies can prevent cookies from being attached to form submissions, they do not prevent the form submission itself. To address this, web developers can implement additional measures.

One approach is to maintain an allow list on the server, specifying the origins that are allowed to submit forms. The browser, when making requests to other origins, includes an origin header that can be checked against the server's allow list. If the origin is not on the allow list, the server can reject the form submission.

Another approach is to use same-site cookies in conjunction with the allow list. While same-site cookies do not prevent form submissions, they can be used to verify the authenticity of the request. By checking if the same-site cookie is present, the server can determine if the form submission is coming from an authorized source.

It is important to note that relying solely on client-side checks or same-site cookies is not foolproof, as clients can manipulate requests or write scripts to bypass these checks. Therefore, it is essential to implement server-side validation and authentication mechanisms to ensure the security of web applications.

In addition to preventing unauthorized form submissions, web developers may also consider preventing the embedding of images from external origins. This can be useful in scenarios where large images are being embedded, causing performance issues or consuming excessive bandwidth. By disallowing the embedding of images from external origins, web developers can have more control over the resources being loaded on their web pages.

Understanding the Same Origin Policy and its exceptions is crucial for web application security. By implementing appropriate measures, such as allow lists, same-site cookies, and server-side validation, developers can enhance the security of their web applications and protect users from potential attacks.

The Same Origin Policy is a fundamental security concept in web applications. It restricts how a web page or script can interact with resources from another origin, which is defined by the combination of the protocol, domain, and port. The policy aims to prevent malicious websites from accessing sensitive data or performing unauthorized actions on behalf of the user.

However, there are some exceptions to the Same Origin Policy that allow limited interaction between different origins. One common exception is hot linking, where a website embeds resources, such as images, from another site. This can lead to bandwidth theft, as the hosting site ends up paying for the requests made by the embedding site. To prevent this, the hosting site can block the loading of embedded images by checking the Referer header. However, this approach is not foolproof, as the header can be manipulated by attackers.

Another exception to the Same Origin Policy is when a logged-in avatar from one site needs to be displayed on another site. In this case, the Referer header can be used to verify the request's origin. Additionally, same-site cookies can be employed to ensure that the request is legitimate and prevent unauthorized access to user data.

Scripts are another resource that can be embedded from one site to another. While the Same Origin Policy does not prevent the embedding of scripts, there may be cases where a website wants to block the inclusion of its scripts on other sites. This could be to prevent bandwidth theft or to maintain control over the usage of certain effects or functionalities. However, it's important to note that this is not a concern for scripts that do not contain private data and are similar to static files like images or CSS.

In some cases, web developers may use content delivery networks (CDNs) to load external libraries or frameworks. CDNs can be used to improve performance and reduce bandwidth usage. However, this can also



lead to the embedding of scripts from third-party sources. To prevent this, the Referer header can be checked to ensure that the request is coming from an allowed origin.

The Same Origin Policy is a crucial security measure in web applications. While there are exceptions that allow limited interaction between different origins, it's important to implement additional measures, such as checking the Referer header and using same-site cookies, to prevent unauthorized access, bandwidth theft, and other potential security risks.

The Same Origin Policy is a fundamental security concept in web applications that restricts the interaction between different origins. An origin is defined by the combination of a protocol, domain, and port. The Same Origin Policy ensures that web pages from different origins cannot access each other's data or execute malicious actions.

Under the Same Origin Policy, web pages can only access resources (such as cookies, DOM elements, or JavaScript objects) from the same origin. Any attempt to access resources from a different origin is blocked by the browser for security reasons. This policy is implemented to prevent cross-site scripting attacks, where an attacker could inject malicious code into a vulnerable website and steal sensitive information from users.

However, there are certain exceptions to the Same Origin Policy that allow web pages to relax these restrictions under specific circumstances. One exception is when web pages intentionally collaborate with each other. For example, if a site wants to make a request to another origin and receive a response, they can use the `postMessage` API. This API allows communication between different origins by exchanging messages securely. By posting a message to another site and receiving a response, web pages can collaborate and share data.

Another exception to the Same Origin Policy is when a web page wants to read data that is not generated by the current page itself but is provided by an arbitrary server response. For instance, if there is an API server that returns a JSON object with relevant information, a web page may want to make a request to that server and retrieve the JSON object for further processing. In this case, the `postMessage` API cannot be used since there is no page to communicate with. Instead, HTTP headers can be used to allow the retrieval of the desired data. By adding specific headers to the server's response, the web page can indicate that it is allowed to access and read the data.

It is worth noting that scraping or proxying, which involves retrieving data from a server and forwarding it to another site, is not a recommended solution to bypass the Same Origin Policy. It can introduce security risks and is generally considered bad practice.

The Same Origin Policy is a critical security measure in web applications that restricts interactions between different origins. However, there are exceptions to this policy that allow intentional collaboration between web pages and the retrieval of data from arbitrary server responses. These exceptions, such as the `postMessage` API and the use of HTTP headers, provide ways to relax the Same Origin Policy under specific circumstances.

The Same Origin Policy is a fundamental security mechanism in web applications that restricts the interaction between different origins (i.e., domains, protocols, and ports) to prevent unauthorized access to sensitive data. However, there are certain exceptions to this policy that allow cross-origin communication in specific scenarios.

One common exception is when a website wants to fetch data from another origin and display it to its users. This can be achieved by making a request from the user's browser directly, instead of proxying the request through the server. By doing so, the website can retrieve different data based on the user's context, such as their login status. However, this approach has limitations, as the server cannot perform this task on behalf of the website.

Another exception to the Same Origin Policy is the ability to make requests to other origins using script tags. Unlike other resources like images and stylesheets, scripts are not subject to the Same Origin Policy. This means that a website can include a script tag pointing to another origin and retrieve the response. However, there are limitations to this approach as well. While the website can make the request, it cannot directly read the data that comes back from the other origin. The response is treated as JavaScript code and can be executed, but the website can only observe the results and cannot access the data directly.

To overcome this limitation, a technique called JSONP (JSON with Padding) can be used. JSONP is a clever

workaround that involves cooperation between the server and the website. The website requests the server to wrap the JSON response in a callback function of its choice. By doing so, the server includes the JSON response within the function call, allowing the website to access the data by implementing the corresponding callback function. This technique is considered a hack, but it effectively enables cross-origin communication by leveraging valid JavaScript syntax.

However, there are downsides to using JSONP. From the server's perspective, it requires writing additional code to support cross-origin requests and handle the wrapping of responses. It also introduces the risk of producing invalid JavaScript files if not implemented carefully due to potential whitespace issues. Additionally, JSONP is vulnerable to callback injection, where a malicious user can provide a callback name that includes malicious JavaScript code, which will be executed automatically by the website.

The Same Origin Policy is a crucial security mechanism in web applications that restricts cross-origin communication. However, there are exceptions to this policy, such as making direct requests from the user's browser and using script tags with JSONP. These exceptions allow websites to fetch data from other origins, but they come with limitations and potential security risks that need to be carefully considered.

The Same Origin Policy (SOP) is a fundamental security concept in web applications. It restricts how web pages or scripts from one origin can interact with resources from a different origin. An origin is defined by the combination of protocol, domain, and port.

However, there are some exceptions to the SOP that can lead to security vulnerabilities. One such exception is the Reflected File Download (RFD) attack. This attack occurs when callback arguments are not properly sanitized. The attacker tricks the user into clicking on a malicious link that includes these unsanitized arguments. As a result, the user's browser downloads a file from a different origin, which can lead to the execution of arbitrary code.

From the perspective of the other site, they were simply trying to fetch data from another origin and perform some actions with it. To achieve this, they had to include a script from the other site. However, this inclusion of a script can be exploited if the server of the other site gets hacked and starts returning code that deviates from the expected format. This gives the attacker remote execution capabilities on the site that included the script.

To mitigate these risks, it is important to limit the execution access granted to other sites. Allowing full execution access can have severe consequences, as demonstrated by the RFD attack. Therefore, it is crucial to implement proper sanitization techniques and restrict the access granted to external resources.

In the next session, we will delve into three methods that can be used to address these security concerns. Additionally, we will explore the Cross-Site Inclusion (XSI) attack, which is another interesting vulnerability that can be exploited.

**EITC/IS/WASF WEB APPLICATIONS SECURITY FUNDAMENTALS DIDACTIC MATERIALS****LESSON: CROSS-SITE SCRIPTING****TOPIC: CROSS-SITE SCRIPTING (XSS)**

Cross-Site Scripting (XSS) is a code injection vulnerability that allows attackers to run malicious code in the context of a targeted website. This can lead to various security issues and compromise the integrity and confidentiality of user data.

In XSS attacks, the attacker takes advantage of a web application's failure to properly validate or sanitize user input. This allows them to inject their own scripts into a website, which are then executed by unsuspecting users' browsers. The injected scripts can perform various actions, such as stealing sensitive information, modifying website content, or redirecting users to malicious websites.

One infamous example of an XSS attack is the MySpace worm, which spread throughout the social networking site in 2005. The attacker, Samy, discovered that he could include extra code in the profile information box on MySpace. This code would be executed when other users viewed his profile, automatically adding him as a friend. Samy then took it a step further and added a script to the profiles of those who viewed his profile, causing a chain reaction where more and more users became infected. This resulted in exponential growth of friend requests and caused significant disruption to the MySpace platform.

To prevent XSS attacks, web application developers must implement proper input validation and output encoding. Input validation ensures that user input is checked against expected formats and patterns, while output encoding ensures that any user-generated content is properly encoded to prevent script execution. Additionally, web application firewalls and content security policies can be implemented to detect and mitigate XSS attacks.

It is important for both developers and users to be aware of the risks associated with XSS attacks. Developers should follow secure coding practices and regularly update their software to patch any vulnerabilities. Users should be cautious when interacting with websites, especially when entering personal information or clicking on suspicious links.

By understanding the fundamentals of XSS and taking appropriate security measures, web applications can be better protected against this common and potentially devastating vulnerability.

**Web Applications Security Fundamentals - Cross-site scripting - Cross-Site Scripting (XSS)**

Cross-Site Scripting (XSS) is a vulnerability that occurs when untrusted user data is combined with code. This code is typically written by the developer to render a webpage. When combining user data with code, it is crucial to ensure that the user's data does not unintentionally become executable code. This vulnerability can be found in any part of the website where user data is inserted into code written by the developer.

Another common vulnerability that can lead to user data becoming executable code is SQL injection. In SQL injection, the user provides data that is added to a SQL command, which is then sent to a database to retrieve data. The user's data can modify the command in a way that allows the attacker to fetch different data or even run additional commands.

The significance of XSS and SQL injection vulnerabilities lies in the fact that once an attacker manages to inject their data as executable code, they can perform any action as if they were the developer of the website. This means they can take actions on behalf of the user who is logged into the site, making it difficult to detect their malicious activities.

For example, an attacker can send a request to modify a user's profile, making it appear as if the actual user initiated the action. Additionally, stealing someone's cookie is another common exploit. By obtaining the user's cookie and sending it to their own server, the attacker can impersonate the user by adding the stolen cookie to their own browser.

To illustrate the potential consequences of XSS vulnerabilities, let's consider a naive example. Suppose we have a search page where users can search for a specific term. In this example, the server retrieves the user's input

and concatenates it with HTML code to display the search results. However, if the server fails to properly sanitize the user's input, an attacker can inject a script into the search query. When the page is assembled, the script will execute, potentially leading to unauthorized actions or data theft.

To mitigate XSS vulnerabilities, it is crucial to properly sanitize and validate user input before including it in code. This can be achieved through input validation, output encoding, and using secure coding practices.

Cross-Site Scripting (XSS) is a vulnerability that allows attackers to inject malicious code into web applications by exploiting the combination of untrusted user data with code written by developers. This vulnerability can lead to unauthorized actions, data theft, and impersonation of users. It is essential for developers to implement proper input validation and output encoding to prevent XSS attacks.

Cross-Site Scripting (XSS) is a common vulnerability in web applications that allows attackers to inject malicious scripts into web pages viewed by other users. In this type of attack, the attacker exploits the trust that a website has in a user's input, allowing them to execute arbitrary code on the victim's browser.

One specific type of XSS attack is known as "reflected XSS". In this attack, the attacker injects malicious code into a URL parameter or form input, which is then reflected back to the user in the website's response. When the victim clicks on the manipulated link or submits the form, the malicious code is executed in their browser.

To understand how this attack works, let's look at an example. Suppose a vulnerable website allows users to search for products, and the search query is reflected in the search results page without proper sanitization. An attacker could craft a URL that includes a malicious script, such as `<script>alert(document.cookie)</script>`, and send it to a victim. When the victim clicks on the link, the script is executed, and an alert box containing their browser's cookie is shown.

It's important to note that the browser cannot distinguish between code injected by the attacker and code generated by the website itself. From the browser's perspective, the malicious code is part of the legitimate website. This allows attackers to steal sensitive information, such as login credentials or session cookies, and perform actions on behalf of the victim.

There are several ways to mitigate the risk of XSS attacks. One approach is input validation, where the web application checks user input for potentially dangerous characters or patterns. For example, if the application detects the presence of angle brackets (`<` or `>`), it can reject the input or sanitize it by encoding the characters.

Another approach is output encoding, which involves converting special characters into their HTML entities. By doing so, the browser interprets the characters as literal text instead of HTML tags. For instance, the angle brackets in the malicious script `<script>` would be converted to `&lt;` and `&gt;`, preventing the script from being executed.

It's important to note that both input validation and output encoding should be used together for effective protection against XSS attacks. Input validation helps prevent the injection of malicious code, while output encoding ensures that user-supplied data is displayed as intended, without being executed as code.

Cross-Site Scripting (XSS) is a critical security vulnerability that allows attackers to inject malicious scripts into web pages viewed by other users. By exploiting the trust between a website and its users, attackers can execute arbitrary code in victims' browsers, leading to the theft of sensitive information or unauthorized actions. To prevent XSS attacks, web applications should implement input validation and output encoding techniques to sanitize user input and prevent the execution of malicious code.

Cross-Site Scripting (XSS) is a type of vulnerability that allows attackers to inject malicious scripts into web pages viewed by users. In this scenario, we will explore how an XSS vulnerability can be introduced into a website and the potential risks associated with it.

Let's consider a hypothetical situation where a bank is partnering with a blog and wants to track users who come to their website through a specific link from the blog. To achieve this, they decide to add a parameter called "source" to the URL. The bank intends to use this parameter to customize the welcome message for users coming from the blog.

To implement this feature, the bank's website includes code that retrieves the value of the "source" parameter from the URL. If the parameter is defined, the website displays a personalized welcome message. However, if the parameter is undefined, the website does not display anything.

The problem arises when the bank fails to properly sanitize the input received from the "source" parameter. This means that a user can manipulate the parameter by injecting malicious scripts into it. For example, an attacker could insert a script that steals sensitive information, such as the user's cookies.

To demonstrate the vulnerability, the attacker can simply enter a script into the "source" parameter, such as `<script>alert(document.cookie);</script>`. When a user with an active session visits the website through this manipulated link, the injected script executes and displays an alert containing the user's cookie information.

It is worth noting that the injected script does not visibly appear in the URL or the rendered page source. This is because the script is treated as code and executed by the browser, rather than being displayed as part of the page content.

This vulnerability poses a significant risk because users may unknowingly click on seemingly trustworthy links, such as those from reputable websites or blogs. If the targeted website is vulnerable to XSS attacks, the attacker's code can execute within the context of the user's session, allowing them to steal sensitive information or perform unauthorized actions on behalf of the user.

It is important to understand that simply avoiding suspicious links is not enough to protect against XSS attacks. Even if a user is cautious and does not click on suspicious links, they can still be vulnerable if the website they are visiting has an XSS vulnerability.

Cross-Site Scripting (XSS) is a serious web application security vulnerability that allows attackers to inject malicious scripts into web pages. By exploiting this vulnerability, attackers can steal sensitive information, compromise user accounts, and perform unauthorized actions. It is crucial for web developers to implement proper input validation and output encoding techniques to prevent XSS attacks and protect user data.

Cross-Site Scripting (XSS) is a prevalent vulnerability in web applications that allows attackers to inject malicious scripts into web pages viewed by other users. This vulnerability arises when user input is not properly validated or sanitized before being included in the web application's output.

To understand XSS, let's take a look at an example. Suppose we have a web page that displays a user's name, which is obtained from a URL parameter. If the web application fails to properly sanitize the user's input, an attacker can inject a script into the URL parameter, leading to potential security risks.

One way to mitigate XSS vulnerabilities is by using HTML escape techniques. HTML escape involves replacing unsafe characters with their safe counterparts. For example, the less than sign (<) can be replaced with the HTML entity "&lt;" to prevent it from being interpreted as part of an HTML tag.

Frameworks and libraries often provide built-in mechanisms for handling user input and preventing XSS attacks. These mechanisms automatically escape user input before it is displayed in the web application. However, it is essential to be aware of the implications and potential risks when opting out of these mechanisms.

It is worth noting that XSS vulnerabilities are still prevalent in web applications. According to research conducted by White Hat Security, approximately half of the websites analyzed were found to have XSS vulnerabilities. This highlights the importance of implementing proper security measures to mitigate the risk of XSS attacks.

The rise of client-side technologies, such as React and other frameworks, has introduced new challenges in preventing XSS attacks. These technologies require additional considerations and techniques to ensure the security of web applications.

XSS is a significant security concern in web applications. It is crucial to adopt proper input validation and sanitization techniques to prevent XSS vulnerabilities. By never trusting user input and implementing appropriate security measures, web developers can protect their applications and users from potential malicious

attacks.

#### Web Applications Security Fundamentals - Cross-site scripting - Cross-Site Scripting (XSS)

Cross-Site Scripting (XSS) is a common vulnerability in web applications that allows attackers to inject malicious scripts into web pages viewed by other users. XSS attacks can occur in different contexts, such as HTML, JavaScript, and user input fields, and understanding the rules and techniques used by attackers is crucial for effective defense.

There are two main types of XSS attacks: reflected and stored. In a reflected XSS attack, the attacker injects malicious code into the HTTP request itself, typically by manipulating the URL or query parameters. The target user is then tricked into visiting a specially crafted URL that includes the attack code. However, this type of attack is limited because the attacker needs to find a way to include the attack code in the URL or query.

On the other hand, stored XSS attacks involve storing the attack code in the server's database. When a user visits a page that retrieves data from the database, the attack code is included in the rendered page, regardless of the URL they came from. This makes stored XSS attacks more powerful, as the attacker can use various means to inject the attack code into the database.

For example, an attacker could exploit a vulnerability in the server that allows them to modify an HTTP header, which is then stored in the database. Alternatively, they could bribe someone with access to the database to add a script containing the attack code. The key point is that the method of injecting the attack code into the database is not important; what matters is that it gets stored and executed when users access the page.

To protect against XSS attacks, web developers must properly sanitize and validate user input. In the case of HTML contexts, all left angle brackets (<) should be replaced with the entity code "&lt;" to prevent them from being interpreted as HTML tags. Additionally, a backslash escape character ("\") should be used to further ensure that the input is treated as plain text and not as code.

However, it is important to note that simply replacing angle brackets is not enough. Developers must also consider other special characters and properly escape them to prevent XSS attacks. For example, if a user enters the string "\"&lt;" (backslash followed by a left angle bracket), the sanitization function should handle it correctly and not interpret it as an HTML tag.

XSS attacks pose a significant threat to web applications, and understanding the different types of XSS vulnerabilities is crucial for effective defense. By properly sanitizing and validating user input, developers can mitigate the risk of XSS attacks and ensure the security of their web applications.

Cross-Site Scripting (XSS) is a common vulnerability in web applications that allows attackers to inject malicious scripts into web pages viewed by other users. In this didactic material, we will discuss the fundamentals of Cross-Site Scripting, focusing on the problem of user input and how to defend against it.

When dealing with user input, it is important to understand the concept of escape sequences. An escape sequence is a combination of characters that represents a special meaning. In the context of web applications, escape sequences are used to indicate that certain characters should be treated as literal characters rather than having their usual interpretation.

One common escape sequence used in web applications is the ampersand followed by a character code and a semicolon (e.g., &lt;). This sequence represents a special character, such as a less than symbol (<). However, if user input is not properly sanitized, an attacker can exploit this by injecting their own escape sequences.

For example, if a user enters the characters "&lt;script>" as input, the web application may interpret it as an opening script tag instead of literal characters. This allows the attacker to inject and execute their own JavaScript code on the web page, potentially compromising user data or performing other malicious actions.

To mitigate this vulnerability, web developers must properly sanitize user input. One approach is to replace certain characters with their corresponding HTML entities. For example, the less than symbol (<) can be replaced with "&lt;". This ensures that the user input is treated as literal characters and not interpreted as HTML tags or escape sequences.



However, it is important to note that the mitigation strategy may vary depending on the context. For instance, the rules for handling user input in style and script tags are different from other types of tags. Therefore, developers should be aware of the specific rules and apply the appropriate sanitization techniques accordingly.

In the case of HTML attributes, it is crucial to sanitize user input to prevent attacks. Attributes, such as the alt attribute in an image tag, can be manipulated by attackers to inject malicious code. By exploiting the lack of proper sanitization, an attacker can inject additional attributes or modify existing ones, leading to potential security breaches.

To defend against such attacks, web developers should carefully validate and sanitize user input. In the case of HTML attributes, special attention should be given to characters that have special meaning, such as quotes. By replacing certain characters, such as single quotes, with their corresponding HTML entities, developers can prevent attackers from injecting malicious code through attribute values.

However, it is important to note that using double quotes as attribute delimiters is also a common practice. Therefore, developers should consider implementing a comprehensive sanitization strategy that covers both single and double quotes.

Cross-Site Scripting (XSS) is a serious security vulnerability that can be exploited by attackers to inject malicious scripts into web pages. By understanding the fundamentals of XSS and implementing proper input sanitization techniques, web developers can effectively defend against this type of attack.

Cross-site scripting (XSS) is a web application security vulnerability that allows attackers to inject malicious scripts into web pages viewed by other users. In this didactic material, we will focus on the fundamentals of XSS, specifically on Cross-Site Scripting.

Cross-Site Scripting occurs when a web application does not properly validate user input and allows attackers to inject malicious scripts into web pages. These scripts are then executed by the victim's browser, leading to various security risks.

There are three main types of Cross-Site Scripting: Stored XSS, Reflected XSS, and DOM-based XSS. In this material, we will focus on Stored XSS.

Stored XSS, also known as Persistent XSS, occurs when an attacker injects a malicious script that is permanently stored on the target server. This script is then served to other users when they access the affected web page.

To understand Stored XSS, let's consider an example. Imagine you are creating a new profile on a website and it asks for your name. Instead of providing a regular name, you enter a string that contains a malicious script. When other users visit your profile, their browsers will execute the injected script, potentially leading to unauthorized actions or data theft.

To mitigate Stored XSS attacks, web developers should implement proper input validation and output encoding. Input validation ensures that user-supplied data is within the expected range and format, while output encoding ensures that any user-generated content is properly encoded before being displayed to other users.

One common vulnerability that allows for Stored XSS is the improper handling of quotes. By converting quotes into HTML entities, developers can prevent the execution of injected scripts. Additionally, developers should be cautious of attributes that have more power, such as "source" or "href," as allowing user input in these attributes can lead to script execution.

It is important to note that the lack of proper input validation and output encoding can have severe consequences. Attackers can exploit Cross-Site Scripting vulnerabilities to steal sensitive information, gain unauthorized access, or perform other malicious activities.

Cross-Site Scripting (XSS) is a web application security vulnerability that allows attackers to inject malicious scripts into web pages. Stored XSS, one of the three main types of XSS, occurs when an attacker injects a malicious script that is permanently stored on the target server. To mitigate Stored XSS attacks, developers should implement proper input validation, output encoding, and handle attributes with caution.

Cross-Site Scripting (XSS) is a common vulnerability in web applications that allows attackers to inject malicious scripts into trusted websites. This can lead to various security risks, such as stealing sensitive information, manipulating website content, or performing unauthorized actions on behalf of the user.

One type of XSS attack is called Cross-Site Scripting via data and JavaScript URLs. Data URLs allow embedding data directly into a URL, while JavaScript URLs execute JavaScript code within the context of the current page. These types of URLs are rarely seen but can be dangerous if not properly handled.

A data URL can be constructed by specifying a MIME type and the content to be displayed. For example, "data:text/html,hi" would display the text "hi" as an HTML page. However, if a website does not properly sanitize data URLs, an attacker could inject arbitrary JavaScript code, potentially leading to unauthorized access or data theft.

Similarly, JavaScript URLs allow executing JavaScript code directly within the current page. By prefixing a script with "javascript:", an attacker can run arbitrary code in the context of the website. This can be exploited to perform actions on behalf of the user or manipulate website content.

It is worth noting that some browsers have implemented protection mechanisms to mitigate the risks associated with JavaScript URLs. For example, if a user copies and pastes a JavaScript URL into the address bar, the "javascript:" prefix may be stripped to prevent unintended execution of malicious code. However, this protection can be bypassed by manually adding back the "javascript:" prefix.

Phishing attacks often exploit the JavaScript URL vulnerability by tricking users into pasting malicious code into their address bar. Attackers may prompt users to execute code that appears harmless but actually performs malicious actions. To combat this, browser vendors have implemented measures to remove the "javascript:" prefix when pasting URLs into the address bar, reducing the effectiveness of such attacks.

Cross-Site Scripting via data and JavaScript URLs is a significant security concern in web applications. Proper input validation and output encoding are crucial to prevent the injection of malicious scripts. Additionally, user awareness and caution when interacting with unfamiliar URLs can help mitigate the risks associated with these types of attacks.

Cross-Site Scripting (XSS) is a web application security vulnerability that allows attackers to inject malicious scripts into web pages viewed by other users. This can lead to the theft of sensitive information, session hijacking, or unauthorized actions on the affected website.

One common type of XSS is called "JavaScript URL" or "JavaScript link". This involves using a URL that starts with "javascript:" followed by the code the attacker wants to execute. When a user clicks on the link, the injected script runs in the context of the affected web page. Different browsers handle this differently, with Chrome and Firefox stripping out the "javascript:" part, while Safari prompts the user to enable JavaScript for search fields.

Another technique used in XSS attacks is the use of data URLs. These URLs allow the embedding of data directly into the web page, eliminating the need for an additional HTTP request. This can be useful for small images or other resources that need to be loaded quickly. However, if not properly validated, data URLs can also be used to execute malicious scripts.

To protect against XSS attacks, it is important to validate and sanitize user input before using it in HTML attributes such as href or source. This involves checking that the input is a well-formed URL and does not contain any potentially dangerous characters or scripts. Additionally, it is important to properly encode user-generated content when displaying it on web pages to prevent script injection.

It is worth noting that XSS attacks can have serious consequences and should not be taken lightly. Websites should implement proper security measures, such as input validation, output encoding, and the use of Content Security Policy (CSP), to mitigate the risk of XSS vulnerabilities.

Cross-Site Scripting (XSS) is a web application security vulnerability that allows attackers to inject and execute malicious scripts on web pages viewed by other users. By understanding the different techniques used in XSS

attacks and implementing proper security measures, website owners can protect their users from potential harm.

Cross-Site Scripting (XSS) is a common web application security vulnerability that allows attackers to inject malicious scripts into web pages viewed by other users. This can lead to various attacks, such as stealing sensitive information, manipulating website content, or performing phishing attacks.

One type of XSS attack is known as DOM-based XSS. In this type of attack, the attacker exploits the way the Document Object Model (DOM) handles user input. The DOM is a representation of the HTML structure of a web page, and it is manipulated by JavaScript to dynamically update the page.

To understand how DOM-based XSS works, let's consider an example. Suppose a web application allows users to input data that is then displayed on the page. The application includes a function that runs when the user hovers over a specific element on the page. This function takes user data and incorporates it into the JavaScript code.

The problem arises when the application fails to properly sanitize the user input. If the user input contains special characters or code that is not properly escaped, it can be interpreted as code by the browser and executed. This allows the attacker to inject malicious scripts into the page.

In the example mentioned earlier, the user input is being used in both a JavaScript context and an HTML context. Simply escaping the single quote or double quote characters is not enough to prevent the attack. The attacker can end the quote early and add additional attributes to the HTML tag, effectively injecting their own code.

Additionally, even if the attacker's code is on a different origin (domain), it can still cause problems. Browsers have a single event loop, so if the attacker's code includes an infinite loop or other resource-intensive operations, it can lock up the entire page, causing a denial of service.

Another issue to consider is when the user is allowed to select the ID of an attribute. IDs are supposed to be unique across the entire page, similar to classes. If the user selects an ID that is already in use, it can lead to unexpected behavior. For example, JavaScript code that relies on selecting elements by their ID may mistakenly select the wrong element, potentially leading to security vulnerabilities.

Furthermore, there is a peculiar feature in the DOM where any element on the page with an ID automatically creates a global variable in JavaScript with that name. This means that if the user controls the ID of an element, they can declare variables at the start of the script, potentially altering the behavior of the code and enabling attacks.

Cross-site scripting (XSS) is a serious web application security vulnerability that can be exploited to inject malicious scripts into web pages. DOM-based XSS attacks take advantage of the way the Document Object Model handles user input. Proper input sanitization and validation are essential to prevent XSS attacks and protect user data.

Cross-Site Scripting (XSS) is a common vulnerability in web applications that allows attackers to inject malicious scripts into websites that are viewed by other users. One type of XSS attack is known as Cross-Site Scripting (XSS) and it occurs when user input is not properly validated and is displayed on a web page without proper encoding.

In this example, the speaker discusses the issue of including dynamic user strings inside a variable in a script. This can be a dangerous practice if not done correctly, as it can allow an attacker to inject malicious code into the script. The speaker demonstrates how putting user data between single quotes can be problematic, as it can interfere with the syntax of the script.

To mitigate this issue, the speaker suggests escaping the single quotes by using backslashes. This ensures that the user data is treated as a string and does not interfere with the script syntax. Additionally, the speaker mentions the importance of also escaping double quotes in case the user includes them in their input.

However, the speaker also highlights a potential vulnerability in this approach. An attacker could potentially use

a backslash followed by a quote to break out of the string and inject malicious code. This is because the server's algorithm for handling backslashes and quotes may interpret the input differently.

To address this vulnerability, the speaker suggests using double backslashes in the input to ensure that the server interprets the backslash as a literal character and not as an escape character. By doing so, the server will not treat the backslash followed by a quote as a control sequence, but as a literal backslash and quote.

Cross-Site Scripting (XSS) poses a significant threat to web applications and can lead to the compromise of user data and the execution of malicious code. Proper input validation and encoding are essential in preventing XSS attacks. Developers should be aware of the potential vulnerabilities associated with including user input in scripts and take appropriate measures to mitigate these risks.

Cross-Site Scripting (XSS) is a common web application vulnerability that allows attackers to inject malicious scripts into web pages viewed by other users. In this lesson, we will discuss the fundamentals of XSS and explore different types of XSS attacks.

One type of XSS attack is called Cross-Site Scripting (XSS). This occurs when an attacker is able to inject malicious scripts into a trusted website. These scripts are then executed by unsuspecting users, allowing the attacker to steal sensitive information or perform unauthorized actions on behalf of the user.

To understand how XSS works, let's consider a scenario where a user submits a form on a website, and the input is displayed on a page without proper sanitization. If an attacker is able to inject a script into the input, it will be executed when the page is loaded by other users.

There are different types of XSS attacks, including Stored XSS, Reflected XSS, and DOM-based XSS. Each type has its own characteristics and potential impact. Stored XSS occurs when the injected script is permanently stored on the server and served to multiple users. Reflected XSS happens when the injected script is included in a URL or form parameter and is only executed when the user interacts with a specific link or form. DOM-based XSS occurs when the injected script manipulates the Document Object Model (DOM) of a web page.

To prevent XSS attacks, it is crucial to properly sanitize and validate user input. This involves removing or encoding any potentially dangerous characters or scripts. One commonly used technique is HTML escaping, which converts special characters into their corresponding HTML entities. By doing so, the browser will treat the characters as plain text and prevent the execution of any embedded scripts.

However, it's important to note that HTML escaping alone is not sufficient to mitigate all XSS vulnerabilities. Attackers can still find ways to bypass these protections, especially in more complex scenarios. Therefore, it is recommended to use a combination of security measures, such as input validation, output encoding, and Content Security Policy (CSP), to effectively defend against XSS attacks.

Cross-Site Scripting (XSS) is a significant security vulnerability that can have severe consequences for web applications and their users. Understanding the different types of XSS attacks and implementing appropriate security measures is essential to protect against this threat.

Cross-Site Scripting (XSS) is a common vulnerability in web applications that allows attackers to inject malicious scripts into web pages viewed by other users. These scripts can be used to steal sensitive information, manipulate web content, or perform other malicious actions.

One way XSS attacks can occur is when a web application includes user-provided data in a web page without properly sanitizing or validating it. This allows attackers to inject their own scripts into the page, which are then executed by the victim's browser.

There are different types of XSS attacks, but one common type is called "Reflected XSS." In this type of attack, the malicious script is included in the URL of a web page. When the victim clicks on the manipulated URL, the script is executed in their browser.

To prevent XSS attacks, it is important to properly sanitize and validate all user-provided data before including it in a web page. One way to do this is by hex encoding the user's data. Hex encoding converts the user's input into a different representation that only uses the characters 0-9 and A-F. This ensures that any possible

combination of these characters is safe and cannot break out of the intended context.

However, it is important to note that the encoded data needs to be decoded at runtime to revert it back to its original form. This ensures that the user's input is correctly displayed and processed by the application.

Another approach to prevent XSS attacks is by using a tag called "template." This HTML tag allows developers to include user data in a web page without executing any scripts. The data can only be accessed and manipulated through JavaScript, making it safer.

To escape the content inside the template tag, developers only need to escape the left angle bracket (<) and the ampersand (&) characters. This is a simple rule that helps prevent XSS attacks.

It is important to be aware of certain contexts that are never safe from XSS attacks. These include placing user input in free-floating parts of a script, using HTML comments, using tag names, and using cell contents. These contexts have specific parsing rules that can be difficult to handle correctly, making it safer to avoid using user input in these contexts altogether.

Preventing XSS attacks is crucial for maintaining the security of web applications. By properly sanitizing and validating user input, using encoding techniques, and being cautious about certain contexts, developers can significantly reduce the risk of XSS vulnerabilities.

Cross-site scripting (XSS) is a common vulnerability found in web applications that allows attackers to inject malicious scripts into web pages viewed by users. This can lead to various security risks, such as stealing sensitive information, manipulating website content, or redirecting users to malicious websites.

One type of XSS attack is called "tag name evasion". Attackers try to bypass the filtering mechanisms implemented by web applications by using different variations of tag names. For example, they may use angle brackets followed by a script tag, thinking that all tags must start with a name followed by a space. However, modern browsers interpret this differently and may still execute the injected script.

There are several ways attackers can exploit this vulnerability. In some cases, certain characters may be ignored by the browser, allowing attackers to insert arbitrary content between the attribute name and the equals sign. This can bypass blacklist approaches where specific tag names are blocked, as the injected script may not match the banned name exactly.

In other cases, attackers may manipulate the structure of the HTML tag itself. They may omit quoting or even reverse the opening and closing tags. Surprisingly, some browsers still execute the script even in these unconventional scenarios. This forgiving behavior of HTML parsers is based on the robustness principle, which suggests that software should be conservative in what it accepts and liberal in what it produces. However, this principle can introduce security risks, as it requires making assumptions about the meaning of malformed input.

The robustness principle is often applied in various contexts, such as TCP implementations, where accepting malformed packets can improve compatibility with different systems. However, in the case of web applications, this principle can lead to unintended consequences. For example, if a browser automatically corrects poorly written HTML, developers may not realize that their code is incorrect and may unknowingly rely on browser-specific behavior. This can result in inconsistencies across different browsers, as each may interpret the code differently.

To mitigate the risks associated with XSS attacks, web developers should adopt a defense-in-depth approach. This involves implementing multiple layers of security controls, such as input validation, output encoding, and proper handling of user-generated content. By carefully validating and sanitizing user input, developers can prevent malicious scripts from being executed on their websites. Additionally, output encoding techniques can ensure that any user-generated content is displayed as plain text, preventing script execution.

Cross-site scripting (XSS) is a significant security concern in web applications. Attackers can exploit vulnerabilities in web pages to inject malicious scripts, potentially compromising user data and website integrity. Understanding the various techniques used by attackers, such as tag name evasion, is crucial for developers to implement effective security measures and protect against XSS attacks.

Cross-Site Scripting (XSS) is a vulnerability that occurs when an attacker injects malicious scripts into a web application, which are then executed by the user's browser. In this didactic material, we will discuss the fundamentals of XSS and how to prevent it.

One important aspect of XSS is the need to properly escape user data. By escaping user data, we ensure that any potentially malicious code is rendered harmless. There are three contexts where user data should be escaped: HTML, JavaScript, and URL.

In the HTML context, user data should be HTML encoded to prevent any malicious code from being interpreted as HTML tags or attributes. This can be done using functions like `htmlspecialchars()` in PHP or the equivalent in other programming languages.

In the JavaScript context, user data should be escaped using JavaScript backslash sequences to prevent any special characters from being interpreted as code. This can be achieved by using functions like `encodeURIComponent()` in JavaScript.

In the URL context, user data should be URL encoded to prevent any malicious code from being interpreted as part of the URL. This can be done using functions like `urlencode()` in PHP or the equivalent in other programming languages.

It is important to note that nesting parsing chains should be avoided. This occurs when user data is nested within multiple layers of code, making it difficult to properly escape. Nesting parsing chains can lead to vulnerabilities and should be avoided.

Additionally, it is crucial to be aware of the different contexts in which user data is used. In some cases, user data may need to be escaped multiple times to ensure it is properly protected. Failure to escape user data correctly can result in XSS vulnerabilities.

To illustrate this, consider the example where user data is used within a string in a JavaScript function. In this case, the user data needs to be escaped as JavaScript code and then again as HTML code to prevent any potential XSS attacks.

Another example is when user data is used in different contexts within the same application. In this case, the user data needs to be escaped for each specific context to avoid vulnerabilities.

It is worth mentioning that XSS vulnerabilities can have serious consequences, including unauthorized access to user data, session hijacking, and the injection of malicious code. Therefore, it is crucial to follow best practices and properly escape user data to mitigate the risk of XSS attacks.

Cross-Site Scripting (XSS) is a significant security vulnerability that can be prevented by properly escaping user data in different contexts. By understanding the fundamentals of XSS and adopting secure coding practices, developers can protect web applications from potential attacks.



**EITC/IS/WASF WEB APPLICATIONS SECURITY FUNDAMENTALS DIDACTIC MATERIALS****LESSON: CROSS-SITE SCRIPTING****TOPIC: CROSS-SITE SCRIPTING DEFENSES**

Cross-site scripting (XSS) is a type of vulnerability that can occur in web applications. It involves the injection of malicious code into a website, which is then executed by the user's browser. Last time, we discussed the concept of XSS and how it works. Today, we will focus on the defenses against XSS attacks.

To begin, it is important to understand the context in which XSS attacks occur. The attacker's goal is to get their code to run in a user's browser while they are visiting a website. This code can potentially access sensitive information or perform unauthorized actions. Therefore, the target of the attack is not the website server itself, but rather the user's browser.

There are two main types of XSS attacks: reflected and stored. Reflected XSS involves manipulating a URL to trick the server into reflecting back some data that is treated as code by the user's browser. Stored XSS, on the other hand, involves adding malicious code into a database, which is then returned by the server on various responses. Stored XSS attacks are generally more powerful than reflected XSS attacks.

To defend against XSS attacks, we need to escape or sanitize user input before combining it with code. The core problem lies in the combination of untrusted user data with code. In fact, the term "cross-site scripting" can be somewhat misleading, as the issue is more accurately described as HTML injection. This aligns with similar vulnerabilities like SQL injection and other server-side injections.

The solution is to ensure that user input is properly escaped or sanitized. This means that special characters and code are treated as plain text and not executed by the browser. We discussed the specific characters that need to be handled carefully in our previous discussion.

User data can come from various sources, including HTTP requests such as query parameters, form fields, headers, and cookies. It is important to be cautious with all data that comes from the user. Additionally, data from databases should also be treated as potentially untrusted, as we cannot always be certain of how the data was obtained.

Defending against XSS attacks involves escaping or sanitizing user input to prevent the execution of malicious code. By properly handling user data and being cautious with all sources of input, we can mitigate the risks associated with XSS vulnerabilities.

When it comes to web application security, one of the major concerns is cross-site scripting (XSS) attacks. XSS attacks occur when an attacker injects malicious code into a web application, which is then executed by unsuspecting users. This can lead to various consequences, such as stealing sensitive information or manipulating the website's content.

To defend against XSS attacks, it is crucial to understand the fundamentals of XSS and the available defenses. One common defense mechanism is escaping user input. Escaping involves converting special characters into their HTML entity equivalents, ensuring that they are treated as literal characters rather than executable code.

There are two main approaches to escaping user input: escaping on the way into the database or escaping on the way out when rendering the page. Escaping on the way into the database ensures that any data added by different individuals on the team is in a clean format. However, this approach may not account for the context in which the data will appear on the page.

On the other hand, escaping on the way out at render time allows for a more accurate decision-making process. Since the context of the data is known during rendering, the correct escaping mechanism can be applied accordingly. This approach eliminates the need to predict all possible contexts in advance.

It is important to note that different contexts within a web page require different escape characters. For example, content placed inside an HTML tag, an attribute, or a script string may have different escape requirements. Therefore, assuming that the data has not been escaped and performing the escaping process at render time is the safest and preferred option.

Web frameworks often provide built-in HTML escaping functionality, which should be utilized. These functionalities have been developed and tested by a large number of users, increasing the likelihood of discovering and fixing any potential bugs. This concept, known as Linus's Law, emphasizes that with enough scrutiny, all bugs can be identified and resolved.

By leveraging the HTML escaping functionality provided by web frameworks, developers can benefit from the collective expertise of the community. This not only saves time and effort but also increases the likelihood of using a secure and reliable solution. It is important, however, to understand how the specific framework handles HTML escaping and to use it correctly.

Defending against XSS attacks requires implementing proper defenses, such as escaping user input. By escaping on the way out at render time and utilizing the HTML escaping functionality provided by web frameworks, developers can mitigate the risk of XSS vulnerabilities and ensure the security of their web applications.

Cross-site scripting (XSS) is a common vulnerability in web applications that allows attackers to inject malicious scripts into web pages viewed by other users. In order to prevent XSS attacks, it is important to understand the different contexts in which untrusted data can appear and to properly escape that data.

One way to defend against XSS attacks is by using HTML escaping. However, it is not enough to simply rely on a framework to handle the escaping for you. You need to understand how the framework is escaping the data and ensure that it is escaping all the necessary characters. For example, if you are using HTML escaping, you cannot simply take the output and put it into a string in a script tag or a comment without ensuring that all the characters are properly escaped.

Let's take a look at an example using a templating language called ejs (embedded JavaScript). In ejs, you can use HTML tags and the syntax "`<%= %>`" to insert user data into a web page. This syntax will evaluate the JavaScript value, escape it, and then insert it into the page. This ensures that any user data is properly escaped and prevents XSS attacks.

Another popular framework, React, also provides a way to insert user data into a web page. However, React automatically escapes any user data that is inserted using angle brackets. This means that if you try to insert user data using the "`<div>`" syntax, React will escape it to prevent XSS attacks. React also prevents you from using certain JavaScript properties as attributes to avoid potential security risks.

If you need to insert a string into a web page using React and you are certain that the string is safe, you can use the "`dangerouslySetInnerHTML`" attribute. However, it is important to note that this attribute should only be used when you are absolutely certain that the string is safe and has already been properly escaped.

It is worth mentioning that the use of the "`dangerouslySetInnerHTML`" attribute and other similar features in React should be done with caution. These features are intentionally designed to look ugly and discourage their use. If you come across code that uses these features, it is recommended to refactor it and find a safer and cleaner alternative.

Understanding how to properly escape untrusted data is crucial in defending against XSS attacks in web applications. By using the appropriate escaping techniques provided by frameworks like ejs and React, you can ensure that user data is properly sanitized and prevent potential security vulnerabilities.

Cross-site scripting (XSS) is a common vulnerability in web applications that allows attackers to inject malicious scripts into trusted websites. These scripts can then be executed by unsuspecting users, leading to various security risks. In this didactic material, we will explore the fundamentals of XSS and discuss some defenses against it.

One of the main causes of XSS vulnerabilities is the improper handling of user input. When user-supplied data is not properly sanitized or validated, it can be used to inject malicious scripts into the web application. These scripts are then executed by other users who visit the affected page.

To illustrate this, let's consider an example using the EJS template language. EJS allows developers to render

dynamic content in HTML templates. However, if not used correctly, it can make it easy to introduce XSS vulnerabilities.

In our example, we have an Express server listening on port 4000. We define a simple route for the home page, where we send back some HTML. The HTML contains a template that includes a variable called "name". We assume that this variable will be safe because it is rendered using the EJS render function, which should escape any potentially malicious content.

However, if we allow the "name" variable to be specified by the user through a query parameter, we introduce a potential vulnerability. If an attacker sets the "name" parameter to a malicious script, it will be rendered without proper escaping, leading to an XSS attack.

To mitigate this vulnerability, it is important to properly sanitize and validate all user input. In the case of EJS, it is crucial to ensure that the template syntax is used correctly, with the equal sign (=) instead of a dash (-) to render variables. Additionally, any user-supplied data should be properly escaped before being rendered in the template.

Assuming that XSS attacks will happen, it is also important to implement additional defenses to minimize the impact. One approach is to implement a Content Security Policy (CSP) that restricts the types of content that can be loaded on a page. This can help prevent the execution of injected scripts by blocking or sanitizing them.

Another defense mechanism is to implement input validation and output encoding consistently throughout the application. By validating and sanitizing user input on the server-side and encoding output properly, we can reduce the risk of XSS vulnerabilities.

XSS vulnerabilities can have serious consequences for web applications. It is crucial to properly handle user input, validate and sanitize it, and encode output to prevent the injection and execution of malicious scripts. Additionally, implementing defenses such as Content Security Policies can further enhance the security of web applications.

In the field of computer security, one key concept is defense in depth. This concept aims to provide redundant security measures so that even if one measure fails, there are other layers of defense to prevent successful attacks. This is particularly important when it comes to web application security, specifically in the case of cross-site scripting (XSS) attacks.

One example of defense in depth in computer systems is the built-in detection feature in web browsers like Safari. When visiting a potentially malicious site, the browser will warn the user and prompt them to reconsider proceeding. Even if the user decides to proceed, additional measures like requiring a user to right-click to open a file instead of double-clicking can add an extra layer of defense.

Another example is the use of two-factor authentication. Even if an attacker manages to obtain a user's password, they would still need physical access to the user's second factor device, such as a phone, to gain access. Additionally, some services send email notifications when a user logs in from a new location, providing an audit trail and allowing the user to respond in case of a security breach.

It's worth mentioning that defense in depth is not foolproof, and it's important to continually evaluate and update security measures to stay ahead of attackers. One example of a flawed implementation of defense in depth was the use of the Data Encryption Standard (DES) algorithm, which was considered insecure. To compensate for its weaknesses, some organizations resorted to encrypting data multiple times, mistakenly believing it would enhance security. However, this approach actually increased the risk of losing access to the encrypted data.

Now, let's focus on a specific aspect of defense in depth: defending user cookies. In the context of web applications, attackers may attempt to steal user cookies to gain unauthorized access. One defense mechanism is the use of HTTP-only cookies. These cookies cannot be accessed by JavaScript running in the browser, preventing attackers from exfiltrating sensitive information. Even if an attacker tries to access the cookie using the "document.cookie" method, it will return as if the cookie does not exist.

To understand how HTTP-only cookies work, it's essential to know that when making HTTP requests to a server,

the cookie is automatically included as a header. This allows the server to validate the user's identity and provide the necessary access rights.

Defense in depth is a fundamental concept in computer security, aiming to provide multiple layers of security measures to protect against attacks. Examples include built-in detection features in web browsers, two-factor authentication, and the use of HTTP-only cookies to defend against cross-site scripting attacks. However, it's crucial to regularly reassess and update security measures to stay ahead of evolving threats.

Cross-site scripting (XSS) is a common vulnerability in web applications that allows attackers to inject malicious scripts into web pages viewed by users. These scripts can be used to steal sensitive information, manipulate content, or perform other malicious actions. To defend against XSS attacks, various security measures can be implemented.

One approach to defending against reflected XSS attacks is to use a feature called the XSS auditor, which is built into some web browsers like Chrome. The XSS auditor works by parsing the HTML of a web page and comparing it to the URL. If it detects that the URL parameter is reflected in the page, it blocks the execution of the script. This helps protect websites from XSS attacks, even if the website itself is insecure.

However, the XSS auditor has limitations. It can produce false negatives, allowing some malicious scripts to bypass the filter. Attackers can encode the query in different ways to evade detection. Additionally, the XSS auditor also suffers from a false positive problem. It cannot distinguish between a truly reflected script and a script intentionally added by the website author. This can lead to legitimate scripts being blocked, impacting the functionality of the website.

To illustrate this, let's consider a scenario where a web page includes a script that the site author intended to run. If a user visits this page without any manipulated query strings, the XSS auditor will not interfere and the script will execute as intended. However, if an attacker attaches a query string that contains the same script, the XSS auditor may mistakenly identify it as a reflected XSS attack and block the script from running.

Despite its limitations, the XSS auditor can still provide a layer of defense against reflected XSS attacks. It is important to note that it should not be relied upon as the sole defense mechanism. Implementing multiple layers of security measures, such as input validation, output encoding, and secure coding practices, is crucial to effectively protect web applications against XSS attacks.

The XSS auditor is a feature in some web browsers that helps defend against reflected XSS attacks by comparing the URL to the HTML of a web page. While it provides a level of protection, it has limitations and should be used in conjunction with other security measures to ensure comprehensive web application security.

#### Web Applications Security Fundamentals - Cross-site scripting - Cross-Site Scripting Defenses

Cross-site scripting (XSS) is a type of security vulnerability that allows attackers to inject malicious scripts into web pages viewed by other users. These scripts can be used to steal sensitive information, manipulate web content, or redirect users to malicious websites. To defend against XSS attacks, various techniques and defenses have been developed.

One common defense mechanism against XSS attacks is frame busting. Frame busting is a technique that prevents a web page from being loaded within an iframe. It works by using JavaScript code that checks if the page is being loaded in a frame and, if so, redirects the browser to the main website. This defense mechanism is effective in preventing attackers from loading their malicious scripts within an iframe on their own website.

Another defense mechanism involves preventing the execution of malicious scripts on the target website. This can be achieved by modifying the URL of the target website and adding the script as a query parameter. By doing so, the script is not executed by the browser, as it is flagged as a potential XSS attack. This defense mechanism relies on the browser's built-in XSS protection feature, which detects and blocks potentially malicious scripts.

However, it is important to note that this defense mechanism may have unintended consequences. Some legitimate websites use query parameters to pass HTML content between pages. If the browser blocks all scripts in query parameters, it may break the functionality of these websites. Therefore, it is necessary to carefully

consider the impact of implementing this defense mechanism.

In the past, there have been debates and changes in the default behavior of XSS defenses. Initially, bypasses were found, leading to the decision to keep the defense mechanism in place despite its limitations. However, when it was discovered that attackers could selectively remove unwanted scripts from target websites, the severity of the issue was taken more seriously. The default behavior was then changed to block the entire page from loading if a script was detected. Eventually, the decision was made to remove the defense mechanism altogether, considering it a flawed approach.

It is worth mentioning that the behavior of XSS defenses can vary across different versions of web browsers. For instance, in Chrome Canary, the latest beta version of Chrome, the XSS defense mechanism has been removed. This means that malicious scripts are no longer filtered out, even when they are included in the URL. However, it is important to note that this is a beta version, and the final release may include different security measures.

Cross-site scripting (XSS) is a significant security vulnerability that can be exploited by attackers to inject malicious scripts into web pages. To defend against XSS attacks, various techniques and defenses have been developed, including frame busting and script blocking. However, it is essential to carefully consider the impact of these defenses and stay updated on the latest browser security features and updates.

Cross-site scripting (XSS) is a common web application security vulnerability that allows attackers to inject malicious scripts into web pages viewed by users. These scripts can then be executed by the victim's browser, leading to various harmful consequences such as data theft, session hijacking, or defacement of the website.

One type of XSS attack is known as cross-site scripting defenses. In this attack, the attacker injects malicious scripts into a vulnerable web page, which are then executed by the victim's browser. The scripts can be used to steal sensitive information, manipulate the website's content, or perform other malicious actions.

To defend against cross-site scripting attacks, web developers can implement various security measures. One common defense mechanism is called frame busting. Frame busting prevents a web page from being displayed within an iframe on another website. This can help protect against clickjacking attacks, where an attacker tries to trick users into clicking on hidden or disguised elements on a web page.

However, some modern web browsers, such as Chrome, have implemented a feature that prevents iframes from redirecting users. This can interfere with the frame busting defense mechanism. The reason behind this feature is to prevent ads from redirecting users to other websites without their consent. While this feature improves user security, it can also break legitimate frame busting code.

A more effective defense against cross-site scripting attacks is the use of HTTP headers. Web developers can include an HTTP header in their web pages that instructs browsers not to display the page within an iframe. This approach is more reliable and less prone to interference from browser features.

It is important to note that cross-site scripting attacks can target not only inline scripts but also external scripts loaded by web pages. By injecting malicious code into a vulnerable web page, attackers can manipulate the behavior of the website and potentially compromise user data or session information.

In addition to the frame busting and HTTP header defenses, there is another attack vector related to cross-site scripting. By manipulating the URL parameters of a web page, an attacker can trigger certain behaviors in the victim's browser. For example, if a specific string is present in the URL, the browser may block the page from loading. This can be used by attackers to detect the presence of certain strings or variables in the web page, potentially revealing sensitive information.

To summarize, cross-site scripting is a serious web application security vulnerability that can lead to various malicious activities. Web developers can defend against cross-site scripting attacks by implementing frame busting, using HTTP headers, and carefully validating and sanitizing user input to prevent script injection.

Cross-Site Scripting (XSS) is a common vulnerability in web applications that allows attackers to inject malicious scripts into trusted websites. To protect against XSS attacks, it is important to understand the fundamentals of web application security and the defenses that can be implemented.

One way to defend against XSS attacks is by using the XSS auditor, a feature in web browsers that detects and blocks potential XSS vulnerabilities. However, the XSS auditor has its limitations and is being phased out. Therefore, it is crucial to explore alternative methods of protecting web applications.

One such method is Content Security Policy (CSP), which is the reverse of the same origin policy. CSP allows website owners to specify which servers their web applications are allowed to communicate with. By limiting communication to a whitelist of trusted servers, CSP effectively prevents attackers from exfiltrating data and restricts their ability to exploit cross-site scripting vulnerabilities.

To enforce CSP, an HTTP header is used. This header is sent by the server along with the HTML response and instructs the browser to enforce the specified policy for all code running on the page. Since the attacker may already have code running on the page, it is essential that the policy cannot be modified by JavaScript. By including the policy in the HTTP header, the server ensures that the policy is set and cannot be altered by the attacker.

CSP works by blocking any requests that violate the specified policy. For example, a simple policy that restricts communication to the same origin can be implemented by adding a specific header to the server's response. This effectively limits an attacker's ability to communicate with other servers and reduces the potential harm they can cause.

However, it is important to note that even with CSP, an attacker's code can still perform actions on behalf of the user within the same site. To mitigate this risk, it is crucial to ensure that sensitive data is protected and that any inline scripts are carefully reviewed. By default, CSP prevents all inline scripts from running, which adds an extra layer of defense against XSS attacks. However, it is essential to thoroughly test and review the web application to ensure that legitimate scripts are not inadvertently blocked.

Content Security Policy (CSP) is a powerful defense mechanism against cross-site scripting (XSS) attacks. By specifying a policy in the HTTP header, web application owners can restrict communication to trusted servers and prevent attackers from exploiting XSS vulnerabilities. However, it is important to carefully review and test the web application to ensure that legitimate scripts are not blocked.

Cross-site scripting (XSS) is a common web application security vulnerability that allows attackers to inject malicious scripts into web pages viewed by other users. These scripts can be used to steal sensitive information, manipulate content, or perform other malicious actions.

To defend against XSS attacks, web developers can implement various defenses. One common defense is to sanitize user input by escaping special characters and validating data before displaying it on web pages. This helps prevent the execution of injected scripts.

Another defense is Content Security Policy (CSP), which allows web developers to specify the types of content that can be loaded and executed on their web pages. By setting a CSP header, developers can define trusted sources for scripts, stylesheets, images, and other resources, effectively blocking unauthorized content from being loaded.

For example, if a web page allows user input to be displayed, the developer can use CSP to explicitly deny the execution of scripts in that context. By doing so, even if an attacker manages to inject a script, it will not be executed.

CSP also allows developers to specify trusted sources for loading scripts. This can be useful when working with third-party scripts or subdomains. By adding trusted sources to the CSP policy, developers can ensure that only scripts from those sources are allowed to run.

It is important to note that deploying CSP requires careful configuration and testing. If a policy is misconfigured, it can potentially break the functionality of a website for all users. To avoid this, developers can use CSP in report-only mode initially. In this mode, violations are reported to a specified URL, allowing developers to review and adjust the policy before deploying it in production.

Additionally, developers can include a report URL in the CSP header to receive notifications whenever content is blocked by the policy. This helps identify any missed sources or potential attacks in real time.



There are several directives available in CSP, including default source, image source, object source, and script source. These directives allow developers to specify trusted sources for different types of content. For example, the script source directive is particularly important for preventing cross-site scripting attacks, and developers should be cautious when configuring it.

Cross-site scripting is a significant security concern for web applications. To defend against XSS attacks, developers can implement various measures, including input validation, sanitization, and the use of Content Security Policy. By following best practices and configuring CSP correctly, developers can significantly mitigate the risk of XSS attacks.

Web applications security is a crucial aspect of cybersecurity. One common vulnerability that attackers exploit is cross-site scripting (XSS). XSS occurs when an attacker injects malicious code into a website, which is then executed by the victim's browser. To protect against XSS attacks, it is important to understand the different types of XSS and implement appropriate defenses.

One type of XSS is called cross-site scripting defenses. In this type, the attacker uses the "base" tag in HTML to manipulate the relative URLs on a website. By setting the base tag to their own domain, the attacker can redirect requests to their own server instead of the intended destination. To prevent this, it is recommended to have a policy that restricts the use of the base tag or prevents it altogether.

Another XSS defense is the "frame ancestors" tag. This tag allows website owners to specify whether their site can be framed by other websites. Although there is already a tag for this purpose, the frame ancestors tag provides a newer way to achieve the same result.

The "upgrade and secure requests" tag is another interesting defense mechanism. It is useful when migrating an old website to a more secure version. By applying this tag, all HTTP requests made by the website will be automatically upgraded to HTTPS. This ensures that sensitive data is transmitted securely. These defenses are part of the content security policy (CSP).

When deploying a CSP, it is important to consider the presence of inline code. Inline code refers to scripts embedded directly within HTML documents. By default, CSP policies do not allow inline code execution, as it poses a security risk. However, this can cause issues when certain scripts, such as those from Google Analytics, rely on inline code. To address this, the "unsafe-inline" directive can be added to the CSP policy, allowing inline scripts. However, this compromises security and should be used with caution.

Another challenge with CSP policies is finding the right balance between security and functionality. Overly restrictive policies can break websites, as certain resources may be blocked. For example, if a CSP policy only allows images to be loaded from the website's own origin, it can prevent Google Analytics from functioning properly. In such cases, specific exceptions need to be made in the policy to allow resources from trusted sources.

Cross-site scripting defenses are essential in protecting web applications from XSS attacks. Understanding the different types of XSS and implementing appropriate defenses, such as managing the base tag, using the frame ancestors tag, and upgrading and securing requests, can greatly enhance the security of web applications. However, it is important to carefully consider the impact of CSP policies on website functionality and make necessary exceptions to ensure a balance between security and usability.

Cross-site scripting (XSS) is a common vulnerability in web applications that allows attackers to inject malicious scripts into trusted websites. These scripts can then be executed by unsuspecting users, leading to various security risks such as data theft, session hijacking, or defacement of the website.

To defend against XSS attacks, web developers can implement Content Security Policy (CSP) as a security measure. CSP is a set of directives that instruct the browser on what types of content are allowed to be loaded and executed on a web page. By defining a strict policy, developers can mitigate the risk of XSS attacks by blocking the execution of any unauthorized scripts.

However, implementing an effective CSP policy can be challenging. The transcript highlights some of the difficulties and limitations associated with CSP. One issue is that the policy is dependent on the behavior of the

scripts running on the website. If the script's behavior changes or if new scripts are added, the policy may break, rendering the website vulnerable to XSS attacks.

To address this challenge, a solution proposed in a research paper called "CSP is dead, long live CSP" by Google researchers is to propagate trust from the initial script to any scripts included at runtime. This means that if the initial script is trusted, any scripts it includes, regardless of their source, are implicitly trusted as well. This approach aims to ensure that the CSP policy remains effective even when new scripts are added to the website.

The research paper also highlights the shortcomings of existing CSP implementations. It reveals that a significant number of websites that have deployed CSP still contain unsafe endpoints, allowing attackers to bypass the policy. Additionally, many CSP policies that attempt to limit script execution are found to be ineffective against XSS attacks.

However, the paper proposes a solution that has been successfully deployed and proven to be effective in mitigating XSS attacks. Although the details of the solution are not provided in the transcript, it offers hope that a reliable defense against XSS can be achieved.

Implementing an effective CSP policy is crucial for protecting web applications from XSS attacks. While there are challenges and limitations associated with CSP, ongoing research and development are continuously improving its effectiveness. By staying informed about the latest advancements in web application security and following best practices, developers can enhance the security posture of their web applications and safeguard sensitive user data.

Cross-site scripting (XSS) is a common vulnerability in web applications that allows attackers to inject malicious scripts into trusted websites. This can lead to various security issues, including unauthorized access to user data and the execution of arbitrary code.

One type of XSS attack is called "echoing," where an attacker exploits a vulnerable input field or parameter that echoes user-supplied data without proper sanitization. If the attacker can input valid JavaScript code, it will be executed on the website, giving them control over the site.

Another example involves AngularJS, a popular JavaScript framework. AngularJS allows developers to include template code that is dynamically executed. However, if an attacker can manipulate the page's content, such as adding braces in a div element, AngularJS may interpret it as code and execute it. This can lead to arbitrary code execution on the site.

A more complex example involves exploiting a server's error page. If the error page echoes user-supplied data without proper validation, an attacker can inject JavaScript code disguised as a URL. By naming a line of code and using a go-to statement, the attacker can effectively execute their malicious code on the site.

Even seemingly harmless data, such as CSV files, can be used to inject JavaScript code. If the attacker can control the data inside the CSV file, they can include JavaScript code that will be executed on the site.

Implementing Content Security Policy (CSP) is a recommended defense against XSS attacks. However, CSP can be challenging to configure correctly. The website "uselesscsp.com" showcases examples of companies, including Apple, that have failed to implement CSP effectively.

One proposed solution to address these challenges is called "strict dynamic." Instead of explicitly listing all trusted domains in the script policy, strict dynamic allows implicit trust of certain domains. These trusted domains can then load any code they want, cascading the trust to other loaded scripts. This approach helps mitigate the risk of XSS attacks.

To ensure the trustworthiness of scripts loaded from trusted domains, a nonce can be used. A nonce is a random value generated by the server and included in the HTTP header. By including the nonce in the script tag, the browser will only execute scripts that have a matching nonce value, preventing the execution of attacker-injected code.

Cross-site scripting (XSS) poses significant security risks to web applications. Attackers can exploit vulnerabilities to inject and execute malicious scripts. Implementing proper defenses, such as Content Security

Policy (CSP) with strict dynamic and nonces, can help mitigate the risk of XSS attacks.

In the context of web application security, one of the common vulnerabilities is Cross-Site Scripting (XSS). XSS occurs when an attacker injects malicious scripts into a trusted website, which then gets executed by the users' browsers. To mitigate this risk, web developers need to implement proper defenses.

One effective defense against XSS is the use of nonces (number used once). Nonces are random values generated by the server and included in both the HTTP header and the HTML page. The browser is instructed to only execute scripts that include the nonce value as an attribute. This means that any injected script without the correct nonce will not be executed.

To implement this defense, the server includes a nonce value in the HTTP header and in the HTML page. The attacker, however, cannot easily access these values. They cannot view the server's response headers, nor can they access the nonce value in an HTML page from a different origin. This makes it challenging for attackers to determine the nonce value.

Furthermore, to make the defense even stronger, the server should generate a different nonce value for each page request. By using a completely random and unpredictable nonce, it becomes impractical for attackers to guess or iterate over all possible nonce values.

It is important to note that HTML responses should not be cached when using this defense, as the nonce value needs to be different for each request. Additionally, it is crucial to choose a nonce value that is long enough to prevent attackers from easily guessing it.

However, it is worth mentioning that some browsers, such as Safari, do not yet support the strict dynamic attribute, which is a part of this defense. In such cases, developers need to decide how to handle Safari users. One option is to include other unsafe properties in the content security policy, allowing scripts to be loaded from any location and letting code run more freely. However, this approach should be carefully considered, as it may introduce additional security risks.

Using nonces as a defense against Cross-Site Scripting attacks is an effective strategy. By generating random and unpredictable nonce values for each page request, web developers can significantly reduce the risk of XSS vulnerabilities. However, it is crucial to consider browser compatibility and make appropriate decisions when certain browsers do not support all aspects of this defense.

Web applications security is a crucial aspect of cybersecurity. One common vulnerability in web applications is cross-site scripting (XSS). XSS occurs when an attacker injects malicious code into a website, which is then executed by unsuspecting users. In this didactic material, we will focus on the fundamentals of XSS and explore various defenses against it.

One defense mechanism against XSS is Content Security Policy (CSP). CSP is an HTTP response header that allows website administrators to control which resources can be loaded and executed by a web page. By specifying a CSP, administrators can restrict the types of content that can be loaded, thereby mitigating the risk of XSS attacks.

One aspect of CSP is the use of nonces and strict dynamic policies. Nonces are random values that are generated by the server and included in the CSP header. When a script is loaded, the browser checks if the nonce matches the one specified in the CSP. If it does, the script is allowed to execute. This prevents the execution of injected scripts that do not have the correct nonce.

Strict dynamic policies, on the other hand, allow scripts to be loaded and executed only if they have been explicitly whitelisted in the CSP. This ensures that only trusted scripts are allowed to run, further reducing the risk of XSS attacks. However, it is important to note that strict dynamic policies may not be supported by all browsers.

Another defense mechanism against XSS is the use of a feature policy. Feature policy allows website administrators to disable certain browser functionalities that may be exploited by attackers. For example, the geolocation API can be disabled to prevent an attacker from accessing a user's location. By selectively disabling unnecessary features, the potential damage that an attacker can cause is significantly reduced.

It is also worth mentioning the concept of dom-based XSS. Unlike reflected or stored XSS, where the attacker modifies the HTML returned by the server, dom-based XSS occurs when the attacker manipulates the DOM (Document Object Model) at runtime. This can be done by tricking a trusted script into adding malicious DOM nodes to the page. To mitigate dom-based XSS, developers should avoid using `innerHTML` to add user-generated content to the page, and instead, use `textContent`, which treats the content as plain text and automatically escapes any HTML tags.

Web applications security is a critical aspect of cybersecurity. Cross-site scripting (XSS) is a common vulnerability that can be mitigated through various defenses, such as Content Security Policy (CSP) and feature policies. By implementing these security measures, website administrators can significantly reduce the risk of XSS attacks and protect their users' data and privacy.

Content Security Policy (CSP) is a security mechanism that protects web applications from cross-site scripting (XSS) attacks. It specifically defends against reflected and stored XSS attacks. However, CSP does not provide protection against DOM-based XSS attacks. To address this, the Trusted Types API can be used in conjunction with CSP.

The idea behind DOM-based XSS defenses is to prevent the execution of potentially malicious code by disallowing the direct assignment of HTML strings to the `innerHTML` property. Instead, the use of the Trusted Types API is enforced, where `innerHTML` assignments should only accept trusted HTML objects and fail if passed a string.

To implement this defense, a template factory is created. The factory includes a `createHTML` function that escapes user input. Although the `createHTML` function appears to return a string, it actually returns a trusted HTML object. This ensures that any HTML code added to the page must flow through this factory, preventing unauthorized DOM-based XSS attacks throughout the codebase.

This approach simplifies code auditing as the focus is on ensuring the security of the template factory function. By verifying the safety of this function, it can be concluded that there are no other ways for DOM-based XSS to occur on the site.

To further enhance web application security, it is important to be vigilant and never trust any data from the client. All user input should be sanitized and appropriately escaped depending on the context in which it is used. Additionally, correctly implementing CSP with strict dynamic and utilizing trusted types can effectively prevent various types of XSS attacks.

By combining these defenses, web applications can significantly reduce the risk of XSS vulnerabilities. It is crucial to maintain a constant state of vigilance and adopt a proactive mindset towards cybersecurity.

**EITC/IS/WASF WEB APPLICATIONS SECURITY FUNDAMENTALS DIDACTIC MATERIALS****LESSON: WEB FINGERPRINTING****TOPIC: FINGERPRINTING AND PRIVACY ON THE WEB**

In this didactic material, we will explore the topic of web fingerprinting and its implications for privacy on the web. We will begin by understanding the reasons why websites track user activity and how this has become a prevalent practice. Then, we will delve into the concept of fingerprinting and how it has emerged as an alternative to traditional cookie-based tracking. Finally, we will discuss countermeasures and solutions to address fingerprinting and its impact on user privacy.

Websites track user activity for various reasons, including targeted advertising, personalization, and analytics. This practice allows websites to gather data on user behavior, preferences, and demographics, which can be used to deliver more relevant content and advertisements. However, this tracking has raised concerns about privacy violations and the potential misuse of personal information.

Fingerprinting is a technique used to track users without relying on cookies or other identifiable information. It involves collecting various data points, such as browser and device characteristics, network information, and installed plugins, to create a unique identifier or "fingerprint" for each user. This fingerprint can then be used to track user activity across different websites, even if cookies are cleared or disabled.

The rise of fingerprinting has posed new challenges for user privacy. Unlike cookies, which can be easily managed and controlled by users, fingerprinting is more difficult to detect and block. Additionally, fingerprinting can be used to identify and track users across multiple devices, further eroding privacy.

To address the concerns raised by fingerprinting, various countermeasures have been developed. These include browser extensions and privacy-focused browsers, such as Brave, which aim to block or limit fingerprinting techniques. Additionally, browser vendors and standardization bodies are working on implementing privacy-enhancing features and guidelines to mitigate the impact of fingerprinting.

In order to better understand the current state of fingerprinting and its impact on user privacy, it is important to identify and address the vulnerabilities and surface areas exposed in the browser. This involves analyzing the different techniques used for fingerprinting and exploring potential solutions to protect user privacy.

Brave, a privacy-focused browser, is actively working towards addressing fingerprinting and protecting user privacy. In addition to its privacy-preserving features, Brave proposes an alternative approach to funding websites through its advertising model. Users are incentivized to view ads, and a portion of the payments received by Brave is used to fund websites, creating a more balanced incentive structure.

Web fingerprinting has become a significant concern for user privacy on the web. Websites track user activity for various purposes, but this practice has raised concerns about privacy violations. Fingerprinting has emerged as an alternative to traditional tracking methods, posing new challenges for user privacy. However, efforts are being made to develop countermeasures and solutions to address fingerprinting and protect user privacy.

Tracking exists in the web for the purpose of understanding user behavior and preferences in order to deliver targeted advertisements. In the past, tracking was primarily focused on tracking content rather than individual users. Advertisers would place ads next to specific content that they believed would attract the target audience. This model involved matching ads to content, not to specific individuals.

Even before the internet, there were efforts to track users. For example, if someone signed up for a magazine subscription, their information could be sold to advertisers who wanted to target people with similar interests. However, with the rise of the internet, tracking users across different websites became more prevalent.

Initially, websites would form alliances called web rings, where they would advertise each other's sites based on shared interests. This allowed for some level of tracking user interests across websites. However, as technology advanced, tracking became more sophisticated and third-party tracking services emerged.

In the current web ecosystem, ads are no longer matched solely based on the content of a website. Instead, they are matched based on the interests and preferences of individual users. When a browser requests content

from a website, information about the user, such as their purchasing history, location, and name, is sent along with the request. The third-party tracking service then decides which ad to display based on this information.

This shift towards tracking individuals rather than just content has led to the creation of massive databases by third-party tracking services. These databases contain information about users' browsing habits and preferences across multiple websites. This allows advertisers to target specific individuals with ads, regardless of the quality of the website they are visiting.

As a result, users may encounter ads that are not relevant or of low quality on websites that they visit. The quality of the ad is no longer directly connected to the quality of the content. This has contributed to the current state of the web, where users often experience a poor browsing experience due to the prevalence of irrelevant or low-quality ads.

Tracking on the web has evolved from tracking content to tracking individuals. This shift has allowed advertisers to target specific users with ads, leading to the creation of large databases containing user information. However, it has also resulted in a decline in the quality of ads and the overall browsing experience.

In the realm of web applications security, one important aspect to consider is web fingerprinting. Web fingerprinting refers to the process of identifying and tracking users based on their unique online characteristics. This technique has significant implications for privacy on the web.

Traditionally, advertisers would place ads on high-quality websites to reach their target audience. However, with the advent of web fingerprinting, advertisers can now target users directly, regardless of the website they visit. This has resulted in a shift of revenues from high-quality websites to low-quality ones, as advertisers no longer have to rely on website quality for effective ad placement.

The impact of this shift is twofold. Firstly, high-quality websites no longer have the same incentive to display ads, as advertisers can reach their desired audience on various low-quality websites. Consequently, these high-quality websites suffer a loss of advertising revenue. Secondly, the content creators on high-quality websites are negatively affected, as the revenue generated from ads decreases.

In theory, advertisers should allocate their ad spend budget proportionally based on the distribution of web usage. For example, if 50% of users visit the Chicago Tribune and 50% visit other websites, advertisers should allocate 50% of their budget to the Chicago Tribune and 50% to other websites. However, in practice, this is not the case. Placing ads on high-quality websites, such as the Chicago Tribune, is significantly more expensive than on low-quality websites. As a result, advertisers find it more cost-effective to target crummy websites, where their advertising dollars can go much further.

This dynamic explains why websites like [breitbart.com](http://breitbart.com), which may have less reporting and original work, can still generate substantial revenue. It is not necessarily a reflection of political preferences, but rather a consequence of the current advertising system and web fingerprinting techniques.

Furthermore, web fingerprinting has led to insidious downstream effects. Advertisers can track users' online behavior and target them with ads on various websites, even if those websites are unrelated to the user's original interests. This tracking can have unintended consequences, such as revealing personal information to unintended parties. For instance, there have been cases where target ads inadvertently exposed a teenager's pregnancy to her father before she had the chance to share the news herself.

It is important to note that web fingerprinting involves numerous middle parties that collect and analyze user data to build profiles. Many of these middle parties are not profitable, except for a few major players like Google and Facebook. This highlights the extensive presence of entities solely focused on gathering user information.

Web fingerprinting and its impact on web applications security and privacy have significant consequences. It has shifted advertising revenues from high-quality websites to low-quality ones, negatively affecting both high-quality content creators and the overall quality of online content. Additionally, the tracking enabled by web fingerprinting techniques has raised concerns about unintended disclosure of personal information. Understanding these dynamics is crucial for comprehending the current state of online advertising and privacy.

Web fingerprinting is a technique used to track users' online activities by collecting information about their web



browsers, devices, and behavior. In a study conducted by Stephen Englehart at Princeton, researchers visited a million websites to measure the extent of web tracking. They categorized the third-party resources requested by these websites and found that even the average website had connections to almost 20 different tracking parties.

These tracking parties can be thought of as databases that record information about users' website visits. While Google Analytics emerged as the most popular tracker, accounting for over 80% of the websites visited, there were other profitable advertising companies that also engaged in tracking. These companies had access to users' information but were not categorized as trackers.

It is important to understand the concept of first-party and third-party in web fingerprinting. The first-party refers to the top-level URL frame, which is usually the domain name entered by the user. The third-party refers to any other page or resource that is loaded within the first-party context. Sometimes, third-party resources can be considered first-party depending on the measurement being taken.

While some organizations explicitly guarantee privacy and refrain from tracking, others have the technological capability to track but are not believed to engage in tracking activities. However, it is worth noting that tracking is a prevalent practice on the web, with numerous parties collecting user information.

This study sheds light on the extent of web fingerprinting and highlights the need for users to be aware of the tracking practices employed by various websites. Understanding web fingerprinting can help users make informed decisions about their online privacy and take necessary precautions to protect their personal information.

Web fingerprinting is a technique used to track and identify users based on their unique browser and device characteristics. It involves collecting information about a user's browser, operating system, plugins, and other attributes that can be used to create a unique identifier, or fingerprint.

One common method of web fingerprinting is through the use of third-party resources, such as Google Analytics. When a website includes a third-party resource like Google Analytics, it can determine the first-party website that the user is visiting and record their behaviors. This means that even though Google Analytics is a third-party resource, it has access to information about the first-party website.

It is important to note that web fingerprinting can be carried out by various means, and the common case is often facilitated by choices made in the history of the web. For example, Brendan Eich, the inventor of JavaScript, made choices that unintentionally enabled certain tracking mechanisms. These mechanisms were not the intended purpose of the web, but rather a result of decisions made under pressure during the early days of the web.

In the early days of the web, requests made to servers did not carry any state information. Each request was independent, and there was no way to carry authentication or user-specific information between requests. To address this issue, the concept of cookies was introduced. Cookies are values that a server gives to a user's browser, which are then returned on subsequent requests. This allows the server to recognize the user and provide personalized content or access to restricted resources.

Cookies serve as a way of transmitting tokens across browser views or page views, allowing the server to track and identify users. When a user visits a website, the server sends a cookie identifier along with the content. The user's browser then returns this identifier on future requests, allowing the server to associate the request with the specific user. If someone else tries to access the same resource without the proper cookie identifier, the server will reject the request.

Another factor that enables web fingerprinting is the use of multiple hosts for hosting resources like images. In the past, hosting resources was expensive, and websites would often host the same image on multiple domains to distribute the load. By referring to images across domains, websites can track users across different websites by linking the requests for those images.

Web fingerprinting is a powerful technique that allows websites and third-party resources to track and identify users based on their unique browser and device characteristics. It has become an integral part of online tracking and has implications for user privacy and security.

Web fingerprinting is a technique used to identify and track users on the web by analyzing unique characteristics of their browsers. Unlike traditional tracking methods that rely on storing values, fingerprinting focuses on gathering information about the browser itself.

One common method of fingerprinting is analyzing the size of the browser window. Each user may have a different window size, allowing websites to differentiate between them. Other characteristics that can be used for fingerprinting include the user's operating system, browser version, installed fonts, and the presence of certain plugins or extensions.

Fingerprinting allows websites to create a unique identifier for each user, even if they clear their cookies or use different devices. This enables tracking across multiple websites and poses a significant threat to user privacy.

To mitigate the risks associated with fingerprinting, several countermeasures have been implemented by different browsers. Safari and Firefox, for example, have taken a strong stance against third-party cookies, limiting their usage. Brave, a privacy-focused browser, also incorporates countermeasures to protect against fingerprinting.

However, as websites become more sophisticated, they find new ways to track users. Some sites have started using URL parameters or alternative storage methods, such as local storage, to avoid relying on cookies. Additionally, techniques like HTTP Strict Transport Security (HSTS) have been co-opted for tracking purposes. Advertisers realized that by creating multiple subdomains with different HSTS instructions, they could uniquely identify users.

Web fingerprinting is a complex and evolving issue in the realm of cybersecurity. It is crucial for users to stay informed about the different tracking methods employed by websites and to take steps to protect their privacy.

Web fingerprinting is a technique used to identify and track individual users on the internet. It involves collecting various pieces of information about a user's web browser and device, and combining them to create a unique identifier. This identifier, or fingerprint, can then be used to track the user's online activities across different websites.

There are several factors that contribute to the uniqueness of a web fingerprint. These include the browser type and version, operating system, installed plugins and fonts, and other browser settings. Individually, these factors may not be very identifying, as many people may have similar configurations. However, when combined together, they create a set of identifiers that can be used to distinguish one user from another.

To understand how fingerprinting works, let's consider an example. Imagine there are billions of internet users, and you are one among them. The challenge for websites and trackers is to identify you specifically, without storing any personal information. They achieve this by collecting and analyzing various pieces of information about your browser and device.

For example, let's look at the user-agent string, which is a part of the HTTP request sent by your browser to a website. This string contains information about the browser type, version, and other details. In most browsers, this string includes multiple identifiers, such as "Mozilla", "KHTML", and "Chrome". While individually these identifiers may not be very identifying, the combination of them, along with other factors, can create a unique fingerprint.

Another example is the use of fonts on a web page. Web browsers use a complex algorithm to match the fonts specified in the CSS with the fonts installed on your system. By analyzing the fonts used on a web page, it is possible to create a fingerprint that can be used to identify you.

To be successful at fingerprinting, trackers need a large number of semi-unique identifiers. These can include browser settings, installed plugins, and other factors that may vary among users. On the other hand, protecting against fingerprinting involves minimizing the number of identifying identifiers that can be collected.

Web fingerprinting is a technique used to identify and track individual users on the internet. It involves collecting various pieces of information about a user's browser and device, and combining them to create a unique identifier. By understanding how fingerprinting works, users can take steps to protect their privacy.

online.

Web fingerprinting is a technique used by websites to gather information about users' devices and browsers. One aspect of web fingerprinting is font fingerprinting, which involves extracting information about the fonts installed on a user's system. Websites can then use this information to uniquely identify users based on their font configurations.

Although there is no direct API to retrieve a list of installed fonts, websites can indirectly extract this information by creating a page element and applying different fonts to it. By observing changes in the size of the text element, websites can determine whether a user has a specific font installed or not. This method relies on the fact that each user's font configuration is unique, as they may have additional fonts installed that are not commonly found on other systems.

Defending against font fingerprinting is challenging because it is difficult to prevent websites from reading the width of elements. Additionally, imposing restrictions on font usage can negatively impact the user experience. Browsers have attempted to mitigate font fingerprinting through various countermeasures, but finding practical solutions remains a complex problem.

Another aspect related to web fingerprinting is cookie syncing, which involves linking multiple cookies to identify the same user. Cookie syncing often utilizes font fingerprinting as one of the fingerprinting methods. This process presents significant challenges in terms of web privacy protection.

There are ongoing efforts to address font fingerprinting through standardization. Fixing the standards would be beneficial, as it would provide a comprehensive solution across different browsers. However, implementing these fixes in browsers like Chromium can be challenging due to resistance from certain teams. Nonetheless, some browsers, such as Mozilla and Apple, prioritize privacy and are more inclined to adopt privacy-preserving measures.

In the meantime, alternative approaches have been proposed. One approach is to consider local fonts as empty, which may work for most English-language and Western websites. However, this approach may break websites that rely on specific fonts, particularly in Asian regions. Another alternative is to establish consistent mappings between requested languages and possible local fonts. By limiting the number of fonts associated with each language, the privacy of users can be better protected. Additionally, requiring users to opt-in for fonts they want to use on the web can also be considered as a solution.

Font fingerprinting is a technique used by websites to gather information about users' font configurations. Defending against font fingerprinting is challenging, and finding practical countermeasures remains a complex problem. Efforts are being made to address this issue through standardization and alternative approaches.

### Web Fingerprinting and Privacy on the Web

Web fingerprinting is a technique used to uniquely identify users based on their browser characteristics. It involves gathering information about the user's browser, device, and network to create a unique identifier, or fingerprint, that can be used to track their online activities. This can be a concern for user privacy, as it allows websites and third parties to gather information about users without their knowledge or consent.

One aspect of web fingerprinting is web fingerprinting through fonts. When a user visits a website, the website can request a list of fonts installed on the user's device. This information can be used to create a unique fingerprint for the user. To mitigate this, some propose using a standard set of fonts across all devices, so every user appears the same. However, this approach has its limitations. If a website uses a large number of fonts, it can significantly slow down the user's browsing experience, especially on mobile devices. Additionally, requesting a large number of fonts can consume a significant amount of data, potentially costing the user money.

Another method of web fingerprinting is through the Canvas API. The Canvas API allows websites to draw and manipulate graphics on the user's browser. By instructing the browser to perform specific drawing operations and then reading the results, subtle differences in how different hardware platforms and browsers render the graphics can be observed. These differences can be used to create a unique fingerprint for the user. This technique is particularly interesting because it exploits an API that was not intended for malicious purposes. The

Canvas API was primarily designed for drawing graphics and was not meant to be used for tracking users. However, some websites use this API for legitimate purposes, such as detecting emoji support.

To better understand how web fingerprinting through the Canvas API works, researchers have conducted experiments to compare the rendering of graphics across different platforms and graphics cards. These experiments reveal subtle differences in how graphics are rendered, which can be used to create unique fingerprints. By analyzing these differences, researchers can identify users with a high degree of accuracy.

Web fingerprinting is a technique used to uniquely identify users based on their browser characteristics. It can be done through various methods, such as font fingerprinting and Canvas API fingerprinting. These techniques raise concerns about user privacy, as they allow websites and third parties to track users without their knowledge or consent. Mitigating web fingerprinting is challenging, as some proposed solutions can negatively impact user experience or have limitations. It is important for users to be aware of these techniques and take steps to protect their privacy online.

Web fingerprinting is a technique used to identify and track users on the web based on unique characteristics of their devices or browsers. One of the most identifying features of a browser is its height and width, which can be used for fingerprinting purposes. Extracting this information can be challenging, especially if JavaScript cannot be executed.

There are a few ways to obtain the height and width of a browser window without using JavaScript. One simple approach is to access the "window" object and retrieve the height and width properties. However, this method can be easily manipulated, making it unreliable for fingerprinting.

To overcome this limitation, alternative methods can be used. For example, determining if an image is being displayed or analyzing the text flow can provide insights into the actual height and width of the browser window. These factors are more difficult to falsify, making them more reliable for fingerprinting purposes.

The Electronic Frontier Foundation (EFF), a nonprofit organization based in San Francisco, has been actively working on privacy-related issues. They have developed tools like Privacy Badger and have contributed to projects like HTTPS Everywhere. In one of their papers called "Panopticon," they explore various identifiers used for web fingerprinting and assess their level of identifiability.

While web fingerprinting techniques can be used for legitimate purposes, such as enhancing user experience and security, they also raise concerns about privacy. Currently, websites can access this information without requiring any permissions, which poses a potential risk to user privacy.

Efforts are being made to address these privacy concerns. Organizations like the EFF are actively working on technical solutions to mitigate the identifiability of browsers and devices. However, finding a balance between privacy and functionality remains a challenge.

Web fingerprinting techniques, including the extraction of height and width information, can be used to identify and track users on the web. While alternative methods can be employed to obtain this information without relying on JavaScript, the reliability and accuracy of such methods are still being explored. Privacy-conscious organizations like the EFF are working on solutions to protect user privacy while maintaining the necessary functionality of web applications.

Web fingerprinting is a technique used to uniquely identify users based on various characteristics of their web browser. This process involves collecting information about the user's browser configuration, such as the installed fonts, screen resolution, and user-agent string, among others, and using this information to create a unique identifier or fingerprint.

One important point to note is that web fingerprinting is often used as a benchmark to measure browser privacy. When websites or browsers claim to improve privacy, they are usually referring to mitigating the effectiveness of fingerprinting techniques.

In practice, fingerprinting code typically hashes the collected information to create a unique value for each characteristic. For example, the number of fonts installed and the screen resolution may be hashed separately and then combined to create a final fingerprint. These fingerprints are often stored in a database for future

reference.

One countermeasure against web fingerprinting is to deliberately introduce noise or false information into the fingerprinting process. By lying about one or more characteristics, it is possible to generate a different fingerprint each time, effectively rendering the fingerprinting technique ineffective. However, it is important to note that this countermeasure may not always be successful, especially against more sophisticated fingerprinting implementations.

To gain a better understanding of web fingerprinting, it is recommended to explore the fingerprintjs library. This library is widely used for fingerprinting and is available in both open-source and commercial versions. By examining the code of this library, it is possible to identify various fingerprinting approaches and mechanisms commonly used in the real world.

Some examples of fingerprinting approaches that can be found in fingerprintjs include:

1. Gyroscope-based fingerprinting: By utilizing the gyroscope sensor in a device, it is possible to determine the orientation of the device, which can be used as a unique identifier.
2. Audio-based fingerprinting: This approach involves generating audio and analyzing the subtle differences in the generated waveform to create a unique fingerprint. Different audio cards may produce slightly different waveforms, leading to distinct fingerprints.
3. WebRTC-based fingerprinting: WebRTC allows websites to access information about the audio and video devices connected to a user's device. This information, such as device labels or unique IDs, can be used to create a fingerprint.
4. Plugin-based fingerprinting: If a user has plugins installed, such as Java or Flash, these plugins can be used as identifiers. Older versions or less common plugins may be particularly effective in uniquely identifying users.
5. IP address-based fingerprinting: Although not directly related to web fingerprinting, IP addresses can also be used as identifiers. Companies that provide IP-to-geographic address mapping services can use this information to create fingerprints based on the user's location. However, this technique has become less reliable with the introduction of IPv6 and the depletion of IPv4 addresses.

Web fingerprinting is a technique used to uniquely identify users based on various characteristics of their web browser. By understanding the different approaches and countermeasures involved in web fingerprinting, it is possible to better protect user privacy and mitigate the effectiveness of these techniques.

Web fingerprinting is a technique used to track users' online activities by collecting unique information about their devices and browsers. One approach to web fingerprinting is IP cookies, which involves tracking users based on their unique IP addresses. However, this method becomes difficult when users employ privacy tools like VPNs or the Tor network.

Another method of web fingerprinting involves identifying users based on inconsistencies in the information provided by privacy-preserving tools. For example, if a user installs a canvas fingerprinting protection extension, it may provide inconsistent information about the user's preferred language or location, making it easier to identify them.

There are several ways to combat web fingerprinting. One approach is to remove the functionality that allows fingerprinting, such as disabling canvas drawing or font reading. Another approach is to make the functionality consistent across different browsers, reducing the uniqueness of each user's fingerprint. A third approach is to limit access to fingerprinting data, allowing only first-party sites to access certain information.

Noise or randomization is another promising method to combat web fingerprinting. By introducing random elements into fingerprinting data, it becomes more difficult for trackers to accurately identify users. Additionally, Google has proposed a concept called privacy budget, which allows users to perform certain actions for a limited time before restricting further access. However, this approach has raised concerns about privacy implications.

While removing certain APIs or making functionality consistent can be effective in combating web fingerprinting, it may also break legitimate use cases and negatively impact user experience. Therefore, finding a balance between preserving functionality and protecting user privacy is crucial.

Web fingerprinting is a complex issue with various methods used to track users' online activities. Different approaches, such as removing functionality, making it consistent, limiting access, introducing noise, or implementing a privacy budget, are being explored to combat this issue. However, finding the right balance between functionality and privacy remains a challenge.

Web fingerprinting is a technique used to track and identify users based on their unique browser characteristics. It involves collecting various data points such as browser version, screen resolution, installed fonts, and plugins to create a digital fingerprint that can be used to identify individuals across different websites. While web fingerprinting has legitimate uses such as fraud prevention and security, it also raises concerns about privacy.

One approach to address privacy concerns is through permission prompts. When a website wants to access certain functionalities or collect specific data, the browser prompts the user for permission. However, this approach can lead to permission prompt fatigue, where users quickly grant permissions without fully understanding the implications. It is also not a scalable solution to cover all fingerprinting endpoints.

User gestures play a crucial role in determining permissions. User gestures refer to standardized actions such as clicking or hovering on a webpage, indicating the user's intention to interact with the page. By giving functionality access only to frames with user gestures, certain privacy risks can be mitigated.

Another consideration is the distinction between first-party and third-party entities. First-party entities are trusted websites that users willingly interact with, while third-party entities, such as analytics services, may not have the same level of trust. Differentiating between these entities can help determine the level of permission granted.

Google's proposal of engagement as a user gesture takes into account various signals like frequency of visits, bookmarks, and homepage additions to decide on granting permission. While this approach may improve the current state of art, it is difficult to predict and reason about functionality availability, making it challenging for developers and users to understand privacy implications.

Steganography, a technique used to hide data within media, offers a promising solution to inject noise into high entropy fingerprinting endpoints. By making each canvas unique, subtle differences can be introduced that are imperceptible to casual observers but disrupt fingerprinting techniques. This approach can be applied to images, audio, and even user agent strings.

Privacy budget is another proposal under consideration. It aims to limit the amount of data that can be collected and used for fingerprinting purposes. This approach is currently being developed by the Blink project. However, some experts argue that it may not be an effective solution and have concerns about its implementation.

Web fingerprinting poses challenges to user privacy. Approaches such as permission prompts, user gestures, differentiating between first-party and third-party entities, steganography, and privacy budget are being explored to address these concerns. Each approach has its advantages and limitations, and further research and development are needed to strike a balance between privacy and functionality.

Web fingerprinting is a technique used to track and identify users based on unique characteristics of their web browser or device. It involves collecting information such as screen resolution, installed fonts, browser plugins, and other attributes that can be used to create a unique profile for each user. This profile can then be used to track the user across different websites and online activities.

One approach to address web fingerprinting is to assign a privacy budget to each user. The idea is to limit the amount of identifying information that can be collected about a user. For example, a website may allow a user to be identified with a precision of one out of a thousand. Once enough identifying information has been collected to surpass this threshold, privacy protections are imposed to prevent further tracking.

However, this approach has several limitations. Firstly, it raises the question of what happens once the privacy budget is exhausted. If a user visits a website multiple times, do they continue to accumulate identifying bits? If



so, there is a risk of breaking the website or compromising privacy. Secondly, the scope of the privacy budget is unclear. If a third-party is included on a page, do they have their own budget or is it shared with the first party? This can lead to a situation where the third party exhausts the first party's budget, undermining privacy protections.

Another challenge is finding practical ways to nullify each fingerprinting method. Even if a user decides to block certain attributes, such as screen resolution, it is important to have mechanisms in place to prevent third parties from accessing this information. This requires effective means of blocking access to sensitive data once it has been delegated or restricted.

In addition, there are other techniques that can be used for web fingerprinting, such as exploiting different error cases in browsers or manipulating HTML parsing errors. These methods can be used to extract identifying information or perform code execution. Addressing these issues requires standardizing error conditions and implementing measures to prevent these techniques from being exploited.

To combat web fingerprinting as a website owner, it is important to consider different strategies. This could include implementing feature policies to control what third-party frames have access to, ensuring that error cases are handled consistently across browsers, and finding ways to nullify fingerprinting methods by blocking access to sensitive information.

While the idea of assigning a privacy budget to limit web fingerprinting is appealing, it poses several challenges and limitations that need to be addressed. It is crucial to find practical solutions to nullify fingerprinting methods and prevent third parties from accessing sensitive information. Standardizing error conditions and implementing effective feature policies can also contribute to mitigating the risks associated with web fingerprinting.

Web fingerprinting is a technique used to track and identify users based on their unique browser configurations. It involves collecting information about a user's browser, operating system, plugins, fonts, and other characteristics to create a unique identifier or "fingerprint". This fingerprint can then be used to track the user's online activities across different websites.

Web fingerprinting poses serious privacy concerns as it allows third parties to track and monitor users without their knowledge or consent. This can lead to targeted advertising, profiling, and potential security risks.

To address this issue, various approaches have been taken to protect user privacy and prevent web fingerprinting. One approach is to make the failure of fingerprinting attempts as catastrophic as possible. This can be done by implementing measures such as spinning loops, throwing up warnings, and requesting additional permissions. The goal is to make the website so intolerable to use that the fingerprinting tool vendor will roll back their efforts.

However, these measures are often seen as ugly and not user-friendly. Some websites and trackers intentionally break websites to discourage users from protecting their privacy. This highlights the need for more effective and user-centric solutions.

One browser that aims to address web fingerprinting and protect user privacy is Brave. Brave incorporates a feature called "shields" which is a collection of privacy tools. These tools are enabled by default for every website visited, but users have the option to disable them if needed. One of the main features of shields is cross-site tracking blocking, which prevents known trackers from loading. Brave utilizes lists of known tracker URLs, including EasyList, EasyPrivacy, uBlock Origin, and its own generated lists, to determine if a URL is associated with tracking.

In addition to blocking cross-site tracking, Brave also takes measures to prevent the sending of third-party cookies to anyone. Exceptions are made only in rare cases to unbreak websites. Brave also disables certain fingerprinting techniques by default, making it harder for websites to gather identifying information.

While Brave's approach is relatively simple, it has shown promising results. However, further research and development are needed to improve and advance privacy protection measures. Brave is open to collaborations and internships for those interested in contributing to this field.

It is important to note that web privacy concerns extend beyond the web itself. IoT systems, for example, also

pose significant privacy risks. Therefore, it is crucial to address privacy issues in various domains and not solely focus on the web.

Web fingerprinting is a privacy concern that allows for the tracking and identification of users based on their unique browser configurations. Brave is a browser that incorporates privacy tools, such as cross-site tracking blocking and disabling fingerprinting techniques, to protect user privacy. However, further research and collaboration are necessary to address privacy concerns in other domains and advance privacy protection measures.

Privacy is a crucial aspect of web applications security. It is important to consider privacy as more than just a feature to be added to a system. Any harm caused to users through the lack of privacy is a significant issue that should not be overlooked. Therefore, it is essential to reject the notion of privacy as a mere sticker on a box.

When choosing an employer, it is crucial to consider the impact of your work. One way to evaluate potential employers is by determining whether the projects you will be working on empower powerful individuals over weaker ones or vice versa. It is important to think deeply about the implications of working on projects that perpetuate power imbalances.

In terms of legislation, Europe has taken a significant step forward with the General Data Protection Regulation (GDPR). This legislation has set a high standard for privacy protection, and efforts are being made to hold international companies accountable to this standard. Brave and similar companies have been successful in challenging advertisers, although it is still early in this process.

Legislation in the United States is also improving, with California passing robust privacy legislation. One notable aspect of this legislation is the concept of dual purpose, which prevents companies from using user data for purposes other than what was initially consented to. For example, if a user provides their data to improve product quality, it cannot be used to enhance advertising. This type of legislation helps protect user privacy and ensures that data is not misused.

However, it should be noted that many companies operate outside the jurisdiction of US and European laws, making it challenging to enforce privacy regulations globally. Therefore, technological solutions will also play a crucial role in addressing privacy concerns.

One such solution is anonymization, which involves treating certain transactions as survey results and using techniques like Chum mixed nuts to mix and obfuscate data. This helps improve privacy guarantees by altering how information is processed and transmitted. Anonymization is a complex topic that requires further exploration, but it is an important aspect of how Brave handles transactions.

Privacy is a fundamental aspect of web applications security. It should be considered beyond a simple feature and should not cause harm to users. Legislation, such as the GDPR and California's privacy laws, plays a significant role in protecting user privacy. However, the global nature of the internet requires technological solutions to complement legal efforts.

**EITC/IS/WASF WEB APPLICATIONS SECURITY FUNDAMENTALS DIDACTIC MATERIALS****LESSON: DOS, PHISHING AND SIDE CHANNELS****TOPIC: DENIAL-OF-SERVICE, PHISHING AND SIDE CHANNELS**

In today's lecture, we will be discussing the fundamentals of web application security, specifically focusing on denial-of-service (DoS) attacks, phishing, and side channels. Before we delve into these topics, I would like to start with a group activity.

Please take out your laptops or iPads, although laptops are preferred. Open up a browser that you don't normally use, ensuring that it doesn't have any unsaved work. We will be visiting a website together, so make sure you are willing to force quit the browser if necessary.

Now, type in the URL "annoyingsite.com" and hit enter. Please refrain from pressing any other buttons at this time. Take a moment to observe what is happening to your browser. Partner up with someone and discuss the surprising things you notice about the website's behavior.

Some of the things you may have observed include the appearance of the print dialog, the website opening in full screen, and the downloads folder being filled with files. These actions demonstrate a UI denial-of-service attack, where the website overrides the browser's default behavior to make it difficult for users to escape the trap it has set.

This type of attack can have different goals. In this case, the website attempted to prevent users from closing the browser window by overriding the keyboard shortcut and displaying persistent messages. Such attacks can be used to create scareware, where users are convinced that their computer is infected with a virus and are prompted to purchase ineffective products. Additionally, there are troll sites that aim to annoy users without causing any harm.

Understanding the different levels of APIs provided by the browser is crucial when considering these types of attacks. Some APIs can be used without restrictions, such as those related to the Document Object Model (DOM). However, there are other APIs that have limitations and require user permission. By exploiting these APIs, attackers can manipulate the user interface and create a disruptive experience.

Denial-of-service attacks, phishing, and side channels are significant threats to web application security. It is important to be aware of these vulnerabilities and take appropriate measures to protect against them.

**Web Applications Security Fundamentals - DoS, Phishing, and Side Channels**

Web applications are an integral part of our online experience, allowing us to interact with websites and perform various tasks. However, these applications can also be vulnerable to attacks that compromise their security. In this material, we will discuss three common threats to web application security: Denial-of-Service (DoS), phishing, and side channels.

Denial-of-Service (DoS) attacks aim to disrupt the availability of a web application by overwhelming it with a flood of requests or by exploiting vulnerabilities in its infrastructure. This can result in the application becoming slow or completely unavailable to legitimate users. Attackers may use various techniques, such as sending a large number of requests simultaneously or exploiting vulnerabilities in the application's code.

Phishing attacks, on the other hand, target users by tricking them into revealing sensitive information, such as passwords or credit card details. Attackers often impersonate legitimate entities, such as banks or social media platforms, and send deceptive emails or create fake websites that closely resemble the original ones. When users unknowingly provide their information, attackers can use it for malicious purposes, such as identity theft or financial fraud.

Side channels refer to unintended channels of communication that can be exploited by attackers to gather information about a web application or its users. These channels are typically unintentional and arise from the design or implementation of a system. For example, an attacker may analyze the timing of responses from a web application to gain insights into its internal workings or to extract sensitive information.

To mitigate these threats, web browsers have implemented security measures in the form of Application Programming Interfaces (APIs). These APIs regulate the interactions between web applications and users, ensuring that certain actions require explicit user consent or engagement. There are different levels of API restrictions, with each level imposing stricter requirements on web applications.

Level one APIs require the user to interact with the web application in some way before certain actions can be performed. This interaction can be as simple as pressing a key or clicking on the application. However, passive interactions, such as scrolling, do not count as valid interactions for accessing these APIs.

Level two APIs involve more invasive actions that trigger permission prompts. These prompts ask for the user's consent to access sensitive resources, such as location, camera, or microphone. This level of API restriction ensures that applications cannot access these resources without the user's explicit permission.

Level three APIs represent a compromise between strict permission prompts and immediate access. The browser dynamically decides whether to allow these APIs based on the user's engagement with the web application. For example, autoplaying sound is an API that requires user engagement. If a user frequently interacts with a site and plays videos, the browser may allow autoplaying sound without a prompt.

It is important to strike a balance between user consent and usability. While it is crucial to protect users from potential threats, overly restrictive measures can hinder the user experience. Web browsers continuously update their security measures to adapt to evolving threats and user expectations.

Web applications face various security threats, including DoS attacks, phishing, and side channels. To mitigate these risks, web browsers enforce API restrictions that require user engagement or explicit permission for certain actions. By understanding these fundamentals, developers and users can better protect themselves and ensure a secure online experience.

#### Web Applications Security Fundamentals - DoS, Phishing, and Side Channels

Web applications are susceptible to various security threats such as denial-of-service (DoS) attacks, phishing attacks, and side channels. These threats can compromise the integrity, availability, and confidentiality of web applications and their users' data. In this didactic material, we will explore these security fundamentals in detail.

Denial-of-Service (DoS) attacks aim to disrupt the normal functioning of a web application by overwhelming its resources. Attackers achieve this by flooding the application with an excessive amount of requests, rendering it unable to respond to legitimate user requests. One common form of DoS attack is the use of infinite loops combined with alert windows. By continuously generating alert windows, the attacker prevents the user from closing the tab or browser window, effectively locking them into the malicious site. To mitigate this, modern browsers now employ multi-process architecture, allowing users to close the tab or window without being blocked by the malicious code.

Phishing attacks involve tricking users into revealing sensitive information, such as login credentials or financial details, by masquerading as a trustworthy entity. Web applications can be vulnerable to phishing attacks through the use of pop-up windows. These windows, opened using the `window.open` API, can be designed to mimic legitimate websites, leading users to unknowingly disclose their information. Browsers have implemented countermeasures by restricting pop-up windows to be opened only in response to user interactions, such as clicking on a button. This helps mitigate the risk of users falling victim to phishing attacks.

Side channels are covert channels through which attackers can gather sensitive information without directly exploiting vulnerabilities in the web application. One example of a side channel is the ability to intercept and respond to user events, such as mouse clicks or keyboard inputs. By intercepting these events, attackers can open additional windows, further complicating the user's ability to close the malicious site. This can lead to a frustrating user experience and potentially expose users to further security risks.

It is important for web application developers and users to be aware of these security fundamentals and take appropriate measures to protect against DoS attacks, phishing attempts, and side channels. Developers should implement secure coding practices, such as input validation and output encoding, to prevent vulnerabilities that can be exploited. Users should exercise caution when interacting with unfamiliar websites, avoiding clicking on suspicious links or providing sensitive information unless they can verify the legitimacy of the site.

By understanding these security fundamentals, we can work towards creating and using web applications that prioritize the protection of user data and ensure a safe online experience.

In the realm of web application security, it is crucial to understand and address potential threats such as denial-of-service (DoS) attacks, phishing, and side channels. These threats can have severe consequences for both users and organizations, making it essential to be equipped with the knowledge and strategies to mitigate them effectively.

Denial-of-service attacks aim to disrupt the availability of a web application by overwhelming it with a flood of requests or exploiting vulnerabilities in its infrastructure. Attackers may use various techniques, including flooding the application with excessive traffic or exploiting weaknesses in the application's code or infrastructure. The result is often a significant degradation in the application's performance or even a complete outage.

Phishing, on the other hand, involves tricking users into divulging sensitive information such as passwords, credit card details, or personal data by impersonating a trustworthy entity. Attackers typically accomplish this through deceptive emails, messages, or websites that mimic legitimate sources. Once users unknowingly provide their information, it can be exploited for malicious purposes, such as identity theft or unauthorized access.

Side channels refer to unintended channels of information leakage that can be exploited by attackers to gain unauthorized access or extract sensitive data. These channels can arise due to flaws in the design or implementation of a web application, allowing attackers to infer information by analyzing variations in response times, power consumption, or other observable behaviors. Side channels can be particularly challenging to detect and mitigate, as they often exploit subtle and unintended interactions between different components of the system.

To defend against these threats, web application developers and security professionals employ various strategies and best practices. These include:

1. Implementing robust access controls and authentication mechanisms to ensure that only authorized users can access sensitive resources or perform critical actions within the application.
2. Regularly updating and patching the application's software and infrastructure to address known vulnerabilities and protect against emerging threats.
3. Employing secure coding practices to minimize the risk of introducing vulnerabilities during the development process. This includes practices such as input validation, output encoding, and secure session management.
4. Implementing monitoring and logging mechanisms to detect and respond to suspicious activities or anomalies in real-time. This enables early detection of potential attacks and allows for timely mitigation.
5. Educating users about the risks and best practices for online security, such as recognizing phishing attempts, using strong and unique passwords, and being cautious when sharing personal information online.

It is important to note that the examples mentioned in the original transcript, such as bypassing pop-up blockers or manipulating window behavior, are not recommended practices and should not be employed for legitimate purposes. These actions can be seen as unethical and may violate legal and ethical boundaries. It is essential to prioritize the security and privacy of users and adhere to industry best practices.

Understanding the fundamentals of web application security is crucial in today's digital landscape. By being aware of potential threats such as denial-of-service attacks, phishing, and side channels, and implementing the appropriate security measures, developers and organizations can protect their applications and users from harm.

#### Web Applications Security Fundamentals - DoS, phishing, and side channels

Web applications are vulnerable to various security threats, including Denial-of-Service (DoS) attacks, phishing attacks, and side channels. In this didactic material, we will explore these threats and discuss their impact on web application security.

Denial-of-Service (DoS) attacks are designed to disrupt the normal functioning of a web application by

overwhelming it with a high volume of requests. This can lead to a significant decrease in performance or even a complete shutdown of the application. Attackers may exploit vulnerabilities in the application's code or infrastructure to launch these attacks. To defend against DoS attacks, web developers can implement measures such as rate limiting, traffic filtering, and load balancing to ensure the application can handle a large number of requests without being overwhelmed.

Phishing attacks are a form of social engineering where attackers trick users into revealing sensitive information, such as passwords or credit card details, by impersonating a trustworthy entity. These attacks often involve sending fraudulent emails or creating fake websites that closely resemble legitimate ones. To protect against phishing attacks, users should be cautious when clicking on links or downloading attachments from unknown sources. Web application developers can also implement security measures such as email validation, SSL/TLS encryption, and two-factor authentication to mitigate the risk of phishing attacks.

Side channels refer to unintended channels of communication that can be exploited by attackers to gather sensitive information. For example, web applications may inadvertently leak information through the browser's history, allowing attackers to track a user's browsing activity. Developers can prevent such information leakage by implementing proper session management techniques, such as using secure cookies and ensuring sensitive data is not stored in the browser's history.

Web application security is a complex field that requires a balance between functionality and security. While many APIs and features provide valuable functionality, they can also be abused by attackers. Developers must carefully consider the trade-offs between power and security when implementing these features. It is crucial to follow best practices, regularly update software, and conduct security audits to ensure the robustness of web applications.

Web applications face various security threats, including DoS attacks, phishing attacks, and side channels. Understanding these threats and implementing appropriate security measures is essential to protect web applications and the sensitive data they handle.

#### Web Applications Security Fundamentals - DoS, Phishing, and Side Channels

Web applications are vulnerable to various security threats, including denial-of-service (DoS) attacks, phishing attacks, and side channel attacks. In this didactic material, we will explore these security threats in detail and understand how they can impact the security of web applications.

Denial-of-Service (DoS) attacks aim to disrupt the availability of a web application by overwhelming it with a high volume of requests. Attackers achieve this by exploiting vulnerabilities in the application's infrastructure or by utilizing botnets to launch coordinated attacks. The result is a significant decrease in the application's performance or complete unavailability. Web application developers must implement robust security measures, such as rate limiting, traffic filtering, and load balancing, to mitigate the risk of DoS attacks.

Phishing attacks involve tricking users into revealing sensitive information, such as login credentials or financial data, by impersonating a trusted entity. Attackers often create deceptive websites or send fraudulent emails that mimic legitimate organizations to deceive users into providing their personal information. Web application developers should educate users about phishing attacks, implement secure authentication mechanisms, and employ email filtering systems to detect and prevent phishing attempts.

Side channel attacks exploit unintended information leakage from a system to gain unauthorized access or extract sensitive data. In the context of web applications, side channel attacks can target various elements, such as the user's cursor or file downloads. For example, attackers can hide the cursor or manipulate its behavior to disorient users and make it harder for them to close windows or interact with the application. Developers must be aware of these vulnerabilities and implement proper security measures, such as validating user input and sanitizing data, to prevent side channel attacks.

Furthermore, web applications must address potential security risks associated with file downloads. Attackers can exploit insecure file download mechanisms to deceive users into downloading malicious files or execute unauthorized actions. Developers should implement secure file download functionalities, including proper validation of file types and names, to ensure the integrity and safety of user downloads.



It is also crucial to consider the security implications of full-screening functionalities in web applications. While full-screening is a useful feature for enhancing user experience, it can also be exploited by attackers to deceive users or launch malicious actions. Developers should be cautious when implementing full-screen functionalities and ensure that proper security measures, such as browser compatibility checks and prefix usage, are in place.

Lastly, we must address the importance of protecting against cross-site request forgery (CSRF) attacks. CSRF attacks exploit the trust between a user's browser and a web application to perform unauthorized actions on behalf of the user. By tricking a user into visiting a malicious website or clicking on a malicious link, attackers can execute actions that can compromise the user's account or steal sensitive information. Web application developers should implement robust CSRF protection mechanisms, such as same-site cookies, to prevent these attacks and ensure the security of user accounts.

Web applications face various security threats, including denial-of-service attacks, phishing attacks, and side channel attacks. Developers must implement appropriate security measures, such as rate limiting, traffic filtering, secure authentication mechanisms, and CSRF protection, to safeguard web applications and protect user data.

#### Web Applications Security Fundamentals - DoS, Phishing, and Side Channels

In the field of cybersecurity, it is important to understand the various threats that can compromise the security of web applications. This didactic material will focus on three specific threats: Denial-of-Service (DoS) attacks, phishing attacks, and side channels.

Denial-of-Service (DoS) attacks are aimed at disrupting the availability of a web application by overwhelming its resources. Attackers achieve this by flooding the target application with a large volume of requests, causing it to become unresponsive or crash. To carry out a DoS attack, attackers often exploit vulnerabilities in the application's code or infrastructure.

Phishing attacks, on the other hand, target users rather than the application itself. In a phishing attack, attackers deceive users into revealing sensitive information such as usernames, passwords, or credit card details. This is typically done through fraudulent emails or websites that mimic legitimate ones. Users are tricked into providing their information, which is then used for malicious purposes.

Side channels refer to unintended channels of communication that can be exploited by attackers to gain unauthorized access or extract sensitive information. One example of a side channel attack is tab nabbing. In tab nabbing, a malicious site is linked from a harmless site, and when the user clicks the link, it opens in a new tab. This technique is often used by chat services to prevent users from leaving their platform. However, it can be exploited by attackers to deceive users and steal their information.

To protect against these threats, web application developers and administrators should implement robust security measures. This includes regularly updating and patching the application's software, implementing strong authentication mechanisms, and educating users about the dangers of phishing attacks. Additionally, security headers such as X-Frame-Options can be used to prevent clickjacking attacks and protect against tab nabbing.

It is important to stay vigilant and keep up-to-date with the latest security practices to ensure the safety of web applications and the users who interact with them.

#### Web Applications Security Fundamentals - Denial-of-Service (DoS), Phishing, and Side Channels

Web applications are vulnerable to various security threats such as denial-of-service (DoS) attacks, phishing, and side channels. In this didactic material, we will explore these threats and discuss measures to defend against them.

Denial-of-Service (DoS) attacks aim to disrupt the normal functioning of a web application by overwhelming it with a flood of requests or by exploiting vulnerabilities in the application's code. This can lead to a temporary or permanent loss of service for legitimate users. Attackers may use techniques such as flooding the application with excessive traffic or exploiting software vulnerabilities to crash the application.

Phishing is a type of attack where attackers deceive users into revealing sensitive information such as login credentials or financial details. One method of phishing is through the use of malicious links. Attackers can send deceptive links that appear to be legitimate but actually lead to fake websites designed to steal user information. These fake websites may mimic popular sites like social media platforms or banking portals.

Side channels are unintended channels of communication that can reveal sensitive information about a web application. In the context of web applications, side channels can be exploited to gain unauthorized access to user data or perform other malicious activities. One example of a side channel attack is tab nabbing, where an attacker manipulates the behavior of browser tabs to deceive users into entering their credentials on a fake website.

To defend against these threats, web application developers can implement various security measures. One approach is to add the "rel=noopener" attribute to links that open in new tabs. This attribute prevents the new tab from having a reference to the previous tab, reducing the risk of tab nabbing attacks.

Additionally, there is a new HTTP header called the "Isolation" header that can be used to request isolation from other sites. This header, although not yet widely supported, can provide enhanced security by isolating a web application in a separate browser process, disabling the window opener functionality, and preventing side channel attacks.

Web applications are vulnerable to denial-of-service attacks, phishing, and side channels. Understanding these threats and implementing appropriate security measures is crucial to protect user data and ensure the integrity of web applications.

Side-channel attacks are a powerful type of attack in cybersecurity. They exploit vulnerabilities in web applications, such as Denial-of-Service (DoS), phishing, and side channels. In a side-channel attack, an attacker can gather information about a target system by analyzing its behavior, even without direct access to the system itself.

One example of a side-channel attack is when a website links to another site. The linked site's window opener becomes known, but it doesn't have a pointer back to the original site. This severs cross-window ties and prevents potential side-channel attack vectors. By using a specific header, the iframe of the linked site can run in a separate process, providing a higher level of isolation. Additionally, this header also breaks post message communication between the sites, further enhancing security.

It is worth noting that annoying features in websites can be used to identify and address vulnerabilities. For example, the site may not function well on mobile devices due to the lack of window manipulation. By identifying and adding more browser APIs that can be abused, the site becomes a test case for browser vendors to assess their UI security. This conflict between the simplicity of a document viewer and the desire for a powerful app platform creates tension for browser vendors. While some users prefer a simple browsing experience, others want the web to compete with native apps on various platforms.

Adding new APIs to browsers can introduce additional fingerprinting vectors, making it easier to identify users. In the past, browser vendors dismissed these concerns, arguing that the problem already existed, and adding more APIs wouldn't worsen the situation. However, there is a shift in mindset among non-Chrome browsers to avoid exacerbating the fingerprinting problem and work towards a less fingerprintable web.

Pete, in his paper titled "Websites Don't Need to Vibrate: A Cost-Benefit Approach to Improving Browser Security," explores the usefulness of different APIs and their impact on site functionality. By disabling certain features and observing how many sites break, the paper evaluates the practicality of these APIs. This approach helps in understanding the trade-offs between powerful features and browser security.

Side-channel attacks pose a significant threat to web application security. Understanding and addressing vulnerabilities related to DoS, phishing, and side channels are crucial for safeguarding web applications. By implementing measures such as using specific headers and evaluating the usefulness of browser APIs, developers and browser vendors can enhance security while balancing the needs of users and the capabilities of the web.

Web Applications Security Fundamentals: DoS, Phishing, and Side Channels

Web applications are an essential part of our online experience, allowing us to perform various tasks and access information. However, they can also be vulnerable to attacks that compromise their security. In this didactic material, we will discuss the fundamentals of web application security, focusing on denial-of-service (DoS) attacks, phishing, and side channels.

#### Denial-of-Service (DoS) Attacks:

A DoS attack is an attempt to disrupt the normal functioning of a web application by overwhelming it with a flood of requests or by exploiting vulnerabilities. The goal is to make the application unavailable to its intended users. Attackers can achieve this by sending a large number of requests, exhausting server resources, or exploiting vulnerabilities in the application's code.

#### Phishing:

Phishing is a type of cyber attack where attackers attempt to deceive users into revealing sensitive information such as passwords, credit card numbers, or personal data. They often do this by impersonating a trustworthy entity, such as a bank or a popular website, and sending deceptive emails or messages. Once the user falls for the deception and provides their information, it can be used for malicious purposes.

#### Side Channels:

Side channels are unintended channels of communication that can be exploited by attackers to gather sensitive information. These channels are typically created due to the implementation of a web application's features or the underlying technology. Attackers can exploit side channels to gain unauthorized access to user data or perform other malicious activities without directly attacking the application's security mechanisms.

To mitigate the risks associated with these threats, web application developers and security professionals employ various techniques. These include implementing secure coding practices, regularly updating and patching software, using encryption to protect sensitive data, and educating users about potential risks and best practices.

It is essential to strike a balance between security and usability when designing web applications. While it may seem tempting to prompt users for every action, research has shown that excessive security prompts can lead to user fatigue and decreased effectiveness. Therefore, it is crucial to carefully design security prompts that are clear, concise, and provide meaningful information to users.

Advanced users, such as Linux users and early adopters, tend to click through phishing warnings at higher rates than average users. This highlights the importance of continuous research and improvement in the design of security prompts and user interfaces to effectively protect users from potential threats.

Web application security is a complex and evolving field. It requires a combination of secure coding practices, regular updates, user education, and thoughtful design of security prompts and interfaces. By understanding the fundamentals of DoS attacks, phishing, and side channels, developers and security professionals can better protect web applications and the sensitive information they handle.

Phishing, Denial-of-service (DoS), and side channels are important concepts in web application security. Phishing involves tricking users into divulging sensitive information by impersonating reputable entities. It is often easier to deceive users than to directly attack a system. Bruce Schneier, a security expert, emphasizes that security is fundamentally a people problem. Therefore, phishing aims to exploit human vulnerabilities rather than technical weaknesses.

DoS attacks, on the other hand, aim to disrupt the availability of a web application by overwhelming it with an excessive amount of traffic or resource requests. This can lead to the website becoming slow or completely inaccessible to legitimate users. DoS attacks can be launched using various techniques, such as flooding the target server with traffic or exploiting vulnerabilities in the application's code.

Side channels refer to unintended channels of information leakage that can be exploited by attackers. These channels provide insights into the internal workings of a system or application, allowing attackers to gain unauthorized access or extract sensitive information. One example of a side channel attack is the manipulation of URLs to deceive users. Attackers can create URLs that appear legitimate but actually lead to malicious websites. This can be achieved by using Unicode characters that look similar but have different byte

representations, making it difficult for users to detect the deception.

To defend against these threats, web developers and users must be aware of best practices. For example, users should be cautious when clicking on links and should verify the legitimacy of websites before entering sensitive information. Developers should implement security measures, such as adding the "rel=noopener" attribute to links, to prevent attackers from exploiting vulnerabilities. Additionally, organizations should educate their employees about the risks of phishing and provide training on how to identify and report suspicious emails or websites.

Understanding and addressing the vulnerabilities associated with phishing, DoS attacks, and side channels are crucial for ensuring the security of web applications. By implementing appropriate security measures and promoting user awareness, organizations can mitigate the risks posed by these threats.

Certain languages have letters that are visually indistinguishable from Latin letters used in English. Instead of reusing the Latin letters from the ASCII set, these languages duplicate them in their own character space. This creates an issue because although the letters may look the same, they are treated as different letters by computers. To address this, if a hostname contains a Unicode character, it can be translated into puny code to make it more obvious. Chrome and Safari use this translation, which is why we see the xn- prefix in URLs.

The puny code translation strips out the Unicode letters that are not in the standard ASCII set, leaving only the remaining letters. The xn- prefix indicates that it is a puny code domain and not a real domain that someone could register. This prevents confusion between puny code domains and regular domains. For example, if someone registered the puny code domain, it would be differentiated from the regular domain.

Chrome intervenes and shows puny code only when it thinks it might be confusing. In cases where it is less likely to be a phishing attempt, Chrome will show the original character directly. Chrome uses rules to determine when to intervene, such as if a domain is made up of letters from two languages and one of the letters looks like an English letter.

Firefox had trouble in the demo because all the letters in the domain were changed to come from Cyrillic, making it appear as if it is one language. This workaround does not work well, and Chrome implements a fix for top-level domains (TLDs) that do not use Unicode letters. If the left side of the domain has any look-alike characters and the TLD itself contains foreign language characters, then it will be rendered as puny code.

This issue is known as an ibn homograph attack, which is similar to domain squatting. It involves registering a domain name that looks visually similar to another domain, potentially leading to phishing attempts or confusion. In 2017, this issue was addressed by most browsers, but Firefox has not fixed it yet.

#### Web Applications Security Fundamentals - DoS, Phishing, and Side Channels

Web applications are susceptible to various security threats, including denial-of-service (DoS) attacks, phishing, and side channels. In this didactic material, we will explore these threats and understand their implications for web application security.

Phishing is a common attack where cybercriminals create fake websites that resemble legitimate ones. They often register domain names that are one letter off from popular companies' domains, hoping that users will make typos when typing the URL. When users visit these phishing websites, they are tricked into entering sensitive information, such as login credentials or credit card details. While users can accidentally mistype a URL, phishing attacks rely on social engineering techniques to deceive users into visiting malicious websites.

Interestingly, even handwriting can be prone to similar issues. For example, the Arabic word "stumped," which means direction, can be mistakenly transcribed as "summit" if the scribe confuses an M with an N. Similarly, certain typefaces can cause confusion, where an M may appear as an RN or an RR. This confusion extends to domain names as well, where some fonts used by browsers do not distinguish Cyrillic letters from English letters. This lack of distinction makes it easier for attackers to create deceptive domain names.

In some cases, intentional use of similar-looking characters can be found. For instance, a Turkish typewriter omitted the number one key due to limited space. To work around this limitation, users can use the lowercase letter L to represent the number one. Similarly, the absence of a semicolon key can be overcome by typing a

colon and then backspacing to replace it with a comma. This use of creative workarounds demonstrates the ingenuity of users in adapting to limited resources.

To address the confusion caused by similar-looking characters in domain names, the use of puny code has been introduced. Puny code is a way to represent internationalized domain names with non-ASCII characters in a readable and unambiguous form. Since 2017, modern browsers like Safari and Chrome automatically display puny code when an entire domain is composed of look-alike letters, and the top-level domain is not an international domain.

Using a password manager can also provide protection against phishing attacks. Password managers are designed to recognize phishing websites and prevent users from entering their credentials. When a user tries to log in to a website that is not recognized by the password manager, it raises a security alert, ensuring that users do not fall victim to phishing attempts.

Additionally, web users should look for the secure lock symbol in their browser's address bar. The secure lock indicates that the connection between the user's device and the website's server is encrypted. However, it is important to note that the presence of a secure lock does not guarantee the legitimacy of the website. In the past, it was believed that certificate authorities would thoroughly verify certificate requests before issuing them. However, with the rise of services like Let's Encrypt, which provide free TLS certificates without human involvement, the process has become more automated, potentially increasing the prevalence of sketchy websites with secure locks.

Web applications face security threats such as DoS attacks, phishing, and side channels. Understanding these threats and implementing appropriate security measures, such as using password managers and being cautious of phishing attempts, is crucial for safeguarding sensitive information online.

#### Web Applications Security Fundamentals - DoS, Phishing, and Side Channels

In the field of cybersecurity, it is essential to understand various threats that can compromise the security of web applications. Three common threats are Denial-of-Service (DoS) attacks, phishing attacks, and side channels.

DoS attacks involve overwhelming a target system or network with an excessive amount of traffic, rendering it unable to function properly. This can be achieved by flooding the target with requests or exploiting vulnerabilities in the system. The goal of a DoS attack is to disrupt the availability of the target, making it inaccessible to legitimate users. It is important for web application developers and administrators to implement measures to mitigate the impact of DoS attacks, such as rate limiting, traffic filtering, and load balancing.

Phishing attacks, on the other hand, aim to deceive users into revealing sensitive information, such as login credentials or credit card details. Attackers often impersonate trusted entities, such as banks or popular websites, by creating fake websites that resemble the legitimate ones. These fake websites are usually hosted on subdomains or similar-looking domains. Users may unknowingly enter their credentials on these fake websites, which are then captured by the attackers. To protect against phishing attacks, users should be educated on how to identify legitimate websites, such as checking the domain name and looking for security indicators like SSL certificates.

Side channels refer to unintended information leakage that can be exploited by attackers. In the context of web applications, side channels can provide attackers with insights into the internal workings of a system, which can be used to launch attacks. For example, attackers may analyze the behavior of a web application under different conditions to gather information about its vulnerabilities. Web application developers should be aware of potential side channels and implement appropriate security measures to prevent information leakage.

To address the issue of subdomains and their potential for misleading users, web browsers have implemented various visual cues to help users identify legitimate websites. For example, some browsers gray out the path in the URL bar to avoid confusion, while others highlight the actual domain name. However, these visual cues are not foolproof and can be manipulated by attackers. Users should exercise caution when entering sensitive information on websites and verify the legitimacy of the domain before proceeding.

Understanding the fundamentals of web application security is crucial in protecting against threats such as DoS



attacks, phishing attacks, and side channels. Web developers and administrators should implement appropriate measures to mitigate these risks, while users should be vigilant and educated about the signs of potential threats.

#### Web Applications Security Fundamentals - DoS, Phishing and Side Channels

Web applications are an essential part of our daily lives, allowing us to perform various tasks online. However, they are also vulnerable to attacks that can compromise our security and privacy. In this didactic material, we will explore three common types of attacks: Denial-of-Service (DoS), phishing, and side channels.

Denial-of-Service (DoS) attacks aim to disrupt the availability of a web application by overwhelming it with a flood of requests or by exploiting vulnerabilities in its infrastructure. This can lead to the application becoming unresponsive or even crashing. Attackers may use various techniques, such as sending a large volume of requests or exploiting vulnerabilities in the application's code.

Phishing attacks, on the other hand, aim to deceive users into revealing sensitive information, such as passwords or credit card details. Attackers often create fake websites or emails that mimic legitimate ones, tricking users into believing they are interacting with a trusted entity. Phishing attacks can be highly sophisticated, making it difficult for users to distinguish between genuine and fake websites or emails.

Side channels refer to unintended channels of communication that can be exploited by attackers to gain unauthorized access to sensitive information. One example is the "picture-in-picture" attack, where a fake browser window is displayed on top of a legitimate one. This can deceive users into entering their credentials or other sensitive information into the fake window, compromising their security.

Another example of a side channel attack is the "cookie jacking" attack. In this attack, the attacker exploits a vulnerability in the browser to extract the user's cookies, which may contain sensitive information. By tricking the user into visiting a malicious website or by being on the same network as the user, the attacker can intercept and steal the user's cookies.

Defending against these attacks requires a multi-layered approach. One effective defense is to use a password manager, which can help prevent users from entering their credentials into fake websites. Password managers can recognize the true domain of a website and only autofill credentials for legitimate domains.

Another defense is to use hardware security keys. These keys establish a secure protocol with the website being accessed, ensuring that sensitive information is not revealed to unauthorized sites. Hardware security keys can be used on both computers and smartphones, providing an added layer of protection.

It is important to stay vigilant and be cautious when interacting with web applications. Always verify the authenticity of websites and emails before entering sensitive information. Regularly updating software and using reputable security tools can also help mitigate the risk of these attacks.

Web applications are susceptible to various security threats, including Denial-of-Service attacks, phishing, and side channels. Understanding these threats and implementing appropriate defense measures is crucial for ensuring the security and privacy of our online activities.

In the field of web application security, there are several fundamental concepts that are important to understand. In this didactic material, we will discuss three such concepts: denial-of-service (DoS) attacks, phishing attacks, and side channels.

Denial-of-service attacks are a common type of cyber attack where the attacker tries to make a service or website unavailable to its intended users. This is typically achieved by overwhelming the target system with a flood of requests, causing it to become slow or unresponsive. One way to carry out a DoS attack is by exploiting vulnerabilities in the target system's code or infrastructure. Another method involves using botnets, which are networks of compromised computers, to launch a coordinated attack. DoS attacks can have serious consequences, as they can disrupt business operations or prevent users from accessing important resources.

Phishing attacks, on the other hand, aim to deceive users into revealing sensitive information such as passwords, credit card numbers, or personal data. Attackers often use social engineering techniques to trick



users into clicking on malicious links or providing their information on fake websites. Phishing attacks can be carried out through various channels, including email, instant messaging, or even phone calls. It is important for users to be aware of phishing techniques and to exercise caution when interacting with unfamiliar or suspicious sources.

Side channels are another aspect of web application security that can be exploited by attackers. Side channels refer to unintended channels of communication that can leak sensitive information. For example, an attacker might use a technique called clickjacking to trick users into unknowingly performing actions on a website. By hiding elements on the page or manipulating the user interface, the attacker can deceive the user into clicking on hidden buttons or links. Similarly, file jacking involves tricking users into thinking they are downloading a file when, in reality, they are uploading their own files to the attacker's server. These attacks exploit human perception and the trust users place in certain indicators, such as the URL bar.

To mitigate these types of attacks, various measures can be taken. For example, Google Safe Browsing is a service that aims to protect users from visiting malicious websites. It maintains a list of known bad sites and provides warnings when users attempt to access them. To protect users' privacy, Google Safe Browsing uses a protocol that allows checking the safety of a URL without revealing the full browsing history to Google.

Understanding the fundamentals of web application security is crucial to protect against cyber attacks. Denial-of-service attacks, phishing attacks, and side channels are all important concepts to be aware of. By implementing appropriate security measures and being vigilant, users can help safeguard their personal information and ensure a safer online experience.

**EITC/IS/WASF WEB APPLICATIONS SECURITY FUNDAMENTALS DIDACTIC MATERIALS****LESSON: INJECTION ATTACKS****TOPIC: CODE INJECTION**

Web applications are vulnerable to various types of attacks, one of which is injection attacks. In this type of attack, an attacker injects malicious code into a web application, which can lead to unauthorized access, data breaches, or other security compromises.

One specific type of injection attack is code injection. Code injection occurs when an attacker is able to insert malicious code into a web application's codebase, which is then executed by the application. This can happen when user input is not properly validated or sanitized, allowing the attacker to inject their own code.

To mitigate the risk of code injection attacks, web developers need to implement proper security measures. One such measure is the use of Google Safe Browsing, a system devised by Google to help protect users from malicious websites.

Google Safe Browsing works by maintaining a list of known malware and phishing URLs. When a user tries to visit a website, their browser can query this list to check if the site has been reported as malicious or suspicious. If the site is flagged, the user is warned before proceeding.

There are two approaches to implementing this security measure. The naive approach involves sending real-time browsing data to Google every time a user visits a website. This approach has two downsides: it compromises user privacy by sharing their browsing history, and it introduces latency as the browser waits for a response from Google.

A better approach is to download a local list of suspicious URLs and have the browser check this list before loading a website. This approach is faster because the list is stored locally, but it has the drawback of being constantly changing and potentially large.

To address these challenges, Google provides an API called the update API. Instead of sending the URL to be checked, the browser downloads a list of hash prefixes from Google. When a suspicious URL is encountered, the browser can compare its hash prefix with the downloaded list to determine if it is safe.

This approach improves privacy as the URL is not sent to Google, and it also reduces latency as the browser can start loading the website while checking the hash prefix locally.

Code injection attacks pose a significant threat to web applications. Implementing security measures such as Google Safe Browsing can help protect users from visiting malicious websites. By downloading a list of hash prefixes and checking them locally, browsers can provide a faster and more privacy-friendly experience.

In the context of web application security, injection attacks are a common and serious threat. One type of injection attack is code injection, which involves malicious code being injected into a web application's codebase. This can lead to various security vulnerabilities and potential exploitation of sensitive data.

To address this issue, Google has implemented a security protocol called Google Safe Browsing. This protocol aims to protect users from visiting malicious websites by checking the safety of URLs before loading them. The protocol utilizes cryptographic hash functions, specifically SHA-256, to determine the safety of URLs.

Here's how the Google Safe Browsing protocol works:

1. The client, which is the user's browser, makes a request to the Google Safe Browsing service to obtain a list of hash prefixes. These hash prefixes represent the beginning part of a hash that corresponds to a suspicious URL.
2. The client receives the list of hash prefixes from the service and stores it locally. It's important to note that multiple suspicious URLs may have the same prefix, and it's also possible for non-suspicious URLs to have the same prefix.

3. When the client wants to check the safety of a specific URL, it first hashes the URL using SHA-256. The resulting hash value is then truncated to match the same prefix length as provided by Google.
4. The client compares the truncated hash value with the locally stored list of hash prefixes. If the truncated hash value is not found in the list, it means that the URL is guaranteed to be safe.
5. In the case where the truncated hash value is found in the list, it indicates that the URL may be unsafe. However, instead of sending the full URL to Google, the client sends only the prefix that matches the truncated hash value. This is done to avoid revealing the exact URL to Google, as there is a possibility that a safe URL may have the same hash prefix.
6. Google receives the prefix from the client and responds with a list of full hash values that share the same prefix. These full hash values correspond to URLs that potentially match the prefix.
7. The client then compares the full hash value obtained from SHA-256 with the list provided by Google. If the full hash value is found in the list, it confirms that the URL is unsafe.

By following this protocol, users can be protected from visiting malicious websites. The Google Safe Browsing service provides an efficient and secure mechanism for checking the safety of URLs before loading them in a web browser.

In the context of web application security, injection attacks are a common and dangerous vulnerability that can lead to unauthorized access, data breaches, and other malicious activities. One specific type of injection attack is code injection, where an attacker injects malicious code into a vulnerable web application.

Code injection attacks occur when an application does not properly validate or sanitize user input and allows the execution of arbitrary code. This can happen in various areas of a web application, such as input fields, URLs, or even HTTP headers. Attackers can exploit this vulnerability to execute arbitrary commands, access sensitive data, or take control of the entire application.

To prevent code injection attacks, it is crucial to implement proper input validation and sanitization techniques. This involves validating and filtering user input to ensure that it does not contain any malicious code or characters that could be interpreted as code. Additionally, using parameterized queries or prepared statements when interacting with databases can help prevent SQL injection attacks, which are a specific type of code injection attack targeting databases.

Furthermore, keeping web application software and frameworks up to date is essential to mitigate code injection vulnerabilities. Developers should regularly apply security patches and updates provided by the software vendors to address any known vulnerabilities.

In addition to these preventive measures, it is important to implement a robust monitoring and logging system. This can help detect any suspicious activities or unusual behavior in the web application, allowing for timely response and investigation.

It should be noted that code injection attacks can have severe consequences, ranging from unauthorized access to sensitive data to complete system compromise. Therefore, it is crucial for developers and security professionals to be aware of this vulnerability and take appropriate measures to protect web applications from such attacks.

Code injection attacks pose a significant threat to web application security. By implementing proper input validation, sanitization techniques, and keeping software up to date, developers can significantly reduce the risk of code injection vulnerabilities. Additionally, monitoring and logging systems play a crucial role in detecting and responding to any potential attacks. By following these best practices, organizations can enhance the security of their web applications and protect sensitive data.

Side-channel attacks are a type of attack where a system is functioning correctly, following its implementation, but still leaks sensitive information. These attacks occur when an attacker can learn information that they shouldn't have access to. Common examples of side channels include timing leaks, where the time it takes for a certain operation to complete can reveal information about the program's state.

To illustrate the concept of side-channel attacks, let's consider an example unrelated to web security. Imagine we want to build a room where we can have secret conversations without anyone outside being able to hear us. We decide to construct the room using a thick glass material, believing that it will prevent sound from escaping. When we enter the room and talk, no sound can be heard from the outside, so we assume our design is successful.

However, since the room is made of glass, an attacker can see through it. This raises the question: Can we use the fact that we can see into the room to learn about the conversations happening inside? In a research paper, MIT researchers demonstrated a side-channel attack that exploits this scenario. By capturing high-speed video of the glass material, they were able to extract subtle visual signals caused by the vibrations of sound hitting the glass. Through suitable processing algorithms, they could partially recover the sounds and turn everyday visible objects into visual microphones.

In their experiments, the researchers recorded video of a potted plant while a nearby loudspeaker played music. Even though the plant's leaves moved by less than a hundredth of a pixel, the researchers were able to combine and filter the tiny motions across the video to recover the sound. They also recovered live human speech from high-speed video of a bag of chips, even when the camera was placed outside behind a soundproof window. The recovered sound was then used with audio recognition software to automatically identify the songs being played.

This example demonstrates how side-channel attacks can exploit seemingly unrelated aspects of a system to leak sensitive information. It highlights the importance of considering all possible attack vectors when designing secure systems, even if they may not appear directly related to the system's primary purpose.

Side-channel attacks are a type of attack where a system leaks information through unintended channels. Timing leaks and visual side channels are just a few examples of how attackers can exploit these vulnerabilities. Understanding and mitigating side-channel attacks is crucial for ensuring the security of web applications and other systems.

Injection attacks, specifically code injection, are a significant concern in web application security. Code injection occurs when an attacker is able to insert malicious code into a web application, which can then be executed by the application. This can lead to various security vulnerabilities, such as unauthorized access, data breaches, and even complete system compromise.

One type of code injection attack is known as a side channel attack. In a side channel attack, an attacker exploits unintended information leakage from a system. For example, an attacker may be able to determine the decryption algorithm being used by monitoring the power usage of a computer. In the context of web applications, side channel attacks can be used to gather sensitive information about users, such as their browsing history.

A classic example of a side channel attack in web applications is the link color attack. This attack takes advantage of the fact that browsers render visited links in a different color than unvisited links. By creating a link and checking the color of the rendered link, an attacker can determine whether a user has visited a specific URL. This can be a serious privacy concern, as it allows an attacker to gather information about a user's browsing habits.

To mitigate this attack, browsers implemented a fix in 2010. Instead of rendering visited links in a different color, browsers now render all links in the same color. This prevents attackers from using link colors as a side channel to gather information about users' browsing history. However, this fix does have some usability drawbacks, as users may become confused when visited links are no longer visually distinguishable from unvisited links.

Alternative solutions were also considered, such as rendering visited links in a different color only if the user had not visited the site before. However, this approach was ultimately discarded due to potential confusion for users. Additionally, attackers could still exploit other visual attributes of links, such as background images, to gather information about visited URLs.

In an effort to address the issue, Mozilla, the organization behind Firefox, focused on making the link color

attack slower and more difficult for attackers. The goal was to reduce the number of URLs an attacker could check per second, thereby making the attack less practical. This approach aimed to provide a partial solution to the problem, acknowledging that it may not be possible to completely eliminate the vulnerability.

Code injection attacks, including side channel attacks, pose a significant threat to web application security. The link color attack is a classic example of a side channel attack that allows attackers to gather information about users' browsing history. Browsers have implemented various measures to mitigate this attack, such as rendering all links in the same color. However, finding a perfect solution remains challenging, as usability and practicality concerns must be carefully balanced.

Web applications are vulnerable to various types of attacks, including injection attacks. One specific type of injection attack is code injection. In code injection attacks, an attacker inserts malicious code into a vulnerable web application, which then gets executed by the application's interpreter or compiler.

To mitigate the risk of code injection attacks, web browsers have implemented certain security measures. One such measure is the restriction of CSS properties that can be applied to visited links. This prevents attackers from using CSS properties to detect whether a link has been visited or not. Additionally, the position and size of visited links cannot be changed, as this could be used to detect changes in the page layout.

Internally, browsers have also made changes to their code paths to prevent timing attacks. Previously, when the href of a link was changed, the code path for visited links was slower, allowing attackers to detect whether a link had been visited or not. Browsers now aim to make the code execution time the same for both visited and unvisited links.

However, these measures are not foolproof. There are still potential vulnerabilities that can be exploited. To demonstrate this, let's consider a scenario where an attacker uses the rendering time of shadows to determine whether a link has been visited. Rendering a large shadow takes more time, so by measuring the rendering time of links with shadows, an attacker can infer which links have been visited.

To exploit this vulnerability, the attacker creates a page with numerous links, including both visited and unvisited ones. By swapping the href attribute of a link from a known unvisited site to a site the attacker wants to query, the browser will redraw the link, even though it will appear the same to the user. By timing the rendering of these links, the attacker can determine which links have been visited based on the differences in rendering time.

In a demonstration, the attacker shows that the rendering time for visited links is generally longer than that for unvisited links. By setting a threshold based on the median rendering time, the attacker can classify links as visited or unvisited. However, it's important to note that this technique is not completely reliable, as it may fail to correctly classify some links.

While browsers have implemented measures to mitigate code injection attacks, there are still potential vulnerabilities that can be exploited. Code injection attacks remain a significant threat to web application security, and developers must continuously update their defenses to protect against these attacks.

### Code Injection in Web Applications Security

Code injection is a type of injection attack that poses a significant threat to the security of web applications. It involves the insertion of malicious code into a vulnerable application, which can then be executed by the application's interpreter or compiler. This can lead to various security vulnerabilities, including data breaches, unauthorized access, and even complete system compromise.

One specific type of code injection attack is known as "injection attacks - code injection." In this type of attack, an attacker exploits vulnerabilities in the application's code execution process to inject and execute their own code. This can be achieved through various means, such as manipulating user input, exploiting insecure coding practices, or taking advantage of poorly implemented security controls.

The transcript discusses a specific scenario where an attacker leverages code injection to perform a browser fingerprinting attack. By manipulating the href attribute of a link, the attacker is able to determine whether the link has been visited or not by timing the browser's redraw process. This technique highlights the challenges in

mitigating code injection attacks, as even seemingly harmless CSS properties can be exploited to leak sensitive information.

To address code injection vulnerabilities, several mitigation strategies can be employed. One approach is to ban CSS properties that affect rendering speed, such as text shadow, even on unvisited links. However, this may not be a comprehensive solution and could impact the user experience. Another suggestion is to double-key the visited link history, which involves considering the site from which the link was visited to determine its visited/unvisited status. While this approach may provide some level of protection, it is not foolproof and relies on accurate tracking of link history.

Web browsers play a crucial role in implementing mitigations for code injection attacks. When making changes to address vulnerabilities, browsers aim to comply with existing specifications. However, if a change is deemed necessary for security reasons and other browsers do not agree, a browser may choose to violate the specification independently. This process is often iterative and involves testing changes in real-world scenarios to assess their impact on websites and user experience.

In extreme cases, completely removing the ability to style certain elements, such as links, can effectively mitigate code injection vulnerabilities. However, this approach may have unintended consequences for web design and user interaction.

It is worth noting that code injection attacks are not limited to manipulating links. Images can also be used to leak sensitive information. For example, an image on a webpage may change its appearance based on the user's login status. By observing the layout changes caused by different image widths, an attacker can infer the user's login status. This highlights the importance of considering all aspects of web application security, including the handling of images and associated cookies.

Code injection attacks pose a serious threat to the security of web applications. The discussed scenario demonstrates the challenges in mitigating such attacks, as even seemingly harmless CSS properties and images can be exploited. Effective mitigation strategies require a combination of secure coding practices, proper input validation, and continuous monitoring for vulnerabilities.

Web applications are vulnerable to various types of attacks, including injection attacks. One specific type of injection attack is code injection. Code injection occurs when an attacker is able to insert malicious code into a web application, which is then executed by the application. This can lead to unauthorized access, data theft, and other security breaches.

One example of code injection is through the use of a phishing page. In this scenario, the attacker creates a page that appears to be a legitimate login page for a service, such as Gmail. By tricking the victim into entering their login credentials on this page, the attacker can gain unauthorized access to the victim's account.

To make the phishing page more effective, the attacker can use a technique that violates the same origin policy. The same origin policy is a security measure that prevents web pages from accessing content from other domains. However, by embedding an image from the target domain, such as Gmail, into the phishing page, the attacker can determine whether the victim is logged into their Gmail account. This is done by checking the size of the image, which changes depending on whether the user is signed in or not.

When the victim visits the attacker's site, the victim's browser automatically attaches cookies to the request. Cookies are small pieces of data that are stored on the user's browser and are used for various purposes, including session management. If the attacker can determine that the victim is logged into Gmail, they can use this information to carry out their attack.

One way to mitigate this type of attack is through the use of same-site cookies. Same-site cookies are cookies that are only sent to the same site that set them. If the cookie used for authentication was a same-site cookie, it would not be attached to the request made to the attacker's site, and the attacker would not be able to determine the victim's login status.

Another interesting example of a side channel attack involves the use of the ambient light sensor API. This API allows web applications to detect the ambient light conditions in a room. While this feature can be useful for adjusting the UI based on lighting conditions, it can also be exploited by attackers.



By manipulating the colors displayed on the screen and reading the sensor's readings, an attacker can infer information about the user's browser history. For example, by displaying a white page and then a black page and comparing the sensor readings, the attacker can determine whether certain links have been visited or not. This can be used to gather information about the user's browsing habits and potentially compromise their privacy.

It is important to note that these side channel attacks are not intentional design decisions, but rather unintended consequences of certain features and APIs. Web application developers need to be aware of these vulnerabilities and take steps to mitigate them, such as using secure authentication mechanisms and properly handling user data.

Injection attacks, including code injection, pose a significant threat to web applications. Attackers can exploit vulnerabilities in web applications to gain unauthorized access and steal sensitive data. It is crucial for developers to understand these attack vectors and implement robust security measures to protect against them.

#### Web Applications Security Fundamentals - Injection attacks - Code injection

Code injection is a type of injection attack that involves inserting malicious code into a vulnerable application. This code is then executed by the application, leading to potential security vulnerabilities and unauthorized access to sensitive information.

One common form of code injection is known as SQL injection. In SQL injection attacks, an attacker exploits vulnerabilities in an application's database query system to manipulate the SQL statements executed by the application. By inserting malicious SQL code, an attacker can bypass authentication mechanisms, extract sensitive data, modify or delete data, and even gain control over the entire database.

Another form of code injection is command injection. In command injection attacks, an attacker exploits vulnerabilities in an application's command execution system to execute arbitrary commands on the underlying operating system. By injecting malicious commands, an attacker can perform unauthorized actions, such as executing malicious scripts, deleting files, or gaining remote access to the system.

Code injection attacks can have severe consequences, including data breaches, unauthorized access to sensitive information, and system compromise. To prevent code injection attacks, it is essential to follow secure coding practices, such as input validation and parameterized queries, to ensure that user input is properly sanitized and validated before being executed by the application.

Additionally, regular security testing, such as penetration testing and code reviews, can help identify and mitigate potential code injection vulnerabilities in web applications. It is also crucial to keep software and frameworks up to date, as vendors often release security patches to address known vulnerabilities.

Code injection is a significant security risk in web applications. By understanding the various forms of code injection attacks, implementing secure coding practices, and regularly testing and updating applications, developers can mitigate the risk of code injection vulnerabilities and ensure the security of their web applications.

#### Code Injection in Web Applications Security

Code injection is a type of attack that occurs in web applications when user-supplied data, which cannot be trusted, is combined with code written by the programmer. This combination can confuse the interpreter and result in the attacker's input being treated as a command that the programmer intended to run. One specific type of code injection is command injection, where the attacker's goal is to execute arbitrary commands on the server's operating system.

To understand command injection, let's consider an example. Imagine a script that reads a file name from the command line and produces a command to display the contents of that file. If the user enters a valid file name, the command works as intended. However, if the user enters a malicious input such as `"; rm -rf /"`, the script will interpret it as two separate commands, resulting in the deletion of all files on the server.

Command injection vulnerabilities arise when untrusted user data is used to construct shell commands without proper validation or sanitization. Attackers can exploit these vulnerabilities by injecting special characters or syntax that breaks out of the expected data context and introduces their own commands.

Mitigating command injection attacks requires careful input validation and sanitization. Developers should ensure that user-supplied data is properly validated and sanitized before using it to construct commands. This includes using parameterized queries, input validation techniques, and avoiding the direct concatenation of user input with command strings.

Code injection attacks, specifically command injection, can be a serious security risk for web applications. By combining untrusted user data with code written by the programmer, attackers can execute arbitrary commands on the server's operating system. To prevent such attacks, developers must implement proper input validation and sanitization techniques.

### Code Injection in Web Applications

Code injection is a type of injection attack that occurs when an attacker is able to insert malicious code into a vulnerable web application. This code can then be executed by the web application, leading to unauthorized actions or data breaches. One common form of code injection is known as "command injection," where an attacker is able to execute arbitrary commands on the server hosting the web application.

In a web application, code injection can occur when user input is not properly validated or sanitized before being used in dynamic code execution. This can happen, for example, when user input is concatenated directly into a command that is executed by the server. If an attacker is able to manipulate this input, they can inject their own commands and potentially gain control over the server.

To understand how code injection can occur, let's consider an example where a web application allows users to view files on the server. The application has a form where users can enter a file name, and when they submit the form, the server retrieves the contents of the specified file and displays it to the user.

In the vulnerable implementation, the server simply concatenates the user-supplied file name into a command that is executed using the `child_process.execSync` function. This function runs the command in the shell and returns the output.

However, this implementation is susceptible to code injection. An attacker can manipulate the file name input to include additional commands or special characters that can be interpreted by the shell. For example, an attacker could enter a file name like `"; rm -rf /"`, which would result in the server executing the `"rm -rf /"` command and deleting all files on the server.

To mitigate code injection vulnerabilities, it is important to properly validate and sanitize user input before using it in dynamic code execution. In the case of the vulnerable implementation described above, the server should use a different function, such as `child_process.spawnSync`, that allows for the proper escaping and handling of user input.

By using `child_process.spawnSync` with an array of arguments instead of a concatenated command string, the server can ensure that user input is properly escaped and treated as data rather than code. This helps prevent code injection attacks by separating the command and its arguments, and by properly handling special characters and escaping user input.

Code injection is a serious security vulnerability that can allow attackers to execute arbitrary code on a web application server. It occurs when user input is not properly validated or sanitized before being used in dynamic code execution. To prevent code injection attacks, it is crucial to validate and sanitize user input, and to use secure coding practices, such as using functions that handle user input properly, like `child_process.spawnSync`.

Code injection is a type of injection attack that occurs when an attacker is able to insert malicious code into a web application. This can lead to serious security vulnerabilities and potential exploitation of the application. One specific type of code injection is known as "injection attacks - code injection".

In code injection attacks, the attacker is able to inject their own code into the application, which is then executed by the server. This can happen when the application does not properly validate or sanitize user input before using it in dynamic code execution. This allows the attacker to manipulate the application's behavior and potentially gain unauthorized access or perform malicious actions.

One example of code injection is command injection. In command injection, the attacker is able to inject their own commands into the application, which are then executed by the server. This can lead to the execution of arbitrary commands on the server, potentially allowing the attacker to gain control over the system.

Another example of code injection is SQL injection. In SQL injection, the attacker is able to inject their own SQL queries into the application, which are then executed by the database. This can lead to unauthorized access to the database, manipulation of data, and even execution of arbitrary commands on the underlying operating system.

To protect against code injection attacks, it is important to implement proper input validation and sanitization. This involves validating user input to ensure that it conforms to expected formats and does not contain any malicious code. Additionally, it is important to use parameterized queries or prepared statements when interacting with databases, as this helps to prevent SQL injection by separating the query from the user input.

Code injection attacks, such as command injection and SQL injection, can pose serious security risks to web applications. It is crucial to implement proper input validation and sanitization techniques to prevent these types of attacks. By following best practices and using secure coding practices, developers can help protect their applications from code injection vulnerabilities.

Injection attacks, specifically code injection, are a common vulnerability in web applications that can lead to unauthorized access and manipulation of data. In code injection attacks, an attacker exploits vulnerabilities in the application's code to inject malicious code that is executed by the server.

One type of code injection attack is SQL injection. SQL injection occurs when an attacker is able to manipulate the SQL queries executed by the server. This can be done by injecting malicious SQL code into user input fields that are not properly validated or sanitized.

To understand how SQL injection works, let's consider an example. Suppose we have a web application that allows users to search for other users by their username. The application constructs an SQL query based on the user's input and retrieves the corresponding user information from the database.

If the application does not properly validate and sanitize the user input, an attacker can exploit this vulnerability by injecting malicious SQL code. For example, the attacker can input a username followed by a single quote, which can cause a syntax error in the SQL query. This error can reveal valuable information about the application's database structure and potentially allow the attacker to extract sensitive data.

To mitigate this vulnerability, developers should employ proper input validation and sanitization techniques. One common technique is to use parameterized queries or prepared statements, which separate the SQL code from the user input and ensure that the input is treated as data rather than executable code.

In addition to SQL injection, code injection attacks can also occur in other contexts, such as command injection. In command injection attacks, an attacker exploits vulnerabilities in the application's command execution mechanism to execute arbitrary commands on the server.

To protect against command injection attacks, developers should carefully validate and sanitize user input before using it to construct commands. It is important to avoid directly concatenating user input into command strings and instead use proper escaping or parameterization techniques.

Code injection attacks, including SQL injection and command injection, are significant security vulnerabilities in web applications. Developers must implement proper input validation and sanitization techniques to prevent these attacks and protect sensitive data.

When developing web applications, it is crucial to consider security measures to protect user data and prevent

unauthorized access. One common vulnerability is code injection, where an attacker inserts malicious code into an application's input fields, leading to potential exploitation.

In the provided material, we see an example of a code injection vulnerability. The application checks for errors in the database and whether a user exists. If no user is found, the application displays a "failed to log in" message. However, there are several issues with this approach.

Firstly, the code does not properly sanitize user input. This allows an attacker to manipulate the input fields and execute arbitrary SQL queries. For example, by entering "Bob' --" in the username field, the attacker gains access to Bob's account without providing a password. This is possible because the code does not validate the input and blindly executes the query.

Additionally, the code uses a vulnerable method called "DB get," which only retrieves one result from the database. This means that if an attacker enters a username like "Bob' OR 1=1 --," the query will return all rows from the database, but the code will only take the first result. Consequently, the attacker gains access to the first user's account in the database, regardless of the username entered.

Furthermore, the code lacks protection against multiple queries. An attacker could exploit this by appending a semicolon and injecting additional queries. For example, by entering "'; UPDATE users SET password='root' WHERE username='Bob'; --" in the username field, the attacker attempts to change Bob's password to "root." However, the code in this case prevents the execution of multiple queries, safeguarding against this specific attack.

It is important to note that not all libraries or APIs have built-in protections against code injection vulnerabilities. Developers must be aware of the limitations of the tools they use and implement proper input validation and sanitization techniques to mitigate these risks. Additionally, logging user login attempts can be beneficial for security auditing purposes, but it should be done carefully to avoid storing sensitive information.

To prevent code injection attacks, developers should follow secure coding practices such as:

1. Input validation and sanitization: Validate and sanitize all user input to ensure it adheres to the expected format and does not contain any malicious code. Use parameterized queries or prepared statements to prevent SQL injection.
2. Principle of least privilege: Ensure that database users have the minimum required permissions to access and modify data. Limiting privileges reduces the potential impact of an attacker gaining unauthorized access.
3. Use secure coding libraries and frameworks: Utilize well-established libraries and frameworks that have built-in security features and protections against common vulnerabilities.
4. Regularly update and patch software: Keep all software components up to date with the latest security patches to address any known vulnerabilities.
5. Conduct security testing and code reviews: Perform regular security testing, including penetration testing and code reviews, to identify and fix any potential vulnerabilities.

By implementing these best practices, developers can significantly enhance the security of web applications and protect user data from code injection attacks.

Injection attacks are a common type of cybersecurity vulnerability that can occur in web applications. One specific type of injection attack is code injection, where an attacker is able to insert malicious code into a web application's codebase.

In code injection attacks, the attacker takes advantage of vulnerabilities in the web application's input validation mechanisms to inject malicious code into the application's codebase. This can happen when the application does not properly sanitize user input or fails to validate input data before using it in code execution.

One example of code injection is SQL injection, where an attacker is able to inject malicious SQL code into a web application's database queries. This can allow the attacker to manipulate the application's database, access

sensitive information, or even execute arbitrary commands on the underlying server.

In a SQL injection attack, the attacker typically tries to exploit vulnerabilities in the application's database queries by injecting SQL code into user input fields. For example, an attacker could try to manipulate a login form by injecting SQL code that alters the application's database or bypasses authentication mechanisms.

To prevent SQL injection attacks, it is important to implement proper input validation and sanitization techniques. This includes using parameterized queries or prepared statements, which separate user input from the actual SQL code and automatically handle proper escaping and encoding of input data.

Additionally, it is important to regularly update and patch web applications to address any known vulnerabilities that could be exploited by code injection attacks. Regular security audits and penetration testing can also help identify and address potential vulnerabilities before they can be exploited.

Code injection attacks, such as SQL injection, pose a significant threat to the security of web applications. By understanding the fundamentals of code injection and implementing proper security measures, developers and system administrators can help protect their web applications from these types of attacks.

Injection attacks, specifically code injection, are a significant threat to web application security. In this context, code injection refers to the insertion of malicious code into a web application's database query, which can lead to unauthorized access or manipulation of data. This didactic material will explore the concept of code injection, its implications, and how it can be exploited.

To understand code injection, let's consider an example. Suppose we have a web application that allows users to log in with a username and password. The application uses a database to store user information, including passwords. Our goal is to determine the first letter of a user's password by exploiting a vulnerability in the application.

The first step in this attack is to craft a specially designed input that triggers the vulnerability. In this case, we want to determine if the first letter of the password is 'P'. If it is, we will execute a slow query, and if it's not, we will execute a fast query. The difference in execution time will allow us to infer the correct answer.

To achieve this, we need to construct a SQL query that incorporates our malicious code. One approach is to create a slow SQL expression that takes a significant amount of time to execute. For example, we can convert a large blob of data into a hexadecimal string, convert it to uppercase, and then compare it to a specific value. This process intentionally introduces a noticeable delay.

To implement this in code, we can use an IF statement in SQL. If the expression evaluates to true, we execute the slow query; otherwise, we execute the fast query. By measuring the execution time, we can determine if the first letter of the password is 'P' or not.

Putting it all together, our query would look like this:

```
IF (SUBSTRING(password, 1, 1) = 'P') THEN
-- Slow query
...
ELSE
-- Fast query
...
END IF;
```

By running this query and observing the timing, we can deduce the first letter of the password. If the slow query takes a noticeable amount of time, it means the first letter is 'P'. Otherwise, it is a different letter.

It's important to note that the success of this attack relies on the vulnerability in the web application that allows the injection of code into the database query. Preventing code injection requires implementing proper input validation and parameterization techniques to ensure that user-supplied data is treated as data, not as executable code.

Code injection attacks pose a significant risk to web application security. By exploiting vulnerabilities in the application, attackers can inject malicious code into database queries, leading to unauthorized access or data manipulation. Understanding the techniques used in code injection attacks is crucial for developers and security professionals to build secure web applications.



**EITC/IS/WASF WEB APPLICATIONS SECURITY FUNDAMENTALS DIDACTIC MATERIALS****LESSON: TLS ATTACKS****TOPIC: TRANSPORT LAYER SECURITY**

TLS attacks, also known as Transport Layer Security attacks, are a significant concern in web application security. In this material, we will discuss the fundamentals of TLS attacks and explore some solutions to mitigate them.

Sequel injection is a common vulnerability that can lead to TLS attacks. In sequel injection, the application server allows users to modify the sequel query, which can result in unauthorized access to data. Even if the user cannot directly control what is displayed on the webpage, there are other ways sequel injection can be exploited.

The core problem with sequel injection is that the application server is responsible for deciding which queries to run on the database. It is crucial for the server to ensure that users can only make queries that align with their permissions. However, if the user can modify the sequel query, they can create any query they want, potentially accessing unauthorized data.

One possible solution to sequel injection is using parameterized sequel. Instead of building sequel queries using string concatenation, developers should use a function that combines untrusted user input with the query. The function will insert the user input in the correct place, ensuring it is escaped correctly for the context where it will be inserted. By using parameterized sequel, developers can prevent sequel injection vulnerabilities.

Another option to mitigate sequel injection is using an Object Relational Mapper (ORM). An ORM allows developers to represent each row of the database as an object in their object-oriented programming language. While the primary purpose of an ORM is to provide an easier way to access data in the database, it also offers the benefit of automatically escaping user input, reducing the risk of sequel injection.

TLS attacks, specifically sequel injection, can pose a significant threat to web application security. By implementing measures such as parameterized sequel and using an ORM, developers can mitigate the risk of sequel injection vulnerabilities and enhance the overall security of their web applications.

In the field of cybersecurity, it is crucial to understand the fundamentals of web application security, including the concept of Transport Layer Security (TLS) and the potential attacks that can occur within this framework. TLS is a protocol that ensures secure communication over a network, commonly used to protect sensitive information transmitted over the internet.

Before delving into TLS attacks, it is important to understand the basics of web application security. In a typical scenario, a web application interacts with a database through objects, such as a "person" object. These objects, although represented as entities within the application, are ultimately mapped to database rows. This mapping allows for the execution of queries on the database, such as sorting the "users" table based on certain criteria like username and password.

When dealing with user input, it is crucial to ensure the security of the application. By utilizing functions that escape untrusted user input, the risk of SQL injection attacks can be mitigated. However, it is important to note that even in a NoSQL database scenario, where queries are represented as objects rather than strings, the same security concerns arise when allowing users to select objects freely.

To address these security concerns, frameworks like Mongoose provide built-in functionalities to handle input sanitization. By relying on these frameworks, developers can minimize the risk of SQL injection attacks and ensure the security of their applications.

Moving on to the topic of TLS, it is essential to understand why HTTP, the unencrypted counterpart to HTTPS, is considered insecure. In HTTP, the entire communication between the client (browser) and the server is visible to anyone on the network, including routers and ISPs. This lack of encryption poses a significant security risk, as sensitive information, such as usernames, passwords, and session cookies, can be intercepted by passive attackers.

Passive attackers observe the traffic passing through the network without actively modifying it. By intercepting HTTP requests and responses, they can gain access to valuable information, compromising the security of the application and the user's data.

In more severe cases, active attackers can actively manipulate the requests and responses they intercept. Although this requires more effort, readily available software allows attackers to modify the traffic passing through a network. This type of attack can be particularly dangerous, as it enables attackers to modify the content of requests and responses, potentially leading to unauthorized access or data manipulation.

To address these security concerns, the implementation of TLS is crucial. TLS provides a secure channel for communication between the client and the server, encrypting the data exchanged during the session. By utilizing TLS, sensitive information, such as usernames, passwords, and session cookies, is protected from interception by passive attackers.

However, it is important to note that TLS itself can be vulnerable to attacks if not implemented correctly. Inadequate configuration or improper usage of TLS can expose the application to various vulnerabilities. Therefore, it is essential to understand the potential pitfalls and best practices associated with TLS implementation to ensure the security of web applications.

Web application security and the implementation of TLS are critical aspects of cybersecurity. By understanding the fundamentals of web application security and the potential attacks that can occur within the TLS framework, developers can take proactive measures to protect their applications and users' data.

In the context of web application security, one of the fundamental aspects to consider is the security of the transport layer, specifically the Transport Layer Security (TLS) protocol. TLS is responsible for providing secure communication between a client and a server over a network, ensuring privacy, integrity, and authentication.

When a client sends a request to a server, it passes through various network infrastructure components, including routers and internet service providers (ISPs). However, these components can be compromised by network attackers who have the ability to intercept, modify, or eavesdrop on the communication.

In a scenario where an attacker is present, they can act as a proxy server intercepting the client's request. The attacker then modifies the HTML response received from the actual server, injecting their own malicious JavaScript code. This modified response is then forwarded to the client, who remains unaware of the attack. Consequently, the client unknowingly executes the attacker's code, leading to potential information theft or other malicious actions.

The threat model in this context revolves around network attackers, which can be anyone who controls the network infrastructure, such as routers or ISPs. These attackers can passively eavesdrop on communication or actively manipulate packets, inject additional packets, modify packet content, or even control the timing of packets. Examples of places where such attacks can occur include wireless networks in cafes or hotels, as well as at national borders where traffic can be tampered with.

The primary goal in web application security is to establish secure communication in the presence of these network attackers. To achieve this, three essential properties need to be ensured: privacy, integrity, and authentication.

Privacy guarantees that communication remains confidential, preventing passive attackers from intercepting and understanding the exchanged data. Integrity ensures that the messages have not been tampered with. Even if an attacker cannot directly view the communication, tampering with the response can lead to the execution of malicious code on the client-side. Lastly, authentication is crucial to verify the identity of the server, ensuring that the client is indeed communicating with the intended party. Without authentication, the other two properties become meaningless, as the client may unknowingly communicate with an attacker.

To achieve these three properties, TLS is employed. TLS provides a secure communication channel by encrypting the data exchanged between the client and server. It uses cryptographic algorithms to ensure privacy and integrity. Additionally, TLS employs certificates to authenticate the server's identity, allowing the client to verify its authenticity.

Ensuring the security of web applications requires protecting the transport layer through the use of TLS. By implementing TLS, we can establish secure communication channels that provide privacy, integrity, and authentication, even in the presence of network attackers.

In the context of web application security, Transport Layer Security (TLS) plays a crucial role in ensuring secure communication between clients and servers. TLS is commonly used with HTTP, resulting in the familiar HTTPS protocol. However, TLS can also be applied to other protocols such as email and instant messaging.

One fundamental aspect of TLS is the anonymous Diffie-Hellman key exchange. In this protocol, there is no authentication of the server, meaning it does not protect against active attackers who can modify packets. Instead, it focuses on preventing passive eavesdroppers from gaining any information.

The key exchange begins with a browser and a server. Both parties agree on a cyclic group, which can be thought of as a set of numbers. This group is defined in a standard and is publicly known to all parties involved, including browser makers, attackers, and servers. Group operations can be performed within this cyclic group, such as multiplication and exponentiation.

To initiate the key exchange, the client and server each randomly select a group element. These elements are denoted as "a" for the client and "b" for the server. The client then sends  $G^a$  ( $G$  raised to the power of  $a$ ) to the server, while the server sends  $G^b$  to the client. It is important to note that  $G^a$  and  $G^b$  still belong to the cyclic group  $G$ .

The objective is to generate a shared key that both the client and server possess. However, an observer in the middle, who intercepts these messages, cannot derive the shared key. The client and server can now agree on this shared key through further steps, which will be discussed later.

It is worth mentioning that TLS not only ensures secure communication but also verifies that responses come from the intended source. This prevents attackers from impersonating the server and sending malicious responses. The process of verifying the source will be explored in subsequent discussions.

TLS is a vital component of web application security, providing encryption and secure communication between clients and servers. The anonymous Diffie-Hellman key exchange is one of the fundamental mechanisms used in TLS to establish a shared key. While it does not authenticate the server, it protects against passive eavesdroppers. The next topic of discussion will delve into the process of verifying the source of responses.

Transport Layer Security (TLS) is a crucial component of web applications security. It ensures secure communication between a client and a server by encrypting the data exchanged between them. However, TLS is not immune to attacks. In this didactic material, we will focus on understanding TLS attacks and how they can compromise the security of web applications.

One type of TLS attack is known as a Man-in-the-Middle (MITM) attack. In a MITM attack, an attacker secretly intercepts and modifies the communication between the client and the server. This allows the attacker to eavesdrop on the data being exchanged or even manipulate it.

To understand how a MITM attack can be carried out, let's consider the Diffie-Hellman key exchange protocol, which is commonly used in TLS. The Diffie-Hellman protocol allows two parties, the client and the server, to securely establish a shared secret key without exchanging it directly.

In the Diffie-Hellman protocol, both the client and the server agree on a common group element, denoted as  $G$ . They each select a random private value, denoted as  $a$  for the client and  $b$  for the server. The client and the server then compute a public value by raising  $G$  to the power of their respective private values. The client sends its public value, denoted as  $A$ , to the server, and the server sends its public value, denoted as  $B$ , to the client.

To derive the shared secret key, each party takes the public value received from the other party and raises it to the power of its own private value. This results in both the client and the server deriving the same shared secret key, denoted as DH key.

In a MITM attack on the Diffie-Hellman key exchange, the attacker intercepts the communication between the client and the server. The attacker can see the public values exchanged between them, but not their private

values. However, the attacker can generate its own private value, denoted as  $C$ , and compute a public value by raising  $G$  to the power of  $C$ .

When the client sends its public value  $A$  to the server, the attacker intercepts it and sends its own public value, denoted as  $G$  to the  $C$ , to the client instead. Similarly, when the server sends its public value  $B$  to the client, the attacker intercepts it and sends its own public value, denoted as  $G$  to the  $C$ , to the server instead.

As a result, the client and the attacker derive a shared secret key, denoted as DH key 1, based on the public value  $G$  to the  $C$ . Simultaneously, the server and the attacker derive another shared secret key, denoted as DH key 2, based on the public value  $G$  to the  $B$   $C$ .

At this point, the attacker can decrypt and manipulate the data exchanged between the client and the server. The client, however, remains unaware of the presence of the attacker and believes it is communicating securely with the server.

This type of attack highlights the importance of authentication in TLS. Without proper authentication, the client cannot be certain if it is communicating with the intended server or an attacker. In the case of anonymous Diffie-Hellman key exchange, where the client and the server do not authenticate each other, the communication lacks authentication.

TLS attacks, such as the Man-in-the-Middle attack on anonymous Diffie-Hellman key exchange, demonstrate the vulnerabilities in web applications security. These attacks exploit weaknesses in the protocols and can compromise the confidentiality and integrity of the data exchanged between the client and the server.

Transport Layer Security (TLS) is a crucial component of web applications security. However, it is susceptible to attacks, such as man-in-the-middle attacks. In a man-in-the-middle attack, an attacker intercepts the communication between the client and the server, posing as the server to the client and as the client to the server. This allows the attacker to eavesdrop on the communication and even modify the data without detection.

To prevent man-in-the-middle attacks, authentication is essential. One way to achieve authentication is by using public key cryptography. In this scheme, a generator algorithm ( $G$ ) generates a public key and a secret key. The public key is widely distributed, while the secret key is kept private. The signing algorithm ( $S$ ) takes the secret key and an input ( $X$ ) and produces a tag ( $T$ ). The verification algorithm ( $V$ ) takes the public key, the input ( $X$ ), and the tag ( $T$ ) to confirm the authenticity of the tag.

To illustrate the concept, consider a scenario where a party wants to announce something to the world and wants everyone to know it is them who said it. The party generates a public key and a secret key using the generator algorithm. They keep the secret key secret and publish the public key. When the party wants to make a statement, they use the signing algorithm with their secret key to produce a tag. They then post the tag along with the statement. Anyone who comes across the statement can use the verification algorithm with the public key, the statement, and the tag to confirm that it was indeed said by the party.

To apply authentication to the Diffie-Hellman key exchange in TLS, the server generates a public key and a secret key using the generator algorithm. The server keeps the secret key private and announces the public key. The client already knows the public key. The client and the server proceed with the Diffie-Hellman key exchange as before. However, before the server sends a response, it creates a transcript of the entire exchange, including the client's message, and signs it using the secret key.

By incorporating signature schemes into the Diffie-Hellman key exchange, we achieve authenticated Diffie-Hellman key exchange. This ensures that the client securely derives a shared key only with the intended server. The signature scheme provides authentication by confirming that the tag was generated by the server using its secret key.

TLS attacks, such as man-in-the-middle attacks, can be mitigated by adding authentication to the Diffie-Hellman key exchange. By using signature schemes based on public key cryptography, we can ensure that the communication between the client and the server is secure and that the server's identity is verified.

Transport Layer Security (TLS) is a crucial component of web application security. It provides secure

communication between clients and servers by encrypting data and verifying the authenticity of the server. However, TLS attacks can compromise the security of web applications. In this didactic material, we will discuss the fundamentals of TLS attacks and how they can be mitigated.

One common attack on TLS is the manipulation of the key exchange process. In a typical TLS handshake, the client and server exchange public keys to establish a shared secret key. This shared key is then used to encrypt and decrypt data during the session. However, an attacker can intercept the key exchange and manipulate the exchanged keys, leading to a compromised session.

To prevent such attacks, TLS employs a technique called digital signatures. During the key exchange, the server signs the exchanged public key with its secret key, producing a tag. The server then sends the signed key and tag to the client. Before deriving the shared key, the client verifies the tag's validity by calling the verification function on the public key with the transcript and tag received from the server. If the tag is valid, it indicates that the server is the owner of the secret key and ensures the integrity of the key exchange.

By verifying the tag, the client can detect if it is communicating with a man-in-the-middle attacker. If the tag is invalid, the client can reject the connection and prevent further compromise. This validation step adds an extra layer of security to the TLS handshake.

It is important to note that this validation process is one-way authentication. The server does not authenticate the client during the key exchange. However, this is not a significant concern because once the shared key is derived, the client can securely send its credentials, such as a username and password, to the server over the encrypted connection. The server can then validate the credentials at the web application layer.

Now, let's discuss how the client obtains the server's public key. Including the server's public key in the browser is not a viable option due to the large number of websites and the constant changes in public keys. Another approach could be for the server to send its public key to the client during the key exchange process. However, this approach poses a security risk. If the server sends the public key to the client, the client cannot verify the authenticity of the key, as an attacker could send their own public key instead.

To address this issue, TLS relies on the use of Certificate Authorities (CAs). CAs are trusted entities that issue digital certificates, which contain the server's public key and other identifying information. The client's browser or operating system has a pre-installed list of trusted CAs. During the TLS handshake, the server sends its digital certificate to the client. The client then verifies the certificate's authenticity by checking its signature against the trusted CAs. If the certificate is valid, the client can extract the server's public key from the certificate and proceed with the key exchange.

TLS attacks can compromise the security of web applications. To mitigate these attacks, TLS employs techniques such as digital signatures and the use of trusted Certificate Authorities. These measures ensure the authenticity of the server and the integrity of the key exchange process, providing a secure communication channel between clients and servers.

### Certificate Authorities in Web Applications Security

In the context of web applications security, one important aspect is the secure exchange of keys between the client and the server. This is where certificate authorities (CAs) come into play. A certificate authority is an entity that issues digital certificates to site owners. These certificates certify that a specific subject is the owner of a particular public key.

The role of a certificate authority is to vouch for the authenticity of the public key. By trusting a small number of certificate authorities, we can establish a secure way to verify the ownership of a public key. When a server sends its public key to a client, it also includes a statement from the certificate authority vouching for the correctness of the key. The client can then trust that the public key received is indeed the correct one.

To see the details of certificates sent by servers, one can click on the lock icon in the browser and then select "more information." In the details, there are several fields, but the most important one for the browser is the common name field. The browser compares the common name field with the URL the user is visiting to determine if the certificate is valid for that site.

Another important field in the certificate is the issuer field, which indicates the certificate authority that issued the certificate. In some cases, the certificate authority is the same entity as the site owner, while in others, it can be a trusted third party.

In addition to the common name field, there is also the alternate subject name field. This field allows for the inclusion of multiple domains without using a wildcard. It provides a way to specify additional domains that the certificate is valid for.

By relying on certificate authorities and their issued certificates, web applications can establish a secure and trusted connection between the client and the server. This ensures that the public key received by the client is indeed the correct one for the site being visited.

Transport Layer Security (TLS) is a fundamental aspect of web application security. It ensures secure communication between clients and servers by encrypting data and providing authentication. However, TLS attacks can compromise this security and allow attackers to intercept and manipulate data.

One type of TLS attack involves compromising the trust in Certificate Authorities (CAs). CAs are entities that issue digital certificates, which are used to verify the authenticity of websites. When a user visits a website, their browser checks if the website's certificate is signed by a trusted CA. If it is, the browser trusts the website and establishes a secure connection.

In some cases, multiple domain names may be associated with a single certificate. For example, the certificate for google.com may also be valid for gmail.com or googlemail.com. Browsers compare the website's domain name against all the alternate subject names listed in the certificate.

To determine which CAs a browser trusts, there is a built-in list in the browser settings. In Firefox, for example, you can view the list of trusted CAs. It is interesting to explore this list as it contains various CAs with the power to issue certificates that your browser will trust. Alongside well-known CAs like Google Trust Services LLC, there are other interesting ones such as Hong Kong Post, which can issue certificates for any site.

Furthermore, some organizations have multiple root certificates that are trusted by the browser. By expanding the details of an organization in the list, you can see the different root certificates associated with it. This can help identify the organization behind the certificate.

Different browsers handle trusted CAs differently. Chrome and Safari, for example, rely on a certificate store built into the operating system. When you receive a computer directly from the factory, it comes with a hard-coded list of trusted CAs. Chrome, Safari, and other Chromium-based browsers refer to this certificate store for trust decisions.

Removing trusted CAs from your browser can have consequences. Any site that has a certificate issued by a removed CA will trigger warnings in your browser. Therefore, it is important to consider the risk and reward before deleting trusted CAs.

Attackers can exploit compromised CAs to conduct man-in-the-middle attacks. If an attacker manages to add their own key as a trusted key in your CA store, they can issue a certificate for a targeted website. When you visit that website, the attacker intercepts your request and sends back a fake page with their own certificate. Since your browser trusts their certificate, you will unknowingly communicate with the attacker instead of the legitimate website.

This highlights the significance of maintaining the integrity of the trusted CA list. Employers, for example, may add a trusted CA to the certificate store on company-issued laptops to monitor network traffic for security purposes. Similarly, some companies use network appliances to actively perform man-in-the-middle attacks for security inspection.

Understanding the trust in CAs and the potential risks associated with compromised trust is crucial in maintaining web application security. By exploring the trusted CA list in your browser and being aware of the implications of removing trusted CAs, you can better protect yourself from TLS attacks.

In the context of web application security, Transport Layer Security (TLS) plays a crucial role in ensuring secure



communication between clients and servers. However, there are instances where organizations may want to inspect the encrypted traffic flowing through their network. This could be to identify potential malware or unauthorized data exfiltration. To achieve this, organizations may consider breaking TLS, although this approach raises concerns about compromising the security provided by TLS.

One of the arguments for breaking TLS is the need for organizations to inspect requests and make security decisions. By decrypting the traffic, organizations can analyze the content and detect any malicious activity or policy violations. For example, they can identify if an employee is uploading company data to an unauthorized server or sharing sensitive information through platforms like Dropbox. Despite these potential benefits, the decision to break TLS for inspection purposes is still debated due to the potential risks it introduces.

When implementing TLS, a key component is the certificate authority (CA). The server, in this case, needs a certificate to establish secure communication. The server generates a public key and a secret key for itself, but the client does not trust this public key. To address this issue, the server sends its public key to the CA for validation. The CA requires proof that the server is the legitimate owner of the domain. This proof can be provided through various means, such as modifying the DNS settings or adding a specific file to the domain's website. Once the CA is satisfied with the proof, it signs a message stating that the domain has a public key and sends it back to the server as a certificate.

When a client connects to the server, the server sends the certificate to the client during the Diffie-Hellman key exchange. The client then validates the certificate using the public key of the trusted CA. The validation process involves verifying if the certificate originated from the trusted CA and if it corresponds to the domain the client intended to communicate with. If the certificate passes these checks, the client can trust the public key provided by the server and proceed with the secure communication.

It is important to note that the certificate exchange process only occurs once, unless the certificate expires. This ensures that subsequent connections between the client and server can be established without repeating the certificate exchange process.

TLS attacks and the decision to break TLS for inspection purposes are complex topics in web application security. While organizations may have legitimate reasons to inspect encrypted traffic, it is crucial to carefully consider the potential risks and trade-offs involved in compromising the security provided by TLS.

Transport Layer Security (TLS) is a crucial component of web applications security. It provides encryption and authentication for data transmitted between a client and a server. However, TLS is not immune to attacks. In this didactic material, we will explore some TLS attacks and their impact on web application security.

One potential attack on TLS is the disruption of the certificate issuance process. The server relies on certificates to prove its identity to clients. If the server fails to obtain a certificate, it cannot establish trust with its users. This disruption can be caused by various factors, such as network issues or malicious interference. In recent years, the introduction of Let's Encrypt, a free certificate authority, has made certificate issuance more accessible. However, if the process is disrupted, the server may be unable to provide certificates to clients, leading to potential security risks.

To mitigate the risks associated with disrupted certificate issuance, Let's Encrypt has implemented a command-line tool that automates the certificate retrieval process. This tool interacts with Let's Encrypt servers and returns the certificate to the user. By automating the process, Let's Encrypt aims to reduce the possibility of human error and make the issuance more efficient. However, the use of automated tools also means that certificates tend to have shorter expiration periods, typically around three months. This approach enhances security by limiting the potential damage caused by compromised certificates. If the automated process is disrupted, such as by a failed cron job, the website may go offline, impacting user experience.

TLS 1.3 is the latest version of the TLS protocol. It has replaced previous versions due to their inherent vulnerabilities. SSL, a predecessor to TLS, is no longer supported by modern browsers. TLS 1.0 and 1.1 are also on their way to being deprecated. Browsers have implemented measures to prevent communication with servers using older TLS versions, as they are considered insecure. TLS 1.3, on the other hand, is currently considered secure and has not exhibited any significant problems.

TLS 1.3 consists of two phases. The first phase involves establishing a shared secret between the client and

server. This process ensures mutual authentication and confidentiality. The second phase focuses on encrypting the data exchanged between the client and server using the established shared secret. This encryption provides confidentiality and integrity for the transmitted data.

TLS attacks can pose significant risks to web application security. Disruptions in the certificate issuance process can lead to compromised trust between the server and clients. TLS 1.3, the latest version of the protocol, addresses many of the vulnerabilities present in previous versions. By understanding these attacks and the measures in place to mitigate them, web application developers and security professionals can better protect their systems and data.

Transport Layer Security (TLS) is a crucial component of web application security. It provides a secure channel for communication between a client and a server, ensuring that data transmitted over the internet remains confidential and tamper-proof. In this context, it is important to understand how TLS attacks can compromise this security.

When implementing TLS, a key is used to encrypt and decrypt data. This key is put through an encryption algorithm to ensure secure communication. When HTTPS is added on top of HTTP, it provides an additional layer of security, indicated by the lock symbol in the browser. However, the browser needs to determine when to display this lock symbol and when not to.

The rules for displaying the lock symbol are more complex than just checking the initial HTML page. This is because web pages often contain references to scripts, images, and resources from other sites. To ensure complete security, the lock symbol should only be displayed if every element on the page is fetched using HTTPS. If even a single element is fetched using HTTP, it opens up the possibility of a man-in-the-middle attack.

In a man-in-the-middle attack, an attacker can modify the HTTP response for a script, allowing them to run their own script on the page. This compromises the security even if the main page was fetched using HTTPS. To prevent this, every element from different sites that the browser connects to must undergo a series of checks. These checks include verifying that the certificate was issued by a trusted Certificate Authority (CA), ensuring that the certificate is not expired, and matching the common name or subject alternate name with the URL of the origin. Only when all these checks pass for every element, the lock symbol is displayed to the user, providing them with a sense of security.

However, there are criticisms regarding the issuance of certificates by Certificate Authorities. Some argue that Certificate Authorities only focus on verifying domain ownership and do not consider if a site may be used for phishing or is similar to another site. To address this, Google Safe Browsing is used to mark known phishing sites, and the browser displays a warning. Additionally, Chrome has an experiment that utilizes machine learning to detect typosquatting, where a site's URL is similar to a popular site, and warns the user if they may have intended to visit a different site.

TLS 1.3 is the latest version of TLS and provides enhanced security features. In the implementation of TLS 1.3, specific messages are exchanged between the client and server. These messages include the "hello" message, "server hello" message, and the key share part, which involves the Diffie-Hellman key exchange. A nonce, a random number used only once, is also included in the messages.

The inclusion of a nonce is important to prevent replay attacks. Without a nonce, an attacker who observes the message sent to the server could replay it, potentially compromising the security. This becomes particularly significant when implementing zero round-trip time (zero RTT) to improve performance. Zero RTT reduces the number of back-and-forth exchanges between the client and server, but it introduces the risk of replay attacks. By including a nonce, the risk of replay attacks is mitigated.

Understanding TLS attacks and the security measures in place is crucial for web application security. By implementing TLS correctly and following the necessary security checks, we can ensure secure communication between clients and servers, protecting sensitive information from unauthorized access.

In the context of web applications security, Transport Layer Security (TLS) plays a vital role in ensuring secure communication between clients and servers. TLS provides encryption, authentication, and integrity for data transmitted over the internet. However, it is important to understand the fundamentals of TLS attacks in order to effectively protect web applications.

One type of TLS attack is known as a replay attack. In a replay attack, an attacker intercepts a message sent from the client to the server and later re-sends that message to the server. This can cause the server to perform an action multiple times, potentially leading to unwanted consequences. To prevent replay attacks, TLS uses a nonce, which is a unique value included in the initial message from the client to the server. If the nonce is missing, an attacker could intercept the message and resend it, causing the server to perform the action twice.

Another important aspect of TLS is the use of certificates. When the server receives a request from the client, it sends a certificate that has been encrypted with a shared secret. This encryption ensures that a passive observer cannot determine the server's identity by inspecting the certificate. By encrypting the certificate, TLS protects the confidentiality of the communication.

In addition to encryption, TLS also utilizes data signing. The data that is signed includes the transcript of all the communication that has taken place so far. By signing the data, TLS ensures that the communication cannot be tampered with or modified by a man-in-the-middle attacker. The data signing process adds an extra layer of security to the communication.

Once the communication is complete, TLS establishes session keys using a key derivation function. These session keys are used to encrypt subsequent HTTP requests. By using session keys, TLS ensures that each session is unique and provides forward secrecy. Forward secrecy is a property of TLS that prevents an attacker from decrypting past communication even if they obtain the secret key used for encryption.

TLS attacks, such as replay attacks, can pose a threat to web application security. By implementing nonces, encrypting certificates, signing data, and using session keys, TLS safeguards against these attacks and provides secure communication between clients and servers.

Transport Layer Security (TLS) is a fundamental aspect of web application security. It provides encryption and authentication for data transmitted between a client and a server. TLS ensures that the communication remains confidential and secure, protecting sensitive information such as passwords and credit card details.

One important feature of TLS is forward secrecy. This means that even if an attacker manages to obtain the private key used for encryption, they cannot decrypt previously recorded traffic. This is because TLS uses a different session key for each session, making it virtually impossible for an attacker to decrypt past communications.

TLS also provides identity protection by encrypting the server's certificate. However, it is important to note that while a passive observer may not be able to determine the specific certificate being used, they can still identify the server and the IP address it is communicating with. This can be revealing, especially if a single IP address hosts only one website.

Additionally, if DNS requests are not encrypted, an observer can see the specific DNS lookup performed before the encrypted communication. Therefore, it is crucial to encrypt DNS requests as well to ensure complete privacy.

TLS also includes server-side authentication, where the client verifies the identity of the server. However, client certificates, which allow the server to verify the identity of the client, are rarely used in practice. Although client certificates provide an additional layer of security, the complexity and lack of familiarity with installing certificates on browsers limit their adoption.

The adoption of HTTPS, which is the secure version of HTTP using TLS, has significantly increased in recent years. According to the Google HTTP Transparency Report, almost all of the top 100 websites (excluding those owned by Google) now support HTTPS. This positive trend can be attributed to browser vendors threatening to label websites without HTTPS as "not secure" and the availability of free certificates from Let's Encrypt.

TLS plays a crucial role in securing web applications by providing encryption, forward secrecy, identity protection, and authentication. The widespread adoption of HTTPS has greatly improved the security of online communication.

Transport Layer Security (TLS) is a crucial component of web applications security. It ensures secure

communication between clients and servers, protecting sensitive data from unauthorized access or tampering. However, TLS itself is not immune to attacks. In this didactic material, we will explore some common TLS attacks and their implications.

One notable attack on TLS is the compromise of Certificate Authorities (CAs). CAs play a vital role in the TLS ecosystem by issuing digital certificates that authenticate the identity of websites. If a CA is compromised, it can issue fraudulent certificates for any domain, undermining the security of all websites on the internet. To mitigate this risk, the concept of intermediate CAs was introduced. Intermediate CAs are authorized by top-level CAs to issue certificates. However, they must be extremely cautious when issuing these certificates, as they effectively grant the power to issue more certificates. This ensures that only trusted entities become intermediate CAs.

To better understand the significance of CAs, let's delve into the process of becoming a CA. The exact process may vary, but it typically involves approaching a recognized CA and paying a fee. There is a mailing list where discussions and petitions take place, allowing different countries to have a say in the decision-making process. This approach aims to ensure a diverse and inclusive representation in the TLS ecosystem.

Now, let's shift our focus to the adoption of TLS across different platforms. Recent data from Google reveals the percentage of web pages loaded over HTTPS, broken down by platform. Windows users exhibit a high adoption rate, with approximately 80-90% of pages being loaded securely. However, Linux users seem to have a lower adoption rate, with a preference for unencrypted pages. Chrome OS, primarily used by Chromebook users, shows the highest adoption rate due to its close integration with Google services.

Despite the overall positive trend towards HTTPS adoption, it is interesting to note that Google's own services still have a significant number of pages loaded over HTTP. While they have made substantial progress, they have not yet achieved 100% HTTPS adoption. This discrepancy can be attributed to older mobile devices that communicate with certain APIs over HTTP. These devices are no longer receiving updates and are awaiting obsolescence. Blocking these requests or finding alternative solutions may be necessary to ensure complete HTTPS adoption.

Additionally, specific Google products, such as Google News and Google Maps, have not yet reached 100% HTTPS adoption. This could be due to factors like outdated maps applications or news sites that have not migrated to HTTPS. It is crucial to address these remaining gaps to ensure a more secure browsing experience.

While TLS provides a robust security framework for web applications, it is essential to be aware of potential vulnerabilities and attacks. The compromise of Certificate Authorities poses a significant threat to the security of all websites. Understanding the process of becoming a CA and the measures taken to ensure trust is vital. Furthermore, analyzing the HTTPS adoption rates across different platforms and addressing any remaining gaps is crucial for a more secure web.

Transport Layer Security (TLS) attacks are a serious concern in web application security. These attacks can compromise the confidentiality and integrity of data transmitted over the internet. In this didactic material, we will explore some real-world examples of TLS attacks and discuss their implications.

One notable case occurred in 2011 when a Certificate Authority (CA) in the Netherlands issued a fake certificate for Gmail. This allowed the holder of the fake certificate to perform man-in-the-middle attacks on Gmail connections. As a result, hundreds of thousands of Iranian users visiting Gmail.com were unknowingly intercepted. The CA responsible for issuing the fake certificate went out of business after major browser vendors removed their trust from their browsers. This incident highlights the severe consequences of mishandling TLS certificates.

Another case involved a company that somehow emailed the private keys of around 22,000 users to a random person. This is puzzling because the company should not have had access to the private keys in the first place. The correct process involves users sending their public keys to the company, which then creates certificates without ever seeing the users' secret keys. The company's possession of secret keys raises questions about their security practices.

Komodo, a well-known CA, has also faced security issues. One of their resellers was hacked, leading to the issuance of fake certificates for popular websites like Google, Yahoo, Skype, Mozilla, and Microsoft. Surprisingly,

the president and CEO of Komodo downplayed the incident, stating that it was just a sequel attack on a Brazilian company selling their products. This case emphasizes the importance of taking TLS attacks seriously, even if they may seem insignificant at first.

When a CA's trust is compromised, browsers may decide to remove their trust from the browser entirely. Symantec experienced this consequence due to multiple violations, even after being acquired by another company. Other browsers typically follow Mozilla's lead in such cases. These decisions are not made lightly and involve extensive discussions within the security community.

Now, let's shift our focus to an attack called TLS strip (formerly known as SSL strip). Many servers implement HTTPS and redirect all incoming requests to the HTTPS version of the site. However, if an attacker intercepts the initial unencrypted HTTP request before the redirect, they can manipulate the traffic and keep the user on the unsecured HTTP version of the site.

To illustrate this attack, let's consider a scenario. A client makes a request for the home page of example.com. The server recognizes that the request arrived over HTTP, triggering a 301 response code for redirection to the HTTPS version. The server also returns some HTML indicating the page has moved. However, the browser ignores this HTML and focuses on the location header, redirecting the user to the HTTPS version of the site. If an attacker intercepts the initial HTTP request and prevents the redirect, the user will remain on the unsecured HTTP version, exposing their data to potential attacks.

TLS attacks pose significant risks to web application security. Real-world examples demonstrate the consequences of mishandling certificates and the importance of maintaining trust in the browser ecosystem. Understanding these attacks and implementing appropriate security measures is crucial to safeguarding sensitive information.

Transport Layer Security (TLS) is a protocol used to secure communication between web applications and servers. It ensures that the data transmitted between the client and the server is encrypted and protected from unauthorized access. However, there are certain attacks that can compromise the security provided by TLS, such as TLS strip attacks.

In a TLS strip attack, there is an active attacker who intercepts the communication between the client and the server. The attacker observes the unencrypted request made by the client and passes it on to the server. The server responds by sending a redirect response to the HTTPS version of the page. Instead of forwarding this response to the client, the attacker handles it and follows up with another request to the server, this time over TLS. The server sends back a response, which the attacker modifies by changing the HTML. Specifically, the attacker modifies all the links in the HTML to point to HTTP versions of the URLs. The modified HTML is then sent to the client.

The client, unaware of the attack, receives the modified HTML. Since the communication is unencrypted, the client does not realize that the HTML has been modified. Any links clicked by the client will now make HTTP requests, which can be intercepted by the attacker. This allows the attacker to force the client to use HTTP throughout their entire session, compromising the security of the communication.

To prevent TLS strip attacks, HTTP Strict Transport Security (HSTS) can be used. HSTS allows the server to instruct the client to always use HTTPS, regardless of the protocol specified by the user. If the user enters an HTTP URL, the browser automatically adds the "s" to the URL and ensures that HTTPS is used. Similarly, if the user clicks on a link from another site or an internal link within the site that uses HTTP, the browser will rewrite the URL to use HTTPS, as instructed by the server.

By implementing HSTS, web applications can protect against TLS strip attacks and ensure that all communication is encrypted and secure.

Transport Layer Security (TLS) is a protocol used to secure communication between web applications and servers. It ensures that data transmitted between the two parties is encrypted and protected from unauthorized access.

One method used to enforce the use of TLS is through the use of the Strict Transport Security (STS) header. When a server sends an HTTP response to the browser, it includes the STS header, which instructs the browser



to only use HTTPS when communicating with that server for a specified duration. This preference is then stored by the browser and followed for the specified period.

However, there is a downside to using the STS header. It needs to be sent to the browser before it can take effect. This means that the very first request made to a site may potentially be unencrypted if the user types in HTTP. The browser cannot know in advance that it needs to use HTTPS until it receives a response from the server that includes the STS header. This is known as a trust on first use model.

A similar concept can be seen when using SSH to log into a server for the first time. The user is asked to trust the fingerprint of the server, which is a way to identify whether they are communicating with the correct server. If the fingerprint is different, it indicates a potential man-in-the-middle attack, and the user should not trust the server.

In practice, most users do not verify the fingerprint and blindly trust the server. While not ideal, this is generally safe if done from a trusted network. Once the fingerprint is trusted, SSH will remember it for future connections and raise an alert if it changes.

Returning to the STS header, once the browser trusts a server for the first time, it is considered secure for future connections. However, if the user clears their browser history, including cookies, there is a question of whether the list of sites that have requested to always use HTTPS should also be cleared. If the list is not cleared, an attacker or someone with access to the user's computer can view this information and gain insights into the user's browsing history.

On the other hand, if the list is cleared, the user will lose the information, and they can potentially be vulnerable to a man-in-the-middle attack on their next connection to any of these sites.

This presents a trade-off between privacy and security. Browsers tend to prioritize privacy in this case and clear the list when the user clears their history.

To address the trust on first use problem, browsers offer a solution called the preload list. This allows website owners to request that their site be hardcoded into the browser, instructing it to always use HTTPS even before receiving the STS header.

To enable this, the server must include the "preload" and "include subdomains" directives in the STS header. By doing so, the server gives permission for the browser to preload the header for all users, even those who have not visited the site before.

This preload list helps ensure that all requests to the site are made using HTTPS, providing an additional layer of security.

Transport Layer Security (TLS) is a crucial component of web application security. The Strict Transport Security (STS) header helps enforce the use of HTTPS by instructing the browser to only communicate with a server using HTTPS for a specified duration. However, there are trade-offs between privacy and security, such as the trust on first use model and the decision to clear the list of sites that have requested HTTPS. Browsers offer the preload list as a solution to the trust on first use problem, allowing website owners to hardcode their site into the browser to always use HTTPS.

Transport Layer Security (TLS) is a crucial aspect of web application security. It ensures the confidentiality and integrity of data transmitted between a client (such as a browser) and a server. In this material, we will discuss TLS attacks and how they can compromise the security of web applications.

One important aspect of TLS is the use of HTTPS, which stands for Hypertext Transfer Protocol Secure. HTTPS encrypts the communication between the client and the server, preventing eavesdropping and tampering with the data. To ensure that a website always uses HTTPS, web developers can opt their domains into a list maintained by browsers. This list is compiled into the browser itself and guarantees that requests to these domains will always be made using HTTPS.

It is worth noting that once a domain is added to this list, it is difficult or impossible to be removed. Therefore, web developers should carefully consider whether they want to enforce HTTPS for their site permanently. If a



domain is removed from the list, users may have difficulty connecting to the site if it is served over HTTP in the future.

To have a domain added to the HTTPS preload list, web developers can visit the HSTS Preload website and follow the instructions. The site will verify that the domain has the necessary HTTP Strict Transport Security (HSTS) header, with appropriate settings such as a minimum max age of one year, and the "include subdomains" and "preload" keywords. Once the domain passes the verification process, it will be added to the list of sites to be preloaded in future browser releases.

The HSTS preload list is maintained by the Chrome team, and other browsers pull the same list from the Chrome repository. This ensures consistency across different browsers. Some top-level domains (TLDs) have opted to add their entire TLD to the preload list. For example, the dev TLD automatically enforces HTTPS for all dev domains. This approach simplifies the process for individual domain owners and helps prevent the list from becoming unmanageable.

In addition to TLS attacks and HTTPS, there are other interesting topics related to web application security that we could explore further. Public key pinning, certificate transparency, and DNS certificate authority authorization are all worth investigating. These topics provide additional layers of security and can enhance the overall protection of web applications.

TLS is a complex and fascinating subject, and there is much more to learn. However, this material provides a solid foundation for understanding TLS attacks and the importance of web application security.

**EITC/IS/WASF WEB APPLICATIONS SECURITY FUNDAMENTALS DIDACTIC MATERIALS****LESSON: HTTPS IN THE REAL WORLD****TOPIC: HTTPS IN THE REAL WORLD**

HTTPS in the real world is an important topic in web applications security. While HTTPS provides confidentiality, integrity, and authentication for web traffic, not all web traffic is encrypted. This is why browsers still support HTTP. In order to address this issue, systems have been developed to enhance the security of HTTP and HTTPS.

One such system is Strict Transport Security (HSTS). HSTS aims to solve the problem of websites still being accessible over HTTP, even if they support HTTPS. For example, when a user types a website URL without specifying the protocol, the browser may default to sending the request over HTTP. Similarly, if a web page loads sub-resources using the HTTP scheme, the traffic may end up being unencrypted. This opens up opportunities for attackers to intercept and tamper with the HTTP requests and responses.

To illustrate the real-life implications of this issue, a researcher named Moxie Marlinspike demonstrated how an attacker can intercept HTTP to HTTPS redirects and insert malicious traffic. In one instance, GitHub experienced an attack where analytics scripts loaded over HTTP were replaced with malicious scripts, resulting in a denial-of-service attack.

Initially, browsers attempted to address this problem by displaying a lock icon to indicate that the traffic is encrypted. However, research has shown that relying solely on positive security indicators is not effective. Users may not notice the absence of the lock icon and continue to use insecure connections.

Therefore, additional measures were needed to ensure that traffic intended for HTTPS would not be sent over HTTP. HSTS was introduced as a solution. It allows websites to instruct browsers to always use HTTPS for future connections. Once a browser receives this instruction, it will automatically upgrade any HTTP requests to HTTPS, preventing the possibility of interception and tampering.

HTTPS adoption has made significant progress, but not all web traffic is encrypted. Systems like HSTS have been developed to enhance the security of HTTP and HTTPS. By enforcing the use of HTTPS, these systems help protect users' data and prevent attackers from intercepting and tampering with web traffic.

Strict Transport Security (HSTS) is a mechanism implemented by browsers to enhance web application security. When a website opts into HSTS, it informs the browser that it should only be accessed over a secure connection. This results in two key behaviors.

Firstly, if the browser detects an HTTP request to a server that has opted into HSTS, it will automatically rewrite the URL to HTTPS. This ensures that the traffic never travels over the network unencrypted, eliminating the need for server redirects.

Secondly, when a website has enabled HSTS, browsers will not allow users to bypass certificate errors. Normally, users have the option to ignore certificate errors, but this option is disabled for HSTS-enabled websites.

To demonstrate this behavior, Chrome's developer tools display an internal redirect when making a request to a website that has opted into HSTS. This means that the request is immediately internally rewritten to HTTPS, even before it reaches the network. Additionally, if an invalid certificate is encountered on an HSTS-enabled website, users do not have the option to bypass the error.

HSTS is implemented using an HTTP response header. The header contains several directives that control its behavior. One important directive is "max-age," which determines how long the information is cached in the browser's HSTS cache. Initially, a site may set a low max-age value for testing purposes. However, mature sites typically set it to a year or more to ensure long-term protection.

Another directive is "includes subdomains," which covers all subdomains of a website if enabled. While this is generally good practice, it can create challenges for large organizations that have one or two subdomains without HTTPS support. In such cases, enabling HSTS for the entire domain becomes problematic.

On the other hand, there is no standardized way for a subdomain to set HSTS for the entire domain. This is particularly relevant for landing pages like "www.example.com," which users often visit. Ideally, such subdomains should be able to set HSTS for the entire domain, but currently, there is no solution for this issue.

HSTS is a powerful mechanism that enhances web application security by ensuring that websites are accessed over secure connections. By opting into HSTS, websites can prevent traffic from ever being transmitted unencrypted and eliminate the option to bypass certificate errors. However, there are still some challenges and limitations associated with HSTS implementation, particularly regarding subdomains and large organizations.

HTTPS in the real world is an important aspect of web application security. When it comes to HTTPS Strict Transport Security (HSTS), there are a few key points to understand.

Firstly, HSTS does not have a built-in mechanism for exceptions. Once HSTS is set and the header is served to users, there is no way to undo it. This means that users will remember and enforce the HSTS policy for the specified time period. Organizations typically roll out HSTS gradually, increasing the max-age value over time to avoid sudden expiration.

Secondly, HSTS does not protect the first visit to a website. This is a trade-off made in practice, as the majority of users' exposure to potential security risks occurs after their initial visit. Upon installation of a browser like Chrome, users will receive HSTS policies for the websites they regularly visit, providing protection from subsequent visits.

Another important consideration is the potential for HSTS tracking, also known as super cookies or fingerprinting vectors. HSTS allows websites to set persistent state that can be queried from third-party contexts. This means that if a website loads an image from a third-party domain, that domain can set and read HSTS state. This capability can be abused for privacy invasion and tracking purposes. Unlike cookies, which can be cleared or controlled, HSTS tracking vectors are not always as easily restricted.

To illustrate how HSTS tracking can be implemented, consider the following high-level idea. When a user visits a website, such as shopping-site.com, a script from ad-network.com assigns a random identifier to the user. This identifier is represented in binary, and the ad network script issues sub-resource requests for each bit that is set in the identifier. These sub-resource loads set HSTS for specific subdomains. When the user visits another site that loads a script from ad-network.com, the script reads the previously set HSTS state by issuing sub-resource requests for each bit in the identifier. By observing which of these requests redirect to HTTP, the script can correlate the user's behavior across different sites.

It is worth noting that the designers of HSTS were aware of this tracking vector and acknowledged it in the specification. However, no concrete action has been taken to address this issue.

HSTS plays a crucial role in web application security by enforcing secure connections. Organizations should carefully consider the implementation of HSTS and gradually roll it out to ensure a smooth transition. Additionally, the potential for HSTS tracking should be recognized and addressed to protect user privacy.

In the world of web applications security, one important aspect to consider is the use of HTTPS to ensure secure communication between clients and servers. HTTPS, or Hypertext Transfer Protocol Secure, is an extension of HTTP that adds encryption and authentication mechanisms to protect data transmitted over the internet. In this didactic material, we will explore the real-world implementation of HTTPS and how it is being used to mitigate tracking and improve security.

Recently, Apple discovered evidence of websites tracking users through a technique called HSTS tracking. As a result, they deployed two mitigations to address this issue. The first mitigation focuses on setting cookies, where only subresources from the top-level domain or the registerable domain of that top-level domain are allowed to set HSTS headers. For example, if you are on "example.com" and a subresource from "bar.example.com" is loaded, the HSTS header will only be honored if it is from "fubar.example.com" or "example.com". This approach helps prevent unauthorized tracking by limiting the scope of HSTS headers.

The second mitigation implemented by Apple involves preventing the reading of HSTS cookies. This is achieved by leveraging Safari's rules for blocking third-party cookies. If a subresource does not allow cookies, HSTS is not applied to that subresource. By combining these two mitigations, both the setting and reading of HSTS cookies

are restricted, enhancing user privacy and security.

In the case of Google Chrome, a different approach is being considered to mitigate HSTS tracking. It aligns with an ongoing effort to eliminate mixed content, which refers to HTTP subresources on HTTPS pages. By disallowing the loading of HTTP subresources on HTTPS pages, the reading of HSTS cookies on HTTPS pages is mitigated. However, a trade-off is made by not applying HSTS to subresources on HTTP pages. This means that if you are on an HTTP page and it loads an HTTP subresource, HSTS will not be applied. While this approach may result in a slight decrease in security, it is deemed acceptable given the already compromised nature of HTTP pages.

It is worth noting that the detection of HSTS tracking can be achieved through various methods. One approach involves serving different resources, such as images, depending on whether the request is made over HTTP or HTTPS. By analyzing the response, it is possible to determine if HSTS is being honored or not.

The deployment of HTTPS in the real world has brought about several mitigations to address HSTS tracking. Both Apple and Google have implemented strategies to limit the setting and reading of HSTS cookies, thus enhancing user privacy and security. While trade-offs between security and privacy must be carefully considered, the goal is to strike a balance that minimizes the risk of tracking while still providing a secure browsing experience.

### HTTPS in the Real World

In the real world, the implementation of third-party cookie blocking has raised concerns regarding web application security. For example, when browsing Facebook, the website loads scripts from various analytics providers. However, if Safari determines that it is inappropriate to send third-party cookies to these providers, the website becomes vulnerable to malicious script injection. Previously, these scripts would have been loaded over an encrypted connection, ensuring security. Additionally, different web browsers have varying approaches to third-party cookie blocking. Chrome, for instance, does not adopt Safari's rules for blocking third-party cookies.

Another issue arises when a web page is accessed via HTTP instead of HTTPS. In this case, if a resource does not support HTTPS, it will not be loaded separately, potentially causing breakage. The process of rolling out HTTPS and phasing out mixed content involves running experiments over two years to measure the extent of breakage. This phase timeline, spanning several months, aims to minimize disruption. Although this transition may cause some HTTP links that load sub-resources to stop working, it is an essential step in shaping the web platform in the desired direction.

Despite concerns and resistance, upgrading to HTTPS has numerous advantages. Contrary to outdated beliefs, HTTPS can outperform HTTP due to the introduction of the HTTP/2 protocol, which significantly improves performance. Additionally, the cost of obtaining certificates has decreased significantly over time. Previously, certificates could cost thousands of dollars, but now they are often free or available at low prices. Despite the availability of data supporting the benefits of HTTPS, some individuals hold onto outdated notions and resist the transition.

However, the numbers indicate a positive trend. The adoption of HTTPS has reached high percentages, particularly for critical aspects of web security. While there will always be a small group of deniers, progress is being made. Eventually, as with FTP, HTTP will be marginalized and potentially removed from web browsers. This process will be guided by measurements and usage statistics, ensuring a smooth transition.

It is important to consider the regional differences in HTTPS adoption. While tech startups in the South Bay may not encounter significant challenges in upgrading to HTTPS, smaller companies in regions like South America may face more difficulties. Understanding the diverse constituency and addressing their specific needs is crucial in promoting widespread adoption of HTTPS.

HTTPS in the real world is an important aspect of web application security. While there were regional differences in the adoption of HTTPS, it has now become a standard practice globally. Getting a certificate and implementing HTTPS has become easier with the availability of documentation and tooling. However, migrating large existing websites to HTTPS can still be challenging for both technically capable organizations and non-experts.

In recent years, there has been a trend towards centralization in the industry, with platforms like WordPress and AWS offering ready-to-go web server solutions. This centralization has led to a significant increase in the number of websites using HTTPS, as these platforms have made it easier for non-experts to secure their sites.

One challenge in implementing HTTPS is the issue of first visit problems. When a user visits a website for the first time, they don't know if it has opted-in to HTTPS or not. To address this, browsers ship with a massive list of websites that have opted-in to HTTPS from the beginning. This list is known as the HSTS preload list. However, adding websites to this list is a complex process, as it requires meeting certain requirements and going through an automated checking process.

Maintaining the HSTS preload list is a delicate task, as binary size is a make-or-break concern for browsers like Chrome. Increasing the binary size can lead to performance issues and affect the user experience. Additionally, in markets where users pay for the data they download, increasing binary size can have financial implications for users.

To get a website onto the HSTS preload list, website owners can visit the preload website and submit their site for inclusion. However, once a website is on the list, it is not easy to get removed quickly. Website owners have to wait for a release cycle, and different browsers may have their own policies and update cycles for the list.

Currently, there is no automated pruning of entries on the HSTS preload list. Websites that are no longer active or don't fulfill the requirements remain on the list. However, there is a possibility of automating the pruning process in the future to remove stale entries.

The first visit to a website without HTTPS does pose a security risk, but it is not necessarily more sensitive than subsequent visits. The vulnerability lies in the lack of protection during the first visit, and similar security properties apply to subsequent visits without HTTPS.

HTTPS adoption has become widespread globally, with the help of documentation, tooling, and centralized platforms. The HSTS preload list plays a crucial role in ensuring websites are opted-in to HTTPS from the beginning, but adding and removing websites from the list can be a complex process. The first visit to a website without HTTPS poses a security risk, but it is not inherently more sensitive than subsequent visits.

HTTPS in the real world is an important aspect of web applications security. In this material, we will discuss the concept of HTTPS and its significance in ensuring secure communication between a user's browser and a website.

HTTPS, or Hypertext Transfer Protocol Secure, is an extension of the HTTP protocol that adds an extra layer of security through the use of encryption. This encryption ensures that any data transmitted between the user and the website is protected from unauthorized access or tampering.

One of the key components of HTTPS is the use of SSL/TLS certificates. These certificates are issued by trusted certificate authorities (CAs) and serve as a digital proof of identity for the website. When a user visits a website secured with HTTPS, their browser checks the validity of the SSL/TLS certificate to ensure that the website is genuine and not an imposter.

However, the process of issuing and managing SSL/TLS certificates is not without its challenges. In the past, there have been instances where certificate authorities were compromised, leading to the issuance of fraudulent certificates. This means that attackers could potentially impersonate legitimate websites and intercept sensitive information.

To address this issue, efforts have been made to establish a preload list. This list includes high-value websites, such as Google, that have been verified and approved to be included in the list. Websites on the preload list are automatically included in the browser's HSTS (HTTP Strict Transport Security) list, which ensures that all future connections to these websites are made securely over HTTPS.

However, the preload list is not a perfect solution. It was initially intended for high-value sites but has grown to include other websites as well. This lack of clear criteria for inclusion has led to concerns about the effectiveness and sustainability of the preload list.

Another challenge in the HTTPS ecosystem is the issue of certificate authority trust. While any certificate authority can issue certificates for any website, this also means that attackers can exploit this system. In the past, there have been instances where certificate authorities were compromised, allowing attackers to issue fraudulent certificates for well-known websites like Google.

To mitigate this risk, various solutions have been proposed, but none have proven to be foolproof. One suggestion is to restrict certificate authorities from issuing certificates for websites based in different countries. However, this approach has its limitations and does not address the underlying issue of certificate authority trust.

HTTPS plays a crucial role in ensuring the security of web applications. It provides encryption and authentication mechanisms that protect user data from unauthorized access or tampering. However, challenges remain in the form of fraudulent certificates and issues with certificate authority trust. Efforts are ongoing to address these challenges and improve the security of HTTPS in the real world.

In the context of web applications security, HTTPS plays a crucial role in ensuring secure communication between clients and servers. However, there are certain challenges and considerations when implementing HTTPS in the real world.

One common misconception is the idea of blocking certain countries or regions from accessing web applications as a security measure. While this may seem like a simple solution, it is not a viable option from a business perspective. Blocking entire regions can lead to missed opportunities and potential loss of customers. Moreover, it is important to note that spam or security threats can originate from any location, and blocking specific regions does not guarantee protection.

Another approach to enhance HTTPS security is through certificate pinning. Certificate pinning allows site operators to specify which certificate authorities (CAs) are trusted to issue certificates for their web applications. By doing so, attackers are limited to compromising only the chosen CAs, reducing the attack surface. It is important to emphasize that pinning should be done at the public key level, rather than the certificate level, to ensure cryptographic integrity. This means that the public key used for pinning is hashed and included in the HTTP response header.

To implement certificate pinning, site operators should include the "Public-Key-Pins" header in their HTTP responses. This header specifies the pins (hashed public keys) that the browser should expect from the server. It is recommended to include multiple pins for redundancy and security purposes. The pins are generated by hashing the subject public key info, which contains information about the type and bits of the public key.

In a certificate chain, which includes the end entity certificate, intermediate CAs, and root CAs, it is possible to pin any of the public keys. For example, in the case of Google, their end entity certificate for google.com is signed by the GTS CA, which is in turn signed by the global sign root CA. Site operators can choose to pin any of the keys in this chain to enhance security.

When a client makes a connection to a web server using HTTPS, the browser performs regular CA-authorized certificate validation. In addition to this, the browser also validates the pins it has previously encountered. This extra step ensures that the server's public key matches the pinned keys, providing an additional layer of security.

HTTPS in the real world requires careful consideration of security measures. Blocking regions or countries is not a recommended approach, as it can have negative business implications and does not guarantee complete security. On the other hand, certificate pinning allows site operators to specify trusted CAs and reduces the attack surface. By pinning the public keys, rather than the certificates themselves, cryptographic integrity is maintained.

In the real world, HTTPS (Hypertext Transfer Protocol Secure) is an essential security measure for web applications. It ensures that the communication between a client (usually a web browser) and a server is secure and encrypted. However, there are practical challenges and limitations when it comes to implementing HTTPS effectively.



One challenge is the complexity of managing the cryptographic keys used in HTTPS. In order to establish a secure connection, the client needs to verify the authenticity of the server's public key. This verification is done by checking if any of the trusted keys intersect with the key provided by the server. If they do, it indicates that the connection is secure. If not, it could mean that the client is communicating with an impersonator or that the server has set the wrong keys.

However, successfully implementing this verification process is not easy. It requires a deep understanding of public key infrastructure and the concept of subject public key info. Most website operators may not possess this knowledge, resulting in potential security vulnerabilities. To address this issue, efforts have been made to simplify the process, such as providing shell scripts that generate a set of trusted keys. Unfortunately, even this approach has proven to be too complex for many website operators.

Another challenge lies in the certificate chain presented by the server during the TLS (Transport Layer Security) handshake. The server sends its certificate chain to the client for validation. However, the client may construct a different chain than what was served by the server. This can happen when the client discovers additional root certificates from its own database. If the client constructs a different chain and the server operator has pinned their trust to a specific root certificate, the pin validation will fail. This creates a situation where the server operator cannot predict or control the chain the client will build, leading to potential failures in pin validation.

Furthermore, the diversity of web browsers and their different behaviors adds to the complexity. Each browser has its own set of trusted root certificates, TLS client libraries, and certificate libraries. These factors can vary across different versions and platforms. Additionally, browser behaviors can change over time due to administrator policies or other factors. Consequently, server operators have no reliable way to anticipate the chain that a client will build during a connection.

This unreliability creates a significant challenge for server operators and makes it difficult to rely on pin validation as a security measure. It would require constant monitoring and adaptation to the changing behaviors of various clients, which is impractical and burdensome.

From a user perspective, there are also challenges in using browsers with pin validation. If pin validation fails, users are unable to access the desired website. There is no option to bypass the error screen, and the message is often non-actionable. Users are left with no clear instructions on what to do next, leading to a frustrating experience. Moreover, if there are issues with the website's pins or certificates, there is no guarantee that the problem will be resolved in the future. This lack of recoverability and actionable solutions further diminishes the user experience.

In addition to these challenges, there is a theoretical concern of hostile pinning. If an attacker gains control of a server, they could set their own pin set for the domain, potentially redirecting clients to their own servers or compromising their security.

While HTTPS and pin validation are important for web application security, there are significant challenges and limitations when implementing them in the real world. The complexity of managing keys, the unpredictability of certificate chains, the diversity of browsers, and the user experience issues all contribute to the difficulty of relying on pin validation as a reliable security measure.

In the realm of web application security, one crucial aspect is the use of HTTPS to ensure secure communication between users and websites. However, there are certain challenges associated with HTTPS implementation, particularly in the real world. This didactic material aims to shed light on the practical aspects of HTTPS and its implications.

One issue with HTTPS is the use of public key pinning, which involves associating a specific cryptographic key with a particular domain. This mechanism helps prevent attackers from impersonating a website by using fraudulent certificates. However, there are limitations to pinning, such as the fact that pins have a limited lifespan. After the expiration date, users may lose the ability to access a website even after it has recovered from an attack.

Unlike other protocols like SSH, which allow users to easily recover from key mismatches, HTTPS lacks a straightforward recovery mechanism. This lack of usability poses a challenge, especially considering the vast number of users and websites on the internet. To address this issue, an alternative solution called "unship

pinning" was introduced. In 2018, after seven years of implementation, the decision was made to remove pinning due to its complexity and limited benefits.

Despite the removal of pinning, concerns were raised by users who had come to rely on its safety features. To mitigate these concerns, alternative mechanisms were proposed. One such mechanism is Certification Authority Authorization (CAA), which involves adding a DNS record to a domain. This record specifies the authorized certificate authority (CA) for issuing certificates for that domain. If a CA receives a request for a domain that does not match the specified issuer, the issuance process is halted. While CAA relies on the good behavior of CAs, it has proven to be effective in practice.

Another approach is static pinning, which involves preloading a list of pins in the browser. This list is carefully curated and serves as a reference for validating certificates. Despite the existence of static pinning in modern browsers, it is important to note that it is not as prevalent as dynamic pinning.

It is worth mentioning that DNS, which is used in the CAA mechanism, is not currently secured. However, efforts are being made to address this issue. In the future, DNS will provide encryption, integrity protection, and authentication, making the CAA mechanism even more robust.

While public key pinning in HTTPS has its limitations, alternative mechanisms such as CAA and static pinning have been introduced to address usability and security concerns. These mechanisms provide additional layers of protection against attacks and help ensure a safer web browsing experience.

HTTPS in the real world is an important topic in web application security. One aspect of HTTPS is pinning, which involves associating a specific cryptographic key with a particular website. This ensures that the browser only accepts a certificate if it matches the pinned key. However, pinning can be challenging because if the key needs to be changed, the site becomes inaccessible until the new key is released. Despite this drawback, pinning provides value by ensuring secure connections to trusted websites.

Another approach to web application security is certificate transparency. This concept involves the use of public logs where all certificates are recorded. These logs can be monitored by domain owners to detect unauthorized certificates for their domains. Researchers can also use the logs to identify certificates that do not adhere to proper practices. Certificate transparency makes it difficult for attackers to use malicious certificates without being noticed.

To implement certificate transparency, one idea is to have the browser check with the public logs before accepting a certificate as valid. However, this approach raises concerns about privacy and performance. An alternative solution is to have the logs provide a receipt or statement confirming that a certificate has been logged. This receipt can be included with the certificate during the validation process. By validating the receipt using the log's public keys, the browser can ensure that the certificate has been logged without directly querying the log during the connection setup.

It is important to note that the receipt provided by the log is a promise to log the certificate within 24 hours. This ensures that the log remains up to date and that any malicious or accidental certificates are discovered and remediated.

Both pinning and certificate transparency are valuable tools in ensuring the security of web applications. While pinning provides a direct association between a website and its cryptographic key, certificate transparency makes all certificates publicly accessible, allowing for detection of malicious or accidental certificates. By implementing these security measures, web applications can enhance their protection against unauthorized access and improve overall cybersecurity.

Certificate Transparency (CT) is a system designed to enhance the security of web applications by providing transparency and accountability in the issuance and use of digital certificates. It addresses the challenge of trusting Certificate Authorities (CAs) and aims to prevent malicious certificate issuance.

The implementation of CT involves the use of logs, which are responsible for recording and making certificates publicly visible. However, the logs themselves need to be trusted, as they can potentially misbehave or provide inconsistent information to different observers.

To ensure the integrity of the logs, a cryptographic construction is employed. This construction generates a short summary of the data observed by the logs, which possesses two important properties. First, if two observers have the same summary, they can efficiently compare it and verify that they have seen the same data. Second, given a certificate and a summary, it is possible to efficiently determine if the certificate was included in the data that produced the summary.

When a server encounters a certificate, it can request a summary and a proof from the log to verify that the certificate was indeed included. Different observers can then compare the summaries from different logs to ensure they are seeing the same view.

It is important to note that CT does not provide all the security properties that key pinning aims to achieve. CT primarily focuses on detection rather than prevention. This means that there may be a period of time during which a malicious certificate can be used before it shows up in CT logs. Additionally, there is a possibility of malicious logs existing, which may take an unspecified amount of time to be discovered.

Despite these limitations, CT has several benefits. It discourages malicious certificate issuance, as the transparency makes it more likely for such actions to be detected. It also helps organizations and researchers identify and address bad practices within the certificate ecosystem.

Certificate Transparency is a system that enhances the security of web applications by providing transparency and accountability in the issuance and use of digital certificates. It utilizes logs to record and make certificates publicly visible, while employing cryptographic constructions to ensure the integrity of the logs. While CT focuses on detection rather than prevention, it serves as a valuable tool in discouraging malicious certificate issuance and improving the overall security of web applications.

In the field of cybersecurity, one important aspect is web application security, specifically the use of HTTPS in real-world scenarios. HTTPS, or Hypertext Transfer Protocol Secure, is a protocol that ensures secure communication over a network. It provides encryption and authentication, making it difficult for attackers to intercept or modify data being transmitted between a web server and a client.

To understand the significance of HTTPS in the real world, it is crucial to highlight the role of Certificate Authorities (CAs). CAs are organizations trusted to issue digital certificates that verify the authenticity of websites. These certificates are essential for establishing a secure connection between a client and a server. It is worth noting that there are numerous CAs, each having their requirements and processes.

Another critical concept related to HTTPS is Certificate Transparency (CT). CT is a system that aims to make the issuance and management of digital certificates more transparent. It has been in development for several years and has reached significant milestones. However, it is still a work in progress, with many challenges and open problems to be addressed.

One of the challenges in CT is the tracking of HTTP Strict Transport Security (HSTS) policies. HSTS is a security feature that forces web browsers to communicate with a website only over HTTPS. Balancing security and privacy concerns remains an ongoing challenge in this area.

Additionally, there are concerns about static analysts, which are tools used to analyze and identify vulnerabilities in software. These tools can be exploited by attackers, posing a significant security risk. Currently, there is no proposed system that can replace the security guarantees provided by key pinning without introducing new vulnerabilities.

Moreover, there are open questions regarding the honesty of CT logs. CT logs are public repositories that store information about issued certificates. Ensuring the integrity and trustworthiness of these logs is a critical concern.

It is important to acknowledge the complexity of these challenges. The conceptual difficulty is compounded when considering the scale of the problem, as it involves a vast ecosystem with billions of users, each with their motivations and goals. Despite the challenges, progress has been made, and the hard problems encountered are a testament to the success achieved in solving the easier parts.

Understanding HTTPS in the real world involves grasping the role of CAs, the development of Certificate

Transparency, and the challenges surrounding HSTS tracking, static analysts, and CT logs. The field of cybersecurity offers exciting and rewarding opportunities to address these challenges and make a positive impact.

**EITC/IS/WASF WEB APPLICATIONS SECURITY FUNDAMENTALS DIDACTIC MATERIALS****LESSON: AUTHENTICATION****TOPIC: INTRODUCTION TO AUTHENTICATION**

Authentication is a fundamental concept in cybersecurity, specifically in the realm of web applications security. It involves verifying the identity of a user, ensuring that they are who they claim to be. This is crucial for building secure systems, even in situations where an attacker has obtained the user's password.

There are several factors that can be used for authentication. The first is something the user knows, such as a password stored in their memory. The second factor is something the user possesses, like a phone, an ID badge, or a cryptographic key. Lastly, authentication can also be based on something the user is, such as their fingerprints, retina, or other biometric data.

Biometric data, in particular, offers unique and interesting possibilities for authentication. For example, gait analysis, which involves analyzing a person's walking pattern, can be used to identify individuals. Another intriguing example is the use of the shape of a person's rear-end as a biometric data point for car theft prevention.

Authentication plays a vital role in ensuring the security of web applications. By implementing robust authentication mechanisms, we can build systems that are resilient even in the face of compromised passwords. This aligns with the concept of defense-in-depth, which emphasizes multiple layers of security to mitigate the impact of a single system failure.

Authentication is the process of verifying a user's identity. It can be based on something the user knows, possesses, or is. By implementing strong authentication measures, we can enhance the security of web applications and protect against potential attacks.

Authentication is a fundamental aspect of web application security. It involves verifying the identity of users to ensure that they are who they claim to be. This is crucial in preventing unauthorized access and protecting sensitive information.

One common method of authentication is the use of biometric data, such as fingerprints or facial recognition. However, there are limitations to this approach. Biometric data is not changeable, meaning that if it is stolen, it cannot be reset. This is why many systems that use biometric authentication keep the data on the device itself, rather than sending it to a remote service. For example, on iPhones, touch ID or face ID data stays on the device and is used to unlock a key store that contains a cryptographic key for authentication.

Another important concept in authentication is the use of multiple factors. By using multiple factors from different categories, such as something you know (e.g., a password), something you have (e.g., an ATM card), and something you are (e.g., biometric data), we can increase the certainty that a user is who they claim to be.

It's important to note the difference between authentication and authorization. Authentication is the process of verifying a user's identity, while authorization is the process of determining what actions a user is allowed to perform. Authentication is typically done through login forms, cookies, or other methods, while authorization is handled through access control lists (ACLs) or capabilities.

When implementing authentication in web applications, there are several common mistakes to avoid. One such mistake is storing usernames in a case-insensitive manner. This can lead to security vulnerabilities, as an attacker could potentially bypass authentication by exploiting case-insensitive comparisons.

Another mistake is failing to properly separate authentication and authorization. These are two distinct processes, and it's important to keep them separate to prevent confusion and ensure proper access control.

Additionally, it's important to avoid issuing long-lived tokens or session IDs for authorization. These tokens should be regularly updated or expired to reflect any changes in user privileges or authorizations.

Authentication is a critical aspect of web application security. By implementing secure and robust authentication mechanisms, we can protect user identities and prevent unauthorized access.

Authentication is a fundamental aspect of web application security. It ensures that users are who they claim to be and grants them access to the appropriate resources. In this context, it is important to consider two key factors: usernames and passwords.

When it comes to usernames, it is crucial to store them consistently in the database. This means that if a user types their username in different cases, such as lowercase or uppercase, it should be stored exactly as entered. Failure to do so can result in multiple user accounts with the same name but different cases, leading to confusion and potential security issues. Additionally, enforcing uniqueness of usernames is essential to prevent multiple users from registering with the same name.

To accommodate user preferences for capitalization, it is possible to create a second column in the user table to store the preferred casing of usernames. This allows for the comparison of lowercase input during authentication while rendering the preferred casing in the user interface.

Moving on to passwords, it is a well-known fact that users often choose weak passwords. Commonly used passwords, such as "password" or "123456," pose a significant security risk. Despite efforts to educate users about the importance of strong passwords, the situation has not improved significantly over the years.

To address this issue, websites and applications often implement password requirements. However, the traditional approach of forcing users to change their passwords regularly and select complex combinations of uppercase and lowercase letters, numbers, and symbols is outdated and ineffective. Users tend to add predictable patterns, such as appending a number to the end of their password, which does not enhance security.

It is important to note that requiring complex passwords with a combination of various character types is not recommended. Instead, a more effective approach is to encourage users to choose longer passwords that are easy for them to remember but difficult for others to guess. Length and uniqueness are key factors in creating strong passwords.

Outdated practices also include saving a history of previously used passwords and preventing users from reusing them. This approach can lead to a false sense of security, as attackers can exploit patterns in password changes. Therefore, it is advisable to avoid implementing such features.

The traditional password requirements based on complexity and regular changes are not effective in improving security. Instead, focusing on longer and unique passwords is recommended. By educating users about the importance of strong passwords and adopting modern practices, we can enhance the security of web applications.

#### Web Applications Security Fundamentals - Authentication - Introduction to authentication

In the realm of web application security, there are certain practices that should never be implemented, as they can compromise the security of user authentication. Some real websites have been found to engage in these practices, which are considered to be awful and highly detrimental to the security of user accounts.

One such practice is the imposition of a maximum length on passwords. This is typically done due to the presence of legacy systems that cannot handle passwords longer than a specific length. As a result, passwords are either capped at a certain number of characters or truncated, giving users a false sense of security. This means that any additional characters beyond the specified limit do not contribute to the overall security of the password.

Another horrendous practice involves allowing users to log in over the phone using a touchpad. The issue with this approach is that touchpads only have numbers, which are mapped to multiple letters. In the backend, the letters chosen by the user for their password are translated into the corresponding numbers on the touchpad. This significantly reduces the entropy of the password and increases the likelihood of unauthorized access. It has been confirmed that by swapping a couple of letters in their password based on the same number on their phone's touchpad, attackers were able to gain access to user accounts.

These are just a few examples of the many horror stories surrounding web application authentication. There



have been instances where websites implement a minimum password age policy, preventing users from changing their passwords too frequently. The intention behind this policy is to deter users from bypassing the requirement to change their passwords regularly by simply changing it and then immediately reverting back to the original password. However, this practice is highly insecure since immediate password changes are necessary in case of password loss or exposure.

Disabling the cut and paste functionality on login forms is also a highly flawed practice. The rationale behind this decision is to prevent users from saving their passwords in a text file. However, this breaks password managers, which rely on the ability to copy and paste passwords securely. It is crucial to allow users to utilize password managers for enhanced security.

Password hints are another aspect that is often implemented incorrectly. Many websites provide password hint questions with limited entropy. For example, some airlines offer a set of hint questions with answers selected from a drop-down menu containing only eight options. This significantly weakens the security of the account, as an attacker has a one in eight chance of guessing the correct answer. Even if multiple questions are combined, the limited set of questions still poses a significant risk.

Lastly, a misguided idea in the realm of authentication is the use of on-screen keyboards to prevent keyloggers from capturing passwords. While this may be effective against keyloggers, it fails to address other methods of capturing sensitive information, such as taking screenshots or intercepting mouse clicks on the on-screen keyboard. This approach also introduces usability issues for users, making the login process more cumbersome.

It is essential to avoid these terrible practices in web application authentication. Despite claims of being implemented for security reasons, they ultimately compromise the security of user accounts. Implementing strong and secure authentication practices is crucial to protect user data and maintain the integrity of web applications.

Authentication is a fundamental aspect of web application security. It involves verifying the identity of users before granting them access to sensitive information or functionalities. In this context, the concept of an ID shield or secure ID has been proposed to help users detect phishing pages. The idea behind this approach is that when users create an account, they select an image that is supposed to be unique to them. Whenever they visit the bank's website and enter their username, the bank displays the selected image as a way to confirm the authenticity of the page.

However, there are several problems with this approach. Firstly, a phishing site can easily replicate the image selection page and display the chosen image, making it ineffective in detecting phishing attempts. Additionally, users often ignore positive user interface indicators, which further diminishes the effectiveness of this method. Furthermore, since only one image is used, it becomes even easier for attackers to exploit this system by simply entering the victim's username and capturing the image dynamically.

Studies have shown that the effectiveness of security images is generally poor. In one study, 72% of participants entered their password even when the security image and caption were removed. This highlights the limitations of relying solely on visual cues for authentication.

In recent years, there has been a shift in password requirements. The complexity of a password, such as including numeric, alphabetic, and special symbols, does not necessarily make it stronger. Instead, choosing multiple words from a large dictionary can result in stronger passwords, even if they are composed of lowercase letters and contain no punctuation. This approach increases the entropy of the password and makes it harder to guess or crack.

To address the issue of users choosing weak passwords, organizations can implement password checking against known breached data. By comparing user-chosen passwords with leaked password databases, organizations can prevent users from selecting passwords that have been compromised in previous breaches. This approach is recommended by the National Institute of Standards and Technology (NIST) as a more effective way to enhance password security.

Furthermore, NIST no longer recommends changing passwords regularly. Instead, the focus is on ensuring that passwords are not easily guessable or compromised. This shift reflects the understanding that frequent password changes can lead to weaker passwords as users tend to choose simpler and easier-to-remember

passwords.

It is important to note that while these recommendations improve security, they also have usability implications. Striking the right balance between security and usability is crucial to avoid overly restrictive policies that may frustrate users. However, by implementing these recommendations, organizations can significantly enhance the security of their authentication systems.

Authentication is a fundamental aspect of web application security. It involves verifying the identity of users to ensure that only authorized individuals have access to sensitive information or services. In this context, the choice of passwords plays a crucial role in maintaining the security of user accounts.

While some users may not consider the security of their passwords for certain services, such as Netflix, it is important to understand the potential risks associated with using weak or breached passwords. Attackers often exploit breached databases or attempt to guess passwords using common patterns and known passwords. Therefore, it is essential to encourage users to choose strong and unique passwords to mitigate these risks.

One popular method for creating strong passwords is through the use of substitutions, numerals, and punctuation. However, it is important to note that the additional entropy gained from these techniques is relatively small. Moreover, assuming users follow a specific pattern of selecting an English word, making substitutions, and adding a symbol and a number at the end, the number of possible passwords is limited. Therefore, it is advisable to consider alternative approaches, such as using password managers, which can generate complex passwords for users.

Another common issue with user passwords is their length. Short passwords are particularly vulnerable to cracking attempts. Research has shown that passwords shorter than twelve characters can be easily cracked, especially when attackers employ offline methods that involve reversing hashed passwords. Therefore, it is recommended to use passwords of sufficient length to enhance security.

To illustrate the impact of password length on cracking time, a mapping of password lengths to cracking times is presented. The assumption is that the password includes upper and lower case letters, as well as numbers. The graph clearly demonstrates that shorter passwords are more susceptible to being cracked quickly, while longer passwords provide significantly greater security.

In order to better understand the implications of password strength, an interactive website allows users to input example passwords and observe the time it would take to crack them. This tool helps users visualize the importance of selecting strong passwords and reinforces the notion that even a seemingly strong password can be compromised under certain circumstances.

It is important to consider different attack scenarios when evaluating password security. The online attack scenario involves an attacker attempting to log in to a server by sending multiple username and password combinations. This scenario is inherently limited by the number of requests that can be sent without being blocked or detected. However, offline attack scenarios, which occur when an attacker has access to a database of hashed passwords, are more concerning. In these situations, attackers can leverage local resources to crack passwords much faster. This is particularly true when using a cracking array of computers.

Authentication is a critical aspect of web application security. By understanding the risks associated with weak or breached passwords, users can make informed decisions to protect their accounts. Choosing strong and unique passwords, utilizing password managers, and considering password length are all important steps in enhancing authentication security.

Authentication is a fundamental aspect of web application security. It ensures that users are who they claim to be before granting them access to sensitive information or functionalities. In this didactic material, we will introduce the concept of authentication and discuss best practices for implementing secure authentication systems.

One common method of authentication is the use of passwords. When users create an account, they choose a password that they will later provide to verify their identity. However, not all passwords are equally secure. Attackers can use various techniques to crack weak passwords and gain unauthorized access to user accounts.

To illustrate this, let's consider an example. Suppose a website allows users to create accounts with passwords. The website's developers have implemented a feature that checks the strength of the chosen password. They have also assured users that their passwords are not sent to the server. However, even with these measures in place, it is still not recommended to use real passwords for demonstration purposes.

In order to create a strong password, it is important to consider its length and complexity. Longer passwords are generally more secure, so it is advisable to allow users to choose passwords with a significant length, potentially up to 64 characters. However, a minimum length requirement should also be enforced, such as a minimum of eight characters.

Another consideration is the limitation on password length. Some popular hashing functions, like bcrypt, have a maximum length limit. For example, bcrypt allows a maximum of 72 ASCII characters. This can be problematic if the recommended maximum password length is higher, such as 64 characters. To address this issue, it is recommended to hash the password with another hash function before using bcrypt.

In addition to length and complexity, it is crucial to check passwords against known breach data. This means comparing the chosen password with a database of previously compromised passwords. If a match is found, the user should be prompted to choose a different password.

To protect against brute-force attacks, where attackers try multiple passwords in rapid succession, it is important to implement rate limiting. This means restricting the number of authentication attempts a user can make within a certain time frame. By doing so, the system can prevent attackers from repeatedly guessing passwords until they find the correct one.

For highly sensitive web applications, it is advisable to require a second factor of authentication, such as a one-time password sent to the user's mobile device. This adds an extra layer of security and makes it more difficult for attackers to gain unauthorized access.

When implementing authentication systems, there are some common mistakes to avoid. One of them is silently truncating long passwords. This can happen with certain hashing functions, like bcrypt, which have a maximum length limit. Developers should ensure that passwords are not truncated and that the chosen hashing function can handle the desired maximum length.

Another mistake to avoid is restricting the characters that users can choose for their passwords. It is important to allow a wide range of characters, including Unicode symbols and emojis. Restricting the character set can weaken the security of passwords and limit user creativity in choosing memorable passwords.

Finally, it is crucial to be mindful of logging practices. Developers should avoid accidentally including passwords in plain text log files. Logging HTTP requests can be useful for debugging and analysis purposes, but passwords should never be logged. This can prevent unauthorized access to user credentials in case of a security breach.

Authentication is a critical aspect of web application security. Implementing secure authentication systems involves considering password length, complexity, rate limiting, and the use of additional factors of authentication. Developers should also avoid common implementation mistakes, such as truncating passwords and logging sensitive information.

Authentication is a crucial aspect of web application security. It ensures that only authorized users have access to sensitive information and functionalities. In this section, we will introduce the concept of authentication and discuss some fundamental techniques to defend against network-based guessing attacks.

Authentication is the process of verifying the identity of a user or system. It is typically achieved by presenting credentials, such as a username and password, to prove one's identity. However, there are various types of network-based attacks that can compromise the authentication process and gain unauthorized access to user accounts.

One common attack is brute force, where an attacker systematically tries all possible passwords to target a specific account. Another attack is credential stuffing, where an attacker reuses username and password combinations obtained from a breach on one site to gain access to another site. This attack exploits the common practice of users reusing passwords across multiple platforms. Lastly, password spraying involves

trying a known weak password on multiple accounts, hoping that some users have chosen that password.

To defend against these attacks, several strategies can be employed. One effective approach is rate limiting, which restricts the number of authentication attempts an attacker can make within a certain time period. For example, using the Express rate limiter package in Node.js, you can limit the number of login attempts per username or IP address.

Another technique is to implement additional verification measures to ensure that the user is a real person. One popular method is CAPTCHA (Completely Automated Public Turing test to tell Computers and Humans Apart). CAPTCHA presents a challenge that is easy for humans to solve but difficult for automated bots. By successfully completing the CAPTCHA, users can prove their authenticity and continue with their login attempts.

Authentication plays a critical role in web application security. To defend against network-based guessing attacks, rate limiting and CAPTCHA can be effective measures to protect user accounts from unauthorized access. Implementing these techniques can significantly enhance the security of web applications and safeguard user data.

Authentication is a fundamental aspect of web application security. It ensures that users are who they claim to be before granting them access to sensitive information or functionalities. In this didactic material, we will introduce the concept of authentication and discuss its importance in securing web applications.

Authentication is the process of verifying the identity of a user or entity. It involves the use of credentials, such as usernames and passwords, to validate the user's identity. The goal of authentication is to prevent unauthorized access to web applications and protect sensitive data from malicious actors.

One common method of authentication is the use of CAPTCHAs (Completely Automated Public Turing test to tell Computers and Humans Apart). CAPTCHAs are designed to distinguish between humans and automated bots. They typically present users with a challenge, such as identifying distorted characters in an image, and require them to provide the correct response.

However, CAPTCHAs have their limitations. Research has shown that certain CAPTCHA implementations have high breakage rates, making them susceptible to attacks. For example, one implementation had a breakage rate of 92%, while another had a breakage rate of 33%. These rates indicate that even a low success rate for attackers can compromise the security of CAPTCHAs.

Furthermore, CAPTCHAs can be inconvenient for users, especially those with visual impairments. To address this, some websites offer alternative CAPTCHAs that use audio instead of visual challenges. However, if the audio-based CAPTCHA is less secure than the main one, attackers can exploit it by parsing the audio file.

Another vulnerability of CAPTCHAs is the use of screenshot attacks. Attackers can detect when they are presented with a CAPTCHA, capture a screenshot of it, and then present it to users on another site. These users unknowingly solve the CAPTCHA, and the attacker uses their answers in real-time to bypass the CAPTCHA on the target site.

To make matters worse, there are dark market services that offer CAPTCHA-solving APIs, allowing attackers to outsource the task of breaking CAPTCHAs. These services are relatively cheap, with prices as low as 50 cents for a thousand CAPTCHA solves.

In response to these vulnerabilities, interactive CAPTCHAs have been introduced. These CAPTCHAs require users to perform actions, such as clicking on specific images, to prove their humanity. The interactivity of these CAPTCHAs makes them harder to bypass, as they require capturing and transmitting user behavior across the browser session.

However, even interactive CAPTCHAs have their limitations. In 2014, researchers from Stanford University developed an algorithm called NMLAlgo that could break text-based CAPTCHAs with high accuracy and speed. They demonstrated that all existing text-based CAPTCHAs were insecure in practice, with breakage rates ranging from 5% to 51%.

As a result, new approaches to authentication, such as the use of trust scores, are being explored. One example

is reCAPTCHA, which analyzes user behavior, such as mouse movements, scrolling patterns, and IP addresses, to create a trust score. If a user is deemed suspicious, a CAPTCHA may be presented. However, most of the time, users can simply click a checkbox to verify their humanity.

It is important to note that even these advanced authentication methods are not foolproof. Attackers continuously evolve their techniques, and new vulnerabilities may be discovered in the future. Therefore, it is crucial for web application developers and security professionals to stay updated on the latest authentication best practices and adapt their strategies accordingly.

Authentication is a critical aspect of web application security. CAPTCHAs have been widely used to authenticate users, but they have their limitations. Researchers have demonstrated vulnerabilities in various CAPTCHA implementations, leading to the development of more advanced authentication methods. However, even these methods are not immune to attacks. It is essential for security practitioners to continuously improve authentication mechanisms to ensure the protection of sensitive data and prevent unauthorized access.

Authentication is a fundamental aspect of web application security. It ensures that users are who they claim to be and prevents unauthorized access to sensitive information or actions. In this didactic material, we will explore different techniques and considerations related to authentication.

One popular method to verify user authenticity is through the use of CAPTCHA, specifically reCAPTCHA. reCAPTCHA is a service that analyzes user interactions with a website to determine if they are a real person or a bot. It does this by logging login attempts and other actions performed on the site. Over time, reCAPTCHA builds a reputation for each user, indicating their trustworthiness. This reputation can be queried by other services to make decisions on allowing user actions. This approach is effective in combating automated login attempts and unwanted bot activities on websites.

However, there are certain limitations to consider. For example, reCAPTCHA relies on IP reputation to determine user authenticity. This can lead to issues for users who utilize the Tor browser for privacy purposes. Since Tor IP addresses are often associated with suspicious activities, these users may frequently encounter CAPTCHA challenges, even for legitimate actions. Finding a solution to address this challenge remains a topic of discussion.

Another technique to enhance authentication security is the implementation of a defense-in-depth approach. This involves requesting users to re-enter their password before performing sensitive actions, such as changing passwords, email addresses, or adding new shipping addresses. This technique provides an additional layer of protection against vulnerabilities like XSS, CSRF, or session fixation. Even if an attacker manages to inject malicious code into a website, they would still require the user's password to perform these sensitive actions. Notable examples of this technique can be observed on platforms like GitHub, where users are prompted to enter their password again before performing critical operations.

Furthermore, response discrepancy information exposure is another important consideration in authentication systems. This refers to the unintended leakage of information to attackers. For instance, providing different responses when a user attempts to log in, such as indicating whether an email address exists or if the password is incorrect, can inadvertently reveal information to attackers. This can help them determine the existence of user accounts on a service. Therefore, it is essential to carefully design authentication systems to avoid such response discrepancies and prevent information exposure.

Authentication plays a vital role in web application security. Techniques like reCAPTCHA, defense-in-depth, and mitigating response discrepancy information exposure contribute to ensuring the authenticity of users and protecting sensitive information. Understanding and implementing these fundamentals are crucial in building secure web applications.

Authentication is a fundamental aspect of web application security. It involves verifying the identity of users before granting them access to sensitive information or functionalities. However, improper implementation of authentication mechanisms can lead to security vulnerabilities, such as information disclosure and timing attacks.

One common mistake in authentication is providing specific error messages that reveal information about the state of the account or the existence of a user. For example, displaying messages like "invalid password,"

"invalid user ID," or "account disabled" can give attackers valuable insights into the system. To mitigate this risk, it is recommended to always respond with generic error messages, such as "login failed" or "we couldn't log you in," without providing specific details about the error.

It is crucial to apply this practice consistently across all possible entry points where an attacker might try to gather information. This includes not only login forms but also password reset forms and account creation forms. Failure to do so might allow attackers to exploit vulnerabilities in these forms to gather sensitive information.

Another consideration in authentication is the timing of responses. Attackers can leverage the time it takes for a system to respond to determine whether a user exists or not. This can be particularly problematic when there are different code paths depending on the existence of a user. By measuring the response time, an attacker can infer whether a user is valid or not. To prevent timing attacks, it is important to ensure that the response time is consistent, regardless of the validity of the user.

In addition to error messages and timing, the HTTP status code can also leak information about the state of the system. It is essential to ensure that the status code remains consistent, even if the error message is the same. Inconsistencies in the status code can still provide valuable information to attackers.

While implementing these security measures is crucial, it is important to consider the trade-off between security and user experience. Generic error messages and consistent response times might be less informative for users, potentially leading to frustration. Therefore, the level of disclosure should be carefully evaluated based on the sensitivity of the information and the potential impact of an attacker enumerating users.

Proper authentication is essential for web application security. By avoiding specific error messages, ensuring consistent response times, and maintaining consistent HTTP status codes, the risk of information disclosure and timing attacks can be significantly reduced.

Authentication is a fundamental aspect of web application security. It ensures that users are who they claim to be before granting them access to sensitive information or functionalities. In this context, it is crucial to understand the potential risks and challenges associated with authentication.

One common concern in authentication is the possibility of introducing timing differences between different code paths. This can occur when validating user credentials against a database. To mitigate this risk, it is recommended to execute the entire code path, regardless of whether the user exists or not. By doing so, the timing difference is eliminated, ensuring a consistent and secure authentication process.

Another important consideration is the need for empirical testing. If authentication timing is critical for a particular service, it is essential to thoroughly test the system. This testing can help identify any potential timing vulnerabilities in the database. By detecting such vulnerabilities, appropriate measures can be taken to mitigate the risks.

Mitigations in authentication have trade-offs, and one such trade-off is the impact on user experience. Using generic error messages can frustrate legitimate users, making it difficult for them to troubleshoot login issues. To address this, it is advisable to provide specific error messages that guide users in identifying the exact cause of authentication failures. This approach enhances user experience and reduces frustration.

To further improve user experience, rate limiting can be implemented for authentication attempts. By limiting the number of attempts within a specific timeframe, potential attackers are deterred from systematically enumerating usernames or passwords. This approach strikes a balance between providing friendly error messages and preventing large-scale enumeration attacks.

Determining the optimal rate limit requires careful consideration. It should be set high enough to accommodate legitimate users, including those sharing the same IP address, such as users in a corporate network. However, it should also be low enough to deter automated bots from overwhelming the system. A rule of thumb is to set the rate limit an order of magnitude higher than the expected maximum number of attempts by legitimate users. Monitoring the system and adjusting the rate limit as needed is crucial to maintain a secure and user-friendly authentication process.



In addition to these considerations, it is important to address the issue of data breaches. Data breaches, such as the Equifax and Yahoo incidents, have become increasingly prevalent, exposing millions of users' sensitive information. These breaches highlight the importance of robust security measures in protecting user data.

To mitigate the impact of data breaches, individuals can take proactive steps, such as locking their credit, to minimize the risk of identity theft. Organizations must also prioritize security measures, including robust authentication mechanisms, encryption, and regular security audits, to safeguard user data and prevent unauthorized access.

Authentication plays a critical role in web application security. By understanding and addressing the challenges and risks associated with authentication, organizations and individuals can enhance security, protect user data, and provide a seamless and secure user experience.

Authentication is a fundamental aspect of web application security. It ensures that only authorized users have access to restricted resources and protects against unauthorized access. In this context, authentication refers to the process of verifying the identity of a user or entity trying to access a system.

One common security issue that can occur is misconfiguration of servers, which can allow unauthorized access to sensitive data. For example, a server might be misconfigured, allowing anyone to connect to it and read the data stored on it. Another issue is command injection or SQL injection attacks, where an attacker can exploit vulnerabilities in a server to gain access and potentially exfiltrate data from a database.

Sometimes, security breaches occur due to simple mistakes, such as leaving an S3 bucket public instead of private. This means that anyone who finds the URL can download all the data stored in the bucket. These breaches can have severe consequences, as attackers can use the stolen information to compromise other servers and systems.

To protect users on a website, proactive measures can be taken. For example, by analyzing data from breaches, it is possible to identify if any users on the site are reusing passwords that were compromised in a breach. If such a case is detected, the user's account can be locked to prevent unauthorized access.

It is important to note that breaches are common, and it is likely that many individuals have been affected by them. Services like "Have I Been Pwned" can be used to check if an email address has been compromised in a breach. This service provides information about the companies that experienced breaches and the type of data that was lost. Additionally, users can set up alerts to be notified whenever a breach occurs, allowing them to change their passwords on other sites if necessary.

When it comes to storing passwords, it is crucial to never store them in plain text. In the event of a data breach, attackers would gain access to all the passwords, which can lead to unauthorized access on other sites where users reuse their passwords. Storing passwords securely, such as using hashing algorithms, is essential to protect user accounts.

Authentication is a critical component of web application security. It helps verify the identity of users and prevents unauthorized access to sensitive resources. Misconfigurations, command injection, and SQL injection attacks can lead to security breaches and data exfiltration. It is important to take proactive measures, such as analyzing breach data, to protect users on a website. Additionally, password storage should be done securely to prevent unauthorized access in the event of a breach.

Authentication is a crucial aspect of web application security. One fundamental concept in authentication is the storage of passwords. Storing passwords in plain text is not recommended as it poses serious security risks. If an attacker gains access to the database, they can easily see all the passwords. Therefore, it is important to hash the passwords before storing them in the database.

Hashing is a process that converts plain text into a fixed-length string of characters. It is a one-way function, meaning that it is computationally infeasible to reverse the process and obtain the original password from the hash. The cryptographic hash function used for password hashing should have certain properties. While speed is generally desirable for hash functions, for password hashing, it is actually better for the function to be slow. This slows down attackers who are trying to crack passwords.

Let's take a look at an example of how password hashing can be implemented in Node.js using the crypto library. We can use the sha-256 hash function to hash the passwords. The code snippet below demonstrates this implementation:

1.	const crypto = require('crypto');
2.	
3.	function sha256(password) {
4.	const hash = crypto.createHash('sha256');
5.	hash.update(password);
6.	return hash.digest('hex');
7.	}
8.	
9.	// When a user provides a password, we hash it and store the hash in the database
10.	const hashedPassword = sha256(userPassword);
11.	// Store hashedPassword in the database

By hashing the passwords, we have improved the security of our application. However, when it comes to authentication, we need to compare the user's input with the stored hash to determine if the password is valid. One property of a hash function is that it is deterministic, meaning that the same input will always produce the same output. This property allows us to compare the hashed password with the stored hash.

1.	// When a user attempts to authenticate, we compare the hashed password with the stored hash
2.	const hashedPasswordFromDatabase = getHashedPasswordFromDatabase();
3.	const userEnteredPassword = getUserEnteredPassword();
4.	
5.	if (sha256(userEnteredPassword) === hashedPasswordFromDatabase) {
6.	// Authentication successful
7.	} else {
8.	// Authentication failed
9.	}

While password hashing provides an extra layer of security, there is still a vulnerability in the system. If two users have the same password, their hashed passwords will also be the same. This information can be exploited by attackers. Additionally, attackers can perform pre-computed lookup attacks. They can generate a database of common passwords and their corresponding hashes, and then compare the hashes in the stolen database to find matches.

To defend against these attacks, we can use techniques such as salting and using stronger hash functions. Salting involves adding a unique random value to each password before hashing it. This ensures that even if two users have the same password, their hashed passwords will be different. Stronger hash functions, such as bcrypt or Argon2, are designed to be slower and more resistant to attacks.

Password hashing is a crucial aspect of web application security. By hashing passwords, we protect user data even if the database is compromised. However, additional measures like salting and using stronger hash functions should be implemented to further enhance security.

Authentication is a crucial aspect of web application security. It ensures that only authorized users can access sensitive information or perform certain actions. One common method used for authentication is the use of passwords. However, passwords alone may not provide sufficient security.

To address this issue, the concept of password salts is introduced. Password salts prevent identical passwords from being easily revealed or identified. They also add entropy to weak passwords, making pre-computer lookup attacks less effective. A salt is a randomly chosen value, typically a short amount of bytes like 16 or 32 bytes. It is concatenated to the password before being hashed.

When a user creates an account, a random salt is generated and combined with their password to produce a hash. This hash, along with the salt, is stored in the database. When the user attempts to log in, their password attempt is combined with the stored salt in the same way, and the resulting hash is compared to the one stored in the database. If they match, the user is authenticated.

It is important to note that the salt itself does not need to be kept secret. It is stored alongside the password in the database. This allows for the validation of passwords during login attempts.

To simplify the implementation of password salts, libraries like bcrypt can be utilized. Bcrypt is a widely-used library designed to handle password hashing. It generates a salt at the time of account creation and includes it in the output hash. This eliminates the need for developers to manually handle salting.

By using bcrypt, developers can simply provide the user's password and specify the desired number of rounds for hashing. The library takes care of generating the salt, combining it with the password, and producing the hash. The resulting hash, which includes the salt, is then stored in the database.

During login attempts, bcrypt is used again to compare the plaintext password provided by the user with the stored hash. The library extracts the salt and other necessary information from the hash and processes the plaintext password in the same way. If the resulting hash matches the stored hash, authentication is successful.

Bcrypt offers additional security features, such as an expensive key setup algorithm. This algorithm ensures that the hashing process takes a significant amount of time, making it more difficult for attackers to crack passwords. Bcrypt is a reliable and widely-used library for password hashing.

Authentication is an essential part of web application security. Password salts provide an extra layer of protection by preventing the easy identification of identical passwords and adding entropy to weak passwords. Libraries like bcrypt simplify the implementation of password salts by handling the generation of salts and the hashing process. By utilizing bcrypt, developers can ensure the secure storage and validation of passwords.

Authentication is a fundamental aspect of web application security. It ensures that users are who they claim to be and prevents unauthorized access to sensitive information. In this context, authentication refers to the process of verifying the identity of a user.

When it comes to storing passwords securely, one commonly used method is bcrypt hashing. Bcrypt is a cryptographic algorithm that adds an additional layer of security to passwords. It uses a combination of salting and multiple iterations to generate a hash. The salt, a predetermined number of bytes, is added to the password before hashing. This salt is stored alongside the hash in the database.

If an attacker gains access to a database containing bcrypt hashes, they face significant challenges. The bcrypt algorithm makes it computationally expensive to crack passwords. In fact, Microsoft published a blog post outlining the costs and efforts required to crack passwords hashed with sha-256, a similar algorithm. They estimated that a machine capable of cracking a hundred billion passwords per second against sha-256 could be built for \$20,000. This highlights the importance of using strong hashing algorithms to protect user passwords.

However, even with strong hashing algorithms, there is still a risk of password compromise. Attackers can exploit breaches where plaintext passwords are exposed. By compiling a list of breached passwords, an attacker can attempt to crack hashed passwords in a database. Statistically, this method has a high success rate, as many users choose weak and easily guessable passwords. Adding song lyrics, news headlines, and other commonly used strings to the list can further increase the success rate.

Given the risks associated with password compromise, it is crucial to implement additional security measures, such as multi-factor authentication (MFA). MFA adds an extra layer of protection by requiring users to provide something they have or something they are, in addition to their password. This can include factors like a fingerprint scan, a one-time password generated by a mobile app, or a hardware token.

Microsoft claims that using MFA can significantly reduce the likelihood of an account being compromised, stating that it is 99.9% less likely to be breached. By implementing MFA, web applications can mitigate the risks associated with password-based attacks.

Authentication is a critical aspect of web application security. Using strong hashing algorithms like bcrypt and implementing additional security measures such as multi-factor authentication can help protect user passwords and prevent unauthorized access.

Authentication is a fundamental aspect of web application security. While passwords are commonly used for

authentication, they are not sufficient to protect against various attacks. Even if a strong password is chosen, it does not guarantee protection against attacks such as credential stuffing, phishing, man-in-the-middle attacks, malware, physical theft of passwords, or brute force attacks.

To enhance security, it is recommended to require a second factor of authentication. This can be done by prompting the user to present a code from their phone or another device. However, to avoid inconveniencing users, this requirement can be selectively enforced based on suspicious behavior. For example, a site can keep track of browsers by using cookies and only require the second factor when a new browser without the cookie is detected. Other factors such as IP addresses, location, or user agent can also be used to prompt for the second factor.

To implement a second factor, a common method is to use time-based one-time passwords (TOTP). This involves using an Authenticator app on the user's phone, such as Google Authenticator. The user scans a QR code provided by the website, which establishes a shared secret key between the server and the phone. The app generates a six-digit code that changes every thirty seconds. This code is then presented to the website as proof of possession of the device.

The server generates a secret key specific to each user and shares it with their phone using a QR code. The phone app initializes a counter to ensure the code changes over time. The counter, along with the secret key, is processed through a function, resulting in a one-time password. The user provides this password to the website, which verifies it against the shared secret key.

By implementing a second factor such as TOTP, web applications can significantly enhance their authentication security. This approach provides an additional layer of protection beyond passwords alone, making it more difficult for attackers to gain unauthorized access.

Authentication is a fundamental aspect of web application security. It involves verifying the identity of users and granting them access to the system based on their credentials. In this context, we will discuss the process of authentication and how it ensures secure access to web applications.

To begin with, authentication requires the use of a shared secret between the user and the server. This secret is typically stored on the user's device and is used to generate a unique code that is sent to the server for verification. The secret can be obtained through various means, such as scanning a QR code or manually inputting it.

Once the secret is obtained, the user's device uses it to generate a code based on a predetermined algorithm. This code is then sent to the server for validation. The server performs the same calculation using the shared secret and compares the generated code with the one received from the user's device. If they match, the user is granted access.

To ensure synchronization between the user's device and the server, a counter is utilized. Both the device and the server maintain a counter that is incremented at regular intervals, typically every 30 seconds. This counter is used as an input in the code generation process, ensuring that both parties are using the same counter value.

To establish the shared secret, the server generates a random set of bytes and stores it in the user's database entry. This secret is then provided to the user via a QR code. When the user wants to generate a code, they take the current time and divide it by 30 seconds to determine the counter value. This counter, along with the shared secret, is used in the code generation process.

The generated code is typically a long hash value. To present it to the user in a more user-friendly format, a few steps are taken. Specific bytes are selected from the hash and then reduced to the desired number of characters, usually a six-digit code. This code is then displayed to the user for authentication purposes.

It is important to note that the server performs the same code generation process to compare the generated code with the one received from the user. If they match, the server confirms that the user is in possession of the shared secret, indicating successful authentication.

Authentication in web applications involves the verification of user identity through the use of a shared secret and a code generation process. By following this process, both the user and the server can ensure secure

access to web applications, even in the absence of an internet connection.

**EITC/IS/WASF WEB APPLICATIONS SECURITY FUNDAMENTALS DIDACTIC MATERIALS****LESSON: AUTHENTICATION****TOPIC: WEBAUTHN**

Authentication is a crucial aspect of cybersecurity, especially when it comes to web applications. It involves verifying the identity of a user to ensure that they are who they claim to be. In this context, the question we need to address is how we can build secure systems even when an attacker has access to a user's password.

One of the key principles in cybersecurity is defense-in-depth, which means that we should have multiple layers of protection to mitigate the impact of a single system failure. With this in mind, we need to explore strategies to enhance authentication security.

Authentication can be based on three factors: something the user knows, something the user has, or something the user is. The first factor, something the user knows, typically involves a password that the user has stored in their memory. The second factor, something the user has, refers to physical items like a phone, an ID badge, or a cryptographic key. The third factor, something the user is, involves physical characteristics such as fingerprints, retina scans, or other biometric data.

It is worth noting that biometric data is becoming increasingly popular for authentication purposes. For example, gait analysis, which analyzes a person's walking pattern, can be used as a unique identifier. Similarly, there are experiments with using the shape of a person's rear-end as a means of authentication in cars. While these may seem unconventional, they highlight the innovative ways in which authentication is evolving.

In the context of web applications, a promising technology for authentication is WebAuthn. WebAuthn allows for the interaction between web applications and various authentication devices, such as fingerprint sensors, face ID sensors, and physical hardware tokens, in a standardized manner. This technology is already being used by platforms like GitHub.

Authentication is a critical aspect of web application security. By implementing robust authentication mechanisms, we can ensure that systems remain secure even when an attacker gains access to a user's password. Exploring innovative authentication methods, such as biometrics and technologies like WebAuthn, can further enhance the security of web applications.

Biometric data is a form of authentication that is not ideal because it is not changeable. Unlike other forms of authentication, such as passwords or cryptographic keys, biometric data cannot be reset if it gets stolen. This is why many systems that use biometric authentication perform the authentication on the device itself, rather than sending the biometric data to a remote service. For example, on iPhone devices, touch ID or face ID data stays on the device and is used to unlock a key store containing a cryptographic key for authentication. WebAuthn, which is a browser authentication mechanism, also follows this approach, ensuring that biometric data is not sent to the server.

In terms of factors used for authentication, the more factors from the three categories (something you know, something you have, something you are) that are utilized, the more confident we can be about the user's identity. For example, when using an ATM, two factors are typically provided: the ATM card (something you have) and the ATM PIN (something you know). Biometric authentication is not commonly used in this scenario.

There is a difference between authentication and authorization. Authentication is the process of verifying a user's identity, while authorization is determining what actions or access privileges the user should have. Authentication can be achieved through various means, such as a login form, cookies, or other types of authentication over HTTP. Authorization, on the other hand, is typically handled using access control lists (ACLs) or capability URLs. ACLs define what a user can do, while capability URLs serve as unique tokens that unlock access to specific resources.

It is important not to confuse or mix up authentication and authorization. They are separate processes with distinct purposes. When providing a session ID to a user, it is used for authentication to verify their identity. However, the actual authorization to perform certain actions is determined separately based on the user's privileges.



When implementing authentication in a web application, there are several common mistakes to avoid. One such mistake is storing usernames in a case-insensitive manner. It is important to store usernames in a case-sensitive manner to ensure accurate authentication.

WebAuthn is a fundamental aspect of web application security that focuses on authentication. One important consideration when implementing authentication is the handling of usernames. It is crucial to ensure that usernames are stored consistently in the database. For example, if a user registers with a lowercase username, it is important to enforce uniqueness so that another user cannot register with the same name but with different capitalization. To accommodate users who prefer their usernames to be capitalized, a second column can be created in the table to store their preferred casing while also saving the lowercase version in another column. This way, when a user enters their username, it is automatically converted to lowercase for comparison, but the preferred casing can be rendered in the user interface.

Passwords are another critical aspect of authentication, and unfortunately, users often choose weak passwords. This is evident from the list of top ten passwords used over the years, which includes common choices like "password" and "123456". To mitigate this issue, password requirements are often implemented on the server side during user registration. However, it is essential to consider what these requirements should be. Outdated advice suggests changing passwords regularly, like every three months, and using a combination of upper and lowercase letters, numbers, and symbols. However, these requirements are not practical and can lead to users simply appending a number to their existing password without providing real security improvements.

Instead of relying on outdated advice, it is important to implement password requirements that actually enhance security. This may include encouraging users to choose longer passwords, as longer passwords are generally more secure. Additionally, it is crucial to educate users about the importance of unique and complex passwords, as well as the risks associated with reusing passwords across different accounts. Implementing strong password hashing algorithms and enforcing multi-factor authentication can also significantly enhance security.

When it comes to web application security and authentication, it is essential to handle usernames consistently and enforce uniqueness. Additionally, password requirements should focus on length, uniqueness, and educating users about the risks associated with weak passwords. By implementing these best practices, web applications can better protect user accounts and data.

WebAuthn is a fundamental aspect of web application security, specifically focused on authentication. It is important to understand the best practices and avoid common pitfalls in order to ensure the security of user accounts.

One common mistake that real sites make is implementing a maximum length on passwords. This is often due to legacy systems that cannot handle longer passwords. However, this practice severely limits the security of user passwords, as any characters beyond the maximum length are truncated or ignored. It is crucial to allow users to create strong and complex passwords without arbitrary limitations.

Another alarming practice is allowing users to log in over the phone using a touchpad. In this scenario, the touchpad only has numbers, which are mapped to three letters each. The backend system then translates the chosen letters into the corresponding numbers. This approach significantly reduces the entropy of passwords, making it easier for attackers to guess or crack them. Additionally, this method allows for more passwords to potentially gain access to user accounts, as different combinations of letters can produce the same numbers on the touchpad.

Minimum password age policies are also seen in some sites, which prevent users from changing their passwords too frequently. The intention behind this policy is to discourage users from repeatedly changing their passwords back to a previously used one. However, this practice is flawed, as it hinders users from promptly changing their passwords in case of compromise or loss.

Disabling the cut and paste feature on password forms is another problematic practice. While the intention may be to prevent users from saving passwords in text files, it also hampers the usability and convenience of password managers. Password managers rely on the ability to copy and paste passwords securely, and disabling this functionality can lead to weaker password management practices.

Password hints are often implemented incorrectly, providing inadequate entropy. For example, some sites use a limited set of hint questions with predefined answers, making it easier for attackers to guess the correct answer. This undermines the security of the account, as the password hint becomes a weaker entry point compared to a strong password.

A misguided approach to combat keyloggers is the use of on-screen keyboards that users can click on with their mouse. While this may protect against keyloggers, it introduces new vulnerabilities. Attackers can simply take screenshots or intercept the clicks on the on-screen keyboard, compromising the user's login credentials. This practice also creates usability issues for regular users, making the login process more cumbersome.

It is essential to avoid these poor practices in web application security. Implementing a maximum password length, allowing login via touchpads, enforcing minimum password age policies, disabling cut and paste, implementing weak password hints, and relying on on-screen keyboards are all detrimental to the security and usability of web applications.

Authentication is a crucial aspect of web applications security. One approach that was used in the past to enhance user security was the implementation of ID shields or secure IDs. The idea behind this approach was to help users detect phishing pages. When creating an account, users would select a unique image that would be displayed whenever they logged into the website. The concept was that a phishing site would not be able to replicate this image, thus allowing users to identify fraudulent pages.

However, there were several problems with this approach. Firstly, if the selected image was not displayed correctly or was replaced with a broken image icon, users may not have been alerted to the presence of a phishing page. Additionally, positive UI indicators, such as the presence of the selected image, were often ignored by users. Another limitation was that only one image could be selected, making it easier for attackers to replicate the login process and trick users into entering their credentials on a phishing page.

Furthermore, a study conducted on the effectiveness of these security images revealed that they were not reliable. In fact, 72% of participants entered their password even when the security image and caption were removed. This study highlighted the ineffectiveness of this approach in preventing phishing attacks.

In recent years, there has been a shift in password requirements. It has been recognized that complexity does not necessarily equate to strength. Forcing users to include numeric, alphabetic, and special symbols in their passwords does not necessarily lead to stronger passwords. Instead, an alternative approach is to choose multiple words from a large dictionary. Even if all the words are common and spelled with lowercase letters and no punctuation, this method can result in stronger passwords compared to using symbols.

To address concerns about users choosing weak passwords when given the freedom to do so, one solution is to check passwords against known breach data. By comparing user-selected passwords with those that have been leaked in previous breaches, it is possible to prevent the use of compromised passwords. This approach is recommended by the National Institute of Standards and Technology (NIST) as an effective way to enhance password security.

Additionally, NIST no longer recommends changing passwords regularly as a security measure. Instead, the focus is on ensuring that passwords are not easily guessable and that they have not been compromised in previous breaches. This shift in perspective acknowledges the reality that breaches are inevitable and encourages a more practical approach to password security.

It is important to note that there is no perfect solution when it comes to balancing usability and security. Implementing strict password policies may restrict some users who do not prioritize security or who have their own reasons for selecting certain passwords. Therefore, it is crucial to find a balance that provides a reasonable level of security without compromising usability.

The use of ID shields or secure IDs for authentication purposes has become obsolete due to their limitations and ineffectiveness in preventing phishing attacks. Password requirements have evolved, with a focus on choosing multiple words from a dictionary rather than relying on complex symbols. Additionally, checking passwords against known breach data is now recommended as a more practical approach to enhancing password security. The emphasis is on preventing the use of compromised passwords rather than regularly changing passwords.

Authentication is a crucial aspect of web application security. In this material, we will discuss the fundamentals of authentication and focus on a specific authentication method called WebAuthn.

One of the main concerns in authentication is the selection of strong passwords. Many users tend to choose weak passwords, even for services that contain sensitive information. For example, people may not prioritize the strength of their Netflix password, despite the potential risks associated with compromised accounts. Attackers often target breached databases or use lists of breached passwords to launch online attacks. Therefore, it is essential to educate users about the importance of selecting strong passwords.

To illustrate the concept of password strength, let's refer to an XKCD comic. The comic emphasizes that common password patterns, such as substituting letters with numerals or adding punctuation, do not significantly increase the entropy or randomness of the password. Instead, it suggests that using a combination of four random words can provide a much stronger password that is also easier to remember.

Password length is another crucial factor in determining its strength. A longer password is generally more secure than a shorter one. Research shows that passwords with fewer than twelve characters are vulnerable to cracking attempts. This vulnerability increases as the password length decreases. Therefore, it is advisable to use longer passwords whenever possible.

To help users understand the impact of password strength, there are various tools available online. These tools allow users to input their passwords and estimate the time it would take to crack them. By using such tools, users can gain insight into the strength of their passwords and make informed decisions about their security.

It is important to consider different attack scenarios when evaluating password strength. Online attacks occur when an attacker systematically tries different username and password combinations by sending requests to the server. The speed of these attacks is limited by the server's response time and security measures. Offline attacks, on the other hand, involve attackers attempting to reverse the hash function used to store passwords in a stolen database. These attacks can be significantly faster, especially if the attacker has access to a powerful cracking array of computers.

Selecting strong passwords and understanding their impact on security is crucial in web application authentication. Users should prioritize longer passwords and avoid common patterns that can be easily guessed or cracked. Educating users about the risks and providing tools to assess password strength can significantly enhance the overall security of web applications.

WebAuthn is a web standard for secure and convenient authentication. It aims to provide a strong and phishing-resistant authentication method for web applications. In this didactic material, we will discuss the fundamentals of WebAuthn and its role in web applications security.

One of the key challenges in web applications security is authentication. Traditional methods, such as username and password, are often vulnerable to attacks like brute force, dictionary attacks, and password reuse. WebAuthn addresses these challenges by introducing a public key cryptography-based authentication mechanism.

WebAuthn relies on the concept of authenticators, which can be hardware devices, like security keys, or software-based solutions, like biometric sensors. These authenticators generate and store cryptographic keys, which are used to verify the user's identity during the authentication process.

The authentication process in WebAuthn involves three main entities: the user agent (typically a web browser), the relying party (the web application), and the authenticator. When a user wants to authenticate, the relying party sends a challenge to the user agent, which in turn prompts the user to select an authenticator. The user then performs an action, such as providing a fingerprint or inserting a security key, to prove their identity. The authenticator generates a public-private key pair and signs the challenge with the private key. The signed challenge, along with the public key, is sent back to the relying party for verification.

WebAuthn provides several advantages over traditional authentication methods. Firstly, it eliminates the need for passwords, reducing the risk of password-related attacks. Secondly, it offers strong cryptographic protection, making it resistant to phishing attacks. Thirdly, it supports multi-factor authentication, allowing users to combine different authenticators for enhanced security.

To implement WebAuthn in a web application, certain best practices should be followed. Firstly, it is recommended to allow users to choose long passwords, potentially up to 64 characters. However, a minimum length requirement, such as eight characters, should also be enforced. Secondly, the length of the password should be limited to avoid denial-of-service attacks. Hashing the password with a fast hash function before storing it can mitigate this risk. Additionally, it is important to check user passwords against known breach data and to implement rate limiting for authentication attempts. Finally, depending on the sensitivity of the web application, the use of a second factor for authentication should be encouraged or required.

When implementing WebAuthn, it is crucial to avoid common implementation mistakes. These include silently truncating long passwords, restricting the characters that users can choose, and accidentally including passwords in plain text log files. Logging mechanisms should be carefully designed to ensure that sensitive information, such as passwords, is not exposed.

WebAuthn provides a secure and user-friendly authentication mechanism for web applications. By leveraging public key cryptography and authenticators, it offers protection against various authentication attacks. Following best practices and avoiding implementation mistakes are essential for ensuring the effectiveness of WebAuthn in enhancing web applications security.

Viewing logs is an essential part of an engineer's daily job. However, in theory, any engineer at Facebook could have seen passwords in plain text. This poses a significant security risk as someone could potentially use these passwords on other services or even on Facebook itself. To mitigate this risk, it is crucial to use Transport Layer Security (TLS) for all traffic and encrypt all data.

Another security concern is network-based guessing attacks. Imagine you have implemented a system where users are selecting strong passwords. However, we need to consider what would happen if an attacker tries to guess passwords for targeted users over the internet. To defend against this, we need to implement certain measures.

Firstly, we need to address the issue of weak password selection by users. When users choose passwords, there is no guarantee of their strength. Attackers can exploit this by iterating over all possible passwords, especially those that are short or obvious. There are three main types of network-based attacks: brute force, credential stuffing, and password spraying.

Brute force attacks involve iterating over a list or dictionary of passwords to target a specific account. Credential stuffing occurs when an attacker uses data from a breach on one site to try the same usernames and passwords on another site, exploiting the common practice of password reuse. Password spraying, on the other hand, involves trying a known weak password on multiple accounts.

To defend against these attacks, rate limiting authentication attempts is a common and effective approach. This can be achieved by using the Express rate limiter package in Node.js. By implementing rate limiting middleware, we can track and restrict the number of authentication attempts for a specific username or IP address within a given time period.

Another defense mechanism is to implement a challenge-response system to verify that the user is a real person. This can be done through methods like CAPTCHA (Completely Automated Public Turing test to tell Computers and Humans Apart). CAPTCHAs require users to complete a task that is easy for humans but difficult for automated bots. If the user successfully completes the challenge, they are unblocked and allowed to continue making authentication attempts.

It is worth noting that some early implementations of CAPTCHAs have been compromised. However, modern CAPTCHAs have evolved to be more secure and effective in distinguishing humans from bots.

To enhance web application security, it is crucial to encrypt all traffic using TLS, implement rate limiting to prevent brute force attacks, and utilize challenge-response systems like CAPTCHA to verify user authenticity.

CAPTCHA (Completely Automated Public Turing test to tell Computers and Humans Apart) is a widely used security measure on websites to distinguish between human users and automated bots. However, there are several vulnerabilities and limitations associated with traditional text-based CAPTCHAs.

One major issue is that the breakage rate of some CAPTCHA implementations is quite high. For example, one implementation has a breakage rate of 92%, while another has a breakage rate of 33%. These rates indicate that even succeeding 33% of the time is considered sufficient for attackers. Researchers typically aim for a breakage rate of 1% or lower, so these implementations are considered completely broken.

Another problem is the user experience of CAPTCHAs. Users often find them time-consuming and frustrating. It takes an average of 10 seconds to complete a CAPTCHA, which is a significant amount of time for users to spend on a single task. Additionally, CAPTCHAs can be challenging for users with visual impairments. While some sites offer alternative CAPTCHAs using audio, these are often less secure and can be parsed by attackers.

A specific vulnerability with CAPTCHAs is the use of image-based challenges. Attackers can detect when they are presented with a CAPTCHA and take a screenshot of it. They can then redirect users from their own site to a victim site and present the captured CAPTCHA to other users. These users unwittingly solve the CAPTCHA, and the attacker uses their answers in real-time to bypass the CAPTCHA on the victim site. This method makes it difficult to prevent CAPTCHA bypass attacks.

To make matters worse, there are dark market services available that offer CAPTCHA solving as a service. For a small fee, attackers can outsource the CAPTCHA solving process to these services, making it even easier for them to bypass CAPTCHAs.

In response to these vulnerabilities, interactive CAPTCHAs have been introduced. These CAPTCHAs require users to interact with elements such as clicking on images of traffic lights. These interactive challenges make it harder for attackers to automate the process, as they need to capture and send all the required user interaction data. However, it is unclear if these interactive CAPTCHAs are completely immune to attacks.

In 2014, a research paper from Stanford University declared that all text-based CAPTCHAs are fundamentally broken. The researchers developed an algorithm called NML which could break any text-based CAPTCHA in a single step, rendering them insecure. They successfully broke various real-world CAPTCHA schemes, including Yahoo, reCAPTCHA, and CNN. This research demonstrated that any breakage rate above 1% is considered a security failure.

To address the limitations of text-based CAPTCHAs, a newer approach called WebAuthn has been developed. WebAuthn analyzes user behavior during a browsing session, including mouse movements, scrolling patterns, and IP addresses, to create a trust score for the user. By building up a reputation for IP addresses over time, WebAuthn can determine if a user is trustworthy or suspicious. If a user is deemed suspicious, a CAPTCHA may be presented. However, for most users, a simple click on a checkbox is sufficient to pass the authentication process.

Traditional text-based CAPTCHAs have several vulnerabilities and limitations that make them unreliable for effective web application security. Researchers have developed alternative approaches, such as interactive CAPTCHAs and WebAuthn, to address these issues. However, it is important to continuously monitor and update security measures as attackers constantly evolve their techniques.

WebAuthn is a web application security feature that focuses on authentication. It provides a secure and user-friendly way to authenticate users on websites. One of the main goals of WebAuthn is to eliminate the use of passwords, which are often weak and easily compromised. Instead, WebAuthn relies on public key cryptography to authenticate users.

When a user interacts with a website that uses WebAuthn, the site logs their login attempts and actions using a service called reCAPTCHA. This service builds up a reputation for the user based on their interactions, determining how trustworthy they are. Other services can then query this reputation and make decisions on whether to allow the user to perform certain actions.

WebAuthn also addresses the issue of automated login attempts and bots. By analyzing client-side code, WebAuthn can detect patterns that indicate a bot-like behavior. This helps prevent unauthorized access and protects websites from malicious activity.

One challenge with WebAuthn is its reliance on IP reputation to determine if a user is real or not. This can lead



to issues for users who value their privacy and use the Tor browser, as their IP addresses may be flagged as suspicious. These users may frequently encounter reCAPTCHA challenges when trying to perform actions on websites.

To further enhance security, websites can implement a defense-in-depth technique called reauthentication. This technique requires users to reenter their password before performing sensitive actions, such as changing passwords or adding new shipping addresses. This provides an additional layer of protection against vulnerabilities like XSS, CSRF, and session fixation.

Another security measure is response discrepancy information exposure. This refers to the practice of revealing different responses to different login attempts. For example, a website might inform a user that there is no account associated with their email address or that the password is incorrect. These different responses can inadvertently disclose information to attackers, such as the existence of an account on the service. It is important to ensure that login responses are consistent to prevent this type of information exposure.

WebAuthn is a powerful authentication mechanism that enhances web application security. By eliminating passwords and leveraging public key cryptography, it provides a more secure and user-friendly authentication experience. However, it is crucial to address challenges such as IP reputation issues and response discrepancy information exposure to ensure the effectiveness of WebAuthn in protecting user accounts and website integrity.

WebAuthn is a fundamental aspect of web application security, specifically focusing on authentication. It aims to provide a secure and user-friendly authentication mechanism for web applications. In this didactic material, we will explore the importance of secure authentication and the potential risks associated with improper implementation.

One common vulnerability in authentication systems is the leakage of sensitive information through error messages. Attackers can exploit this information to gain insights into the state of user accounts. To mitigate this risk, it is recommended to always respond with generic error messages, regardless of whether the user entered an incorrect username, password, or if the account doesn't exist or is in a special state (e.g., locked or disabled). By doing so, we prevent attackers from obtaining specific information about the account. It is crucial to implement this approach consistently across all possible avenues an attacker might use to access this information, including password reset forms and account creation forms.

Let's consider some examples to illustrate the proper implementation of generic error messages. The incorrect approach would be to reveal specific information about the account's state, such as "login for user foo invalid password" or "log-in failed account disabled." Instead, it is best to provide a generic message like "we couldn't log you in; your username or password was incorrect." Although this approach may be less informative for the user, it significantly reduces the risk of disclosing sensitive information.

Another scenario where generic messages should be used is during the account reset process. Rather than stating "we sent you a password reset link" or "this email doesn't exist," it is better to say "if this email exists, we'll send you a link." Similarly, during account creation, avoid revealing whether a user ID is already in use or providing a message like "welcome, you've successfully signed up." Instead, assume the process has worked and avoid disclosing unnecessary information. However, it is essential to consider the user experience, as this approach may be frustrating for users who already have an account and receive an email stating otherwise.

In addition to error messages, the HTTP status code can also leak information about the account's state. It is crucial to ensure that the state remains consistent, even if the error message is the same. Returning different HTTP status codes for the same message can still provide attackers with valuable information.

Timing is another critical aspect to consider in authentication systems. In the provided code example, there is a vulnerability where the execution time of the authentication process can be exploited by an attacker to determine the existence of a user. By measuring the time difference between different code paths, an attacker can deduce whether the user exists or not. To address this issue, it is necessary to design authentication systems that have consistent execution times, regardless of the user's existence.

Implementing secure authentication mechanisms is crucial for web application security. By utilizing generic error messages, ensuring consistent HTTP status codes, and mitigating timing vulnerabilities, we can enhance



the security of our authentication systems. However, it is essential to strike a balance between security and user-friendliness, as overly generic messages may frustrate users. The level of concern for disclosing user information depends on the specific service being provided. Therefore, it is crucial to assess the potential impact of information disclosure and tailor the implementation accordingly.

One important aspect of web application security is authentication, which ensures that only authorized users can access certain resources or perform specific actions. In this context, WebAuthn (Web Authentication) is a fundamental technology that provides a secure and user-friendly way to authenticate users on the web.

When implementing authentication in web applications, it is crucial to consider potential timing attacks. Timing attacks exploit differences in the time taken to perform certain operations to gain unauthorized access. For example, an attacker could use the response time of a login request to determine if a user exists in the system, potentially revealing sensitive information.

To mitigate timing attacks, it is recommended to eliminate any timing differences between code paths. This can be achieved by executing the same code regardless of whether the user exists or not. By hashing the user's password and looking it up in the database, the authentication process becomes consistent and secure. If the password is valid, the user is granted access, otherwise an error is thrown.

Empirical testing is also important to ensure the effectiveness of authentication mechanisms. In real-world scenarios, the database could take varying amounts of time to respond, potentially leaking information about the existence of user accounts. By testing the system and identifying such vulnerabilities, appropriate measures can be taken to mitigate the risks.

However, it is essential to strike a balance between security and user experience. Applying mitigations may introduce trade-offs that impact the user. For instance, using generic error messages can frustrate legitimate users, especially when they are unable to determine the cause of the error. To avoid this, it is recommended to provide specific error messages that help users understand and resolve the issue.

Another approach to enhance security without compromising user experience is to implement rate limiting. By limiting the number of authentication attempts within a certain time frame, attackers are prevented from systematically enumerating accounts. This allows for friendlier error messages while reducing the risk of brute-force attacks.

Determining the optimal rate limit requires careful consideration. Setting it too low may inadvertently block legitimate users, such as those sharing the same IP address within a corporate network. A general rule of thumb is to set the rate limit slightly higher than the most extreme usage scenario. Monitoring and adjusting the rate limit based on observed behavior can help strike the right balance between security and usability.

Data breaches are a prevalent concern in the realm of web application security. High-profile breaches, such as the Equifax and Yahoo incidents, highlight the importance of safeguarding user information. In the Equifax breach, the personal data of 143 million US customers, including Social Security numbers, was compromised. Similarly, Yahoo experienced a breach where one billion user accounts were initially reported as compromised, but later revealed to be three billion.

These incidents serve as a reminder that even well-established companies can fall victim to data breaches. It is essential for organizations to prioritize robust security measures to protect user data and mitigate the potential consequences of such breaches. Users can also take proactive steps, such as locking their credit, to minimize the impact of data breaches.

Authentication is a critical aspect of web application security. Implementing secure authentication mechanisms, addressing timing attacks, conducting empirical testing, and considering trade-offs between security and user experience are essential steps in ensuring the integrity and confidentiality of user data.

WebAuthn is a fundamental aspect of web application security that focuses on authentication. It is important to understand the vulnerabilities that can exist in web applications and the potential consequences of these vulnerabilities. For example, misconfigured servers can allow unauthorized access to data or be susceptible to command injection or SQL injection attacks. Once an attacker gains access to a server, they may be able to pivot to other servers and ultimately exfiltrate sensitive data.

One common mistake that can lead to data breaches is leaving S3 buckets public instead of private. S3 buckets are used for hosting static files, and if anyone can access them, it becomes easy for attackers to download all the data. This highlights the need for proper security measures to protect users on a website.

One proactive approach to protecting users is to analyze data from breaches. By checking if any of the users on a service are reusing passwords that were compromised in a breach, it is possible to identify vulnerable accounts and take appropriate action, such as locking the account. This can help prevent attackers from using stolen credentials on the website.

To check if you have been a part of any breaches, you can use a service like "Have I Been Pwned." This website allows you to enter your email address and see if your information has been compromised in any known breaches. It also provides details about the specific data that was lost in each breach. Additionally, you can set up alerts to be notified whenever a new breach occurs, allowing you to take immediate action to secure your accounts.

When it comes to storing passwords, it is crucial to never store them in plain text. In the event of a data breach, attackers would gain access to all user passwords, which can have severe consequences as users often reuse passwords across multiple sites. Storing passwords securely is essential to protect user accounts and prevent further compromises.

WebAuthn is a critical component of web application security, specifically focusing on authentication. Understanding the vulnerabilities that exist in web applications and implementing proper security measures, such as password breach analysis and secure password storage, is vital to protect user accounts and prevent data breaches.

In web applications, it is crucial to ensure the security of user passwords. Storing passwords in plain text is not only insecure but also highly embarrassing if a hack occurs and the information becomes public. To address this issue, it is recommended to hash the plaintext passwords and store only the hash values in the database.

A cryptographic hash function is used to perform the hashing process. The properties of a hash function are important to consider in this context. While speed is desirable for some use cases, it is actually better for password hashing to be slow. This is because it makes it more difficult for attackers to crack passwords by slowing down their attempts. Other properties, such as determinism and uniqueness, are also important for a hash function used in database authentication.

To implement password hashing in a web application, one can use the crypto library in Node.js. The createHash function can be used to generate a hash object, which has update and digest methods. The update method allows for the progressive input of data, while the digest method returns the final hash value. By using the sha-256 hash algorithm, a function called sha-256 can be defined to simplify the process.

When a user provides a password, it can be hashed using the sha-256 function and then stored in the database. This approach improves security compared to storing passwords in plain text. Later, during the authentication process, the user's provided password can be hashed and compared to the stored hash value in the database to determine its validity.

It is important to note that the hash function is deterministic, meaning that the same input will always produce the same output. While this may seem concerning in terms of password security, it is not a problem as long as users choose unique passwords. The main objective is to ensure that the function returns the same hash value for a given user's password during subsequent login attempts.

However, there is a limitation to consider in this approach. Since the sha-256 function is deterministic, if two users have the same password, their hash values will also be the same. This can be observed in the database, even without knowing the actual password. Additionally, precomputed lookup attacks can be performed by attackers who know the hash function used. By computing and saving the hash values of common passwords in a separate database, an attacker can easily match the hash values from the target database to identify the corresponding passwords.

To defend against these vulnerabilities, additional measures are required. One common solution is to use a

technique called "salting." Salting involves adding a unique random value, known as a salt, to each user's password before hashing it. This ensures that even if two users have the same password, their hash values will be different due to the different salts. Salting effectively prevents precomputed lookup attacks and makes it more difficult for attackers to crack passwords.

To enhance web application security, it is essential to hash passwords using a cryptographic hash function and store only the hash values in the database. While deterministic hashing may seem concerning, it is not an issue as long as users choose unique passwords. However, to further strengthen security, the use of salting is recommended to prevent precomputed lookup attacks and increase the complexity of password cracking.

Passwords are commonly used for authentication in web applications, but they can be vulnerable to attacks. One solution to this problem is the use of password salts. Password salts prevent users with identical passwords from being revealed or identified, and they also add entropy to weak passwords, making pre-computer lookup attacks ineffective.

A password salt is a randomly chosen value, typically a short amount of bytes like 16 or 32 bytes. It is concatenated to the password before it is put through the hash function. The resulting output, which includes the salt, is stored in the database. The salt does not need to be kept secret and is stored alongside the password.

When a user provides their password, the salt is used to combine with the password attempt in the same way as before, producing a new hash. This new hash is then compared to the stored hash in the database. If they match, the user is authenticated.

To implement password salts, the following code can be used:

1.	# Generate a random salt
2.	salt = generate_random_salt()
3.	
4.	# Combine the salt with the user's password
5.	salted_password = salt + user_password
6.	
7.	# Hash the salted password
8.	hashed_password = hash_function(salted_password)
9.	
10.	# Store the hashed password and salt in the database
11.	store_in_database(hashed_password, salt)

To validate a password later, the same steps are repeated:

1.	# Retrieve the salt and hashed password from the database
2.	stored_hashed_password, stored_salt = retrieve_from_database()
3.	
4.	# Combine the stored salt with the password attempt
5.	salted_password_attempt = stored_salt + password_attempt
6.	
7.	# Hash the salted password attempt
8.	hashed_password_attempt = hash_function(salted_password_attempt)
9.	
10.	# Compare the hashed password attempt with the stored hashed password
11.	if hashed_password_attempt == stored_hashed_password:
12.	# Authentication successful
13.	authenticate_user()
14.	else:
15.	# Authentication failed
16.	deny_access()

If you prefer not to implement password salts manually, you can use a library like bcrypt. Bcrypt is a widely used library that automatically handles password salting and hashing. It has been around for many years and has a strong track record of security.

To use bcrypt, you can simply call the bcrypt library functions instead of manually implementing the salt and

hash steps. Here's an example:

1.	import bcrypt
2.	
3.	# Generate a salt and hash the password
4.	hashed_password = bcrypt.hashpw(user_password, bcrypt.gensalt())
5.	
6.	# Store the hashed password in the database
7.	
8.	# Validate a password later
9.	if bcrypt.checkpw(password_attempt, stored_hashed_password):
10.	# Authentication successful
11.	authenticate_user()
12.	else:
13.	# Authentication failed
14.	deny_access()

In this example, bcrypt automatically generates a salt and includes it in the hashed password. When comparing passwords later, bcrypt handles the parsing and comparison of the stored hashed password.

Using bcrypt has additional benefits, such as an expensive key setup algorithm that slows down attackers and the elimination of the need to manage salts manually.

Password salts are an important technique to enhance the security of web application authentication. They prevent password-related vulnerabilities and add entropy to weak passwords. Implementing password salts manually or using libraries like bcrypt can greatly improve the security of user authentication.

WebAuthn is a web application security protocol that focuses on authentication. It provides a secure way for users to authenticate themselves on websites. In this didactic material, we will discuss the fundamentals of WebAuthn and its role in ensuring the security of web applications.

One important aspect of WebAuthn is the concept of hashing. When a user creates an account or changes their password, the password is not stored in plain text. Instead, it is hashed using a cryptographic algorithm called bcrypt. The hashed password is then stored in a database. The bcrypt hash includes several components: a number that represents the number of iterations used to hash the password, a dollar sign (\$), a salt value, and the actual password hash. The salt value is a predetermined number of bytes that adds an extra layer of security to the hashing process.

However, even with the use of bcrypt, it is essential to consider the potential risks if an attacker gains access to the database. Microsoft conducted research and found that a machine capable of cracking a hundred billion passwords per second against sha-256, a different hashing algorithm, could be built for \$20,000. This means that even strong hashing algorithms like bcrypt are not foolproof, and there is always a risk of password cracking.

To mitigate the risks associated with password cracking, it is crucial to implement additional security measures, such as multi-factor authentication (MFA). MFA adds an extra layer of protection by requiring users to provide something they have or something they are, in addition to their password. This could be a physical device like a security key or a biometric identifier like a fingerprint.

Microsoft claims that using MFA significantly reduces the likelihood of an account being compromised, stating that it is 99.9% less likely. By implementing MFA, web applications can enhance their security posture and provide users with an added level of protection.

WebAuthn is a web application security protocol that focuses on authentication. It utilizes bcrypt hashing to store passwords securely in databases. However, it is important to recognize the potential risks associated with password cracking. Implementing additional security measures like multi-factor authentication can significantly enhance the security of web applications and protect user accounts.

Passwords alone are not enough to ensure security in web applications. Even if a strong password is chosen, it does not protect against various attacks such as credential stuffing, phishing, man-in-the-middle attacks,

malware, physical theft, or simple negligence. Therefore, it is important to require a second factor for authentication.

One approach to implementing a second factor is to prompt the user to present a code from their phone or another device. However, this may inconvenience users if required every time they log in. To address this, it is possible to only require the second factor in certain situations, such as when suspicious behavior is detected. For example, a site can keep track of browsers by using cookies and only require the second factor when a browser without the cookie shows up, indicating a new device. Alternatively, the site can prompt the user if they are logging in from a new location, country, or IP address. Suspicious IP addresses can be identified by referring to publicly available lists maintained by the community of site operators. Another method is to monitor login attempts and detect abnormal behavior, such as one person trying to log into multiple accounts within a short period of time. Additionally, examining the user agent or other behavioral aspects of browsing can help identify scripts rather than real users.

One common second factor is time-based one-time passwords (TOTP). This method involves using an Authenticator app on a user's phone, such as Google Authenticator. The user scans a QR code provided by the website, which establishes a shared secret key between the server and the phone. The phone generates a six-digit code that changes every thirty seconds. When logging in, the user provides this code to prove possession of the device and verify their identity. The secret key is used in combination with an HMAC function and other steps to generate the code. This process ensures that the server can verify the user's possession of the shared secret.

To implement TOTP, the server creates a secret key specific to each user. This key is then shared with the user's phone app, typically through a QR code. The phone app initializes a counter to track time and uses it, along with the secret key, to generate the one-time password. The user can provide this password to the site for authentication. The process is repeated every thirty seconds to generate a new code.

Passwords alone are not sufficient for web application security. Requiring a second factor, such as time-based one-time passwords, adds an extra layer of protection against various attacks. By implementing measures to prompt for the second factor only in certain situations and monitoring suspicious behavior, web applications can enhance security and protect user accounts.

WebAuthn is a fundamental aspect of web application security, specifically in the realm of authentication. It provides a secure and reliable way to verify the identity of users accessing web applications. In this context, WebAuthn utilizes a shared secret, stored on the user's device, to establish trust between the user and the server.

The process begins with the server generating a random value, which serves as the secret key. This key is then stored in the server's database, associated with the user's account. To convey this key to the user, a QR code is generated. The user scans the QR code and retrieves the key.

When the user wishes to generate a code for authentication, they take the current time and divide it by thirty seconds. This calculation yields a counter value. The counter value, along with the secret key, is passed through an HMAC (Hash-based Message Authentication Code) function, which produces a hash. To create a user-friendly code, a few bytes are selected from the hash and modulated by the desired number of characters, typically six digits. This resulting code is then presented to the user.

The server follows a similar process to verify the user's authenticity. It calculates the counter value based on the current time and the same thirty-second interval. The server then applies the HMAC function to the counter value and the secret key. If the resulting hash matches the code provided by the user, the server confirms the user's possession of the key and allows access to the web application.

It is important to note that hashing and salting passwords or using bcrypt are essential practices to ensure password security. By implementing these measures, the risk of unauthorized access to user accounts is significantly reduced. Additionally, developers should consider additional layers of protection for users, even in scenarios where attackers have acquired passwords and other sensitive information.

WebAuthn is a powerful tool for web application security, offering robust authentication capabilities. By leveraging shared secrets, QR codes, and HMAC functions, users can securely access web applications, while

servers can verify their identities. Implementing best practices, such as password hashing and salting, further enhances the security of user accounts.



**EITC/IS/WASF WEB APPLICATIONS SECURITY FUNDAMENTALS DIDACTIC MATERIALS****LESSON: MANAGING WEB SECURITY****TOPIC: MANAGING SECURITY CONCERNS IN NODE.JS PROJECT**

Today, we will be discussing security concerns in Node.js projects. Our guest lecturer, Miles Boren, is a member of the Node.js technical steering committee and is responsible for ensuring the quality and security of Node.js releases. He is also a member of PC 39, the technical steering committee for the JavaScript language specification. Miles has extensive experience in the industry, working as a developer advocate for Google Cloud Platform, with a focus on the JavaScript ecosystem.

In this lecture, Miles will cover various aspects of web security in the context of Node.js projects. He will begin by providing an overview of the Common Weakness Enumeration (CWE) and the Common Vulnerabilities and Exposures (CVE) system, which are maintained by the organization MITRE. These systems categorize and track different types of vulnerabilities and weaknesses in software.

Miles will then discuss the importance of understanding the expectations of the language community when building secure applications. He will explain that while a piece of code may be technically correct, it may not align with the community's expectations. As a developer advocate, Miles often works with product teams to ensure that technical decisions align with community expectations and best practices.

Additionally, Miles will touch on his role in helping to address security vulnerabilities within Google's own products. He assists in conducting security audits and ensuring that vulnerabilities are addressed and releases are properly managed. This includes working with various teams and advising on technical decisions and release cadence.

Throughout the lecture, Miles will provide real-life examples and practical advice on managing security concerns in Node.js projects. He will also be open to questions and discussions on topics such as open source, working in the industry, and more.

Note: The views expressed in this lecture are solely those of the speaker and do not necessarily reflect the views of Google or its security practices.

In the field of cybersecurity, managing web security is of paramount importance. One aspect of managing security concerns in a Node.js project involves understanding common vulnerabilities and exposures (CVEs) and common weakness enumerations (CWEs).

CVEs are lists of found CWEs in projects and are maintained by MITRE, a non-profit organization. By referring to the CVE list, developers can identify vulnerabilities specific to their project, such as those related to Node.js. Each CVE is assigned a unique number, which helps in tracking and reporting vulnerabilities.

CWEs, on the other hand, provide a common language for researchers to describe specific types of vulnerabilities. For example, CWE 435 refers to improper interaction between multiple correct behaving entities. Understanding these CWEs helps in accurately reporting and addressing vulnerabilities.

The Node.js project has taken the initiative to become its own CVE Numbering Authority (CNA) to expedite the release process and ensure that vulnerabilities are promptly addressed. As a CNA, the Node.js project assigns CVE numbers to vulnerabilities before reporting them to MITRE. Although it may take some time for these vulnerabilities to be updated in international databases, the project is not dependent on MITRE for releasing security fixes.

To assess the severity of vulnerabilities, the Common Vulnerability Scoring System (CVSS) is used. The CVSS provides a framework for evaluating the risk associated with a vulnerability. Factors such as the attack vector (network adjacent, network local, or physical), attack complexity, and privilege requirements are taken into account to determine the base score. By using the CVSS calculator, developers can assess the risk level of a vulnerability and prioritize their efforts accordingly.

Threat modeling is another important aspect of managing security concerns in a Node.js project. By analyzing the environment in which the project operates, developers can identify potential threats and prioritize security

measures. Depending on the environment, certain vulnerabilities may pose a higher risk than others.

Managing security concerns in a Node.js project involves understanding and utilizing CVEs, CWEs, and the CVSS. By staying informed about vulnerabilities specific to the project and assessing their severity, developers can take proactive measures to ensure the security of their web applications.

Web applications security is a crucial aspect of cybersecurity, especially when it comes to managing security concerns in Node.js projects. In order to effectively manage web security in Node.js, it is important to understand the different security concerns and their impacts.

One of the key considerations in web application security is the exposure of data that shouldn't be accessed. This raises concerns about confidentiality. For example, side-channel attacks can lead to data exfiltration without compromising data integrity. On the other hand, availability impact refers to the possibility of a denial-of-service (DoS) attack affecting the availability of the machine.

To assess the severity of these security concerns, various metrics can be used. These include temporal score metrics and environmental score metrics, which provide a comprehensive evaluation of the potential risks. It is worth noting that the scoring system is based on an algorithm and is designed to serve as a general threat model. Therefore, a high score on this system doesn't necessarily mean the same level of risk for every system.

An important concept in web application security is the notion of zero-day vulnerabilities. These are vulnerabilities that have not yet been patched or disclosed. Zero-days pose a significant risk because they can be exploited by malicious actors without the knowledge of the software developers. White hat security researchers play a crucial role in discovering and reporting zero-day vulnerabilities to projects, allowing them to take necessary actions to mitigate the risk.

However, managing security concerns in Node.js projects presents unique challenges. Node.js is a volunteer-run organization, including its security team. While some contributors may have the support of their employers to dedicate time to the project, everyone involved is a volunteer. This introduces risks as the project relies on the trustworthiness and expertise of the contributors.

To mitigate these risks, the project carefully vets individuals before granting them access to sensitive information. Contributors who work for reputable organizations like Google or Microsoft are preferred due to the accountability and trust associated with their employment. Additionally, having colleagues who can vouch for them further strengthens the trust.

Despite these precautions, there are still vulnerabilities that remain unpatched or undisclosed due to complexity or resource limitations. This leaves the project susceptible to zero-day attacks if security researchers disclose them before the project can address them adequately.

One example of a vulnerability is the hash wick vulnerability, which involved a timing attack against the Node.js server to determine the hash seed used for randomization. This vulnerability also exposed the version of Node.js being used. It serves as a reminder of the potential risks and the importance of timely patching.

Managing security concerns in Node.js projects requires a thorough understanding of web application security fundamentals. By assessing the impacts of confidentiality, integrity, and availability, and by addressing zero-day vulnerabilities and managing trust within the contributor community, the project can work towards maintaining a secure web environment.

Web applications security is a critical aspect of managing security concerns in Node.js projects. In this didactic material, we will explore some fundamental concepts related to managing web security in Node.js projects.

One important vulnerability to be aware of is the hash lookup vulnerability. This vulnerability can occur when an attacker exploits the use of a hash seed to cause a hash lookup vulnerability. By repeatedly making requests with the same key, the attacker can overload the lookup table, causing the server to run slow. This vulnerability was fixed in V8, the JavaScript engine used in Node.js and Chrome browser. However, it was disclosed before it could be fully patched.

Node.js is a server-side runtime environment that allows developers to write and execute JavaScript code. One

key difference between Node.js and other JavaScript runtimes is the inclusion of its own system interface called libuv. While browsers heavily sandbox system APIs, Node.js implements its own system APIs, which are not standardized other than through implementation. This can lead to interesting edge cases when providing programs with system access.

When writing code in JavaScript for Node.js, the code is executed by the V8 engine in a virtual machine. Node.js creates manual bindings to give V8 system access. It is important to note that the JavaScript language itself, as specified by TC39, does not have any system interfaces. Efforts are being made to create a standardized system interface called WASI (WebAssembly System Interface) for JavaScript runtimes.

Confidential information in the context of web security is often placed under embargo. Embargoes mean that the information is not publicly disclosed until a certain date. As a member of the Node.js security triage team, all information received is under embargo until it is publicly disclosed. However, managing embargoes can be challenging, especially when wearing multiple hats such as an Node.js release engineer, Node.js security triage member, and a Google employee working on cloud runtimes. It is important to adhere to the embargo policy to avoid leaking information that could give certain companies a competitive advantage.

Triage is a crucial process in managing security vulnerabilities. Vulnerabilities in Node.js core can be reported through a platform called HackerOne. HackerOne serves as a management system for these vulnerabilities, allowing users to report and track vulnerabilities. The triaging process involves assessing and prioritizing reported vulnerabilities to ensure they are addressed in a timely manner.

Managing web security in Node.js projects involves addressing vulnerabilities such as hash lookup vulnerabilities, understanding the unique system interface provided by Node.js, and adhering to embargo policies when handling confidential information. Triage processes, such as those facilitated by platforms like HackerOne, play a crucial role in managing and addressing reported vulnerabilities.

#### Managing Security Concerns in Node.js Project

In a Node.js project, it is crucial to manage security concerns effectively to ensure the safety and integrity of web applications. This didactic material will provide an overview of the various measures and programs in place to address security concerns in Node.js projects.

One important aspect of managing security concerns is utilizing tools such as the Common Vulnerabilities and Exposures (CVE) system. This system allows developers to track vulnerabilities and their impact on projects. Additionally, the Common Weakness Enumeration (CWE) system helps in categorizing and addressing specific weaknesses in the project.

To assist in managing security concerns, the Node.js project collaborates with HackerOne, an organization that provides resources and support for handling security reports. When new reports are received, the HackerOne team reviews them first to ensure that the project team's time is not wasted on non-issues. This pre-triage process helps in efficiently addressing genuine security problems.

Not only does HackerOne assist with the Node.js project, but it also manages security concerns for the ecosystem as a whole. This includes popular modules like Express and Torrent Stream. A separate team is responsible for triaging vulnerabilities reported in these modules. By centralizing the reporting process, the Node.js project can effectively address security concerns across the entire ecosystem.

Another program that aids in managing security concerns is the Internet Bug Bounty (IBB). This program, run by HackerOne, offers monetary rewards to security researchers who discover bugs in core internet infrastructure projects. Node.js participates in the IBB program, but as of now, no bounties have been awarded. The IBB primarily focuses on vulnerabilities related to remote code execution.

It is essential to understand the different types of vulnerabilities that can affect a Node.js project. Firstly, vulnerabilities in the core Node.js platform pose a significant risk as they can affect every application running on Node.js. Secondly, vulnerabilities can also be present in the Node.js ecosystem, which consists of numerous packages available through NPM. Developers must be vigilant about the security of the packages they use in their applications. Finally, vulnerabilities can exist within the applications themselves, which may or may not be related to the packages used.

In the Node.js core, various threats are addressed, including buffer overflow attacks, denial of service attacks, data exfiltration, remote code execution, hostname spoofing, and vulnerabilities in dependencies. It is not just the ecosystem that needs to be concerned about dependencies; Node.js itself has dependencies that must be regularly updated to address vulnerabilities.

Understanding the lifecycle of a vulnerability is crucial for effective management. When a researcher discovers a bug, they report it through HackerOne. The report goes through a triage process to determine its validity and severity. From there, the Node.js project team takes appropriate action to address the vulnerability, including releasing patches and updates.

Managing security concerns in a Node.js project requires a comprehensive approach. By leveraging tools like the CVE and CWE systems, collaborating with HackerOne, participating in the IBB program, and addressing vulnerabilities in the core platform, ecosystem, and applications, developers can ensure the security and reliability of their web applications.

When managing security concerns in a Node.js project, it is important to have a clear process in place. One of the first steps is triaging reported vulnerabilities. This involves reviewing the vulnerability report and determining if it is a genuine security issue. If it is confirmed as a vulnerability, it is considered triaged and work begins on addressing it. However, sometimes reported issues are not actually vulnerabilities but rather spec bugs or other non-security-related problems. In such cases, the person reporting the issue is informed and pointed to the appropriate information.

In the case of a confirmed vulnerability, the next step is to identify a solution. Ideally, this process should be completed quickly, but it can sometimes take months, depending on the complexity of the issue. It is worth noting that cases where vulnerabilities sit in triage for over a year are extremely rare. However, the Node.js team strives to address security concerns as promptly as possible. Communication with the person who reported the vulnerability is crucial, especially for high severity vulnerabilities. Generally, the team aims to stay in touch with the person on a weekly basis, or even every 24 or 48 hours if necessary.

Once a solution has been identified, a security release is created. This is a critical part of the process but can be challenging due to the volunteer nature of the Node.js organization. The build team, which is also run by volunteers, plays a significant role in this step. The team manages the CI infrastructure, which supports various platforms and architectures. Node.js supports multiple flavors of Linux, BSD, Solaris, Windows, OSX, ARM systems, Z system, PowerPC, and s/390. The CI infrastructure consists of different build bots, each representing a specific platform or architecture. These build bots are persistent, meaning they are not killed after each run like containers in some other CI systems. This persistence allows for efficient testing and debugging of security releases.

It is worth mentioning that system differences can sometimes lead to vulnerabilities. For example, in the past, there was a vulnerability related to the implementation of symbolic links and real path resolution in Node.js. To address this vulnerability, a patch was applied, but it inadvertently caused issues on specific platforms due to inconsistencies in the implementation of Unix APIs. A workaround was implemented to ensure compatibility with older versions of Unix that still had the vulnerability.

Managing security concerns in a Node.js project involves a thorough triaging process, prompt communication with the reporter, identifying and implementing solutions, and creating security releases. The volunteer nature of the Node.js organization and the diverse range of supported platforms and architectures add complexity to the process. However, the team strives to address security concerns efficiently and maintain regular communication with the community.

Web applications security is a critical aspect of cybersecurity, and managing security concerns in Node.js projects requires careful attention. One important consideration is the potential exposure of sensitive information through the console output. If the Continuous Integration (CI) system is public during the testing of security patches, unauthorized individuals could gain access to the system and potentially exploit any identified vulnerabilities. To mitigate this risk, it is necessary to restrict access to the CI system during testing. The CI should only be accessible to the triage team and the release team, ensuring that sensitive information remains confidential.

Shutting down the entire CI system for all collaborators, except for the authorized teams, may cause disruptions. However, it is a necessary step to safeguard the project's security. Collaborators must be informed of the temporary unavailability of the CI system and should refrain from running any processes on it. Additionally, a separate sandbox CI environment should be set up specifically for releases. This sandbox CI should be configured to access embargoed repositories and handle the necessary tasks for releasing updates. It is crucial to ensure that the sandbox CI is properly connected to the relevant repositories, even if the team is working on different forks.

Once the necessary precautions are taken, the vulnerability can be disclosed. This typically involves issuing a security advisory, assigning a Common Vulnerability Scoring System (CVSS) score, publishing a blog post, and notifying relevant message boards. Following the disclosure, it is common for users to update their applications to address the vulnerability. More information about the security process can be found on the official website of the organization, specifically on the security tab. The website provides details on reporting bugs, the bounty program, and third-party bug disclosure policies.

To report vulnerabilities, the organization employs the services of HackerOne, a platform that facilitates responsible disclosure. The landing page of the organization's HackerOne program provides information about the reporting process and the response times users can expect. The organization's security team aims to respond to vulnerability reports within 24 hours and provide an update within 48 hours. The team monitors its performance using metrics provided by HackerOne. The service level agreement (SLA) ensures that the team meets its response standards. The organization's bounty program offers rewards starting at a minimum of \$500 for identified vulnerabilities, with no Remote Code Execution (RCE) vulnerabilities reported so far.

The HackerOne page also displays the activity related to reported vulnerabilities. While some reports may still be under embargo, others are publicly available. An interesting example of a security incident involved a domain takeover of the "registry.nodejs.org" subdomain. An abandoned subdomain was found pointing to a service that allowed new distributions without proof of domain ownership. The organization worked with the reporter to resolve the issue, deleting the erroneous CNAME record and conducting a comprehensive audit of their DNS infrastructure.

Managing security concerns in Node.js projects requires careful attention to protect sensitive information and promptly address vulnerabilities. By following established security procedures, such as restricting access to the CI system, disclosing vulnerabilities, and leveraging platforms like HackerOne, organizations can effectively manage security concerns and ensure the integrity of their web applications.

In managing web security for Node.js projects, it is crucial to consider all possible attack vectors and threat modeling. One such vulnerability occurred when an attacker gained control over a portion of the Node.js domain. While the project itself was not at fault, the reliance on a compromised website posed a security concern.

To address security concerns, a submit report feature was implemented, allowing users to report any suspicious activity. This feature guides users through a login flow, providing a secure means of communication. Additionally, it serves as a platform for security researchers to build their profile and potentially earn rewards.

To further understand web application vulnerabilities, the CVE (Common Vulnerabilities and Exposures) database is a valuable resource. CVE-2017-14919 is an example of an improper input validation vulnerability. In this case, Node.js versions prior to 4.8.6, 6.x prior to 6.11.5, and 8.x prior to 8.8 allowed remote attackers to cause a denial of service by exploiting a change in the zealnd module. This change made an invalid value for the window bits permit parameter, resulting in an uncaught exception and a crash.

By exploring the CVE, users can gain insights into the associated CWE (Common Weakness Enumeration) and its relevance to research concepts. The CVE also provides an extended description, examples, and references to relevant standards.

In the case of CVE-2017-14919, the vulnerability was introduced due to a security update to the ZLib library. Upgrading the dependency introduced a flaw in the window bits parameter, causing Node.js to crash when creating a raw deflate stream. Prior to the update, window bits 8 was a valid value, but it was replaced with window bits 9. This change had a significant impact on the interface of ZLib, potentially causing crashes in applications that manually set the window bits.



Exploiting this vulnerability could lead to a denial of service attack, particularly if an application allowed users to specify the window size. By passing an invalid window bits argument, an attacker could disrupt the availability of the service.

To address this vulnerability, a patch was created. The patch, released in Node.js versions 4.8.2 and 6.10.2, updated the ZLib library from version 1.2.8 to 1.2.11. This patch resolved the issue by addressing the low severity CVE-2017-14919. It is important to note that a pull request was submitted to implement an 8-bit window deflate stream in ZLib, but it did not receive a response.

Managing security concerns in Node.js projects requires a proactive approach, considering potential attack vectors and staying informed about vulnerabilities and patches. By utilizing resources like the CVE database, developers can better understand and address web application security.

In the context of web application security, managing security concerns in Node.js projects is crucial. This didactic material will explore two specific security issues in Node.js and provide insights into their impact and mitigation strategies.

The first security concern discussed is related to a vulnerability in the raw deflate stream initialization with window bits set to eight. This invalid value caused an error to be raised, leading to a crash in Node.js. This crash was irreversible in some versions, rendering the service unavailable. The per message deflate library, up to version 0.1.5, did not handle this error gracefully, resulting in service disruptions. To address this issue, a commit was made to revert the behavior of the Zlib library, changing window bits to nine. This fix was accompanied by a test to ensure that the problem would not reoccur. However, it is important to note that the fix is pending review and has not been merged into the main repository yet. Therefore, the temporary solution remains in place.

The second security concern pertains to an authentication bypass and spoofing vulnerability, known as CDE 2018 71 60. This vulnerability affects the Node.js inspector in versions six and later. It is susceptible to a DNS rebinding attack, which can lead to remote code execution. This attack can be initiated from a malicious website running on the same computer or another computer with network access to the Node.js process. By exploiting the DNS rebinding attack, the malicious website can bypass the same-origin policy checks and establish HTTP connections to local hosts or hosts on a local network. If a Node.js process with an active debug port is running on localhost or a local network host, the attacker can connect to it as a debugger and gain full code execution capabilities.

To illustrate the impact of this vulnerability, consider a scenario where a developer is running Node.js on their personal computer for testing or debugging purposes. If the debugger is active and a malicious website tricks the browser into thinking it is running on localhost, the attacker can execute arbitrary code on the developer's machine. This vulnerability highlights the importance of securing the debug port and implementing proper authentication mechanisms to prevent unauthorized access.

Mitigating this vulnerability requires careful consideration. It may no longer be possible to debug a remote computer by using a hostname due to the potential risks associated with DNS rebinding attacks. Instead, connecting using the IP address or employing an SSH tunnel is recommended as a workaround. Although this change may impact some remote debugging scenarios, it is essential to prioritize security and prevent potential exploits.

Managing security concerns in Node.js projects is crucial to ensure the integrity and safety of web applications. By addressing vulnerabilities such as the raw deflate stream initialization issue and the DNS rebinding attack in the Node.js inspector, developers can enhance the security posture of their applications. It is essential to stay updated with the latest security patches, follow secure coding practices, and implement robust authentication mechanisms to mitigate potential risks.

Node.js is a popular platform for building web applications, and managing security concerns is an important aspect of any Node.js project. In Node.js, releases are managed through a versioning system called SemVer (Semantic Versioning). Every six months, a new major version is released, and there is a master release line where all changes are merged. Different release channels are maintained for each major version.



Currently, there are four major versions being maintained: version 13, version 12, version 10, and version 8. Each version has its own independent branch, and when new releases are made, changes from the master branch are backported to the specific release lines. However, semver major changes are not backported to these release lines.

Semver minor changes refer to non-breaking changes that add new features. For example, adding a new API or extending the capabilities of an existing API without changing behavior or expectations would be considered a semver minor change. On the other hand, semver patch changes do not add or break anything, but can include documentation fixes, new tests, bug fixes, or changes to experimental APIs.

One important aspect to note is that breaking changes can still be considered semver minor if they are necessary for security updates. For example, if a security vulnerability is discovered and needs to be patched, even if it breaks existing functionality, it can be considered a semver minor change. In such cases, it is important to prioritize security over compatibility.

An interesting example of managing security concerns in Node.js is the HTTP parser. Node.js has a legacy parser for HTTP called HTTP underscore parser, written in C++. It has been rewritten in TypeScript, which is easier to manage and maintain. The default parser in Node.js version 12 is the new TypeScript parser, while versions 10 and 8 still use the old C++ parser. Although it is believed that the two parsers are semantically equivalent, changes are not made to the LTS (Long Term Support) branches to avoid potential performance or compatibility issues.

Maintainability is a key factor in managing security concerns. While stability, reliability, and security are important, making wide changes to LTS branches can introduce unexpected issues. Refactoring for speed improvements or other changes that could potentially affect performance are avoided in LTS branches to ensure stability and avoid breaking large systems.

Managing security concerns in a Node.js project involves balancing stability, reliability, security, and maintainability. It is important to prioritize security updates, even if they result in breaking changes. Semver minor changes are used for non-breaking feature additions, while semver patch changes are for non-breaking fixes and improvements. The HTTP parser in Node.js is an example of how security concerns are managed while ensuring compatibility and maintainability.

Managing security concerns in a Node.js project is crucial for ensuring the overall web security of an application. In this didactic material, we will discuss some important security vulnerabilities and concerns related to Node.js and how to manage them effectively.

One of the key aspects of managing web security in a Node.js project is understanding the importance of timely patching and updates. Due to time constraints and limited resources, it is essential to prioritize and be pragmatic when deciding which vulnerabilities to fix. For example, when a version of Node.js reaches its end-of-life, it is recommended to stop patching it, even if new vulnerabilities are discovered. This is done to avoid using a version of Node.js that relies on an outdated and unsupported version of OpenSSL, which can lead to security issues.

Now, let's discuss some specific security concerns and vulnerabilities in Node.js. One vulnerability, CVE-2018-12115, involves out-of-bounds writes. In all versions of Node.js prior to a certain release, when using certain encoding formats, it was possible to write outside the bounds of a single buffer. This can lead to various security escalations, such as denial of service attacks or remote code execution. Attackers can exploit this vulnerability by combining it with other attacks, making it even more dangerous.

Another notable vulnerability, discovered in November 2018, was related to the legacy debugger in Node.js. After fixing a vulnerability in the new inspector-based debugger, it was realized that the legacy debugger had the same vulnerabilities. This meant that older versions of Node.js, which were not patched because they didn't have the inspector, were still vulnerable. This highlights the importance of thoroughly assessing all components of a project for potential security vulnerabilities.

In addition to these vulnerabilities, there were other security concerns in Node.js. One involved large HTTP headers, where carefully timed requests with maximum-sized headers could cause the HTTP server to abort due to heap allocation failures. Another concern was the Slowloris attack, where slow requests with a specific timing

could lead to memory leaks and degrade the performance of the Node.js process over time. These attacks could be difficult to detect and could significantly impact the availability of the service.

Furthermore, there were specific vulnerabilities related to the URL parser for the JavaScript protocol, HTTP request splitting, and timing attacks. These vulnerabilities highlight the importance of thorough testing and continuous monitoring of the security of a Node.js project.

To effectively manage security concerns in a Node.js project, it is essential to follow best practices, such as keeping the Node.js version up to date, regularly applying security patches, and conducting thorough security assessments. Additionally, implementing security measures like input validation, proper authentication, and authorization mechanisms can help mitigate potential security risks.

Managing security concerns in a Node.js project is crucial for ensuring the overall web security of an application. By understanding and addressing specific vulnerabilities and following best practices, developers can enhance the security posture of their Node.js applications.

Web applications security is a crucial aspect of cybersecurity, especially when it comes to managing security concerns in Node.js projects. In this context, it is important to understand the potential threats that can affect both the application and the ecosystem.

One such threat is supply chain attacks, where attackers exploit vulnerabilities in the dependencies used by the application. This highlights the importance of using trusted and secure dependencies in Node.js projects. Weak cryptography is another concern, as it can expose sensitive data to potential attackers. It is essential to use strong encryption algorithms and keep them up to date.

Another significant concern is the developer experience. While it may be tempting to implement strict security measures, it is important to strike a balance between security and pragmatism. Creating a poor developer experience by imposing overly restrictive processes can lead to developers finding workarounds, compromising security. One approach to address this is to allow developers to work in isolated development clusters where they can experiment freely, while conducting a thorough review of dependencies before promoting them to production.

Malicious third-party code is also a potential problem, often associated with supply chain attacks. It is crucial to carefully vet and monitor the dependencies used in the project to mitigate this risk. Query injections are another common vulnerability in web applications that can lead to security breaches. Developers should be aware of this risk and implement proper input validation and sanitization techniques to prevent such attacks.

To assist in understanding the threat environment and establishing a security process for Node.js applications, a comprehensive document called the Node SEC Roadmap is available. This document provides insights into threat modeling, differentiating between server-side and client-side code, and classifying threats based on frequency and severity. It also emphasizes the need to tailor the threat model to the specific context of the project.

Managing web security in Node.js projects requires a holistic approach that addresses various security concerns. By understanding the potential threats, ensuring the use of secure dependencies, creating a balanced developer experience, and implementing proper security measures, organizations can enhance the security of their web applications.

Web Applications Security Fundamentals - Managing web security - Managing security concerns in Node.js project

When it comes to managing security concerns in a Node.js project, there are several important factors to consider. One of the key considerations is the use of cloud functions. While cloud functions offer benefits such as easy deployment and scalability, they also introduce new security concerns. For example, the node runtime and operating system can be updated underneath you, which means that you may not have control over the latest version of Node.js or the system libraries included in the base container image.

Another important consideration is the scaling model used in cloud functions. Google, for example, runs cloud functions with one tenant per one function, meaning that each function will spin up a new instance for each

request. This can help mitigate denial of service (DoS) attacks, as an attack on one instance will not affect the availability of other instances. However, it opens up the possibility of resource attacks, where an attacker can cause your cloud function to consume excessive resources and result in increased costs.

To address these security concerns, it is crucial to implement proper threat modeling and security measures. This includes setting maximum scaling limits and implementing observability warnings to detect traffic spikes beyond a certain threshold. Additionally, understanding the dynamic nature of Node.js and managing dependencies is essential. It is recommended to follow a workflow that includes working in a sandbox environment and gradually increasing security measures as you move towards production.

Another potential threat vector in Node.js projects is the execution of remote code during the npm install process. Post-install scripts can run on every module installation, essentially providing a remote code execution environment. It is important to carefully consider where npm install is being run and whether it is within a sandbox or isolated environment to prevent compromising the main system.

Managing security concerns in Node.js projects requires a comprehensive approach that includes threat modeling, understanding the dynamic nature of Node.js, managing dependencies, and implementing proper security measures. By following best practices and staying informed about potential vulnerabilities, developers can ensure the security of their web applications.

Node.js is a popular platform for building web applications, but it is not immune to security vulnerabilities. In fact, there have been over 180 vulnerabilities disclosed in Node.js. Some of these vulnerabilities are quite unique and require a creative mindset to exploit.

One vulnerability that stands out is the Bower arbitrary file write through Imperva of siblings package extraction. This vulnerability allows an attacker to write arbitrary files on the server, which is obviously a serious security concern. Another vulnerability is the command injection in the ASCII art package. This vulnerability allows an attacker to inject malicious commands into the application, potentially leading to unauthorized access or data leakage.

It is important to note that these vulnerabilities are not limited to high-profile packages. Even smaller modules, like the one with only 1400 downloads per month, can have security flaws. For example, the module mentioned above does not properly sanitize the target command line argument, which can lead to command injection attacks.

Understanding the language of Common Weakness Enumeration (CWE) and Common Vulnerabilities and Exposures (CVE) can help in identifying and mitigating such vulnerabilities. For instance, if a vulnerability is labeled as a SQL injection attack, one can search for CWE SQL injection to learn more about the history and characteristics of this type of vulnerability.

Supply chain attacks are another significant concern in the Node.js ecosystem. These attacks exploit the trust we place in the developers of the modules we use. One notable incident involved the event stream module, which was hacked to target Bitcoin wallets. The attacker managed to publish a vulnerable version of the module, which would secretly exfiltrate the keys to users' Bitcoin wallets. This attack was sophisticated, as the malicious code was hidden in the published package but not in the source code on GitHub. Users had to extract the package from npm to discover the attack.

Supply chain attacks are not new, and they can have severe consequences. It is crucial to be aware of the potential risks and take appropriate measures to secure your application. Regularly updating dependencies, reviewing code, and staying informed about the latest security vulnerabilities are some of the best practices to mitigate supply chain attacks.

Managing security concerns in Node.js projects requires a proactive approach. It is essential to be aware of the vulnerabilities that exist in the Node.js ecosystem, regardless of the package's popularity. Understanding the language of CWE and CVE can aid in identifying and addressing vulnerabilities effectively. Additionally, supply chain attacks pose a significant threat, and developers should be vigilant in reviewing and securing their dependencies.

Web security is a critical concern for any project, especially those built on Node.js. The distributed nature of

ecosystems like Node.js makes them vulnerable to supply chain attacks. While Node.js is particularly susceptible due to its broad ecosystem and numerous dependencies, it is important to recognize that supply chain attacks can occur in any software development environment.

No matter how much research and vetting you do, it is likely that vulnerabilities may still be missed. Simply looking at source code is not enough to protect against these attacks. One approach to mitigating this risk is to implement network policies within the container or runtime environment. By restricting network access to specific domains, the attack would be unable to break out of the sandbox and access other domains.

However, it is not necessary to run Node.js exclusively in a Docker container on your computer. This is just one avenue to consider for protecting against supply chain attacks. It is important to acknowledge that it is impossible to solve every security problem. Instead, focus on creating processes that will protect your project as much as possible.

There are several steps you can take to enhance the security of your Node.js project. First, always use the latest version of Node.js. This ensures that you have the latest security patches and updates. Additionally, tools like NPM audit can help you identify and update vulnerable packages. Greenkeeper, a tool recently acquired by GitHub, can also assist in protecting your packages by monitoring for updates and vulnerabilities. It is also important to regularly audit all dependencies to identify any potential security risks.

Sandboxing is another consideration for enhancing security. By implementing sandboxing techniques, you can isolate and control the execution environment of your code. However, it is crucial to strike a balance between security and maintaining development velocity. Being too strict with auditing and security measures can hinder productivity and discourage developers from following the necessary processes.

Finally, the location where you run your code can also impact security. While the Node.js project itself benefits from a pre-triage process by HackerOne, vulnerabilities can still be found. It is important to report any vulnerabilities to the appropriate channels, such as the ecosystem triage team or the security researchers involved in the project.

Managing security concerns in a Node.js project requires a multi-faceted approach. It involves using the latest version of Node.js, leveraging tools like NPM audit and Greenkeeper, considering sandboxing techniques, and being aware of the location where your code is executed. By implementing these measures, you can enhance the security of your web applications and protect against supply chain attacks.

**EITC/IS/WASF WEB APPLICATIONS SECURITY FUNDAMENTALS DIDACTIC MATERIALS****LESSON: SERVER SECURITY****TOPIC: SERVER SECURITY: SAFE CODING PRACTICES**

## Server Security: Safe Coding Practices

In this session, we will be discussing server security and safe coding practices. We will explore different ways in which servers can be compromised and the measures we can take to defend against such attacks. The focus will be on safe coding practices to prevent security issues from arising.

One important aspect of server security is handling user authentication. We have previously discussed SQL injection and rate limiting as methods to protect against unauthorized access. Today, we will delve deeper into these topics and explore additional ways in which servers can be vulnerable to attacks.

Before we begin, I would like to mention a hands-on workshop hosted by Pete Snyder from Brave, our guest lecturer. The workshop will be held on Thursday from 6 to 7 PM in Gates 174. Dinner will be provided. Additionally, we have some new students in class today, so please extend a warm welcome to them.

Now, let's discuss some recent security findings. Three students have discovered security issues and reported them, earning extra credit. One student found a cross-site scripting (XSS) vulnerability in Access, a part of the Stanford bug bounty program. This student reported the bug and received a \$100 reward from Stanford. Another student discovered an XSS vulnerability in a CS course website, which could potentially allow unauthorized access to change assignment scores. Finally, a student found an insecure design in a coding challenge for a job interview, which exposed test cases. While this student did not receive a bug bounty reward, it highlights the importance of secure coding practices.

Moving on, I would like to share a story that demonstrates the severity of server security vulnerabilities. A security researcher was able to bypass the GitHub authentication flow and gain unauthorized access to user data. Typically, users are prompted to grant permissions to an application before accessing their data. However, this researcher found a way to send a request to GitHub with the user's cookies attached, granting the application unrestricted access to the user's GitHub account. This vulnerability affected all GitHub deployments, including enterprise deployments used by companies for sensitive code. GitHub awarded the researcher the highest bounty ever paid out, emphasizing the severity of the issue.

This story highlights the importance of thorough security testing and the potential impact of even small vulnerabilities. It also showcases the value of bug bounty programs, where individuals can make a living by discovering and reporting vulnerabilities.

Server security is crucial to protect against unauthorized access and data breaches. By adhering to safe coding practices and staying vigilant, we can mitigate the risks associated with server vulnerabilities.

Cross-Site Request Forgery (CSRF) is an attack where an attacker forces a user to execute actions against a web application that they are currently authenticated with. This attack takes advantage of the ambient authority model of cookies used in the browser's authentication mechanism. Ambient authority refers to the fact that once a user logs into a site and proves their identity, all future requests to that site automatically have the same authority as that user.

To understand how CSRF works, let's consider a pictorial representation. We have a client and a server, with the server being the victim in this case. The user visits the site and logs in by sending a login request with their username and password. Assuming the credentials are valid, the server sends back a response containing a set cookie header that sets a cookie session ID. This cookie is stored by the browser.

Later, when the user is still logged into the site, they may open another tab or click on a link that leads them to an attacker's website. The attacker's website sends a request to the victim server, including the parameters required for a sensitive endpoint, such as transferring money. The browser automatically attaches the victim server's cookies to this request. From the server's perspective, it appears to be a legitimate request since it contains the necessary information and the attached cookies. This allows the attacker to perform actions on behalf of the user without their consent.

To mitigate CSRF attacks, same-site cookies can be used. Same-site cookies specify that cookies should only be attached when the request is initiated by the same site. If the request is coming from the victim site to the victim site itself, the browser includes the cookie header. However, if the request is coming from an attacker's site to the victim site, the browser does not include the cookie header. Same-site cookies provide a reliable way to prevent CSRF attacks.

It's worth mentioning that before the introduction of same-site cookies, CSRF tokens were used to achieve a similar outcome. CSRF tokens were necessary when browsers did not support the same-site cookie attribute. They provided a way for websites to prevent unauthorized form submissions. However, with the availability of same-site cookies, CSRF tokens have become less relevant.

CSRF is an attack that takes advantage of the ambient authority model of cookies in the browser's authentication mechanism. By forcing a user to execute actions against a web application they are authenticated with, an attacker can perform actions on behalf of the user without their consent. Same-site cookies provide an effective mitigation strategy against CSRF attacks by ensuring that cookies are only attached when the request is initiated by the same site.

A crucial aspect of server security in web applications is safe coding practices. One important practice is the implementation of CSRF tokens. A CSRF token, which stands for Cross-Site Request Forgery token, is a nonce - a secret and unpredictable value generated by the server and transmitted to the client. The client must include this token in all subsequent HTTP requests to the server for them to be recognized as valid.

To include a CSRF token in a request, an input element with the type "hidden" is added to the page. This input element has the name "CSRF token" and its value is set to the actual token. When the form is submitted, all the inputs, including the CSRF token, are sent to the server as part of the form.

There are two approaches to generating CSRF tokens. One approach is to randomly pick a value and use it as the token. The other approach is to generate the token based on information in the request, such as the session ID.

The CSRF token protects against Cross-Site Request Forgery attacks. When a client interacts with a page that includes the token, any subsequent requests made by the client will include the token. The server checks if the token in the request matches the one it provided earlier. If they match, the request is considered valid and the server responds accordingly.

This protection mechanism prevents attackers from exploiting the trust between a user and a website. Even if a user visits an attacker-controlled site, the attacker cannot provide the correct CSRF token, making their requests invalid.

CSRF tokens are an essential part of server security in web applications. By including and validating these tokens in HTTP requests, web developers can protect against Cross-Site Request Forgery attacks and ensure the integrity of their server's operations.

When it comes to server security in web applications, safe coding practices are crucial to prevent attacks. One important aspect of server security is protecting against Cross-Site Request Forgery (CSRF) attacks.

In a CSRF attack, an attacker tricks a victim into performing an unwanted action on a web application. This is done by exploiting the fact that web browsers automatically include cookies in requests to the server. These cookies are used to authenticate users and maintain their sessions.

To protect against CSRF attacks, web developers can implement a technique called CSRF tokens. A CSRF token is a unique value that is generated by the server and included in each form submission or request that modifies data on the server. The token is then checked by the server to ensure that the request is legitimate and not coming from an attacker.

When a user visits a web page that contains a form, the server includes a CSRF token in the HTML code. This token is typically stored as a hidden field in the form. When the user submits the form, the token is sent back to the server along with the rest of the form data.



The server compares the CSRF token received from the client with the token it originally sent to the client. If the tokens match, the server knows that the request is legitimate and can proceed with the requested action. If the tokens do not match, the server rejects the request.

By using CSRF tokens, web developers can prevent attackers from tricking users into performing unwanted actions. The attacker cannot read the CSRF token from the victim's page due to the Same Origin Policy, which prevents JavaScript code from accessing resources from a different domain.

To implement CSRF tokens effectively, servers should generate a unique token for each user session and store it on the server side. When the server receives a request, it compares the CSRF token in the request with the stored token for the user's session. This ensures that even if an attacker manages to obtain a valid session cookie, they cannot generate a valid CSRF token without knowledge of the server's secret.

A common approach is to use the session ID as part of the CSRF token generation process. The server combines the session ID with a secret known only to the server using a cryptographic function like HMAC. This ensures that the CSRF token is unique to each session and cannot be easily guessed by an attacker.

By implementing CSRF tokens and following safe coding practices, web developers can significantly enhance the security of their web applications and protect against CSRF attacks.

CEO surf tokens are an effective defense mechanism against attackers. By choosing tokens randomly from a large range, it becomes nearly impossible for attackers to guess the correct token. This prevents them from sending a large number of requests with all possible tokens, which would be unfeasible.

Same-site cookies provide a simpler solution for server-generated nonces. Instead of having separate logic for generating and checking nonces, a same-site attribute can be added to the cookie when it is set. This eliminates the need for additional logic and reduces the chances of errors.

The attacker may not have full control over the refer header, but they can cause it to be omitted. In such cases, the server needs to determine whether the omission is due to malicious reasons or if the browser is trying to enhance privacy.

Now that we understand CEO surf tokens, let's discuss the attack process. When a user is prompted to grant an app permission to access their GitHub account, they are redirected to a specific URL that includes information about the app. The user is then presented with an authorization page where they can authorize the app. If the user grants access by clicking the authorize button, they are redirected back to the third-party application. During this redirect, a GitHub token is included in the query string, which allows the application to access the user's data by sending requests to GitHub.

The authorize button is implemented as an HTML form with a CSRF token embedded in a hidden form field. This ensures that an attacker cannot simulate a button click by sending a POST request to the URL. When the server receives the POST request, it validates the CSRF token to ensure that the button click originated from the page itself and not from an attacker's site.

It is worth noting that the form submits to the same URL that the page is loaded from. The only difference is that the initial request is a GET request, while the form submission is a POST request. The server can differentiate between the two by examining the HTTP method and perform different actions accordingly.

CEO surf tokens and same-site cookies are effective measures to enhance server security. The flow of authorizing an application involves redirecting the user to an authorization page, obtaining their consent, and redirecting them back to the application with a token for accessing their data. The authorize button is implemented as an HTML form with a CSRF token to prevent unauthorized button clicks.

When it comes to server security in web applications, safe coding practices are crucial to prevent vulnerabilities and potential attacks. One important aspect of server security is the implementation of authorization flows, which allow users to log in and access their accounts securely.

In the context of GitHub, the login and authorization flow involves several steps. When a user clicks on the

"Authorize" button, a form is submitted via a post request to the same URL. This request includes a CSRF token, which is a security measure to prevent cross-site request forgery attacks. The server then checks the validity of the token and, if it is valid, sends a successful response to the user. The user is then redirected back to the application's page with a GitHub token attached, which allows the application to access the user's account.

This flow seems secure and free from issues as long as the server properly checks the CSRF token. However, a security researcher was able to obtain a copy of the GitHub source code and discovered a potential problem in the implementation of the authorization flow.

The source code revealed that requests to a specific URL, "login/authorize," are handled by a controller function. If the request is a GET request, the server serves an HTML page with a green button. If the request is a POST request, the server checks the CSRF token and grants permission to use the application if it is valid.

The potential problem lies in the fact that the server treats HEAD requests as GET requests. A HEAD request is similar to a GET request, but it omits the body of the response, only sending the headers. This is useful for checking information about a resource without actually downloading it. However, HEAD requests are relatively niche and not commonly used by most developers.

The Ruby on Rails framework, used by GitHub, assumes that developers will not bother implementing HEAD requests and automatically handles them as GET requests. While this may seem convenient, it can lead to a security vulnerability. Since GET requests are not supposed to modify any data, developers may not implement proper security measures in the controller function for GET requests. By treating HEAD requests as GET requests, the server may inadvertently expose sensitive information or perform unintended actions.

Server security in web applications requires safe coding practices. Authorization flows, such as the one used by GitHub, play a crucial role in providing secure access to user accounts. However, it is important to be aware of potential vulnerabilities, such as mishandling of HEAD requests, which can lead to unintended consequences and compromise the security of the application.

#### Server Security: Safe Coding Practices

In web application development, it is crucial to ensure the security of the server to protect against potential attacks. One aspect of server security involves safe coding practices. This didactic material will focus on a specific issue related to safe coding practices in server security and how it can be exploited by attackers.

When developing web applications using frameworks like Ruby on Rails or Express, it is important to understand how certain HTTP requests are handled. In the case of a HEAD request, which is used to retrieve header information from a server, it is typically automatically handled by the framework. This means that the same controller code that handles GET requests will also handle HEAD requests.

While this automatic handling of HEAD requests can be a time-saving feature for developers, it can also lead to a potential security vulnerability. This is because the abstraction provided by the framework may not perfectly hide the complexity from the developer, resulting in what is known as a "leaky abstraction."

In the specific case of Ruby on Rails, there is a function called `request.get?` that returns `false` for HEAD requests. This can be unexpected for developers who assume that their controller code will only run for GET and POST requests. If they have an if statement in their code that checks for `request.get?`, it may inadvertently include HEAD requests as well.

This unintended inclusion of HEAD requests in the controller code can lead to a security vulnerability. For example, if the code includes a check for a Cross-Site Request Forgery (CSRF) token, an attacker can exploit this vulnerability by sending a HEAD request without a valid CSRF token. Since the code assumes that the token has already been checked, it will authorize the application, granting the attacker access to the user's account.

To understand how this vulnerability can be exploited, let's consider the typical CSRF token handling in Ruby on Rails. Normally, there is a function that runs before the controller code and checks if the request is a POST request. If it is, the CSRF token is validated. However, a HEAD request does not trigger this CSRF token checking code, bypassing the security measure. As a result, the controller code assumes that the token has already been checked and can be omitted, allowing the attacker to gain unauthorized access to the user's account.

It is important for developers to be aware of these potential vulnerabilities and ensure that their code properly handles different types of requests. In the case of HEAD requests, it may be necessary to explicitly exclude them from certain code paths or implement additional security measures to prevent unauthorized access.

Safe coding practices are essential for maintaining server security in web applications. Developers should be cautious when handling different types of requests and ensure that their code properly handles potential vulnerabilities. By understanding the intricacies of server security and implementing appropriate measures, developers can protect their applications and users from potential attacks.

In web applications security, server security plays a crucial role in ensuring the safety of the application and protecting user data. One important aspect of server security is safe coding practices. In this material, we will discuss a specific vulnerability that could have been exploited to compromise user accounts on GitHub and explore ways to prevent such attacks.

The vulnerability in question involved a lack of Cross-Site Request Forgery (CSRF) protection in the server code. The attack scenario begins when a user visits a malicious server and receives HTML code that triggers a request to GitHub. This request includes the user's cookies but lacks a CSRF token. Surprisingly, the server does not verify the presence of the CSRF token and simply authorizes the request, redirecting the user to the attacker's server with a GitHub token.

To prevent such attacks, it is essential to implement safe coding practices. One approach is to use a package like "C surf" to add CSRF protection to the server code. By configuring this package to only check against POST requests, developers can mitigate the vulnerability. However, relying solely on this check for POST requests may not be sufficient.

A more defensive coding practice is to avoid making assumptions about the type of request being made. Instead of assuming that the request will be either a GET or a POST, developers should include an "else if" statement to handle any other type of request and consider throwing an exception. This defensive coding paradigm ensures that unexpected requests do not bypass the necessary checks and helps identify potential vulnerabilities.

Crashing the server process in case of unexpected requests is actually a desirable outcome in terms of security. While it may temporarily affect the availability of the site, it quickly alerts developers to the presence of an attacker and allows them to investigate and fix the root cause. In a production environment, such crashes trigger alerts, enabling immediate response and resolution.

In addition to these measures, there may be other ways to enhance server security and prevent similar attacks. For example, implementing stricter input validation and sanitization techniques can help prevent injection attacks. Regularly updating server software and libraries to patch known vulnerabilities is also crucial.

Ensuring server security in web applications requires implementing safe coding practices. Specifically, protecting against CSRF attacks is essential. By using packages like "C surf" and adopting defensive coding practices, developers can significantly reduce the risk of unauthorized access and protect user accounts and data.

When it comes to server security and safe coding practices, there are several important considerations to keep in mind. One of the key aspects is the trade-off between explicit and magical behavior in coding. While relying on magical behavior may seem convenient, it can also lead to unexpected issues and vulnerabilities. On the other hand, being explicit in coding can help prevent mistakes and make the code more understandable for new developers.

In the context of web applications, one of the common practices is to use different URLs for different functionalities. For example, having separate URLs for authorization and form submission can help prevent security issues. Additionally, using separate controllers for different functionalities can also be a good practice. By having separate functions for GET and POST requests, developers can ensure that the code behaves as intended and prevent unexpected behavior.

In the case of Express, a popular web application framework, it is common to register methods for specific HTTP

methods like GET and POST. However, it is not possible to mix different HTTP methods in a single registration. An exception to this is the use of the `app.use` function, which allows handling all HTTP methods in a single function. However, when using this approach, developers need to handle every method explicitly to avoid any unexpected behavior.

Another important consideration is the use of CSRF tokens for security. CSRF tokens are used to protect against cross-site request forgery attacks. However, there are alternative approaches, such as using same-site cookies, which can simplify the implementation and reduce complexity. By using same-site cookies, developers can avoid the need for CSRF tokens and the associated complexity.

To ensure the security of server-side code, it is recommended to be explicit in checking the HTTP method used in requests. By explicitly checking for GET or POST methods, developers can prevent unexpected behavior and ensure that the code behaves as intended. In cases where unexpected methods are encountered, it is advisable to throw an exception and crash the application, as trying to recover from such situations can be risky.

In terms of preventing security issues like the one discussed, frameworks like Rails could implement measures to mitigate the risk. For example, requiring a token check for HEAD requests could ensure that only authorized requests are allowed. Additionally, designing the framework in a way that does not trap users into potential vulnerabilities is crucial.

Server security and safe coding practices are essential in web application development. Being explicit in coding, using separate URLs and controllers for different functionalities, and implementing measures like token checks can help prevent security issues. Additionally, considering alternative approaches like using same-site cookies can simplify the implementation and reduce complexity.

When it comes to server security in web applications, safe coding practices are of utmost importance. One aspect that developers need to consider is the handling of head requests. In some cases, it may be tempting to send a head request to retrieve the homepage of a site without prior knowledge of its details. However, this approach can be problematic if the site requires a CSRF token. In such cases, developers would first need to send a get request to obtain the token and then send the head request, which defeats the purpose of sending just the head request.

While it is unclear how common it is to directly check get requests, there are potential solutions to prevent such issues. One approach is to force developers to handle head requests themselves instead of handling them automatically. However, this may lead to a decrease in support for head requests on sites. Another idea is to set the request method to "get" even for head requests. Although this may seem like lying to the developer, it can be justified by considering that the developer explicitly stated their preparedness to handle get requests in the controller. By pretending that the head request is a get request, the developer's code can run correctly, and the framework can handle the response and delete the body, ensuring that the abstraction remains intact.

A term that is relevant in this context is "leaky abstraction." The purpose of abstractions is to hide unnecessary complexity from developers, allowing them to focus on solving the problem at hand. However, a leaky abstraction occurs when certain implementation details or complexities are exposed, violating the intended simplicity of the abstraction. In the case of handling head requests, the abstraction was meant to hide this complexity, but it ended up leaking out, causing developers to have to consider these edge cases, which defeated the purpose of the abstraction.

When it comes to server security in web applications, safe coding practices are crucial. Handling head requests requires careful consideration to avoid potential issues. By following safe coding practices and ensuring that abstractions are not leaky, developers can enhance the security and reliability of their web applications.

One of the reasons why issues like this occur is because there is no way to enforce that every possible case is handled in the developer's code. In strongly typed languages, a compiler error would occur if a case was missed. However, this can slow down development. One possible solution is to manually inspect the code and check if it accesses certain variables or functions. For example, in Ruby, you could use introspection to check the function, while in JavaScript, you could convert the function to a string and parse it. However, this approach is not recommended.

Based on the different solutions proposed, there are some lessons we can learn to prevent these problems in

our code or libraries. One common theme in security is to reduce complexity, as complexity often leads to security issues. Abstractions that hide complexity from developers can be leaky if they have many edge cases. Additionally, when introducing a new component, it's important to consider how many other components it could potentially interact with. Explicit code is better than clever code, as clever code can become difficult to understand over time. It's important to write code that is easy to comprehend, even if it may look less elegant. Finally, failing early is another important concept. If something is in an unexpected state, it's better to crash or throw an exception rather than trying to handle it.

To improve server security in web applications, it is important to reduce complexity, write explicit code, and fail early when necessary. By following these practices, developers can minimize the risk of security issues in their code or libraries.

In the context of web application security, server security plays a crucial role in ensuring the safety and integrity of the server-side code. Safe coding practices are essential to minimize vulnerabilities and protect against potential attacks. In this didactic material, we will discuss the fundamentals of server security and highlight some key safe coding practices.

When developing web applications, it is common to encounter errors or exceptions. These can occur during multi-step processes, making it challenging to identify the exact location of the error and the state of the application. In such cases, it is advisable to crash the entire process instead of attempting to resume. By doing so, we eliminate the risk of continuing with an unknown state, ensuring the integrity of the system. Additionally, crashing the process allows for a fresh start when the server process is rebooted, minimizing downtime for users.

Code defensive programming is another crucial aspect of server security. It involves assuming that our assumptions will be violated and verifying them upfront. For example, when writing functions, we should not assume that they will always be called with the correct arguments. It is essential to validate inputs and handle potential errors gracefully. By anticipating and addressing potential issues proactively, we can enhance the security and reliability of our code.

While discussing safe coding practices, it is important to note that relying solely on obscurity for security is not recommended. Obscurity refers to hiding the inner workings of the code or system to deter potential attackers. However, this approach is not reliable as it does not address the underlying vulnerabilities. Instead, we should focus on implementing robust security measures and following best practices to protect our applications.

In certain scenarios, errors can be expected, such as when querying a database for a user with a specific ID. If the user does not exist, it is acceptable to handle the error gracefully and return an appropriate response, such as a 404 error. In such cases, crashing the server is unnecessary as the error was expected, and the code was designed to handle it. However, when an unexpected exception occurs, it is crucial to address it appropriately rather than simply logging it and continuing. Unhandled exceptions can lead to unpredictable behavior and compromise the security of the system.

Another aspect of server security is API design. Poorly designed APIs can mislead developers and increase the likelihood of security vulnerabilities. One common issue is when default parameters in an API are insecure, requiring additional options to be passed for secure usage. This design flaw can lead to developers unintentionally using the API in an insecure manner. It is crucial to ensure that default parameters are secure and encourage safe usage.

Polymorphic function signatures are another concern in API design. When a function accepts multiple types of parameters, it can become challenging to determine the intended behavior of the function. This ambiguity can lead to errors and vulnerabilities. APIs should have clear and consistent interfaces, making it easier for developers to understand and use them correctly.

Server security and safe coding practices are vital for ensuring the integrity and security of web applications. By crashing processes when unexpected errors occur, implementing defensive coding practices, avoiding reliance on obscurity, and designing APIs with security in mind, developers can enhance the overall security posture of their applications.

In the context of web application security, server security plays a crucial role in ensuring the overall safety and



integrity of the system. One important aspect of server security is safe coding practices. By following certain guidelines and best practices, developers can minimize the risk of vulnerabilities and potential exploits.

One key principle in safe coding practices is to avoid bundling too much functionality into one function. This is particularly important in loosely typed languages like JavaScript. By keeping functions focused on specific tasks based on the type of parameters passed, it becomes easier for users to understand and use them correctly.

An interesting concept related to safe coding practices is function arity. Function arity refers to the number of arguments a function can accept. In some cases, functions can behave differently based on the number of arguments passed. For example, in jQuery, a widely used JavaScript library, a function called ``$`` exhibits polymorphic behavior. Depending on the type of argument passed, such as a CSS selector, an HTML element, a jQuery object, or a function, the ``$`` function performs different actions. It can return a jQuery object, clone an existing object, parse HTML, or execute a function when the page finishes loading.

However, it is important to note that such polymorphic functions can introduce security risks. For instance, in the case of jQuery, the function's behavior changes based on whether the argument looks like HTML or not. This can lead to potential vulnerabilities if the input is not properly validated.

Another important concept in server security is the use of middleware. Middleware is a generalized version of request handling functions in frameworks like Express.js. Middleware functions are executed for every request, regardless of the method or URL. They provide a way to perform common operations, such as logging, before passing the request to the appropriate handler. Middleware can also be used for tasks like CSRF token validation.

In Express.js, if a fourth argument is passed to a middleware function, it becomes an error handling middleware. This means that if an exception occurs in any of the request handlers, the error handling middleware will be called to handle the error. However, it is essential to ensure that the error object and the ``next`` function are used correctly in the error handling middleware. If these arguments are mistakenly removed, the middleware may no longer function as intended, leading to potential issues in error handling.

Server security and safe coding practices are vital for ensuring the security and reliability of web applications. By following guidelines such as avoiding bundling too much functionality into one function and understanding concepts like function arity and middleware, developers can minimize the risk of vulnerabilities and enhance the overall security of their applications.

The buffer class in Node.js is used to allocate memory for various purposes in a server. Before the introduction of buffer, JavaScript did not have a native way to allocate memory. The buffer class allows you to create buffers containing byte values, strings, numbers, or even copy another buffer. The browser also has similar functionality with typed arrays and array buffers.

To create a buffer with byte values, you can pass an array of byte values to the buffer constructor. This will create a buffer with the specified byte values. Alternatively, you can pass a string to the constructor, and it will parse the string and fill the buffer with the ASCII values of each character. If you pass a number, it will create a buffer with the specified length. Finally, if you pass another buffer, it will create a new buffer and copy the contents of the original buffer.

It is important to note that there are multiple ways to handle binary data in Node.js, which can be confusing. However, this is the way it is due to the evolution of the JavaScript language and the introduction of native support for binary data.

Now, let's move on to a demo that shows how things can go wrong with server security. In the demo, we have a server created using Express and listening on port 4000. By default, if you don't specify an IP address, the server will listen for connections from any device on port 4000. This means that anyone who knows your computer's IP address can connect to your server. This is a major security risk and should be avoided.

Next, we will create an API endpoint called `/api/convert` that takes a string as input and converts it to either hex or base64 encoding. The user can specify the input string and the desired encoding by visiting the URL `/api/convert?data={inputString}&type={encodingType}`. For example, if the user visits `/api/convert?data=hello&type=hex`, the server will convert the string "hello" to hex encoding.



To implement this functionality, we need to extract the data and type parameters from the request query object. We can then parse the data parameter as JSON to convert it into a JSON object. We also need to handle cases where the input string is not valid JSON.

It is important to note that this demo is for illustrative purposes and does not include proper security measures. In a real-world scenario, you would need to implement proper input validation and sanitization to prevent security vulnerabilities.

In this material, we will discuss safe coding practices for server security in the context of web application security fundamentals. We will focus on a specific code snippet and analyze potential security issues associated with it.

The code snippet provided is part of a server-side implementation that converts a given string to a specified type (hex, base64, or utf-8). The code attempts to handle various scenarios and provide appropriate error messages when necessary. However, there are several security vulnerabilities present in this code that we need to address.

One issue is the lack of proper error handling. The code does not utilize try-catch blocks to handle potential exceptions, but instead allows the program to fail and display an error message to the user. This approach can expose sensitive information and should be avoided. It is recommended to implement proper exception handling mechanisms to prevent the leakage of sensitive data.

Another issue lies in the validation of user input. The code checks if the provided type is either hex, base64, or utf-8. However, the validation is incomplete, as it does not handle other possible types. This can lead to unexpected behavior or security vulnerabilities. It is crucial to validate all user input thoroughly and only accept trusted and expected values.

Furthermore, the code uses the "convert" function, which is not implemented. This can lead to potential security risks if the function is not properly implemented. It is essential to ensure that all functions used in the code are well-tested, secure, and free from vulnerabilities.

Additionally, the code utilizes the "new buffer" function to convert the string to the desired type. However, this function is deprecated and should not be used in modern code. It is recommended to use alternative methods or libraries that provide secure and up-to-date functionality.

Finally, the code exhibits a vulnerability known as "memory disclosure." When the provided type is not a string but a number, the code reveals raw server memory, potentially exposing sensitive user data or any data that was previously stored in the process's memory. This can lead to severe consequences and should be addressed immediately.

The code provided has several security issues that need to be addressed. It is crucial to implement proper exception handling, thoroughly validate user input, use secure and up-to-date functions and libraries, and prevent memory disclosure vulnerabilities. By following these safe coding practices, we can enhance server security and protect sensitive data from potential threats.

Server security is a critical aspect of web application security. In this lesson, we will discuss safe coding practices that can help enhance server security.

Safe coding practices involve implementing measures to prevent common vulnerabilities and protect the server from potential attacks. By following these practices, developers can minimize the risk of unauthorized access, data breaches, and other security incidents.

One important practice is input validation. It is crucial to validate and sanitize all user input before processing it on the server. This helps prevent injection attacks, such as SQL injection or cross-site scripting (XSS). By validating user input, developers can ensure that only expected and safe data is accepted by the server.

Another key practice is secure password handling. Passwords should never be stored in plain text. Instead, they should be hashed and salted before being stored in the database. Hashing is a one-way process that converts

the password into a fixed-length string, making it nearly impossible to reverse-engineer the original password. Salting adds an extra layer of security by appending a unique value to each password before hashing.

Secure communication is also essential for server security. HTTPS should be used to encrypt data transmitted between the client and the server. This protects sensitive information, such as login credentials or financial data, from being intercepted or tampered with during transmission.

Regular software updates and patches are crucial for maintaining server security. Developers should keep the server's operating system, web server, and other software components up to date to address any known vulnerabilities. Additionally, strong access controls and permissions should be implemented to restrict unauthorized access to sensitive files and directories.

Safe coding practices play a vital role in server security. By implementing input validation, secure password handling, secure communication, regular updates, and strong access controls, developers can significantly enhance the security of web applications and protect against potential attacks.

**EITC/IS/WASF WEB APPLICATIONS SECURITY FUNDAMENTALS DIDACTIC MATERIALS****LESSON: SERVER SECURITY****TOPIC: LOCAL HTTP SERVER SECURITY**

In the previous material, we discussed API design and mentioned some examples of suboptimal design decisions. These included insecure defaults, confusing function signatures, and behaving differently based on function arity. We used jQuery and Express as examples to illustrate these concepts.

We then moved on to the buffer class in Node.js. The buffer class is used to represent binary data in Node.js. Although JavaScript now has built-in features to manipulate raw data, Node.js still maintains the buffer class due to existing code that relies on it. The buffer API allows us to work with different types of data, including other buffers, arrays of single-byte elements, and strings. When working with strings, each character is interpreted as a byte and stored in the buffer.

We demonstrated the use of the buffer class by implementing a server that converts data into different representations such as hex, base64, and utf-8. The server accepts a JSON string, converts it into a JSON object, and extracts the necessary properties for conversion. The conversion is performed by passing the string to the buffer constructor and calling the `toString` method with the desired type.

During the demonstration, we noticed that if a number is passed instead of a string, uninitialized memory from the Node.js process is returned. This is because the buffer API allocates memory without zeroing it out for performance reasons. To use the buffer class responsibly, it is important to zero out the memory before passing it back to the user. This can be done using the `fill` method with the value zero.

By understanding the buffer class and its API, we can ensure the secure and correct handling of binary data in web applications.

Local HTTP server security is an important aspect of web application security. In this context, it is crucial to understand the potential vulnerabilities that can arise from user-controlled data input.

One common vulnerability is the lack of type enforcement in JavaScript. This means that the type of data being passed to the server is not strictly enforced, allowing for potential misuse. For example, if a user passes a number instead of a string, the server may treat it as a number instead of a string.

To mitigate this vulnerability, one possible solution is to ensure that the data being passed is of the correct type. In the case of a string, we can check if the type of the data is not a string and then convert it into a string. This can be done using conditional statements in the server code.

However, it is important to note that this is just one aspect of server security. There are other potential vulnerabilities that need to be considered. For instance, the use of the "new buffer" constructor can lead to unsafe behavior if a number is passed as an argument. This can result in the creation of a buffer filled with sensitive information.

It is also worth mentioning the use of the ws package, which is a popular NPM package for implementing WebSocket servers. WebSocket servers allow for long-lived socket connections between browsers and servers. While this can provide low latency communication, it is essential to ensure that the data being transmitted is secure.

Local HTTP server security is a critical component of web application security. Ensuring type enforcement and being aware of potential vulnerabilities, such as unsafe buffer creation and WebSocket security, can help protect against potential attacks.

A local HTTP server is an essential component in web application development. It allows developers to test their applications locally before deploying them to a live server. However, it is crucial to ensure that the server is secure to prevent any potential vulnerabilities.

One aspect of local HTTP server security is server-side validation of data received from clients. In a typical scenario, the client sends data to the server, and the server processes the data and sends a response back to

the client. However, if the server fails to validate the data properly, it can lead to security vulnerabilities.

For example, let's consider a scenario where the client sends an object to the server containing a "type" and "data" field. The server processes the data and sends back a response based on the type of data received. In a simple "hello world" example, if the client sends an object with the type "eko" and data "hello world," the server would respond with the string "hello world."

At first glance, this seems safe and straightforward. However, a vulnerability arises when the client can manipulate the "data" field to be a number instead of a string. In the example mentioned earlier, if the client sends the data as 100, the server internally calls a function that creates a buffer using the number. This buffer contains uninitialized memory, which is then sent back to the client. Consequently, the client receives a chunk of memory instead of the expected response.

This vulnerability allows an attacker to exploit the server's memory by sending arbitrary numbers as data. By repeatedly sending such requests, an attacker can dump the server's memory, potentially obtaining sensitive information.

This issue was discovered while working on a node package that implemented the BitTorrent protocol, a file-sharing protocol that involves connecting to random peers. The package had a similar vulnerability where a number instead of a string would result in the server sending back chunks of memory to the client. Although exploiting this vulnerability was challenging, it was still a significant security concern.

Upon discovering this vulnerability, the developers fixed it in their package and reported it to the appropriate authorities. They also realized that this issue might be prevalent in other packages as well. Further investigation revealed similar vulnerabilities in the WS package, a popular WebSocket library, and several other packages.

To address this vulnerability, the developers made a simple fix in their package. They modified a helper function called "ID to buffer" to ensure that the "type" field is always a string before processing it. If the "type" is not a string, the function returns null, preventing any potential memory exposure.

This fix highlights the importance of validating data received from clients before processing it. By ensuring that only expected data types are accepted, developers can mitigate the risk of memory exposure vulnerabilities.

Following the discovery of this vulnerability, other developers started actively searching for similar issues in various packages. One developer even created a static analysis tool, an ESLint plugin, to scan packages for potential vulnerabilities. This initiative led to the identification of several issues, including one in the popular request package.

Local HTTP server security is crucial in web application development. Validating data received from clients is an essential step to prevent vulnerabilities like memory exposure. By implementing proper data validation techniques, developers can ensure the security and integrity of their applications.

A crucial aspect of web application security is ensuring the security of the server. In this context, local HTTP server security plays a significant role. In this material, we will discuss a vulnerability related to sending HTTP requests and explore potential solutions to prevent this vulnerability.

When using the "request" package to send HTTP requests, it is possible to attach a file as an attachment. However, if the attached file is a number, a security issue arises. In this case, instead of attaching the file, a thousand bytes of the user's own memory are sent to the server. This vulnerability highlights a problem in the API design.

To address this issue, a fix was implemented by changing the code to convert numbers into strings before sending the request. This simple fix prevents the vulnerability from occurring. Another package, used to handle binary data, had a similar vulnerability and was fixed in the same manner.

To prevent such vulnerabilities, several ideas can be considered. One approach is to reject numbers when using the buffer. By explicitly disallowing numbers, this type of vulnerability can be avoided. Alternatively, JSON validation can be employed using packages like JSON schema. By defining the expected shape of the JSON data, it is possible to check if the received data matches the specified structure. However, it is important to note that

JSON validation may not detect extra properties in the JSON object, which could lead to potential issues if not handled correctly.

A more robust solution involves defining a class that takes in the JSON object as a constructor parameter. This class can then validate the properties and types, ensuring that only the desired data is exposed. By using this approach, the risk of overlooking updates to the validation schema is minimized.

Lastly, it is crucial to address the design of the buffer itself. The issue arises from the fact that by default, uninitialized memory is returned when a number is passed to the buffer. To mitigate this risk, it is suggested to modify the buffer behavior to zero out the memory, providing a safer default option. However, the argument against this modification is that it would result in a performance decrease of approximately 25%, which is considered unacceptable by the Node project.

The main problem discussed in this material is the exposure of sensitive information when untrusted user input is passed into the buffer as a number. To prevent this vulnerability, it is important to implement appropriate API designs, such as converting numbers to strings before sending requests, employing JSON validation, defining classes to validate and expose data, and considering modifications to the buffer to enhance safety.

In the context of server security, it is important to understand the fundamentals of web application security. One aspect of this is the security of local HTTP servers. In this didactic material, we will discuss the concept of local HTTP server security and some common issues related to it.

One common issue in local HTTP server security is the improper handling of memory initialization. When creating buffers in a local HTTP server, it is crucial to ensure that the memory is properly initialized to prevent potential security vulnerabilities. In the past, there was a design flaw in the API for creating buffers, which resulted in uninitialized memory being allocated under certain conditions.

To address this issue, a new approach was proposed and implemented. The new design separates the functionality of creating buffers into three different APIs: `bufferedUpFrom` for converting any type to a buffer, `Alec` for allocating safe memory (zero-filled), and `AlecUnsafe` for allocating uninitialized memory. This separation allows developers to choose the appropriate API based on their specific needs and the level of performance sensitivity.

This change was not without its challenges. Initially, there was a significant discussion and debate surrounding the proposed design change, with thousands of comments being exchanged. However, in the end, the new design was considered a significant improvement and was implemented. The old buffer designs were deprecated, and a warning message was introduced to discourage their use.

However, the transition to the new buffer design posed additional challenges. Many existing NPM packages relied on the old behavior, making it difficult to switch to the new design without breaking compatibility. To address this, a shim library called "safe buffer" was created. This library allowed developers to use the new APIs even in older versions of Node.js, by simulating the behavior of the new buffer design.

It took time for the changes to propagate throughout the ecosystem, as library authors needed to update their packages to support the new buffer design. Additionally, the decision was made not to display deprecation warnings at runtime, as it would have resulted in an overwhelming number of warnings for every program running on Node.js. Instead, the deprecation warnings were documented, and the transition was allowed to happen gradually.

This story serves as an example of the importance of good API design and the challenges involved in making changes to widely adopted libraries and frameworks. It highlights the need to carefully consider the implications of design decisions, especially in security-sensitive contexts.

Local HTTP server security is a critical aspect of web application security. Proper memory initialization when creating buffers is essential to prevent security vulnerabilities. The introduction of separate APIs for creating buffers and the deprecation of the old designs were significant steps towards improving server security. However, the transition to the new design required careful planning and the creation of a shim library to ensure compatibility with older versions of Node.js.

The function discussed in this material is used to hash passwords before storing them in a database. The function takes two variables: the user's password and the number of times the hash function should be applied (referred to as "hash rounds"). The purpose of applying the hash function multiple times is to slow down attackers who might try to reverse engineer the hash function and gain access to user passwords.

It is important to note that if the "hash rounds" variable is accidentally set as a string instead of a number, the function will behave differently. When "hash rounds" is a number, the function generates a salt, adds it to the password, and then applies the hash function the specified number of times. However, if "hash rounds" is a string, the function assumes that a salt has already been selected and uses it to hash the password only once. This means that if the variable is mistakenly set as a string, all users' passwords in the database will have the same salt, rendering the salt ineffective in providing individual user protection.

The question arises as to why a string would be passed as the "hash rounds" variable. While this may not be directly related to untrusted user input, it could occur when important parts of an application are configurable through a config file or environment variables. Environment variables, in particular, are commonly used to store sensitive information like connection keys to external services. However, it is important to note that environment variables do not have types and are always treated as strings. Therefore, if the "hash rounds" value is obtained from an environment variable, it will be passed as a string, leading to unintended consequences.

To mitigate this issue, it is recommended to separate different functionalities into distinct functions. This separation helps ensure that variables are correctly typed and eliminates the possibility of accidental string inputs causing unexpected behavior.

Additionally, it is advisable to hide detailed error messages and stack traces from users, especially in production mode. Attackers can gain valuable information about the application's structure and dependencies by analyzing stack traces. By suppressing this information in production mode, the risk of exposing sensitive details is minimized. This can be achieved by implementing a conditional statement in the error handler, where the stack trace is displayed only in development mode, while in production mode, it is suppressed.

It is crucial to handle variables with care, ensuring their correct types and preventing unintended consequences. Separating functionalities into separate functions can help minimize the risk of accidental errors. Furthermore, hiding detailed error messages and stack traces from users is a good practice to protect sensitive information about the application.

Local HTTP server security is an important aspect of web application security. One reason why storing sensitive information, such as keys, in files is not recommended is because these files can end up on public platforms like GitHub, leading to potential security breaches. To mitigate this risk, it is advisable to store keys as environment variables. Hosting providers typically provide an interface where developers can input the key, and the hosting provider then makes the environment variable available. This ensures that the key is only accessible in production and not by local developers. By using different keys for different databases, the security of user data is maintained, even if a developer's laptop is stolen.

Another security consideration is the revealing of server information. For example, when using Express, the server attaches a header called "X-Powered-By," which reveals the server being used. While this may not seem like a big deal, it is generally recommended to limit the information disclosed to attackers. The less they know about the server's software, the better. For instance, in the case of nginx, the default configuration includes the version of nginx being used. This can be problematic if the version is vulnerable to exploits. Attackers can search for websites running the vulnerable version and target them. To prevent this, it is advisable to turn off the display of server tokens in nginx using the parameter "server tokens off." Additionally, it is important to ensure that the underlying application, such as PHP or Node.js, does not reveal its identity either.

Fingerprinting is another technique used by attackers to gather information about a server. It involves identifying the operating system running on the server. Even if server information is hidden, attackers can use network scanning tools like nmap to send packets to the server and analyze the responses. By examining subtle differences in whitespace, headers, or other network protocol behaviors, attackers can make educated guesses about the operating system. Unfortunately, there is no foolproof method to prevent fingerprinting.

Securing a local HTTP server involves avoiding the storage of sensitive information in files, using environment



variables instead. It is also important to limit the information disclosed by the server, such as the server software and operating system. By implementing these security measures, the risk of unauthorized access and attacks can be significantly reduced.

Local HTTP server security is an important aspect of web application security. When running a server locally on your computer, there are several security risks that arise. One example of this is the potential for unauthorized access if the server is not properly configured.

One common mistake is forgetting to pass a hostname as a second argument to the server's listen function. This can result in the server accepting connections from all interfaces, meaning anyone on the same Wi-Fi network can connect to your machine if they know the port your server is running on. This can be problematic, especially if you are working on a sensitive application and unintentionally expose it to others before it is ready to be released.

Additionally, if there are any vulnerabilities in the server, an attacker on the same network can exploit them and compromise your machine. This is particularly concerning when developing a local app, as you may not have implemented all the necessary security checks yet. It is surprising that this issue hasn't been more prevalent, considering the potential risks involved.

To mitigate this risk, it is crucial to configure your server to only allow connections from applications on your own computer. By doing so, you ensure that only the browser and other applications running on your machine can connect to the server, while preventing unauthorized access from other machines on the network.

One notable incident involving local server security is the case of the popular video conferencing application, Zoom. Earlier this year, a security researcher discovered a zero-day vulnerability in the software. This vulnerability allowed any website on the internet to remotely turn on a user's webcam without their interaction, simply by visiting a website.

The vulnerability was possible because the Zoom software was running a server on the user's computer. This server accepted requests from any site on the internet, and one particular request could activate the webcam and join the user into a meeting controlled by the attacker. The security researcher gave Zoom 90 days to fix the issue, but when they failed to do so, the researcher published a blog post explaining the vulnerability and providing a link that allowed anyone to join a chatroom with other affected users.

This incident highlighted the importance of properly securing local servers. Even widely used applications like Zoom can have vulnerabilities that can be exploited if server security is not taken seriously.

Local HTTP server security is a critical aspect of web application security. Configuring servers to only allow connections from trusted applications on the same machine is essential to prevent unauthorized access. The incident involving Zoom serves as a reminder of the potential risks associated with local server security and the need for thorough security measures.

A local HTTP server is an essential component in web applications, allowing communication between the client and the server. However, it is crucial to ensure the security of this server to prevent unauthorized access and potential vulnerabilities.

One example of a security issue related to a local HTTP server is the case of Zoom, a popular video conferencing application. In a blog post, a security researcher highlighted a vulnerability that allowed any website to forcibly join a user to a Zoom call with their video camera activated, without the user's permission. Additionally, the vulnerability allowed any web page to install the Zoom client on the user's machine without requiring any user interaction. Even if the user had uninstalled Zoom, a localhost web server would still be running on their computer, making them vulnerable to this issue.

Having an installed app running a web server on a local machine with an undocumented API raises concerns about security. The fact that any website can interact with this web server is a significant red flag for security researchers. This situation creates a potential target for attackers, as the web server accepts HTTP GET requests that trigger code execution outside of the browser's sandbox.

The security of a local HTTP server depends on its design and configuration. While it is possible to restrict

connections to only apps on the same computer, there are still potential risks. For example, web browsers visiting untrusted websites can send GET requests to the local server, bypassing the same origin policy. This means that any site on the internet can send requests to the server, potentially triggering code execution or other actions.

To demonstrate the potential impact of such vulnerabilities, an example was provided. A server was created to listen for connections and execute a command when receiving a GET request to the home page. Although this example is relatively safe as it does not take user input, it highlights the need for secure server design to prevent unauthorized access and potential exploits.

Ensuring the security of a local HTTP server is crucial to prevent unauthorized access and potential vulnerabilities. The example of the Zoom vulnerability serves as a reminder of the importance of secure server design and configuration. By implementing proper security measures, such as restricting access and validating user input, the risk of unauthorized access and potential exploits can be minimized.

A local HTTP server is a server that runs on your computer and listens for incoming requests. In this context, we will discuss the security aspects of local HTTP server and how it can be vulnerable to attacks.

When you start a local HTTP server, it opens up a port on your computer that can be accessed by other applications on your computer or even from the internet. For example, if you visit "localhost:4000" in your web browser, it will send a GET request to the local server and open up a dictionary on your computer.

Now, you might wonder who can send these requests to the server on your computer. The answer is that any application or website that knows the address and port of your local server can send requests to it. For example, if you visit "example.com" and open up the console, you can run JavaScript code that sends a request to your local server using the "fetch" function.

However, there is a security mechanism called Cross-Origin Resource Sharing (CORS) that restricts the access to resources on different origins. When the request from "example.com" is sent to your local server, the browser blocks the response from being read due to CORS policy. But, the request is still sent to the server and the server responds with a success message.

To allow access to the response, you can add an additional header called "Access-Control-Allow-Origin" with the value of "\*" (star). This header tells the browser that any website can read the response from the request. After adding this header, the request from "example.com" will be able to read the response from your local server.

It's important to note that even if you can't read the response, you can still send requests to the server and trigger its behavior. For example, embedding an image with the source pointing to your local server will cause the request to be sent and the server to respond accordingly.

To demonstrate the vulnerability of local servers, you can use the "curl" command to send a request to your local server. This shows that any application on your computer that is allowed to use the network can send HTTP requests to your local server and cause it to behave in a certain way.

Local HTTP servers can be vulnerable to attacks if not properly secured. The CORS mechanism helps in restricting access to resources, but it's important to be aware that sending requests to a local server can still trigger its behavior even if the response cannot be read.

To check for vulnerable servers running on your computer, you can use the command "lsof -i -P | grep LISTEN" which lists open file descriptors that are listening for connections. This will show you any web servers running on your machine that are accessible to anyone.

Local HTTP server security is a critical aspect of ensuring overall server security. In this context, it is important to understand the potential risks associated with running local HTTP servers and the measures that need to be taken to mitigate these risks.

Local HTTP servers are used to handle various functionalities, such as voice data forwarding, phone call support, and running specific processes. These servers are often designed to be accessible within the local network, allowing any device on the network to connect to them.

However, this accessibility also poses security concerns. If not properly secured, anyone on the local network can connect to these servers and potentially exploit vulnerabilities. Therefore, it is crucial to implement robust security measures to protect these servers from unauthorized access and malicious activities.

One example of a vulnerable local HTTP server is the antivirus software developed by Trend Micro. In this case, the server installed by the antivirus software was susceptible to remote code execution (RCE) from any website on the internet. This vulnerability allowed attackers to send GET requests to the server, which could execute arbitrary code on the user's computer. This issue was discovered by security researchers at Google's Project Zero, who promptly notified the company.

To demonstrate the severity of the vulnerability, the researchers provided a simple piece of code that could open the calculator application on the victim's computer. This code could have been embedded in an advertisement or on any website visited by users running the vulnerable antivirus software. As a result, the user's computer would have been compromised.

This example highlights the inherent dangers associated with local HTTP servers. When developing such servers, it is essential to implement strict security measures to prevent unauthorized access and ensure the validation of incoming messages. Failure to do so can lead to severe consequences, compromising the security of users' systems.

To gain further insights into security vulnerabilities and the response of companies to such issues, Google's Project Zero maintains an issue tracker where researchers can report vulnerabilities and interact with company representatives. This platform provides valuable information about how seriously companies take security and their internal processes for resolving these issues.

Local HTTP server security is a crucial aspect of overall server security. It is essential to implement robust security measures to protect these servers from unauthorized access and potential vulnerabilities. The example of Trend Micro's vulnerable antivirus software demonstrates the potential risks associated with improperly secured local HTTP servers. By learning from such examples and understanding the importance of security, developers can ensure the integrity and safety of their server systems.

The policy regarding the default camera settings for participants joining a conference room was a key aspect of the security vulnerability in the local HTTP server of Zoom. By default, participants' cameras would be turned on without any permission prompt when they joined the room. This allowed attackers to exploit the vulnerability by tricking users into joining a room where their camera would be automatically turned on.

Uninstalling Zoom did not remove the local server from the user's computer. This was done intentionally by Zoom to ensure that if a user uninstalled Zoom and later received a Zoom link, clicking on the link would trigger the local web server to quickly reinstall Zoom. This meant that even if the user uninstalled Zoom, clicking on a Zoom link would open it as if it was never uninstalled.

Furthermore, the local HTTP server was vulnerable to a denial-of-service attack, where a malicious site could render the user's computer completely unusable. This was achieved by repeatedly sending GET requests to a non-existent conference room number, causing the computer to freeze. As a result, the user would lose control over their browser and other applications, making their computer effectively unusable.

Initially, Zoom defended their decision and claimed that the vulnerability was not a serious issue. However, they later acknowledged their mistake and decided to remove the local web server after the researcher who discovered the vulnerability made it public. It seems that Zoom did not fully understand the severity of the vulnerability until they received feedback from their customers and others affected by it.

To address the issue, Zoom released an updated version of their application that automatically uninstalled the local web server when installed. They also added a user interface prompt to confirm joining a meeting before actually joining. However, there were still some limitations to this solution. Users who did not open the app for a while would not receive the update, leaving them vulnerable until they eventually opened the app and allowed the installation. Additionally, users who had uninstalled Zoom would have no way of knowing about the vulnerability unless they happened to come across the news.

The situation became even more concerning when another security team discovered a remote code execution (RCE) vulnerability in a similar part of the Zoom codebase. Although the initial researcher did not find the RCE, the combination of the camera vulnerability and the newly discovered RCE could potentially allow attackers to execute code on users' computers remotely. This posed a significant risk to the millions of people who were using Zoom.

The local HTTP server security vulnerability in Zoom's camera settings and the subsequent discovery of a remote code execution vulnerability raised serious concerns about the security of the application. The initial vulnerability allowed attackers to turn on users' cameras without their consent, while the RCE vulnerability could potentially enable attackers to run arbitrary code on users' computers. Zoom took steps to address these issues by removing the local web server and releasing an updated version of the application. However, there were still challenges in ensuring all users received the necessary updates and informing those who had uninstalled Zoom about the vulnerabilities.

Apple has a malware removal tool built into all Macs, which periodically pings an Apple server to fetch a list of bad executable files. If a program on a user's computer matches the fingerprint of a file on the list, the tool will kill it and prevent it from running again. This functionality does not require an operating system update or restart. Apple recently used this tool to uninstall a program called zum-zum server from everyone's Macs, in response to the discovery of a critical vulnerability. This action was taken to prevent the spread of the malware.

When joining a conference on Zoom, the flow of communication between the browser and the local server on the user's computer is as follows:

1. The user visits a join URL on the Zoom website.
2. The Zoom page sends a request to the local server, asking it to launch the conference room.
3. The local server launches the Zoom app and sends a response to the browser.
4. The browser receives the response but cannot read its contents due to the lack of a Cross-Origin Resource Sharing (CORS) header.
5. The Zoom app is successfully launched, fulfilling the user's intent.

However, there was an issue with the local server indicating to the Zoom page whether the app was successfully launched or not. Due to the lack of a CORS header, the page could not read the response. As a workaround, the response was implemented as an image with varying sizes. The page would analyze the image's impact on the layout to determine if the app launched successfully or if it needed to be downloaded. This workaround could have been avoided if the developers were aware of the CORS mechanism, which allows the local server to specify that the Zoom page is allowed to read the response.

To address this issue properly, the developers could have used CORS. They would include a header in the response indicating that the Zoom page is allowed to read it. This would eliminate the need for the image-based workaround. However, it is important to note that this solution does not address the main issue of any site being able to send requests to the local server. It only removes the previous workaround and relies on the browser to enforce the CORS policy by comparing the header with the current page's origin.

Apple's malware removal tool silently uninstalled zum-zum server from all Macs to prevent the spread of a critical vulnerability. The flow of communication between the browser and the local server when joining a Zoom conference involves requests and responses. The lack of a CORS header led to a workaround using image sizes to indicate the success of launching the Zoom app. The issue could have been resolved by implementing CORS properly, but this would not address the problem of any site being able to send requests to the local server.

In the field of web application security, it is important to consider server security, specifically local HTTP server security. In this context, we will discuss the flow of an attacker and explore potential solutions to secure the local HTTP server.

When an attacker wants to compromise a server, they typically send back HTML code to the victim's browser. The attacker's page loads, and they send a request to the local server. In response, the local server launches the Zoom app and sends an opaque response. It is important to note that even if the response contains an access control origin header specifying "zoom.us," the browser will still execute the attacker's code. Therefore, using Cross-Origin Resource Sharing (CORS) does not solve the problem.

To address this issue, the best solution is for websites to register a protocol handler. For example, Zoom can

register the URL "zoom://". When a user clicks on a link starting with this URL, the browser will prompt them to open the Zoom app on their computer. This method ensures a secure way to open an app without compromising server security.

However, in cases where a local HTTP server must be used, it becomes crucial to secure it. One approach is to require user interaction before joining a call. For example, when a user clicks on a "zoom.us" link, a button can be displayed asking if they want to join. Additionally, the local server should only allow communication with "zoom.us" to prevent random ads or sites from accessing it.

To implement this, the local server can inspect the origin header. The origin header indicates the page or origin that triggered the request. By checking if the origin is "zoom.us," the local server can launch the Zoom app and send a response back to the page. This approach ensures that only legitimate requests from "zoom.us" are processed.

It is important to note that the origin header cannot be controlled by JavaScript code running on a site. The browser sets the origin header to the true origin that initiated the request, preventing any manipulation. However, if an app running on the user's computer wants to communicate with the Zoom server, it can do so since it is already installed on the computer.

In terms of security, if an attacker tries to join a call using this approach, their request will have an origin header set to the attacker's origin. When the server checks this header, it will not match and will simply close the connection without sending a response. This ensures that the attacker's page cannot gain any information about the local server.

Although this approach provides some level of security, there are some limitations. The browser does not always add the origin header, particularly in simple requests triggered by image or iframe tags. This can be frustrating but should be taken into consideration when implementing local server security measures.

Securing a local HTTP server involves implementing measures such as requiring user interaction and inspecting the origin header. While registering a protocol handler is the recommended approach, in cases where a local server is necessary, these security measures can help protect against unauthorized access.

A web application can be vulnerable to various security threats, including attacks on the server. In this context, it is important to understand the concepts of local HTTP server security and the different types of HTTP requests.

HTTP requests can be classified into two categories: simple requests and preflighted requests. Simple requests include GET, HEAD, and POST requests that do not have any custom HTTP headers set. These requests are considered safe and can be made without JavaScript. For example, a site can embed an image from another site or submit a form without the need for JavaScript.

On the other hand, preflighted requests are more complex and require permission from the server before they can be sent. These requests include HTTP methods such as DELETE, PUT, and PATCH. Preflighted requests are necessary for security reasons, as they involve potentially destructive actions that can't be allowed without prior approval. For example, if a server receives a DELETE request, it will delete the specified item in the URL. To prevent unauthorized requests, the browser needs to ask for permission before sending such requests.

To ensure server security, it is important to consider the origin of the request. Simple requests do not include the origin header, which means that any site can send these requests without the server distinguishing the origin. This can lead to potential security issues. One solution is to block requests that do not have an origin header and force the inclusion of the origin header by making the request a complex request. Another option is to change the endpoint to require a preflighted request, which guarantees that the origin header will always be sent.

It is crucial to understand the types of requests that can be made without JavaScript. Generally, sites can embed images or submit forms without JavaScript. These actions fall under the category of simple requests and are considered safe. The browser does not attach an origin header or perform any additional checks for these requests.



In contrast, preflighted requests require permission from the server before they can be sent. This is necessary for security reasons, as certain requests can be destructive. The browser needs to prevent these requests from being sent until it has received confirmation from the server that they are allowed.

Server security in the context of local HTTP servers involves understanding the distinction between simple requests and preflighted requests. Simple requests, such as GET, HEAD, and POST, do not require JavaScript and are considered safe. Preflighted requests, on the other hand, require permission from the server before they can be sent and are necessary for security reasons.

In the context of web application security, server security plays a crucial role in protecting sensitive data and preventing unauthorized access. One common vulnerability is the lack of proper handling of HTTP methods, particularly when dealing with non-simple requests such as PUT or DELETE. In this didactic material, we will explore the concept of preflighted requests or options requests, which can be used to enhance server security and mitigate potential attacks.

When a server receives a non-simple request, such as a PUT request, the browser first checks if the server supports preflighted requests. This is done by sending an OPTIONS request to the server, asking for permission to send the actual request. The OPTIONS request includes information about the origin of the request and the desired HTTP method. If the server does not support preflighted requests or does not respond to the OPTIONS request, the browser interprets it as a denial of the request.

To illustrate how preflighted requests can be used to protect a local server, let's consider an example where a user joins a Zoom call. The Zoom HTML page sends a request to the local server, asking it to open the app and join the meeting. In this case, the HTTP method used is PUT, which is a non-simple request. The browser, upon encountering this request, realizes that it may need to ask for permission before sending it. Since the request is coming from the Zoom origin and the server is localhost (a different origin), the browser identifies it as a cross-origin request and initiates the preflight process.

The browser automatically sends an OPTIONS request to the server, indicating the origin (Zoom) and the desired HTTP method (PUT). If the server recognizes the OPTIONS request and approves the PUT request from the specified origin, it responds with an empty response, specifically mentioning that PUT is allowed. This response serves as confirmation to the browser that the server permits the PUT request. Subsequently, the browser proceeds to send the actual PUT request to the server. The server, having received the preflighted OPTIONS request earlier, confidently processes the PUT request, knowing that it has been validated by the browser.

By utilizing preflighted requests, the local server effectively protects against unauthorized PUT requests. The browser acts as a gatekeeper, ensuring that non-simple requests are only sent after obtaining permission from the server. This mechanism significantly reduces the risk of malicious actors exploiting server vulnerabilities.

Preflighted requests or options requests are an essential component of server security in the context of web applications. By validating non-simple requests before they are sent, servers can effectively protect against unauthorized access and potential attacks. This technique enhances the overall security posture of web applications, safeguarding sensitive data and ensuring the integrity of server operations.

In the field of cybersecurity, it is crucial to understand the fundamentals of web application security, particularly server security. In this material, we will focus on the topic of local HTTP server security.

When it comes to server security, one important aspect to consider is the protection against unauthorized requests from different origins. To achieve this, a server must perform a check before responding to a request. This check involves verifying if the request is coming from the same origin or if it has the necessary permissions. If the request is not from the same origin and lacks the required permissions, the server will respond with an "access forbidden" message, denying the request.

To illustrate this process, let's consider an example. Suppose a request is made to a local server. Before responding, the server needs to determine if the request is originating from a trusted source, such as Zoom. To do this, the server initiates an options request, commonly known as a pre-flight request, to confirm the origin. If the request is indeed coming from Zoom, the server will allow it to proceed. However, if the request is coming from an unauthorized source, such as an attacker's domain, the server will respond with an "access forbidden" message, indicating that the request is denied.



It is important to note that once the server denies a request, the browser will not attempt to resend it. This is because the browser recognizes that it did not receive permission from the server. Therefore, it is crucial to establish proper authorization and permissions to ensure the security of the server.

While this mechanism provides protection for the local server against unauthorized requests from websites, it is worth mentioning that it does not prevent native applications running on the computer from accessing the server. Native applications, such as curl or games running in the terminal, can set their own headers and bypass the browser's enforcement of the origin header. This means that they can make requests to the local server and pretend to be originating from a trusted source like Zoom. Therefore, it is essential to be aware of this vulnerability and implement additional security measures when dealing with native applications.

Although the concept of pre-flight requests may seem complex, it is a valuable tool for protecting websites in the real world. However, when it comes to local servers, there is another vulnerability to consider. This vulnerability is known as DNS rebinding, which allows any website on the internet to pretend to be the same origin as the local server. In other words, an attacker's domain can masquerade as Zoom and communicate with the local server. It is important to understand the implications of DNS rebinding and how to mitigate this specific type of attack.

Local HTTP server security is a multifaceted topic that requires a comprehensive understanding of various vulnerabilities and protective measures. While pre-flight requests provide a level of protection against unauthorized requests from websites, it is essential to remain vigilant and consider additional security measures, especially when dealing with native applications. In the next session, we will delve into the topic of DNS rebinding and explore ways to mitigate this particular vulnerability.

**EITC/IS/WASF WEB APPLICATIONS SECURITY FUNDAMENTALS DIDACTIC MATERIALS****LESSON: DNS ATTACKS****TOPIC: DNS REBINDING ATTACKS**

DNS rebinding is a type of attack that was discovered by Dan Kaminsky and presented at a security conference. Despite its initial underappreciation, DNS rebinding remains a pervasive issue even today. It is a subtle and tricky attack to describe, which is why this lecture is dedicated to explaining it in detail.

DNS rebinding attacks exploit vulnerabilities in devices connected to the internet, such as IoT devices, local servers, printers, IP cameras, phones, Smart TVs, switches, routers, and access points. These attacks take advantage of the fact that these devices often have weak security measures or outdated software.

The attack works by tricking a victim's web browser into making requests to a malicious website. The attacker first sets up a DNS server that resolves the victim's domain name to the attacker's IP address. When the victim visits a legitimate website, the attacker's server responds with a short time-to-live (TTL) value, causing the victim's browser to make subsequent requests to the attacker's IP address. This allows the attacker to execute arbitrary code on the victim's device, potentially gaining control over it.

The consequences of a successful DNS rebinding attack can be severe. Attackers can intercept and manipulate network traffic, steal sensitive information, or even take control of the victim's device. This is particularly concerning when it comes to routers, as compromising them allows attackers to modify network settings and potentially gain access to all connected devices.

To mitigate the risk of DNS rebinding attacks, it is crucial to keep devices and software up to date, including routers. Additionally, network administrators should implement strong security measures, such as firewall rules and DNS response validation. Users should also be cautious when visiting unfamiliar websites and avoid clicking on suspicious links.

DNS rebinding is a persistent and potentially damaging attack that targets various internet-connected devices. Understanding its mechanics and implementing appropriate security measures are essential to protect against this threat.

**DNS Rebinding Attacks**

DNS rebinding attacks are a type of cyber attack that can exploit vulnerabilities in web applications, specifically those that utilize the Domain Name System (DNS). In these attacks, an attacker takes advantage of the trust established between a user's browser and a web application to gain unauthorized access or control over devices connected to the user's local network.

The first step in a DNS rebinding attack is for the attacker to lure the user to a malicious website or ad. This can be achieved through various means, such as buying ad space or tricking the user into visiting a compromised website. Once the user visits the attacker's website, the attacker's JavaScript code running in the browser initiates a request to a device on the user's local network, such as a thermostat, using the device's IP address.

The crucial aspect of DNS rebinding attacks is that these requests are unauthenticated, meaning they bypass any security measures that would typically require authentication before allowing access to the device. The attacker's request to the device is a simple GET request, which can be used to manipulate the device's settings or retrieve sensitive information.

For example, in the case of a vulnerable thermostat, the attacker could send a GET request to the thermostat and change the temperature setting to an extreme value, such as 95 degrees. This could create discomfort or even pose a danger, especially in certain climates or for vulnerable occupants, such as the elderly or disabled.

Furthermore, a successful DNS rebinding attack could also result in financial consequences. If a user is targeted while on vacation, for instance, and the attacker changes the thermostat settings to extreme levels, the user's energy bill could skyrocket by the time they return.

Additionally, in some cases, a vulnerable device may not only allow changes to its settings but also execute the

code included in the attacker's request. This could potentially lead to more severe consequences, depending on the device's capabilities and the nature of the code executed.

To better understand how DNS rebinding attacks work, let's review a simplified example using the Zoom video conferencing software. When a user installs Zoom on their computer, a local server is set up to listen for connections on a specific port. When the user visits the Zoom server's webpage, the server responds by launching the Zoom app and returning a response to the browser.

In a legitimate scenario, this communication between the server and the browser is secure. However, in a DNS rebinding attack, an attacker can follow a similar set of steps. By getting the user to visit their webpage, the attacker's JavaScript code initiates a request to the local Zoom server. The server launches the Zoom app as before, but this time, the response is blocked from being read by the browser due to the presence of specific headers.

The attacker's objective in this case is to join the user into a Zoom call without their consent. By exploiting the DNS rebinding vulnerability, the attacker can bypass the browser's security measures and execute the necessary steps to achieve their goal.

To protect against DNS rebinding attacks, web application developers need to implement proper security measures. These may include validating and authenticating requests, implementing secure communication protocols, and ensuring that sensitive operations require explicit user consent.

By understanding the fundamentals of DNS attacks, such as DNS rebinding attacks, users and developers alike can take the necessary precautions to mitigate the risks and protect against potential vulnerabilities.

In the field of web application security, one of the important aspects to consider is DNS attacks, specifically DNS rebinding attacks. DNS rebinding attacks can be a serious threat to the security of web applications, as they can allow attackers to bypass the same-origin policy and access sensitive information.

To understand DNS rebinding attacks, let's first discuss the concept of protocol handlers. Protocol handlers are a way for native applications to register and handle custom protocols. For example, a native app can register a protocol called "zoom://" and specify that whenever a URL with this protocol is encountered, the app should be launched.

One possible solution to prevent DNS rebinding attacks is for web applications like Zoom to not have a local HTTP server at all. Instead, they can rely on protocol handlers to handle the URLs. This way, when a user clicks on a link with the "zoom://" protocol, the browser will redirect the request to the Zoom app, which is capable of handling such URLs.

However, if for some reason the application still needs to have a local HTTP server, there are ways to mitigate the risks associated with it. One approach is to require user interaction before joining a user into a call. By displaying a dialog box asking the user to confirm their intention to join the call, the application can ensure that the user is aware of the action they are taking.

Another approach is to ensure that only the legitimate origin, in this case, Zoom, is able to communicate with the local server. This can be achieved by using techniques such as simple HTTP requests and preflight HTTP requests. Simple requests are requests that can be created solely with HTML, without the need for JavaScript or custom headers. On the other hand, preflighted requests involve additional headers or JavaScript code.

By carefully implementing these security measures, web applications can protect themselves against DNS rebinding attacks and enhance the overall security of their systems.

DNS rebinding attacks pose a significant threat to web application security. By understanding the concept of protocol handlers and implementing appropriate security measures, such as avoiding local HTTP servers or requiring user interaction, web applications can mitigate the risks associated with DNS rebinding attacks and ensure a safer user experience.

A preflighted request is a type of request that is performed when attempting to send an HTTP method that is not listed as a simple request. In such cases, the browser needs to check if it is safe to proceed with the request

because it doesn't know in advance if the server expects such requests. To perform this check, the browser sends an options request to the server, which is essentially a request to ask for permission to send another request later.

In the options request, the browser includes all the relevant information that distinguishes this request from a simple request. For example, if the request is a put method instead of a post or get method, the browser will include this information in the request. The browser also includes the origin that is making the request. Depending on how the server responds to this options request, the request will either be allowed or denied.

It is important to consider how older servers on the internet, which were created before this standard was established, will respond to options requests. If these servers do not support options, they will respond with a different kind of response that may not include the necessary headers to inform the browser that the request is allowed. In such cases, the browser will interpret any response that does not explicitly allow the request as a denial. Therefore, the default assumption is to deny the request.

The concept of preflighted requests may seem arbitrary, but it is a pragmatic decision made to address potential security concerns. In the early days of the web, it was assumed that web pages could send certain types of requests. However, as new types of requests became possible through JavaScript, servers that had made assumptions about what types of requests could be generated from a page might be surprised by the new requests. To prevent servers from being compromised, the decision was made to introduce the preflighted request method. This method requires asking for permission before sending requests that the server may not be expecting.

To illustrate the process, let's consider an example where a get request caused an application to launch. To solve this problem, we can change the method to put and specify that the local server only expects put requests. If a get request is received, the server will not take any action. When the page is loaded, the browser will call fetch with the put method. At this point, the browser recognizes that this is not a simple request but a preflighted request. The browser then checks with the server by sending an options request. The options request includes the desired method (put) and informs the server that the request is being made from JavaScript code running on Zoom.us. The server, in this case, would consider the origin and make a decision based on that. If the origin is Zoom, the request will be allowed.

Preflighted requests are performed when attempting to send HTTP methods that are not considered simple requests. The browser checks with the server by sending an options request to ensure that it is safe to proceed with the desired method. This mechanism was introduced to address potential security concerns and prevent servers from being compromised by unexpected requests.

DNS rebinding attacks are a type of cyber attack that can bypass the security measures put in place to protect web applications. In this attack, the attacker tricks the browser into believing that a request to a local server is actually a request from the attacker's own domain. This allows the attacker to send requests to the server without going through the necessary security checks.

To understand how DNS rebinding attacks work, let's first review how the browser handles requests to a server. When a request is made from a web application to a server, the browser first sends an "options" request to the server to check if the request is allowed. This is known as a pre-flight request. The server responds with headers that indicate whether the request is allowed or not. If the request is allowed, the browser proceeds with sending the actual request.

To prevent unauthorized requests, the browser enforces a same-origin policy. This means that requests can only be made from the same domain as the server. In the case of a DNS rebinding attack, the attacker manipulates the DNS (Domain Name System) resolution process to make the browser believe that the request is coming from the attacker's domain. This tricks the browser into thinking that the request is a same-origin request, bypassing the security checks.

To mitigate DNS rebinding attacks, web developers should implement server-side defenses. One approach is to include a header in the server's response that instructs the browser to allow requests from specific domains. This can be done by setting the "Access-Control-Allow-Origin" header to the desired domain. By explicitly specifying the allowed domains, the server can prevent requests from unauthorized sources.

It's important to note that DNS rebinding attacks can still occur even if the local server is secure. This is because the attack relies on tricking the browser, not compromising the server itself. Additionally, it's worth mentioning that while browsers enforce the same-origin policy, other types of applications running on the same device, such as native apps, may not follow the same rules. Therefore, it's crucial to be aware of the limitations of browser-based security measures.

DNS rebinding attacks are a technique used by attackers to bypass security measures and send unauthorized requests to web applications. By manipulating the DNS resolution process, attackers can trick the browser into thinking that the request is coming from a trusted source. To protect against these attacks, web developers should implement server-side defenses and be aware of the limitations of browser-based security measures.

### DNS Rebinding Attacks

DNS rebinding attacks are a type of attack that tricks the browser into thinking that a request is not a cross-origin request. This allows the attacker to bypass cross-origin rules and also bypass the victim's firewall. The attacker can use the victim's browser as a proxy to communicate directly with vulnerable servers on the same network. This means that not only local servers on the victim's computer can be attacked, but also vulnerable IoT devices on the same network.

It's important to note that DNS rebinding attacks do not violate the same origin policy. The same origin policy states that one origin can only talk to another origin, and origins are defined by the protocol, hostname, and port. In this attack, the same origin policy is followed perfectly. The issue lies in the fact that the same origin policy does not consider the IP address of the server in question. This is the crux of the DNS rebinding attack.

To understand how this attack works, let's consider a demo. In the demo, there are two servers. The server on the left is a vulnerable server that can be exploited by sending it a request. The server is running on port 8080. Any site can send a GET request to this server, causing the dictionary on the server to open. This was demonstrated by pointing an image tag to port 8080.

The attacker's website represents the second server in the demo. The attacker can write code that triggers the vulnerable server by making a request to localhost:8080. This can be done by adding an image to the attacker's website that sends the request. When the attacker's website is visited, the request is made to the local server, causing the dictionary to open.

It's worth mentioning that during the demo, there was a confusion with the ports, but it was eventually resolved. The attacker's website was actually making a request to the local server, as intended.

DNS rebinding attacks can have serious implications as they allow attackers to bypass security measures and access vulnerable servers or devices on the victim's network. Understanding this attack is crucial for implementing effective security measures to protect against it.

### Web Applications Security Fundamentals - DNS Attacks - DNS Rebinding Attacks

In this material, we will discuss DNS attacks, specifically focusing on DNS rebinding attacks. DNS (Domain Name System) is a crucial component of the internet infrastructure that translates domain names into IP addresses. It plays a vital role in ensuring that users can access websites by typing in easy-to-remember domain names instead of complex IP addresses.

DNS rebinding attacks exploit the inherent trust between a web browser and the DNS system. These attacks take advantage of the fact that web browsers allow JavaScript code to make requests to remote servers. By manipulating the DNS resolution process, attackers can trick browsers into sending requests to malicious servers under their control.

One common scenario involves an attacker setting up a malicious website and enticing a victim to visit it. The attacker's website contains JavaScript code that initiates a DNS rebinding attack. The attack begins by resolving the attacker's domain name to a legitimate IP address, allowing the victim's browser to establish a connection. However, after a certain period of time, the attacker changes the DNS record for their domain name to a different IP address, one controlled by the attacker.

When the victim's browser attempts to make subsequent requests to the attacker's domain, it unknowingly sends them to the new IP address controlled by the attacker. This allows the attacker to execute arbitrary code on the victim's browser, potentially leading to various security vulnerabilities and attacks such as cross-site scripting (XSS) attacks.

To demonstrate the impact of DNS rebinding attacks, we will use a code example. The code snippet uses the fetch API to make a request to a local server at localhost:8080. The server's response is then displayed on the webpage, potentially introducing a cross-site scripting vulnerability. Additionally, error messages are also displayed on the webpage for debugging purposes.

To mitigate the risk of DNS rebinding attacks, web developers can implement several measures. One approach is to enforce the same-origin policy, which restricts web browsers from making requests to different domains. By disallowing requests to domains other than the one hosting the webpage, the risk of DNS rebinding attacks can be significantly reduced.

Another defense mechanism is to use a different HTTP method, such as PUT, instead of GET. By changing the method of the request, attackers would need to send a PUT request explicitly, making it more difficult to exploit DNS rebinding vulnerabilities.

It is important for web developers and administrators to stay updated on the latest security best practices and regularly patch any vulnerabilities in their web applications. By implementing proper security measures, including protecting against DNS rebinding attacks, organizations can significantly enhance the security of their web applications and protect their users' sensitive information.

#### DNS Rebinding Attacks in Web Applications Security

DNS (Domain Name System) rebinding attacks are a type of cyber attack that target web applications. In these attacks, attackers exploit vulnerabilities in a web application's DNS resolution process to bypass security measures and gain unauthorized access to sensitive information.

During a DNS rebinding attack, the attacker tricks the victim's browser into making a request to a malicious website, which then returns a response that appears to come from a trusted source. This allows the attacker to bypass the same-origin policy, which normally prevents requests from one domain to another.

In the transcript, the speaker discusses an example of a DNS rebinding attack on a local server. The attacker, using the domain "attacker.com", tries to access a dictionary hosted on "localhost" by sending a PUT request. However, due to the server's configuration, the request is blocked.

The speaker explains that the browser sends an OPTIONS request before the PUT request as a preflighted request. The server, by default, allows all methods, including PUT. However, the server's response does not explicitly allow the origin "attacker.com", leading to the browser blocking the request.

To defend against this attack, the speaker suggests requiring the request to be a PUT in order to be treated as a valid request. By doing so, the server can prevent unauthorized access to sensitive resources.

The speaker also mentions an error message related to CORS (Cross-Origin Resource Sharing) policy. The error message suggests setting the request mode to "no-cors", but the speaker clarifies that this does not help in this situation because a PUT request is not considered a simple request.

To make the attack work, the speaker suggests adding a new handler to handle OPTIONS requests and explicitly allowing the origin and the PUT method in the response headers. This would enable the browser to send the PUT request to the server and access the desired resource.

DNS rebinding attacks exploit vulnerabilities in web applications' DNS resolution process to bypass security measures and gain unauthorized access. By understanding how these attacks work and implementing proper security measures, web application developers can protect their applications from such threats.

Firewalls play a crucial role in network security by protecting our devices from unauthorized access and potential attacks. They act as a barrier between our internal network and the external internet, filtering



incoming and outgoing network traffic based on predefined security rules.

When a browser initiates a connection to the internet, the firewall recognizes it as a request from a browser on the network and allows it to proceed. However, incoming connections from random computers on the internet are not allowed by default. This is because unsolicited communication can pose a security risk, as there is no indication that the user is expecting such a connection.

This is particularly important because many devices on our network, such as IoT devices, are often insecure. Devices like refrigerators, printers, webcams, and smart TVs are now equipped with computers and can be vulnerable to attacks. These devices often assume that other devices on the same network are trusted and may allow incoming connections from them. However, allowing connections from random computers on the internet would be a significant security risk.

The firewall's role is to block such unsolicited incoming connections, preventing potential attacks on insecure devices. It ensures that only authorized devices on the local network can establish connections with vulnerable devices, providing an additional layer of protection.

Furthermore, the fact that these devices rarely receive updates makes them even more susceptible to attacks. IoT devices are often manufactured with fixed firmware and are not regularly updated by users. This means that vulnerabilities discovered after the device's production may remain unpatched, leaving them exposed to exploitation.

Firewalls are essential for network security. They regulate incoming and outgoing network traffic, allowing authorized connections while blocking unsolicited ones. This protection is particularly important for devices on our network that are often insecure and rarely receive updates, such as IoT devices.

In the realm of cybersecurity, one of the fundamental aspects of web application security is protecting against DNS attacks. Specifically, DNS rebinding attacks pose a significant threat to the security of internet-connected devices.

DNS rebinding attacks exploit vulnerabilities in devices on a network, such as printers or IoT devices, by leveraging the browser as a proxy. When a user visits a website, the browser recognizes the request and allows the website to load in the browser. However, the attacker's website, running JavaScript code, can then make a request to a local IP address, such as the IP address of a printer.

This is concerning because many IoT devices assume that requests can only be sent by devices on the same network. Consequently, a random JavaScript code running in an ad can send requests to insecure IoT devices. If the attacker can send a request that infects the printer, for example, the device becomes compromised and can potentially perform malicious actions, such as capturing video frames or launching attacks.

To mitigate the risk of DNS rebinding attacks, several protective measures can be implemented. The same origin policy applies, limiting attackers to sending simple requests, such as GET and POST, to the printer. This restriction prevents more complex requests that could further compromise the device. However, if the attacker can successfully infect the printer, the network is once again vulnerable to potential attacks.

It is important to note that DNS rebinding attacks are not limited to IoT devices but also extend to local servers running on a computer. Similar to the printer scenario, if an attacker's code on a website identifies the user's local IP address or uses localhost, it can make requests to local servers on the user's computer. If these servers have vulnerabilities, the attacker can exploit them, leading to potential security breaches.

DNS rebinding attacks pose a significant threat to the security of web applications and internet-connected devices. By leveraging the browser as a proxy, attackers can exploit vulnerabilities in devices and local servers, compromising network security. Implementing measures such as the same origin policy can help mitigate the risk, but it is crucial to remain vigilant and ensure the security of all connected devices and servers.

DNS rebinding attacks are a type of cyber attack that exploit the way web browsers handle DNS resolution. In this attack, an attacker tricks a victim's browser into making a request to a malicious website, which then sends a different response once the victim's browser has established a connection.

To understand how DNS rebinding attacks work, let's take a look at the code involved. The attacker modifies the code on their website to include a button labeled "send a put two". When this button is clicked, the code sends a request to "attacker.com:8080" instead of the expected "localhost:8080". This is the first step in the attack.

Next, the attacker adds an event listener to the button, so that the code only runs when the button is clicked. This ensures that the attack is not immediately triggered upon loading the page. When the button is clicked, the modified code sends a PUT request to "attacker.com:8080".

To further demonstrate the attack, the attacker sets up a real internet server and modifies their hosts file to override the DNS resolution for "attacker.com". This allows the attacker's server to handle requests to "attacker.com:8080" instead of the actual DNS response. By doing this, the attacker can simulate the behavior of a real DNS entry pointing to their server.

When a victim visits "attacker.com:8080", their browser will make a request to the attacker's server. The modified code on the attacker's server will then run, potentially executing malicious actions.

It is important to note that DNS rebinding attacks are complex and require careful setup to demonstrate. The steps described here provide an overview of the attack process, but actual implementation and execution may involve additional considerations and techniques.

DNS rebinding attacks exploit the way web browsers handle DNS resolution to trick victims into making requests to malicious websites. By modifying code and setting up a simulated DNS resolution, attackers can execute malicious actions on victims' browsers.

In the context of web application security, a type of attack known as DNS rebinding attacks can pose a serious threat. DNS, or Domain Name System, is responsible for translating human-readable domain names into IP addresses that computers can understand. DNS rebinding attacks exploit the way DNS works to trick a victim's browser into making unauthorized requests to an attacker's server.

To understand how DNS rebinding attacks work, let's consider an example. Suppose an attacker creates a malicious website and convinces a victim to visit it. When the victim accesses the attacker's site, the attacker's code running in the background changes the DNS entry for their domain. Instead of pointing to the actual server where their code is hosted, the attacker points the DNS entry to a localhost IP address.

Here's how the attack unfolds: the victim remains on the attacker's page, unaware of the DNS change. When the victim interacts with the page, such as clicking a button, the victim's browser sends an HTTP request. However, due to the altered DNS entry, the request is directed to the victim's own computer, specifically to the localhost IP address. This allows the attacker to execute unauthorized actions on the victim's machine, even though the victim initiated the request.

It's important to note that from the browser's perspective, the request appears to be a same-origin request. Same-origin requests are requests made to the same domain, port, and protocol as the page the user is currently on. In the case of DNS rebinding attacks, the domain and port remain the same, tricking the browser into considering the request as same-origin and allowing it to proceed.

One key aspect of DNS rebinding attacks is that the same-origin policy, which is a security mechanism implemented by browsers, does not consider IP addresses when determining whether two entities are of the same origin. Thus, the altered DNS entry does not violate the same-origin policy, enabling the attacker to execute their malicious actions.

To carry out a successful DNS rebinding attack, the attacker needs to have a target device in mind. They may target a specific device with a known port or attempt to scan all possible ports to find the one being used by the victim's application. While there are 65,000 ports available, the attacker can systematically try all of them, given enough time.

As a potential victim, you may experience a DNS rebinding attack by simply visiting a compromised website. As soon as you access the attacker's initial page, the attacker becomes aware of your visit and proceeds to change the DNS entry. The DNS entry is then pointed to the localhost IP address, and the attacker continuously sends unauthorized requests to that address in the background. This ongoing activity allows the attacker to maintain

control over your machine without requiring any further interaction from you.

To protect against DNS rebinding attacks, it is crucial to keep your web applications and browsers up to date. Additionally, implementing proper security measures, such as secure coding practices and input validation, can help mitigate the risk of such attacks.

DNS rebinding attacks exploit the DNS system to deceive browsers into making unauthorized requests to an attacker's server. By altering DNS entries and leveraging the same-origin policy, attackers can execute malicious actions on victims' machines. Understanding the mechanics of DNS rebinding attacks and implementing appropriate security measures is essential to safeguard web applications and protect against this type of threat.

DNS rebinding attacks are a type of cyber attack that exploit vulnerabilities in the Domain Name System (DNS) to deceive users and gain unauthorized access to their devices or networks. In this attack, the attacker registers a domain name and specifies a DNS server that they control. When a user visits a website associated with the attacker's domain, the DNS server returns an IP address that points to the attacker's server instead of the intended server.

One of the key aspects of DNS rebinding attacks is the manipulation of DNS responses. Normally, DNS responses are cached by the user's browser or device for a certain period of time, typically a few minutes. However, the attacker can set the cache expiration time to be very short, even on the order of seconds. This means that the user's device will repeatedly send requests to the attacker's server instead of the intended server until the cache entry expires.

From the user's perspective, the attack is invisible and automatic. They may visit a website and within seconds, their device or vulnerable IoT device sends a request to the attacker's server. The user may not even notice any unusual behavior or experience any visible effects of the attack.

To carry out a DNS rebinding attack, the attacker does not need to modify the DNS settings on the user's device. Instead, they specify their own DNS server when registering the domain. This gives them control over the DNS responses and allows them to change the IP address that gets returned in the answers. They can switch between the attack IP and the local host IP as frequently as they want, potentially every second.

The only unpredictable part for the attacker is how long the DNS entries will be cached by the user's browser. However, the cache duration is typically relatively short, minimizing the impact on the success of the attack.

There are some mitigations that can be implemented to protect against DNS rebinding attacks. One approach is for DNS resolvers, such as those running on the user's device or provided by their ISP, to refuse responses that point to localhost. This would effectively prevent the attack from succeeding. Some enterprises already implement this measure to protect their networks and devices from potential attacks.

It's worth noting that blocking only the 127.0.0.1 IP addresses, which are mapped to local devices, may not be sufficient. The attacker can also iterate through other IP ranges, such as the commonly used 192.168.x.x range, which are also associated with local devices. Therefore, it is important to block all relevant IP ranges to ensure comprehensive protection against DNS rebinding attacks.

DNS rebinding attacks exploit vulnerabilities in the DNS system to deceive users and gain unauthorized access to their devices or networks. By manipulating DNS responses and caching mechanisms, attackers can redirect user requests to their own servers and carry out invisible and automatic attacks. Implementing measures to refuse responses pointing to localhost and blocking relevant IP ranges can help mitigate the risk of DNS rebinding attacks.

DNS attacks, specifically DNS rebinding attacks, are a significant concern in web application security. In this type of attack, an attacker manipulates the DNS resolution process to bypass the same-origin policy and gain unauthorized access to sensitive information or perform malicious actions on a victim's computer.

One example of a DNS rebinding attack is the case of Spotify, where a local server is built into the Spotify app on a user's computer. The server listens on a high-numbered port and is associated with the domain "spottylocal.com". When a DNS lookup is performed on this domain, it resolves to the localhost IP address. This

setup allows Spotify to control the music player through requests sent to "spottylocal.com".

The DNS rebinding attack occurs when an attacker registers a domain, such as "attacker.com", and manipulates the DNS records to point to the IP address of the target server, such as a bank's server. When a user visits the attacker's website, the attacker updates their DNS records to point to the bank's IP address. Subsequent requests made by the user, such as clicking a button, are then sent to the bank's server.

However, from the browser's perspective, the origin remains the attacker's domain, "attacker.com", due to the same-origin policy. Therefore, while the requests reach the bank's server, they are treated as originating from the attacker's domain. This prevents the attacker from directly accessing or manipulating sensitive information on the bank's server.

The damage in this situation is limited because the same-origin policy restricts the attacker's ability to interact with the bank's server. While the attacker can send requests to the bank's server, the browser's security mechanisms prevent the attacker from accessing or modifying the server's response. Additionally, modern browsers enforce strict certificate validation, making it difficult for the attacker to impersonate the bank's website.

It is important to note that the success of a DNS rebinding attack depends on various factors, including the caching behavior of DNS servers and the specific implementation of the victim's browser. Some DNS servers may cache DNS records for longer periods, potentially mitigating the effectiveness of the attack. Additionally, browser vendors continuously improve their security mechanisms to detect and prevent such attacks.

DNS rebinding attacks exploit vulnerabilities in the DNS resolution process to bypass the same-origin policy and gain unauthorized access to web applications. While these attacks can pose a threat, modern browser security mechanisms and strict certificate validation help mitigate the potential damage. It is crucial for web application developers and administrators to stay informed about DNS attack techniques and implement appropriate security measures to protect against them.

DNS rebinding attacks are a type of attack that targets web applications by exploiting the way DNS resolution works. In a DNS rebinding attack, an attacker tricks a victim's browser into making requests to a malicious website, while the victim believes they are interacting with a legitimate website.

The attack starts with the victim visiting a website controlled by the attacker. The victim's browser performs a DNS lookup to resolve the domain name of the attacker's website to its corresponding IP address. The browser then establishes a TCP connection with the server at that IP address and sends an HTTP request.

At this point, the attacker's website includes JavaScript code that triggers a fetch request to the same origin, which is the attacker's website. The browser, considering it a same-origin request, allows the request to proceed without any additional checks.

Here's where the attack takes advantage of DNS resolution. The attacker's DNS server, which is under their control, responds to the browser's DNS query with the IP address of a local server instead of the original IP address. Since the browser does not have the local server's IP address cached, it makes a new TCP connection to the local server and sends the request there instead of the attacker's server.

The local server, which typically assumes that requests come from trusted sources on the same network, processes the request as if it were legitimate. This allows the attacker to potentially exploit vulnerabilities in the local server or gain unauthorized access to resources on the victim's network.

To prevent DNS rebinding attacks, both servers and browsers can take measures. Servers can implement defenses such as:

1. Implementing proper input validation and sanitization to prevent code injection attacks.
2. Enforcing strict access control policies to ensure that only authorized requests are processed.
3. Implementing rate limiting or request throttling mechanisms to detect and mitigate suspicious or malicious activity.

Browsers can also play a role in mitigating DNS rebinding attacks by:

1. Implementing same-origin policy, which restricts the execution of scripts from different origins.
2. Implementing DNS pinning, which ensures that DNS responses are not tampered with or spoofed.
3. Implementing security mechanisms such as Content Security Policy (CSP) to restrict the execution of untrusted scripts.

It is important for users to be aware of the risks associated with insecure IoT devices and take steps to secure their networks. This includes keeping devices up to date with the latest firmware, using strong passwords, and isolating IoT devices on separate networks if possible.

DNS rebinding attacks exploit the way DNS resolution works to trick browsers into making requests to malicious websites. By implementing proper security measures on servers and browsers, and by practicing good network security hygiene, users can protect themselves from these types of attacks.

DNS rebinding attacks are a type of cyber attack that exploit vulnerabilities in the Domain Name System (DNS) to deceive web browsers and gain unauthorized access to sensitive information or perform malicious actions. In this type of attack, the attacker tricks the victim's web browser into making requests to a malicious website disguised as a trusted site.

To defend against DNS rebinding attacks, it is crucial to understand the attack itself. One way to mitigate this threat is by checking a specific header in the server's response. This header is called the "host header," and it informs the server about the intended site of the request.

When a user opens a browser and tries to access a website, the host header contains information about the desired site. If this header is omitted, the server would not know which site the request is intended for. This is particularly important when a server serves multiple websites. The host header helps the server identify the correct site to serve.

To protect against DNS rebinding attacks, the server needs to examine the host header. If the header indicates an origin other than the expected one (e.g., localhost), it signifies a potential DNS rebinding attack. In such cases, the server should reject the request or take appropriate action to prevent further exploitation.

To implement this defense mechanism, a middleware can be used. A middleware is a piece of code that runs before other handlers and intercepts incoming requests. By placing the DNS rebinding check in the middleware, it ensures that the check is performed before any other handlers are executed. This approach avoids the need to duplicate the code in every handler and simplifies maintenance.

The middleware function checks if the host header in the request matches the expected origin (e.g., localhost). If it does not match, an error is thrown, indicating a DNS rebinding attack. If the header is valid, the middleware allows the request to proceed to the next handler using the "next" function.

By implementing this simple check, the server can effectively defend against DNS rebinding attacks. If an attacker attempts to exploit the vulnerability, the server will refuse to process the request, thus protecting the system from potential harm.

It is important to note that DNS rebinding attacks are complex and require a deep understanding of the underlying mechanisms. When developing local servers, it is crucial to exercise caution and carefully follow best practices. Even with precautions, mistakes can happen, so it is essential to be vigilant and continuously update knowledge on potential vulnerabilities.

DNS rebinding attacks are a significant threat to web applications' security. By implementing a check on the host header, servers can defend against these attacks effectively. However, it is crucial to remain cautious and continuously update knowledge on emerging threats in the field of cybersecurity.

**EITC/IS/WASF WEB APPLICATIONS SECURITY FUNDAMENTALS DIDACTIC MATERIALS****LESSON: BROWSER ATTACKS****TOPIC: BROWSER ARCHITECTURE, WRITING SECURE CODE**

## Browser Architecture and Writing Secure Code

In the previous lecture, we discussed DNS rebinding attacks. These attacks involve a browser being directed to an attacker's website, where a DNS lookup is performed for the attacker's domain. The attacker manipulates the DNS response to point to their own server. The browser then makes an HTTP request to this server, which sends back a page. At this point, the attacker changes their DNS entry to return a localhost IP address. The code running on the attacker's page then makes a request to the same URL as before, triggering another DNS lookup. This time, the localhost IP is returned, and the subsequent HTTP request is sent to the server running locally on the victim's machine. The browser, considering this a same-origin request, allows it to proceed, enabling the attack.

To defend against DNS rebinding attacks, we can add a few lines of code to check the host header. If the host header contains a localhost value, it indicates that a DNS rebinding attack has not occurred. This simple check can help prevent such attacks.

Now, let's discuss the attack surface of local servers in general. When starting a local server, it is common to bind it to the local IP address. By specifying the local IP address as the second argument when calling the listen function, we restrict connections to the server from other machines on the same network. This reduces the attack surface. However, even if this step is overlooked, incoming connections to the local server are often blocked by the operating system's built-in software firewall. Therefore, relying on the operating system alone for protection may not be sufficient.

Furthermore, DNS rebinding attacks can still occur even with these defenses in place. It is crucial to take proactive measures, such as checking the host header, to mitigate the risk of such attacks. The idea behind DNS rebinding attacks is that the attacker's code runs in the victim's browser, acting as a proxy. This allows the attacker to make requests to the victim's IoT devices or local servers. Therefore, solely relying on the operating system's defenses is not enough.

In Mac OS, the software firewall can be enabled to provide additional protection. By default, the firewall allows incoming connections to any software that has a valid signature from a developer. This means that if a locally running server is signed by a recognized developer, incoming connections from other devices on the network or the internet will be allowed. This default policy can be customized based on the user's preferences.

In Windows, a similar firewall feature exists, allowing users to control incoming connections to locally running servers.

It is important to understand the browser architecture and the potential vulnerabilities of local servers to ensure the security of web applications. By implementing secure coding practices and being aware of potential attack vectors, developers can reduce the risk of browser attacks.

When using a Windows computer, you may have encountered a prompt asking you to allow or deny access to a network. This prompt is meant to protect you from local servers that the software you are running has started up. By clicking "allow" with the private box checked, you are indicating that you trust the network you are connected to, such as your home network. This allows other devices on the same network to make connections to your computer and the specific app that is listening for connections. On the other hand, if you are on a public network, it is recommended not to check the private box, as you may not want other devices on the network to connect to your computer and the app. This is particularly important when it comes to defending against DNS rebinding attacks. However, it is worth noting that these settings do not protect against DNS rebinding attacks, as the browser acts as a proxy and bypasses the firewall settings.

Now, let's discuss IoT (Internet of Things) devices. IoT devices, such as Chromecast or Google Home, are designed to expect connections from other devices on the same network. These devices typically do not have a firewall configured, as they want incoming connections to go through. As a result, they are vulnerable to attacks from malicious devices. Additionally, any webpage can make connections to these IoT devices if they are



listening for HTTP requests. DNS rebinding allows us to upgrade these requests from simple requests to more complex ones, potentially opening up new avenues for attack.

It is important to note that IoT devices are often difficult to update once they are shipped to customers. This is due to price pressures and the fact that consumers prioritize cost over security when purchasing these devices. This lack of security awareness has led to incidents like the Mirai botnet, where teenagers took over a massive number of IoT devices worldwide. Their motivation was a feud with other teenagers running a competing Minecraft server. They aimed to take down the hosting of the Minecraft server and even attacked the DNS provider associated with it.

Understanding browser architecture and writing secure code is crucial in ensuring web application security. By being aware of network access settings and the vulnerabilities of IoT devices, we can better protect ourselves and our systems from potential attacks.

### Browser Attacks: Browser Architecture and Writing Secure Code

Browser attacks are a common threat in the realm of cybersecurity, particularly when it comes to web applications. Understanding the architecture of browsers and implementing secure coding practices is crucial in order to protect against these attacks.

One aspect of browser attacks involves the use of Internet of Things (IoT) devices. These devices, such as toasters or thermostats, can be utilized by attackers to launch large-scale attacks on web servers. In one notable case, teenagers managed to amass a significant army of IoT devices to direct traffic towards Minecraft servers, causing disruption. The attackers were able to connect directly to these devices, which were often vulnerable to takeover due to weak default passwords or lack of security measures. Additionally, the attackers attempted to divert attention by including Russian language strings in their malware.

Another attack vector to consider is DNS rebinding. This attack exploits the way browsers handle DNS resolution, allowing attackers to bypass security measures and gain unauthorized access to resources. For example, the Blizzard update agent, a piece of software installed alongside Blizzard games, was found to be vulnerable to DNS rebinding. This allowed any website on the internet to send requests to the update server, resulting in the execution of arbitrary code on users' computers. Similarly, the popular BitTorrent client, Transmission, had a vulnerability that allowed attackers to exploit the separation of its user interface and protocol logic, potentially compromising user systems.

To mitigate these risks, it is essential to implement secure coding practices. This includes ensuring that IoT devices have strong, unique passwords and are not listening to the entire network. Additionally, developers should be aware of the potential for DNS rebinding attacks and check the host header to prevent unauthorized access. Regular security audits and updates are crucial to address vulnerabilities promptly.

Understanding browser architecture and implementing secure coding practices are essential in defending against browser attacks. By addressing IoT device vulnerabilities, considering the risks of DNS rebinding, and adopting secure coding practices, individuals and organizations can significantly enhance their web application security.

### Browser Architecture and Writing Secure Code

In web applications, there are typically two processes involved: one for the user interface (UI) and another for communicating with the torrent network and connecting to peers. The UI process sends messages to the client, which acts as the server in a client-server model. However, in some cases, the server can be vulnerable to attacks if not properly secured.

One common vulnerability is known as DNS rebinding, where an attacker can exploit the server by sending requests from any site on the internet that the user is browsing. This can allow the attacker to execute code on the user's computer, potentially leading to unauthorized access and control.

To mitigate this vulnerability, it is crucial to implement proper security measures in the code. One example of a vulnerable application is Transmission, which had a DNS rebinding vulnerability. This vulnerability allowed any site on the internet to send requests to the Transmission server, potentially executing code on the user's

computer. This vulnerability was discovered and reported by a researcher at Google Project Zero.

Similarly, another application called WebTorrent, developed by the speaker, also had a DNS rebinding vulnerability. However, in this case, the vulnerability did not allow the attacker to execute arbitrary code. Instead, it allowed the attacker to send requests to the local server and gather information about the user's activities, such as the content they were downloading.

To fix the vulnerability in WebTorrent, the developer implemented a check in the code. This check ensured that the host header in the incoming requests matched a specific option defined by the developer. If the host header did not match, the connection was closed, preventing DNS rebinding attacks.

However, the initial fix turned out to be ineffective. It appeared to work during testing, but it did not provide the desired protection. The vulnerability persisted for an additional six months before it was finally resolved.

The root cause of the ineffective fix was a flaw in the code that handled incoming HTTP requests. The code did not properly check the host header, allowing attackers to bypass the protection implemented in the previous fix.

To address this issue, a revised fix was implemented. The new fix correctly checks the host header and closes the connection if it does not match the expected value. This ensures that DNS rebinding attacks are effectively mitigated.

Browser architecture plays a crucial role in web application security. Writing secure code is essential to protect against vulnerabilities such as DNS rebinding. Properly checking the host header in incoming requests is a fundamental step in defending against these attacks.

One important aspect of web application security is browser architecture. Browsers play a crucial role in defending users as they browse the web, protecting them from various threats such as malicious sites and untrusted code. In this section, we will explore how browsers secure the execution of untrusted code and what happens when things go wrong.

To begin with, browsers employ a security mechanism known as the same-origin policy. This policy ensures that websites from different origins (domains, protocols, and ports) are isolated from each other, preventing interference and unauthorized access to sensitive information. We have discussed the same-origin policy extensively in previous lessons.

However, there is still the question of how browsers securely execute untrusted code. JavaScript, a widely used scripting language on the web, runs within a sandboxed environment, limiting its capabilities and preventing it from performing actions that could compromise user security. For instance, JavaScript is not allowed to read files on a user's computer, open unauthorized network connections, or execute native code.

Despite these security measures, there is always the possibility of vulnerabilities in the browser's code or the JavaScript interpreter. In such cases, malicious JavaScript could exploit these vulnerabilities to bypass the security restrictions and perform unauthorized actions.

To mitigate this risk, browsers have implemented various security measures. For example, they employ a layered architecture that separates the rendering engine, JavaScript interpreter, and other components. This isolation helps contain potential vulnerabilities and limit the impact of any successful attacks.

Additionally, browsers utilize techniques such as sandboxing and privilege separation. Sandboxing involves running untrusted code in a restricted environment, isolating it from the rest of the system. Privilege separation ensures that different components of the browser run with different levels of access privileges, preventing unauthorized actions.

When a browser encounters a security issue or vulnerability, it is crucial to address it promptly. Browser vendors regularly release updates and patches to fix these vulnerabilities and improve security. It is essential for users to keep their browsers up to date to benefit from these security enhancements.

Browser architecture plays a vital role in ensuring web application security. Browsers employ various mechanisms, such as the same-origin policy, sandboxing, and privilege separation, to protect users from

malicious sites and untrusted code. Despite these measures, vulnerabilities can still occur, and it is crucial for browser vendors to promptly address and patch these issues to maintain a secure browsing experience.

Browsers play a crucial role in protecting our computers while we browse the internet. However, they face challenges due to their complexity and the constant addition of new features. In this didactic material, we will discuss the high-level architectural decisions made by browsers to ensure our security while browsing.

Browsers are complex software with large codebases, and they are also network-visible, meaning they are connected to the internet. This combination makes them vulnerable to attacks. Attackers exploit the complicated code and send untrusted input to break the browser's security. Moreover, browsers are often written in unsafe languages like C++, which further increases the risk.

To mitigate these risks, browsers have improved their architecture over time. Before these improvements, browsers were frequently targeted with remote code execution attacks, where visiting a malicious website could lead to the complete compromise of the user's computer. However, thanks to advancements in browser architecture, such attacks have become less common.

Let's take a look at an example of JavaScript code that demonstrates the potential dangers. In 2014, a vulnerability was discovered in JavaScript that allowed an attacker to manipulate memory and gain control over the browser. The attacker used a buffer overflow attack, where they allocated a small buffer but manipulated the byte length property to access memory beyond the buffer's bounds. This allowed them to read and write arbitrary memory, leading to the compromise of the browser and potentially the entire computer.

While this example showcases the severity of browser vulnerabilities, it is worth noting that achieving full remote code execution requires additional code and techniques. However, the fact remains that attackers can read and write arbitrary memory, which can be a stepping stone for further exploitation.

To address these issues, browser vendors have implemented various security mechanisms. These include sandboxing, where the browser isolates web content from the underlying operating system, and strict code execution policies, which limit the capabilities of web code. Additionally, browsers regularly release updates to patch vulnerabilities and improve security.

Browser architecture plays a vital role in protecting users from browser-based attacks. Despite the complexity and challenges associated with browser security, advancements in architecture have significantly reduced the risk of remote code execution attacks. However, it is crucial for users to keep their browsers up to date and follow best practices to ensure their online safety.

Web browsers are complex software applications that allow users to access and interact with websites and web applications. However, due to their complexity, they can also be vulnerable to various types of attacks. One such type of attack is known as a browser attack, which targets the security vulnerabilities present in web browsers.

To understand browser attacks, it is important to first understand the architecture of a browser. A browser consists of different modules and components that work together to provide the functionality and user experience we are familiar with. These modules include the rendering engine, JavaScript engine, networking stack, and more.

The JavaScript engine is responsible for interpreting and executing JavaScript code on web pages. It takes strings of JavaScript and updates a state machine that represents the state of the JavaScript world. However, if the JavaScript engine is compromised, an attacker can manipulate the code and take control of the entire process.

This is a serious concern because when a browser is launched, it runs under the user's account on the operating system. As a result, the compromised process has the ability to read and write all the files that the user interacts with, including sensitive documents. This makes browser attacks a significant threat to user privacy and security.

One common type of vulnerability that allows browser attacks is memory safety issues. These issues occur when there are errors or vulnerabilities in the code that handles memory management. Attackers can exploit

these vulnerabilities to execute malicious code and gain control over the browser process.

In fact, memory safety issues are the main type of vulnerability that is still prevalent today. A study conducted by Microsoft revealed that 70% of the vulnerabilities they addressed through security updates each year were memory safety issues. Similarly, Chrome, a popular web browser, reported that out of all the critical severity bugs they have encountered, 130 of them were related to memory corruption.

To mitigate the risks associated with browser attacks, browser developers have adopted the "rule of two." According to this rule, there are three criteria that must be considered when designing the browser architecture. These criteria are:

1. Processing untrustworthy input: Web browsers deal with complex data formats like HTML and CSS, which are often received from untrusted sources. Since it is difficult to ensure the correctness of the code that processes these inputs, this criterion is met.
2. Code written in an unsafe language: Unsafe languages like C++ and assembly are prone to memory safety issues. Since web browsers are typically written in C++ and assembly, they meet this criterion as well.
3. Any code that is wanted to run on the user's computer: This criterion refers to any code that is executed within the browser process. Since the browser process runs under the user's account, any code executed within it can potentially access and manipulate sensitive files.

By considering these criteria, browser developers aim to minimize the risks associated with browser attacks. They focus on writing secure code, implementing proper memory management techniques, and regularly releasing security updates to address any vulnerabilities that may arise.

Browser attacks pose a significant threat to user privacy and security. Understanding the architecture of web browsers and the vulnerabilities they may have is crucial in developing secure web applications. By following secure coding practices and staying updated with the latest security patches, users and developers can minimize the risks associated with browser attacks.

In the context of web application security, it is important to understand the architecture of web browsers and how to write secure code to prevent browser attacks. Browsers have different levels of privilege, with some components having higher privilege than others. For example, the bootloader, firmware, and operating system kernel have the highest privilege on a computer. On the other hand, the browser runs as a user account and has privileged permissions within the browser context.

To ensure security, code in a browser should only have the ability to perform two out of three actions: delete files, read files, or cause damage. If code has the ability to do all three, it poses a significant security risk. In the case of Google Chrome, the browser separates the complicated parsing tasks, which involve untrusted input, from other components. This separation is not entirely precise, but it helps identify the components where security vulnerabilities are more likely to occur. For example, image decoding is a common source of bugs due to the complexity of image formats and the fact that images can be crafted by users to exploit buffer overflow issues.

To mitigate these risks, Google Chrome implements a process separation model. The browser kernel and the rendering engine operate as separate processes. The rendering engine handles the potentially error-prone components, while the browser kernel handles sensitive tasks such as reading files, network communication, and managing user cookies. These two processes communicate through an inter-process communication (IPC) channel, which acts as a socket between them.

By separating these processes, the browser can maintain its functionality while minimizing the impact of security vulnerabilities. If an issue arises in one process, it can be isolated and denied access to sensitive resources. The only way for the process to perform these actions is by communicating through the IPC channel. This approach assumes that the rendering process can be compromised, but even in such cases, the potential damage is limited.

The architecture of Google Chrome can be visualized as follows: the browser kernel, which controls the browser's user interface, interacts with the network, and manages file operations, communicates with the

rendering engine through the IPC channel. The rendering engine, operating as a separate process, downloads all the necessary resources for rendering a webpage, such as HTML, CSS, and JavaScript. It then sends the rendered image back to the browser kernel, which displays it on the page. Each tab in the browser corresponds to a separate rendering process, while the browser kernel handles the overall control and interaction with the user.

This architecture provides several benefits. Firstly, it limits the potential damage that can be caused by an attacker who compromises the rendering process. Secondly, if a rendering process crashes due to a bug, it does not affect the entire browser, ensuring a more stable browsing experience.

Understanding the architecture of web browsers and writing secure code is crucial for web application security. Google Chrome's process separation model, with distinct browser kernel and rendering engine processes, helps mitigate security risks and provides enhanced stability. By separating potentially vulnerable components from sensitive operations, the browser can maintain its functionality while minimizing the impact of security vulnerabilities.

### Browser Architecture and Writing Secure Code

In the realm of web application security, understanding browser architecture and writing secure code are essential. One fundamental aspect of browser architecture is the implementation of a multi-process model. This model ensures that each tab in a browser runs in its own separate process. This means that if one tab crashes, it does not affect the entire browser or other open tabs. In contrast, older browsers like Internet Explorer used a single process for the entire browser, so if one tab had a bug or crashed, it would bring down the entire browser, resulting in the loss of all open tabs and work.

The multi-process architecture not only improves security but also enhances robustness. Renderer processes, responsible for rendering web pages, are isolated and cannot access the disk, network, or devices. This isolation is achieved through running these processes in a restricted sandbox. By limiting the capabilities of renderer processes, the potential damage caused by attackers is significantly reduced. Additionally, this architecture allows web pages to run in parallel, leveraging the benefits of parallelism through multiple processes. Furthermore, if a tab crashes, the browser can continue running, and the user does not lose their work.

Another crucial innovation in browser security is the introduction of auto-updates. Initially pioneered by Chrome, all modern browsers now adopt this practice. Auto-updates ensure that the browser is always up to date with the latest security patches and enhancements. In the past, software updates often relied on user prompts or notifications, leading to delayed or neglected updates. Chrome took a different approach, automatically updating the browser without giving users the option to opt out. Although controversial at first, this approach is now considered standard practice, as it ensures users have the most secure browsing experience.

Interestingly, Chrome's development began with an auto-updater as its sole functionality. The initial version of the browser was a blank window that contained only the auto-updater code. Over time, as the project evolved, the browser gradually gained additional features and functionality. This unique approach allowed developers and open-source contributors to follow the project's progress from its early stages.

While the multi-process architecture and auto-updates greatly enhance browser security, there are still limitations. One challenge is the desire to place pages from different origins into separate renderer processes. However, due to the complexity of achieving this, tabs often consist of content from multiple sites running within the same process. This means that within a single tab, there can be multiple sites and their associated JavaScript code running together. While the multi-process architecture isolates the browser kernel from the tabs, it does not completely segregate different sites within a tab.

Understanding browser architecture and writing secure code are crucial aspects of web application security. The multi-process model ensures that each tab operates independently, preventing crashes from affecting the entire browser. Isolating renderer processes through sandboxing limits the potential damage caused by attackers. Auto-updates keep the browser up to date with the latest security patches, ensuring a secure browsing experience. Despite some limitations, these advancements have significantly improved browser security, providing users with a safer online environment.

Web browsers play a crucial role in ensuring the security of web applications. However, they can also be

vulnerable to attacks if not properly secured. One potential vulnerability is the browser architecture itself, which consists of multiple processes, including the renderer process responsible for drawing web pages.

In a typical scenario, when a user visits a website, the renderer process needs to load various resources, such as cookies, to properly display the page. This process is isolated from the browser and other websites through a security mechanism called the same-origin policy. This policy prevents malicious websites from accessing sensitive information from other websites.

However, if an attacker finds a bug in the browser and manages to break out of the renderer process, they gain control over the entire process. This means they can bypass the same-origin policy and access sensitive information, such as cookies, stored in the renderer process's memory. While they still can't access files on the user's disk, this breach of security can have significant consequences.

Another vulnerability that can be exploited is Specter, an issue that allows attackers to read memory from other processes. In the case of the renderer process, there is no process boundary between the attacker's website and the victim's website. Instead, the boundary is maintained by the code written to separate them. However, an attacker can use Specter to read memory from other parts of the process, violating the process model of modern operating systems.

To address these vulnerabilities, web browsers, like Chrome, have implemented a solution called site isolation. Instead of isolating each tab in a separate process, site isolation assigns each different site to its own process. This means that if a tab loads content from multiple sites, each site will have its own dedicated process. These processes work independently to load and render their respective parts, and the final result is stitched together to display the complete page.

With site isolation, the browser can now differentiate between different sites and their respective renderer processes. This allows the browser to make informed decisions when it comes to loading resources or providing cookies. For example, if a renderer process asks the browser to download content from victim.com or provide cookies for victim.com, the browser can now determine if this request is legitimate or if the process has been compromised. This enhanced security measure helps prevent attackers from accessing sensitive information even if they manage to compromise a renderer process.

Browser architecture and writing secure code are fundamental aspects of web application security. Understanding the vulnerabilities associated with browser processes and implementing security measures like site isolation can help mitigate the risks of browser attacks and protect user data.

In the context of web applications security, it is important to understand the architecture of web browsers and how to write secure code. Web browsers consist of multiple processes, including the browser process and the renderer process. The browser process is responsible for managing the user interface, while the renderer process handles rendering and executing web pages.

When a web page is loaded, the browser process collects the results from each renderer process and stitches them together to create the complete page. Each renderer process is isolated and restricted to only accessing content from its own site. For example, if a renderer process belonging to the site "evil.com" tries to access cookies from the site "victim.com," the browser process will deny the request.

However, the renderer process for "evil.com" is still allowed to load images and other content from external sites. This is enforced by checking the headers of the response. If the headers indicate that the content is allowed to be readable by anyone, the browser process will allow it.

To illustrate this concept, let's consider the scenario where "evil.com" wants to load images from "victim.com." This is allowed because it falls under the same origin policy. However, if "evil.com" tries to load an HTML page and embed it as an image, the browser process will reject the request. This is because the HTML page cannot be loaded as a valid image, and allowing it could potentially compromise the security of the renderer process.

This security feature is known as Cross-Origin Read Blocking (CORB). It prevents cross-origin requests from accessing sensitive data and helps protect against attacks like Spectre.

Now, let's shift our focus to secure coding practices. JavaScript, the language commonly used in web



development, has some quirks and pitfalls that developers should be aware of. One such issue is the behavior of the double equals (`==`) operator. It performs a coercion using the abstract equality comparison algorithm, which can lead to unexpected results. To avoid this, it is recommended to always use the triple equals (`===`) operator, which performs a strict equality check by comparing both the value and the type.

For example, if you want to check if an argument is equal to 0, using the double equals operator may yield unexpected results. However, using the triple equals operator will give you the expected behavior by checking both the value and the type of the argument.

Understanding the browser architecture and writing secure code are essential in web applications security. By following best practices like isolating renderer processes and using proper equality operators, developers can mitigate security risks and ensure the integrity of their applications.

In the context of web application security, understanding browser attacks is crucial to ensure the integrity and security of the code we write. One aspect of browser attacks is related to the browser architecture itself and how it handles certain scenarios. This didactic material will cover some common issues and provide solutions to write more secure code.

One issue that developers may encounter is when duplicate function arguments are used. In JavaScript, if two arguments have the same name, the second one will take precedence over the first one. This behavior can lead to unexpected bugs in the code. To address this, developers can enable strict mode in their code by including the string `"use strict"` at the top of their file. This activates a stricter version of JavaScript that fixes some of the language's rough edges. With strict mode enabled, duplicate parameter names will result in a syntax error, providing an early indication of the issue.

Another approach to improving code security is by using a linter, such as ESLint. A linter is a tool that analyzes code and warns developers about potential issues based on configurable rules. ESLint, for example, offers hundreds of rules that can be customized according to the developer's preferences. To simplify the process, developers can install preset configurations, such as the popular `"standard"` preset. Running the linter with the chosen configuration will highlight formatting issues and potential errors in the code. Additionally, the linter often includes an automatic fixer that can address some of the identified issues automatically.

It is important to note that not all issues can be caught by enabling strict mode. One such example is having an object with duplicated key names. JavaScript allows dynamically setting key names at runtime using variables, making it impossible to detect duplicates during static analysis. To identify and prevent this issue, a linter is necessary.

Another common issue that can have severe consequences is when users provide JSON objects with property names that clash with existing object properties. For example, if a user includes a property named `"hasOwnProperty"` in their object, it will override the inherited `"hasOwnProperty"` function from the object superclass. This can lead to unexpected behavior and even server crashes. Using a linter can help catch such issues. Alternatively, developers can directly call the `"hasOwnProperty"` function from the object superclass, ensuring that the correct implementation is used.

Finally, it is worth mentioning the concept of automatic semicolon insertion in JavaScript. This feature was introduced to help beginner programmers by automatically inserting semicolons at the end of lines. However, it can lead to unexpected behavior in certain cases. Developers should be aware that relying on automatic semicolon insertion can introduce bugs. It is recommended to explicitly terminate lines with semicolons to avoid any ambiguity.

By understanding these browser architecture and code writing best practices, developers can enhance the security of their web applications and reduce the risk of browser-based attacks.

In web application development, understanding browser architecture and writing secure code are crucial for ensuring the security of the application. One aspect to consider is the use of semicolons in JavaScript code. While it may seem like inserting semicolons at the end of every line would protect against errors, it is not a foolproof solution. Automatic semicolon insertion can sometimes lead to unexpected behavior, such as returning an object after a semicolon. To avoid such issues, it is important to understand the rules of semicolon usage in JavaScript.

Using a linter, a tool that analyzes code for potential errors, can help catch these issues. A linter can detect unused code and warn about unreachable code, regardless of whether semicolons are inserted or not. Therefore, it is recommended to always use a linter when writing JavaScript code to ensure its correctness.

Another common mistake is adding leading zeros to numbers, which can unintentionally convert them into octal numbers instead of decimal numbers. For example, adding a leading zero to the number 9 would result in 09, as 9 is not a valid digit in octal representation. To avoid such confusion, it is advisable to use a linter or enable strict mode, which disallows octal numbers.

Declaring variables properly is another important consideration. Forgetting to include a declaration keyword, such as "let" or "const," before defining a variable can result in unintentionally creating global variables. This can lead to unexpected behavior, as the variables will persist outside the scope of the function. To prevent this, strict mode or a linter can be used to catch such errors.

JavaScript has had its fair share of peculiarities in the past. For instance, assigning a value to "undefined" used to break programs. However, such issues have been addressed, and using strict mode or a linter can help identify and prevent such errors.

Browser architecture and writing secure code in web applications require attention to detail. Understanding the nuances of JavaScript, such as semicolon usage, preventing octal number conversions, and proper variable declaration, is crucial for ensuring code correctness. Utilizing tools like linters can greatly assist in catching potential errors and improving the overall security of web applications.

In web application development, it is crucial to prioritize security to protect user data and prevent unauthorized access. One area of concern is browser attacks, where attackers exploit vulnerabilities in the browser architecture or insecure code to compromise the application.

To ensure browser security, developers should understand the browser architecture and write secure code. The browser architecture consists of various components, including the rendering engine, JavaScript engine, and DOM (Document Object Model). Each component plays a specific role in rendering and executing web pages.

When writing secure code, it is important to follow best practices and avoid common pitfalls. One common mistake is omitting semicolons at the end of lines in code. While some programming styles discourage the use of semicolons, omitting them inconsistently can lead to syntax errors. To address this, developers can either add semicolons or use a linter, such as Standard or ESLint, to catch such errors.

Another pitfall is relying on clever programming techniques that sacrifice code readability. For example, taking advantage of short-circuit boolean evaluation to skip executing certain code blocks can make the code harder to understand, especially for beginners. It is recommended to prioritize readability over brevity, as code should be written for humans, not just computers.

Furthermore, it is important to avoid overly complex code that unnecessarily complicates logic. Clever code may showcase a programmer's knowledge, but it can be difficult to maintain and understand for other team members, especially beginners. It is advisable to write code that is straightforward and easily comprehensible by the entire team.

When it comes to browser attacks and writing secure code, developers should prioritize browser security by understanding the browser architecture and following best practices. This includes avoiding inconsistent use of semicolons, prioritizing code readability over cleverness, and keeping code straightforward and comprehensible for the entire development team.

When writing code for web applications, it is important to consider the long-term implications and potential lack of context. Code that may seem clever or efficient at the moment may become difficult to understand or maintain in the future. Therefore, it is advisable to avoid using esoteric features or relying on obscure edge cases in the programming language unless there is a significant benefit.

Similarly, writing secure code is akin to writing in a clear and concise manner. Using overly complex or flowery language may hinder communication and understanding. It is important to consider the audience and the

message being conveyed. While some individuals may view coding as an art form and enjoy experimenting with unconventional techniques, it is crucial to balance creativity with practicality.

Another essential aspect of writing robust code is the inclusion of tests. Testing code ensures that it functions as intended and helps identify any potential issues or bugs. Untested code is inherently flawed, and relying solely on personal judgment without proper testing is not recommended.

In the process of developing web applications, it is common to utilize code written by others through open-source libraries and dependencies. These external contributions can significantly impact the functionality and security of the application. It is crucial to acknowledge that by incorporating these dependencies, developers implicitly rely on the work of thousands of individuals they have never met. This reliance introduces potential risks, particularly in terms of security.

The open-source supply chain concept highlights the dependence on various vendors and contributors who provide the raw materials (code) used in the development process. These individuals, often volunteers, may make accidental mistakes or lack the necessary resources to actively maintain their packages. Consequently, developers may encounter unmaintained code that poses security vulnerabilities. It is important to be aware of these risks and take appropriate measures to mitigate them.

### Browser Attacks: Browser Architecture and Writing Secure Code

In the world of web applications security, browser attacks pose a significant threat. To understand how to protect against these attacks, it is essential to have a solid understanding of browser architecture and the importance of writing secure code.

Browser architecture plays a crucial role in determining the security of web applications. Browsers are complex software systems that consist of multiple components, including rendering engines, JavaScript interpreters, and user interface elements. These components work together to render web pages and execute scripts. However, this complexity also makes browsers vulnerable to attacks.

One common type of attack is the accidental vulnerability. This occurs when developers unintentionally introduce vulnerabilities into their code. For example, a simple mistake in an install script could result in the deletion of critical system files. In one case, an install script for a package called Bumblebee had a space in a command, causing it to delete the entire user folder, resulting in widespread damage to users' machines. It is important to note that this was not a malicious act but an honest mistake by the package maintainer.

Another significant vulnerability is the under-maintained package. These are packages that are still available but lack regular maintenance and updates. This lack of attention leaves them vulnerable to attacks. In the open-source ecosystem, it is not uncommon for maintainers to become overwhelmed or lose interest in their projects. This can lead to vulnerabilities going unnoticed and unpatched, leaving users at risk.

In some cases, malicious actors may target open-source projects by compromising the accounts of maintainers. This can result in intentional code changes that introduce vulnerabilities or even the complete deletion of code, causing widespread disruptions. For example, the Heartbleed vulnerability in the OpenSSL library, which is used to negotiate TLS connections, was discovered to be maintained by just one person who received minimal funding. This incident highlighted the importance of adequate resources and support for critical open-source projects.

It is also worth mentioning the impact of individual maintainers on widely-used software. For instance, Curl, a widely-used tool for making HTTP requests, was maintained by a single hobbyist for 20 years. This individual recently transitioned to full-time work on Curl after accepting donations from the community. This example emphasizes the reliance on individual maintainers and the need for sustainable support for critical software.

To mitigate these vulnerabilities, it is crucial to prioritize secure coding practices. Developers should follow best practices, such as input validation, output encoding, and proper error handling, to prevent common attack vectors like cross-site scripting (XSS) and SQL injection. Additionally, regular code reviews, testing, and timely updates are essential to address any vulnerabilities that may arise.

Understanding browser architecture and writing secure code are fundamental to ensuring web application

security. Accidental vulnerabilities, under-maintained packages, and malicious attacks all pose significant risks. By implementing secure coding practices and providing adequate support to open-source projects, we can enhance the security of web applications and protect users from potential threats.

In the world of open-source software, trust is a fundamental aspect. Developers who contribute to open-source projects are given the responsibility and permission to publish new versions of the software based on the trust they have earned. However, this trust can sometimes be exploited by malicious actors.

One example of this is a case where a developer, who had not actively worked on a package for four years, gave permission to someone to publish new versions of the package. Unbeknownst to the developer, the person who published the new versions added malware to the package. The malware was cleverly hidden as an encrypted string that would be decrypted and evaluated at runtime. The decryption key was based on the parent package, making it difficult to detect the malicious code.

This attack was specifically targeted at a Bitcoin wallet. The attacker knew the exact package that was above the wallet in the dependency tree and used it as the decryption key. As a result, the attacker was able to steal a significant amount of Bitcoin from the targeted wallet. It is important to note that the developer did not intend to give their package to a malicious actor, but rather, they were overworked and overwhelmed with other responsibilities.

Similar dynamics can be observed in the browser extension ecosystem. Many browser extensions are developed by individuals, and sketchy actors often approach the creators of popular extensions with offers to buy the rights to the extensions. Once the ownership of the extension changes hands, the behavior of the extension can change dramatically. Malware can be added, and user browsing activity can be collected.

This issue arises due to the difficulty in monetizing browser extensions. Developers who have put their time and effort into creating extensions may be enticed by offers of financial compensation and willingly hand over control of their extensions. This desperation to monetize their work can lead to unintended consequences for users.

To address these challenges, it is crucial to adopt a mindset of thinking like an attacker. When writing code, it is important to consider potential vulnerabilities and how the code can be exploited. This mindset should also extend to reviewing and analyzing other people's code.

Additionally, user input should never be trusted and should always be sanitized to prevent attacks like code injection or cross-site scripting. Defense-in-depth should be employed, assuming that one security control may fail and planning for such scenarios. Passwords should be salted and hashed to protect user data, and bcrypt is recommended as a secure hashing algorithm.

Furthermore, developers should be aware of ambient authority, where browsers attach cookies to requests destined for specific sites. This can be mitigated by using same-site cookies. It is also advised to avoid overly clever code and prioritize explicitness over magic. Being explicit in code helps improve readability, maintainability, and reduces the risk of introducing vulnerabilities.

Trust is a crucial aspect in the world of open-source software and browser extensions. However, it is important to be cautious and aware of potential security risks. By adopting a mindset of thinking like an attacker, sanitizing user input, employing defense-in-depth, salting and hashing passwords, using same-site cookies, and prioritizing explicitness in code, developers can enhance the security of their web applications and protect users' data.

### Explicit Code and Security

When it comes to writing secure code, it is important to prioritize explicit code over magical code. While magical code may be fun at times, explicit code is generally better for security purposes. Dangerous code should be easily identifiable by having a scary name and proper documentation. It is also crucial to avoid combining functions that perform benign tasks with risky code. Keeping these two separate is essential for maintaining security.

In the field of cybersecurity, having a paranoid mindset can be a valuable asset. Being constantly vigilant and

aware of potential threats is key to ensuring the security of web applications. It is better to err on the side of caution and take necessary precautions to protect against attacks.

### Browser Architecture and Writing Secure Code

Browser architecture plays a significant role in web application security. Understanding the structure of browsers can help developers write more secure code. Browsers consist of different components, such as the rendering engine, JavaScript interpreter, and networking stack. Each component has its own vulnerabilities that can be exploited by attackers.

To write secure code, developers should follow best practices, such as input validation, output encoding, and proper handling of user sessions. Input validation ensures that user inputs are checked for malicious content before being processed. Output encoding helps prevent cross-site scripting (XSS) attacks by encoding user-generated content. Proper handling of user sessions involves securely managing session tokens and implementing measures to prevent session hijacking.

### Changes in Web Application Development

Over time, there have been significant changes in web application development that impact security. One major change is the size of packages produced by developers. Modern programming languages and package managers, such as NPM and Rust's package manager, have addressed the issue of dependency conflicts. This has allowed developers to produce smaller units of work, leading to smaller packages and larger dependency trees.

Additionally, collaboration has become easier with platforms like GitHub. This has resulted in faster code development and a larger number of contributors. The scale of open-source projects has increased significantly, with more code coming from various sources.

### Recommended Classes for Further Study

If you are interested in pursuing further studies in cybersecurity, there are several classes that you may find beneficial. CS 255, Introduction to Cryptography, offers a comprehensive introduction to cryptographic principles and techniques. CS 155 focuses on the implementation aspects of security and is more similar to this class. CS 355 is the continuation of CS 255 and delves deeper into cryptography.

For those interested in blockchain technology and its security implications, CS 251, Cryptocurrencies and Blockchain, is a recommended class. However, please note that certain classes may have application deadlines, so it is important to plan accordingly.

Writing secure code involves prioritizing explicit code, maintaining a paranoid mindset, understanding browser architecture, and following best practices. By staying informed and continuously learning, individuals can contribute to the development of secure web applications.

Web applications are an integral part of our daily lives, as they allow us to perform various tasks and access information through our web browsers. However, the widespread use of web applications also makes them a prime target for cyber attacks. In order to ensure the security of these applications, it is crucial to understand the fundamentals of browser attacks and the importance of writing secure code.

Browser architecture plays a significant role in the security of web applications. Browsers are complex software systems that consist of multiple components, including the rendering engine, JavaScript engine, and various plugins. Each component has its own vulnerabilities that can be exploited by attackers. For example, the rendering engine is responsible for displaying web content, but it can also be manipulated to execute malicious code. Similarly, the JavaScript engine can be targeted to execute unauthorized actions.

To mitigate the risks associated with browser attacks, developers must follow secure coding practices. Writing secure code involves implementing measures to prevent common vulnerabilities, such as cross-site scripting (XSS) and cross-site request forgery (CSRF). XSS attacks occur when malicious scripts are injected into web pages, allowing attackers to steal sensitive information or perform unauthorized actions on behalf of the user. CSRF attacks, on the other hand, trick users into unknowingly performing malicious actions on authenticated

websites.

One approach to writing secure code is to sanitize user input and validate it before using it in web applications. This helps prevent the execution of malicious scripts and protects against XSS attacks. Additionally, developers should implement mechanisms to ensure the integrity and authenticity of data exchanged between the browser and the server. This can be achieved through the use of secure communication protocols, such as HTTPS, and by applying encryption and digital signatures to sensitive data.

Understanding browser architecture and writing secure code are essential in maintaining the security of web applications. By being aware of the vulnerabilities present in browsers and following secure coding practices, developers can significantly reduce the risk of browser attacks and protect users' sensitive information.



**EITC/IS/WASF WEB APPLICATIONS SECURITY FUNDAMENTALS DIDACTIC MATERIALS****LESSON: PRACTICAL WEB APPLICATIONS SECURITY****TOPIC: SECURING WEB APPLICATIONS WITH MODERN PLATFORM FEATURES**

Welcome to this educational material on securing web applications with modern platform features. In this material, we will discuss the vulnerabilities that web applications can face and the mechanisms that can be used to protect them.

Web applications are susceptible to various attacks that can compromise user data. At Google, we have run reward programs for over 8 and a half years, where we pay security researchers for disclosing vulnerabilities to us. These vulnerabilities can be found in web applications, server-side infrastructure, Chrome extensions, mobile apps, and more.

The majority of vulnerabilities we have encountered are related to web issues. Specifically, we have seen two main classes of flaws: cross-site scripting (XSS) and injections. XSS occurs when user-controlled data is not properly escaped in server-side responses or when unsafe JavaScript APIs are used. These vulnerabilities are not always obvious and can be introduced unintentionally during the coding process.

When an XSS vulnerability exists in your application, an attacker can navigate a user to their own website, which then redirects the user to your application with a cross-site scripting payload in the URL. When your application server responds to the user's request, it includes the script injected by the attacker. The user's browser will mistakenly believe that the script comes from your application and execute it with full privileges, potentially compromising the user's session and data.

To protect your web applications, modern web browsers have adopted four new web platform mechanisms. These mechanisms provide defenses against common vulnerabilities, including XSS. While we cannot go into detail about these mechanisms in this material, it is important to be aware of their existence and understand how they can be utilized to enhance the security of your applications.

Securing web applications is crucial to protect user data from potential attacks. Understanding the vulnerabilities, such as XSS, and the mechanisms available to mitigate these risks is essential for developers and security engineers. By implementing the appropriate defenses, you can ensure the safety of your users' data and maintain the integrity of your applications.

Web applications security is a crucial aspect of cybersecurity, as it involves protecting sensitive user data and preventing unauthorized access. One of the main concerns in web applications security is the potential for attackers to gain access to user data, modify it, and perform actions that the application normally allows. This can be extremely dangerous and poses a significant threat to user privacy and security.

Insufficient isolation on the web is another major vulnerability that can be exploited by attackers. The web environment has numerous gaps and ways in which applications can interact with each other, without strict boundaries to protect them from interference. One common vulnerability in this category is cross-site request forgery (CSRF), where an attacker tricks a user into performing an unintended action on a website.

To understand CSRF, consider the example of a legitimate form that allows transferring \$10 to a user named Lukas. However, without proper protection against CSRF, an attacker can create a form that interacts with the same application server, leading to the transfer of an infinite amount of money to an account controlled by the attacker. This vulnerability arises because when a browser sends a request to an application server, it includes all the cookies of that application, regardless of who initiated the request. This lack of distinction makes it difficult for the server to determine the request's origin, enabling attackers to exploit the server's confusion and potentially compromise user data.

These examples highlight the importance of securing web applications against such vulnerabilities. Security engineers analyze these bugs to better understand and address them. However, for the purpose of this discussion, it is sufficient to know that these vulnerabilities are increasing in number and severity. New web APIs and changes to the web security model have made these vulnerabilities more likely to occur and more damaging.

Contrary to the assumption that only specific applications face these problems, data from HackerOne, a bug bounty provider, reveals that vulnerabilities related to web applications security are prevalent across various application ecosystems. HackerOne collects vulnerabilities for thousands of companies, and their data shows that the most common vulnerabilities are related to web bugs, including cross-site scripting (XSS) and cross-site request forgery (CSRF).

Similarly, Mozilla's reward program data also highlights the prevalence of web-related vulnerabilities such as cross-site scripting (XSS). These findings demonstrate that web vulnerabilities are not exclusive to specific applications but are a widespread concern across the industry.

To defend against these vulnerabilities, it is essential to implement defensive mechanisms. These mechanisms can be categorized into injection defenses and isolation mechanisms. Injection defenses focus on mitigating vulnerabilities like cross-site scripting. One promising defense mechanism is the implementation of a content security policy (CSP). When properly configured, CSP allows developers to specify which scripts and plugins are allowed to execute on a website. By fine-tuning these settings, developers can significantly reduce the risk of stored and reflected cross-site scripting vulnerabilities.

It is important to note that content security policy serves as a defense-in-depth mechanism. Even with other primary security features in place, there is always a possibility of a bug slipping through. In such cases, a well-implemented content security policy acts as a safety net, protecting users from the consequences of potential vulnerabilities.

Web applications security is a critical aspect of cybersecurity. Insufficient isolation and vulnerabilities like cross-site scripting and cross-site request forgery pose significant threats to user data and application security. However, through the implementation of defensive mechanisms such as content security policy, developers can mitigate the risks associated with these vulnerabilities and enhance the overall security of web applications.

Content Security Policy (CSP) is an important security feature for web applications. It is an HTTP response header that allows you to specify a policy for the content of the response. By applying this policy, you can protect your website from cross-site scripting (XSS) vulnerabilities.

Enabling CSP on your site is simple. You just need to send a response header named "content security policy" with the specified policy. The browser will then read this header and apply the policy to the content of the response. CSP is a client-side security feature and also supports a report-only mode, which is useful for testing CSP without breaking anything.

There are different types of CSP policies, but we recommend using a nonce-based CSP. This type of policy does not require customization for your application and is always the same, except for the nonce value. A nonce is a random token that is not guessable. In a nonce-based CSP, you need to set a nonce attribute on every script tag and ensure that the value of this attribute matches the nonce in the content security policy. Only then will the browser allow the script to execute.

The advantage of a nonce-based CSP is that even if an attacker injects a script tag into your site, they won't know the nonce value. As a result, the browser will block the execution of that script tag, effectively mitigating the consequences of an XSS vulnerability.

However, using a nonce-based CSP can cause issues if you source scripts from CDNs or have third-party JavaScript code. To address this problem, CSP3 introduced a new keyword called "strict-dynamic". This keyword allows already trusted scripts (those with a nonce attribute) to execute code normally and load child scripts without explicitly setting the nonce attribute on them. This unblocks the entire approach of using a nonce-based CSP, even if you don't control all the JavaScript code.

To roll out a nonce-based CSP on your application, follow these three simple steps:

1. Remove CSP blockers: A strong CSP may disable certain dangerous patterns in HTML. Refactor your site to remove these patterns, such as inline event handlers and JavaScript URIs. Instead, register event handlers programmatically through a JavaScript API and remove or refactor JavaScript URIs.
2. Set the nonce attribute on legitimate scripts: Ensure that all the legitimate scripts on your site have the

nonce attribute set. Update your server-side templates to include the nonce attribute.

3. Test and monitor: Test your CSP implementation thoroughly to ensure it doesn't break any functionality. Monitor your application for any security issues and make adjustments as necessary.

By following these steps, you can effectively secure your web applications with modern platform features like CSP.

In web application security, securing web applications with modern platform features is crucial to protect against various vulnerabilities. One important aspect is content security policy (CSP), which helps mitigate cross-site scripting (XSS) attacks.

A nonce-based CSP is a powerful defense against stored and reflected XSS vulnerabilities. Nonces are unique numbers generated by the server and used to validate requests and responses. By including the nonce in the script tag and the response header for the content security policy, the server can enforce the policy effectively. It is recommended to use a nonce-only CSP if there are no third-party JavaScript dependencies, as it offers the highest level of security. However, adopting a nonce-only CSP can be challenging.

To implement CSP, it is essential to set the content security policy string. Google applications commonly use a recommended policy that can be easily adopted. It is advisable to remove the unsafe-eval directive if possible, but caution must be taken to ensure the application does not use JavaScript eval.

There are additional tips and tricks for CSP adoption, such as using CSP hashes instead of nonces for static applications like single-page apps, improving the debug ability of CSP violation reports, and using fallbacks for older browsers. These topics are covered in more detail at [csp.withgoogle.com](https://csp.withgoogle.com), where you can find a guide on deploying CSP on your website.

When deviating from the provided CSP templates, it is crucial to run your CSP through a content security policy evaluator. This tool helps identify any bypasses in the policy, ensuring its effectiveness. The content security policy evaluator is a free and open-source tool.

Nonce-based CSPs offer consistent and application-agnostic protection against XSS attacks. They are more secure than whitelist-based approaches and do not suffer from whitelist bypasses. By adopting nonce-based CSP and following best practices, web applications can defend against stored and reflected XSS vulnerabilities.

While nonce-based CSPs address stored and reflected XSS, DOM-based XSS remains a concern. DOM-based XSS vulnerabilities occur when user-controlled strings are converted into code using dangerous DOM APIs like innerHTML. These vulnerabilities typically occur on the client-side and can be challenging to detect due to the large number of dangerous DOM APIs and the distance between the source and sink of the vulnerability.

To address DOM-based XSS, trusted types are introduced. Trusted types provide a mechanism to defend applications against these vulnerabilities. They help developers identify and validate user-controlled strings, reducing the risk of introducing DOM-based XSS vulnerabilities. With the increasing popularity of client-side code and modern frameworks, trusted types play a crucial role in enhancing web application security.

The concept of securing web applications with modern platform features involves the use of trusted types. This approach allows for the assignment of typed objects instead of strings to dangerous DOM APIs. By using trusted types, developers are required to use trusted HTML objects when assigning values to sinks like innerHTML. This concept can be enforced by various means such as compilers, linters, pre-submit checks, and even by the browser itself through the use of content security policies.

When a content security policy with the trusted types directive is set, the browser will disallow string assignments to dangerous DOM APIs. For example, the previous vulnerability of DOM-based cross-site scripting will no longer work because string assignments are blocked. Instead, developers must use typed objects that can only be created for a trusted types policy. Additionally, there is a report-only mode available for testing applications with trusted types, where string assignments are still allowed but console warnings and CSP violation reports are generated.

To create a trusted types object, developers can utilize the trusted types API, which is supported by browsers or

can be polyfilled if necessary. The process involves invoking the create policy method and specifying the name of the policy and the policy function. For instance, a custom sanitizer can be used to create a trusted HTML object from a string. Once the trusted HTML object is created, it can be assigned to the DOM sink, such as innerHTML. However, it is crucial to list all the policies in the content security policy's trusted types directive to allow the creation of these trusted types objects.

There is a special case called the default policy, where string assignments are not blocked but instead piped through the default policy. This allows for the application of custom logic, such as using a sanitizer or logging string assignments. This approach helps identify insecure string assignments in the application, which can then be converted into trusted object assignments.

Trusted types significantly reduce the attack surface of web applications by channeling all data flows through policies when assigning values to dangerous DOM APIs. This concentration of security critical code simplifies security reviews and minimizes the trusted code base. Trusted types can be checked at compile time, and runtime checks can be enforced by setting the trusted types content security policy. By ensuring secure and restricted access to policies, developers can effectively eliminate DOM-based cross-site scripting vulnerabilities.

Currently, trusted types are available as Chrome Origin Trials, but they can also be polyfilled using the Trusted Types GitHub page. This page provides detailed explanations and instructions for implementing trusted types. Developers are encouraged to explore this feature and try implementing a default policy to log insecure DOM assignments in their applications.

Securing web applications with modern platform features involves the use of trusted types to assign typed objects instead of strings to dangerous DOM APIs. This approach reduces the attack surface, concentrates security critical code, and simplifies security reviews. By enforcing trusted types through content security policies, developers can effectively mitigate DOM-based cross-site scripting vulnerabilities.

Web applications security is a critical aspect of cybersecurity. One of the key vulnerabilities that web applications face is cross-site scripting (XSS). In this didactic material, we will explore practical ways to secure web applications using modern platform features.

One effective approach to mitigating XSS vulnerabilities is by implementing a trusted types policy in combination with a content security policy (CSP). These two technologies work well together and can be set through a single CSP. By implementing these measures, web applications can prevent and mitigate the majority of XSS vulnerabilities.

However, securing web applications involves more than just protecting against XSS. There are other high-risk bugs that need to be addressed. One such bug is Cross-Site Request Forgery (CSRF), where an attacker can make requests to an application using the user's cookies. This can lead to information leakage or unauthorized actions. Another type of attack involves an attacker opening a new window to the application, limiting direct access but still allowing interaction.

To effectively defend against these attacks, it is important to understand the concepts of origins and sites. In the context of security, two URLs are considered to be same origin if they have the same scheme, host name, and port. URLs that share the same scheme and registerable domain are considered same site, implying some level of trust. URLs that are neither same origin nor same site are considered cross-site, indicating a lack of trust.

With this understanding, let's explore some defensive mechanisms. One such mechanism is fetch metadata request headers. These headers provide additional context to HTTP requests, allowing servers to make security decisions. For example, the SEC-Fetch-Site header indicates which site made the request, distinguishing between same origin and cross-site requests. The SEC-Fetch-Mode header helps differentiate between sub resource requests and navigations, while another header indicates if there was an explicit user interaction leading to a navigation.

By leveraging these fetch metadata request headers, servers can differentiate between same origin and cross-site requests, enabling them to make informed security decisions.

Securing web applications involves implementing measures to protect against XSS vulnerabilities, such as

trusted types policies and content security policies. Additionally, understanding the concepts of origins and sites is crucial for defending against other high-risk bugs. By utilizing fetch metadata request headers, servers can enhance their security by gaining more context about incoming requests.

On the topic of securing web applications with modern platform features, one important aspect to consider is the use of fetch metadata request headers. When a request is made from your own application to your own application server using the Fetch API, the header will indicate "same origin," which means that the request can be considered trusted by the server. However, if an external web page makes the same request, the browser will identify it as a cross-site request, giving the server the opportunity to reject it. In the past, without fetch metadata request headers, these two requests would have been indistinguishable to the server, making it difficult to differentiate between trusted and untrusted requests. This simple module allows requests that are same origin and same site, while banning cross-site requests, providing protection against cross-origin attacks on sub-resources.

The module operates by allowing all requests without the headers for backward compatibility. Then, it checks if the request comes from a trusted site and allows it if it does. Additionally, for potentially dangerous cross-site requests, it verifies if the request is a navigation, as cross-site links should still be allowed for users to access the application from other sites. By implementing logic similar to the module described, you can adopt these protections based on fetch metadata headers. Initially, it is recommended to implement the logic in reporting mode without enforcing the restrictions. Once any compatibility issues are resolved, you can start enforcing the security measures. This feature is already implemented in Chrome and will be available in the next stable version.

Another related aspect is SameSite cookies. These cookies are an important tool for security, but it can be challenging to understand the consequences of setting the SameSite attribute in your cookies. Fetch metadata headers can help you test the deployment of SameSite cookies in your application, providing information to ensure compatibility.

Lastly, protecting windows from being referenced by cross-site pages is crucial. When an evil web page opens a window to your application, it holds a reference to your window, allowing it to perform unexpected actions. To mitigate this risk, the cross-origin opener policy response header can be set to either "same origin" or "same site." This causes windows that are not same origin or same site to lose the reference, preventing unwanted actions such as sending messages, navigating to attacker-controlled sites, or accessing frames within your application. This security measure is simple yet powerful in protecting against potential vulnerabilities.

Implementing the cross-origin opener policy may also provide additional benefits. Browsers can put your application in a separate process, even if they don't have full site isolation, providing protection against speculative execution attacks.

The web platform now offers powerful security mechanisms that can significantly enhance the security of web applications. These mechanisms include protection against injections, Content Security Policy (CSP), trusted types, as well as isolation mechanisms like fetch metadata request headers and the cross-origin opener policy. By adopting these security features on your sites, you can greatly improve the safety of your users against a wide range of vulnerabilities commonly found on the web.

Web applications are an integral part of our digital lives, enabling us to perform various tasks and access information online. However, these applications can also be vulnerable to cyber attacks if not properly secured. In this didactic material, we will explore the fundamentals of web application security, with a focus on securing web applications using modern platform features.

Web application vulnerabilities can pose significant risks to the confidentiality, integrity, and availability of sensitive data. Therefore, it is crucial to understand and address these vulnerabilities to protect against potential cyber threats.

One effective approach to securing web applications is to leverage modern platform features. These features provide built-in security mechanisms that can help mitigate common vulnerabilities. By utilizing these features, developers can enhance the security of their web applications without relying solely on external security measures.

Some of the modern platform features that can be utilized for web application security include:

1. Content Security Policy (CSP): CSP allows developers to define the trusted sources of content that a web application can load. By specifying the allowed sources for scripts, stylesheets, and other resources, CSP can help prevent cross-site scripting (XSS) attacks and other code injection vulnerabilities.
2. HTTP Strict Transport Security (HSTS): HSTS ensures that web browsers only communicate with a web application over secure HTTPS connections. This feature helps protect against man-in-the-middle attacks and ensures that sensitive data is transmitted securely.
3. Cross-Origin Resource Sharing (CORS): CORS enables controlled access to resources on a web page from different domains. By defining the allowed origins and HTTP methods, developers can prevent unauthorized cross-origin requests and protect against cross-site request forgery (CSRF) attacks.
4. SameSite Cookies: SameSite cookies allow developers to specify whether cookies should be sent in cross-site requests. By setting the SameSite attribute to "Strict" or "Lax," developers can prevent CSRF attacks that exploit the browser's automatic inclusion of cookies in cross-site requests.
5. Subresource Integrity (SRI): SRI allows developers to ensure the integrity of externally hosted resources, such as JavaScript libraries or stylesheets. By including a cryptographic hash of the resource in the HTML code, SRI can detect and block any modifications or tampering attempts.

In addition to these platform features, developers should also follow secure coding practices, such as input validation, output encoding, and proper authentication and authorization mechanisms. Regular security assessments and penetration testing are essential to identify and address any vulnerabilities that may exist in web applications.

By implementing these modern platform features and adopting secure coding practices, developers can significantly enhance the security of web applications, reducing the risk of cyber attacks and protecting sensitive user data.