

CEH – Certified Ethical Hacker Exam Review

Contents

Chapter 1: Getting Started: Essential Knowledge.....	2
1. Residual Risk.....	2
2. CEH methodology is laid out this way:.....	2
3. BIA.....	3
4. Incident Response Phases.....	3
5. Gray box test.....	4
6. MAC.....	4
7. The three-way handshake: “SYN, SYN/ACK, ACK”	4
8. ALE.....	5
9. A white hat is attempting a white-box test.	5
10. Audit trails.....	5
11. The Privacy Act.....	5
12. Four terms make up the Common Criteria Process.....	6
13. Risk Management – Risk Mitigation	7
14. Scanning and enumeration	7
15. SOX and other Laws	7
16. Logical or technical control.....	8
17. PCI-DSS.....	8
18. Test Types – Gray Box	9
19. Maintaining access.....	10
20. Policy, Standards, Procedures, and Guidelines.....	10
21. Incident Management.....	10
22. Response to child porn during a Pen test	11
23. An Intranet	11
24. Threats and Vulnerabilities Externally	12
25. Operating system (OS) attacks target common mistakes.....	12
Chapter 2: Reconnaissance: Information Gathering for the Ethical Hacker.....	12
1. Nslookup	12
2. Message an invalid email address	12
3. Email header information	12

4. Google hacks - allintitle:CEH V10	13
5. Traceroute vs. Tracert.....	13
6. Active vs. passive Footprinting.....	13
7. Nslookup	14
8. Split DNS.....	14
9. CNAME and other record types	14
10. Protecting against DNS enumeration	15
11. Passive Footprinting.....	15
12. SOA Record	16
13. DNS Zone Trans – TCP 53	16
14. Nslookup	16
15. DNS Poisoning.....	17
16. SOA.....	17
17. Active Directory–integrated DNS server.....	17
18. EDGAR Database	17
19. Traceroute.....	18
20. Improving DNS Security	18
21. A zone file consists of which records? (Choose all that apply.)	19
22. OSRFramework	20
23. Google operators	21
24. Archive.org.....	21
25. Computer Security Incident Response Team (CSIRT)?.....	21
26. Regional Registries.....	21
References	22

Chapter 1: Getting Started: Essential Knowledge

Key Points

1. Residual Risk

Ensure that any remaining risk is residual or low and accept the risk.

2. CEH methodology is laid out this way:

- Reconnaissance (Footprinting),
- Scanning and enumeration,
- Gaining access,

- Escalating privileges,
- Maintaining access, and
- Covering tracks.

3. BIA

A business impact analysis (BIA) the organization looks at all the systems and processes in use and determines which ones are critical to continued operation. Additionally, the assessor (the person or company conducting the analysis) will look at all the existing security architecture and make an evaluation on the likelihood of any system or resource being compromised. Part of this is assigning values to systems and services, determining the maximum tolerable downtime (MTD) for any, and identifying any overlooked vulnerabilities.

4. Incident Response Phases

In the preparation phase, your IR (incident response) team should be preparing for an incident. Preparation includes lots of things—some of which are mentioned here. But virtually anything you can think of that does not involve actions taken during the incident belongs here. Training, exercises, and policies are all examples.

IR phases can be different depending on whom you ask and what the moon phase is, but generally IR is broken down into six phases:

1. Preparation,
2. Identification,
3. Containment,
4. Eradication,
5. Recovery, and
6. Lessons learned.

Preparation we already covered.

Identification refers to the steps taken to verify it is actually an incident, and all the information surrounding that—source, destination(s), exploit used, malware used, and so on.

Containment is the step used to cordon off the infected system(s) and to prevent any further spread of infection or attack.

Eradication refers to steps taken to remove the malware (or other attack-related residuals, such as backdoors).

Recovery involves the steps taken to rebuild and restore the system(s) and network to pre-attack status (with better security, I might add).

Finally, lessons learned is exactly what it sounds like, and should feed right back into your organization's preparation phase.

5. Gray box test

A gray-box test is designed to replicate an inside attacker. Otherwise known as the partial knowledge attack, the idea is to simulate a user on the inside who might know a little about the network, directory structure, and other resources in your enterprise.

You will probably find this one to be the most enlightening attack in out-briefing your clients in the real world—it is amazing what you can get to when you are a trusted, inside user.

You will often find in the real world that gray-box testing can also refer to a test where any inside information is given to a pen tester—you do not necessarily need to be a fully knowledgeable inside user. In other words, if you have usable information handed to you about your client, you are performing gray-box testing.

6. MAC

Access control is defined as the selective restraint of access to a resource, and there are several overall mechanisms to accomplish this goal.

- Mandatory access control (MAC) is one type that constrains the ability of a subject to access or perform an operation on an object by assigning and comparing “sensitivity labels.” Suppose a person (or a process) attempts to access or edit a file. With MAC, a label is placed on the file indicating its security level. If the entity attempting to access it does not have that level, or higher, then access is denied. With mandatory access control, security is centrally controlled by a security policy administrator, and users do not have the ability to override security settings.
- This should not be confused with role-based access control (RBAC) systems, which may actually use MAC to get the job done. The difference is in whether the information itself has a labeled description or whether the person accessing it has their own label. For example, in a classified area, the information classified as Top Secret will have a label on it identifying it as such, while you, as an auditor, will have your own clearance and need-to-know label allowing you to access certain information.
- MAC is a property of an object; RBAC is a property of someone accessing an object.
- Discretionary access control (DAC) allows the data owner, the user, to set security permissions for the object. If you are on a Windows machine right now, you can create files and folders and then set sharing and permissions on them as you see fit.

7. The three-way handshake: “SYN, SYN/ACK, ACK”

In step 1, the host sends a segment to the server, indicating it wants to open a communications session. Inside this segment, the host turns on the SYN flag and sets an initial sequence number (any random 32-bit number).

When the recipient gets the segment, it crafts a segment in response to let the host know it is open and ready for the communications session. It does this by turning on the SYN and ACK flags, acknowledging the initial sequence number by incrementing it, and adding its own unique sequence number.

Lastly, when the host gets this response back, it sends one more segment before the comm channel opens. In this segment, it sets the ACK flag and acknowledges the other’s sequence number by incrementing it.

For example, suppose Host A is trying to open a channel with Server B. In this example, Host A likes the sequence number 2000, while Server B likes 5000. The first segment would look like this: SYN=1, ACK=0, ISN=2000. The response segment would look like this: SYN=1, ACK=1, ISN=5000, ACK NO=2001. The third and final segment would appear this way: SYN=0, ACK=1, SEQ NO=2001, ACK NO=5001.

8. ALE

When performing business impact analysis (or any other value analysis for that matter), the annualized loss expectancy (ALE) is an important measurement for every asset.

To compute the ALE, multiply the annualized rate of occurrence (ARO) by the single loss expectancy (SLE).

The ARO is the frequency at which a failure occurs on an annual basis. In this example, servers fail once every five years, so the ARO would be 1 failure / 5 years = 20 percent.

9. A white hat is attempting a white-box test.

Start with what kind of hacker he is. He is hired under a specific agreement, with full knowledge and consent of the target, thus making him a white hat. Second, to address what kind of test he is performing, simply look at what he knows about the system. In this instance, he has no prior knowledge at all (apart from the agreement), thus making it a black-box test.

10. Audit trails

A detective control is an effort used to identify problems, errors, or (in the case of post-attack discovery) cause or evidence of an exploited vulnerability—and an audit log or trail is a perfect example. Ideally, detective controls should be in place and working such that errors can be corrected as quickly as possible. Many compliance laws and standards (the Sarbanes-Oxley Act of 2002 is one example) mandate the use of detective controls.

11. The Privacy Act

As part of a pen test on a U.S. government system, you discover files containing Social Security numbers and other sensitive personally identifiable information (PII). You are asked about controls placed on the dissemination of this information. Which of the following acts should you check?

The Privacy Act of 1974 protects information of a personal nature, including Social Security numbers. The Privacy Act defines exactly what “personal information” is, and it states that government agencies cannot disclose any personal information about an individual without that person’s consent. It also lists 12 exemptions for the release of this information (for example, information that is part of a law enforcement issue may be released).

Keep in mind that the Privacy Act generally will define the information that is not available to you in and after a test.

Dissemination and storage of privacy information needs to be closely controlled to keep you out of hot water. As a side note, how you obtain PII is oftentimes just as important as how you protect it once discovered. In your real-world adventures, keep the Wiretap Act (18 U.S. Code Chapter 119—Wire and Electronic Communications Interception and Interception of Oral Communications) and others like it in mind.

The Federal Information Security Management Act (FISMA) is not designed to control the dissemination of PII or sensitive data. Its primary goal is to ensure the security of government systems by promoting a standardized approach to security controls, implementation, and testing. The act requires government agencies to create a security plan for their systems and to have it “accredited” at least once every three years.

The PATRIOT Act is not an effort to control personal information. Its purpose is to aid the U.S. government in preventing terrorism by increasing the government’s ability to monitor, intercept, and maintain records on almost every imaginable form of communication. As a side effect, it has also served to increase observation and prevention of hacking attempts on many systems.

The Freedom of Information Act was not designed to tell you what to do with information. Its goal is to define how you can get information—specifically information regarding how your governments work. It does not necessarily help you in hacking, but it does provide a cover for a lot of information. Anything you uncover that could have been gathered through the Freedom of Information Act is considered legal and should be part of your overall test.

12. Four terms make up the Common Criteria Process

What term contains seven levels used to rate the target?

Common Criteria is an international standard of evaluation of Information Technology (IT) products. Per the website (<https://www.commoncriteriaportal.org/>) Common Criteria ensures evaluations and ratings “are performed to high and consistent standards and are seen to contribute significantly to confidence in the security of those products and profiles.”

The EAL (Evaluation Assurance Level) is made up of seven levels, which are used to rate a product after it has been tested.

The current EAL levels are as follows:

- EAL1: Functionally tested
- EAL2: Structurally tested
- EAL3: Methodically tested and checked
- EAL4: Methodically designed, tested, and reviewed
- EAL5: Semi-formally designed and tested
- EAL6: Semi-formally verified, designed, and tested
- EAL7: Formally verified, designed, and tested

TOE is the target of evaluation—the system or product actually being tested.

ST is the security target—the documentation describing the target of evaluation and any security requirements.

PP is the protection profile—a set of security requirements for the product type being tested.

13. Risk Management – Risk Mitigation

An organization's leadership is concerned about social engineering and hires a company to provide training for all employees. How is the organization handling the risk associated with social engineering?

When it comes to risks, there are four different methods of attempting to deal with them.

In **risk mitigation**, steps are taken to reduce the chance that the risk even will occur, and in this example that is exactly what is happening. Training on social engineering should help reduce the likelihood an employee will fall victim (real-life concerns on this notwithstanding—we are talking about test questions here).

The **acceptance of risk** means the organization understands the risk is there, but they do not do anything about it. Why would a company take this action? Perhaps the chance a threat agent will (or even can) exploit the risk is so low it makes the effort to mitigate it pointless. Or it could be the cost to mitigate simply costs more than any damage or recovery from exploitation in the first place. In any case, if the organization does nothing, they are accepting risk.

Avoidance of risk means the organization takes steps to eliminate the service, action, or technology altogether. In other words, the risk is deemed so great the company would rather do without the asset or service in the first place. In the case of social engineering, unless the organization can work without employees, avoiding this risk is nearly impossible.

Transferring risk occurs when the organization puts the burden of risk on another party. For example, the company might hire an insurance company to pay off in the event a risk is exploited.

14. Scanning and enumeration

The scanning and enumeration phase is where you will use things such as ping sweeps to discover available targets on the network. This step occurs after reconnaissance. In this step, tools and techniques are actively applied to information gathered during recon to obtain more in-depth information on the targets. For example, reconnaissance may show a network subnet to have 500 or so machines connected inside a single building, whereas scanning and enumeration would discover which ones are Windows machines and which ones are running FTP.

15. SOX and other Laws

Which of the following was created to protect shareholders and the general public from corporate accounting errors and fraudulent practices, and to improve the accuracy of corporate disclosures?

The Sarbanes-Oxley Act (SOX; <https://www.sec.gov/about/laws.shtml#sox2002>) introduced major changes to the regulation of financial practice and corporate governance in 2002 and is arranged into 11 titles. SOX mandated several reforms to enhance corporate responsibility, enhance financial disclosures,

and combat corporate and accounting fraud, and it created the “Public Company Accounting Oversight Board,” also known as the PCAOB, to oversee the activities of the auditing profession.

The Gramm-Leach-Bliley Act (GLBA; <https://www.ftc.gov/tips-advice/business-center/privacy-and-security/gramm-leach-bliley-act>) requires financial institutions—companies that offer consumers financial products or services such as loans, financial or investment advice, or insurance—to explain their information-sharing practices to their customers and to safeguard sensitive data. Under the Safeguards Rule, financial institutions must protect the consumer information they collect. GLBA protects the confidentiality and integrity of personal information collected by financial institutions.

The Health Insurance Portability and Accountability Act (HIPAA; www.hhs.gov/hipaa/) was designed to protect the confidentiality of private health information. HIPAA contains privacy and security requirements and provides steps and procedures for handling and protecting private health data.

16. Logical or technical control

A logical (or technical) control is one used for identification, authentication, and authorization. It can be embedded inside an operating system, application, or database management system. A security token (such as RSA’s SecureID) can provide a number that changes on a recurring basis that a user must provide during authentication, or it may provide a built-in number on a USB device that must be attached during authentication.

A physical control is something, well, physical in nature, such as a lock or key or maybe a guard.

17. PCI-DSS

The Payment Card Industry Data Security Standard (PCI-DSS) is a security standard for organizations that handle credit cards. A council including American Express, JCB, Discover, MasterCard, and Visa developed standards for the protection and transmission of card data to reduce credit card fraud. It is administered by the Payment Card Industry Security Standards Council. Validation of compliance is performed annually.

The standard is composed of 12 requirements:

- Requirement 1: Install and maintain firewall configuration to protect data.
- Requirement 2: Remove vendor-supplied default passwords and other default security features.
- Requirement 3: Protect stored data.
- Requirement 4: Encrypt transmission of cardholder data.
- Requirement 5: Install, use, and update AV (antivirus).
- Requirement 6: Develop secure systems and applications.
- Requirement 7: Use “need to know” as a guideline to restrict access to data.
- Requirement 8: Assign a unique ID to each stakeholder in the process (with computer access).
- Requirement 9: Restrict any physical access to the data.
- Requirement 10: Monitor all access to data and network resources holding, transmitting, or protecting it.
- Requirement 11: Test security procedures and systems regularly.
- Requirement 12: Create and maintain an information security policy.

The Trusted Computer System Evaluation Criteria (TCSEC), also known as the Orange Book, was created by the Department of Defense (DoD) and defines and provides guidance on evaluating access controls within a system. TCSEC defines four levels of validation:

- verified protection
- mandatory protection
- discretionary protection, and
- minimal protection

ISO 27002 (www.iso27001security.com/html/27002.html) is an “information security standard published by ISO and the International Electrotechnical Commission (IEC) that recommends security controls based on industry best practices.” This standard includes 13 objectives, ranging from structure, risk assessment, and policy to access controls, human resources security, and compliance.

18. Test Types – Gray Box

As part of the preparation phase for a pen test you are participating in, the client relays their intent to discover security flaws and possible remediation. They seem particularly concerned about internal threats from the user base. Which of the following best describes the test type the client is looking for?

There are three types of tests—**white**, **black**, and **gray**—with each designed to test a specific threat.

- **White** tests the internal threat of a knowledgeable systems administrator or an otherwise elevated privilege level user.
- **Black** tests external threats with no knowledge of the target.
- **Gray** tests the average internal user threat to expose potential security problems inside the network.

19. Maintaining access

In which phase of the attack would a hacker set up and configure “zombie” machines?

Zombies are basically machines the hacker has commandeered to do his work for him. If the attacker is really good, the owners of the zombie machines do not even know their machines have been drafted into the war. There is a bajillion method for maintaining access on a machine you’ve already compromised and maintaining that access does not necessarily mean the system will be used as a zombie—you could, for example, simply want to check in from time to time to see what new juicy information the user has decided to leave in a file or folder for you, or to check on new logins, credentials, and so on.

20. Policy, Standards, Procedures, and Guidelines

Which of the following should not be included in a security policy?

Policy is a high-level document that does not get down and dirty into technical details/specifications and is intended to improve awareness. Policies are mandatory, generally short, and easy to understand, providing everyone with the rules of the road.

Standards are mandatory rules designed to support a policy, and they must include one or more specifications for hardware, software, or behavior.

Procedures are step-by-step instructions for completing a task.

Guidelines are not mandatory, but rather are recommendations for accomplishing a goal or on how to act in each situation.

21. Incident Management

Which of the following is best defined as a set of processes used to identify, analyze, prioritize, and resolve security incidents?

Incident management is the process of dealing with incidents and generally always has the same features/steps:

- identify the problem or root cause,
- analyze and research the issue,
- contain the malicious effort,
- eradicate the effort, and

- resolve any damage caused.

ECC defines the process as having eight steps:

1. Preparation
2. Detection and Analysis
3. Classification/Prioritization
4. Notification
5. Containment
6. Forensic Investigation
7. Eradication and Recovery, and
8. Post-incident Activities.

The incident response team (IRT) is charged with handling this process.

22. Response to child porn during a Pen test

During an assessment, your pen test team discovers child porn on a system. Which of the following is the appropriate response?

First and foremost, in the real world, discovery of something that you think might be illegal activity puts you and your team in a very, very tricky spot. Should you accuse fill-in-the-blank of a crime and involve the authorities, you could be setting yourself up for lawsuits and all sorts of trouble. On the other hand, if you ignore it, you might be found complicit, or at the very least negligent.

In the real world, the answer is to make sure your scope agreement advises you and the client of your duty regarding potential criminal activity found during the scope of your investigation.

No guessing is allowed—it better be iron-clad evidence, obvious to all, or you are in a world of hurt.

Lastly, what potentially illegal activity you discover may determine your response regardless of ROE (Rules of Engagement). If you discover child porn, you could be guilty of a crime for not reporting it, which is not necessarily true for many other crimes. For example, if you witness someone breaking into a house across your street or were performing a pen test and reasonably suspected someone had already compromised the network, you are not compelled by law, in most states, to notify authorities.

However, if you witness bodily harm, you likely would be compelled by law in most states. Speaking purely academically, it is clear cut and will be so on your exam. In the real world the true answer is to know the laws regarding your testing very well, and make sure your team has a good lawyer.

23. An Intranet

An intranet can be thought of, for testing purposes, as your own happy little networking safe space. It is protected from outside attacks and interference by the DMZ and all the layers of security on the outside. Internally, you do not assign loads of heavy security restrictions, because as security increases, usability, and functionality decrease. If your organization's users are on the intranet, you want them as productive as possible.

24. Threats and Vulnerabilities Externally

A machine in your environment uses an open X-server to allow remote access. The X-server access control is disabled, allowing connections from almost anywhere and with little to no authentication measures. Which of the following are true statements regarding this situation? (Choose all that apply.)

- An external vulnerability can take advantage of the misconfigured X-server threat.
- An external threat can take advantage of the misconfigured X-server vulnerability.

A threat is any agent, circumstance, or situation that could potentially cause harm or loss to an IT asset. In this case, the implication is the threat is an individual (hacker) either inside or outside the network.

A vulnerability is any weakness, such as a software flaw or logic design, that could be exploited by a threat to cause damage to an asset. In both these answers, the vulnerability—the access controls on the X-server are not in place—can be exploited by the threat, whether internal or external.

25. Operating system (OS) attacks target common mistakes

While performing a pen test, you find success in exploiting a machine. Your attack vector took advantage of a common mistake—the Windows 7 installer script used to load the machine left the administrative account with a default password. Which attack did you successfully execute?

Operating system (OS) attacks target common mistakes many people make when installing operating systems (for instance, accepting and leaving all the defaults). Examples usually include things such as administrator accounts with no passwords, ports left open, and guest accounts left behind. Another OS attack you may be asked about deals with versioning. Operating systems are never released fully secure and are consistently upgraded with hotfixes, security patches, and full releases. The potential for an old vulnerability within the enterprise is always high.

Chapter 2: Reconnaissance: Information Gathering for the Ethical Hacker

1. Nslookup

```
> server ADNS_SERVER
...
> set type=HINFO
> ATARGET_SYSTEM
...
```

2. Message an invalid email address

A bogus internal address has the potential to provide more information about the internal servers used in the organization, including IP addresses and other pertinent details.

3. Email header information

From the partial e-mail header provided, which of the following represents the true originator of the e-mail message?

☑ **C.** E-mail headers are packed with information showing the entire route the message has taken, and I can guarantee you'll see at least one question on your exam about them. You'll most likely be asked to identify the true originator—the machine (person) who sent the e-mail in the first place (even though in the real world with proxies and whatnot to hide behind, it may be impossible). This is clearly shown in line 9: Received: from SOMEONEComputer [217.88.53.154] (helo=[SOMEONEcomputer]). But don't just study and rely on that one section. Watch the entire trek the message takes and make note of the IPs along the way.

4. Google hacks - allintitle:CEH V10

You are looking for pages with the terms CEH and V10 in their title. Which Google hack is the appropriate one?

The Google search operator allintitle searches for pages that contain the string, or strings, you specify. It also allows for the combination of strings in the title, so you can search for more than one term within the title of a page.

The operator inurl looks only in the URL of the site, not the page title. In this example, the search might bring you to a page like this: http://anyplace.com/apache_Version/pdfs.html.

5. Traceroute vs. Tracert

You are on a Cisco router and want to identify the path a packet travels to a specific IP. Which of the following is the best command choice for this?

The **tracert** command will work on a Windows system, but not on a Cisco device.

6. Active vs. passive Footprinting

Active

- Calling the company's help desk line
- Employing passive sniffing

Passive

- Dumpster diving
- Reviewing financial sites for company information
- Clicking links within the company's public website

7. Nslookup

Examine the following command sequence:

```
C:\> nslookup
Default Server:  ns1.anybiz.com
Address:  188.87.99.6
> set type=HINFO
> someserver
Server:  resolver.anybiz.com
Address:  188.87.100.5
Someserver.anybiz.com CPU=Intel Quad Chip OS=Linux 2.8
```

Which of the following statements best describes the intent of the command sequence?

The operator is enumerating a system named **someserver**.

To perform a zone transfer, you would use the **set type=any** command and then **ls -d anybiz.com**

Checking for name servers in the domain would require the **set type=NS** command.

8. Split DNS

An organization has a DNS server located in the DMZ and other DNS servers located on the intranet. What is this implementation commonly called?

The idea behind split DNS is pretty simple: create two zones for the same domain, with one just for the internal network while the other is used by any external networks. Internal hosts are directed to the internal domain name server. Separating the domain servers greatly restricts the Footprinting an attacker can perform from the outside.

9. CNAME and other record types

CNAME records provide for aliases within the zone on that name. For instance, your server might be named mattserver1.matt.com. A sample DNS zone entry to provide HTTP and FTP access might look like this:

NAME	TYPE	VALUE
ftp.matt.com.	CNAME	mattserver.matt.com
www.matt.com	CNAME	mattserver.matt.com
mattserver1.matt.com.	A	202.17.77.5

A. NS

B. SOA

C. MX

D. PTR

E. CNAME

10. Protecting against DNS enumeration

Ensuring there are no A records for internal hosts on the public-facing name server

If your company has a publicly facing website, it follows that a name server somewhere has to answer lookups in order for your customers to find the site. That name server, however, does not need to provide lookup information to internal machines. Of the choices provided, as silly as it seems to point out, ensuring there are no A records (those used to map hostnames to an IP address) on the external name server is a good start.

11. Passive Footprinting

An ethical hacker searches for IP ranges owned by the client, reads news articles, observes when bank employees arrive and leave from work, searches the client's job postings, and visits the client's dumpster. Which of the following is a true statement?

All of the actions are passive Footprinting.

12. SOA Record

12. Examine the following SOA record:

```
@    IN  SOA  DNSRV1.somebiz.com.  postmaster.somebiz.com. (
200408097      ; serial number
                        3600      ; refresh   [1h]
                        600       ; retry    [10m]
                        86400     ; expire   [1d]
7200 )          ; min TTL   [2h]
```

If a secondary server in the enterprise is unable to check in for a zone update within an hour, what happens to the zone copy on the secondary?

The zone copy is unchanged.

In this question, the key portion you are looking for is the TTL (Time-To-Live) value at the bottom, which is currently two hours (7200 seconds). This sets the time a secondary server has to verify its records are good. If it cannot check in, this TTL for zone records will expire, and they will all be dumped. Considering, though, this TTL is set to two hours and the question states it has been only one hour since update, the zone copy on the secondary will remain unchanged.

13. DNS Zone Trans – TCP 53

TCP 53 is the default protocol and port number for zone transfers. DNS actually uses both TCP and UDP to get its job done, and if you think about what it is doing, they make sense in particular circumstances.

A name resolution request and reply? Small and quick, so use port 53 on UDP. A zone transfer, which could potentially be large and requires some insurance it all gets there? Port 53 on TCP is the answer.

14. Nslookup

Examine the following command-line entry:

```
C:\>nslookup
  Default Server:  ns1.somewhere.com
  Address:  128.189.72.5
> set q=mx
>mailhost
```

Which statements are true regarding this command sequence?

- Nslookup is in interactive mode. And
- The output will show all mail servers in the zone somewhere.com.

Nslookup runs in one of two modes—interactive and noninteractive. Noninteractive mode is simply the use of the command followed by an output. For example, `nslookup www.google.com` will return the IP address your server can find for Google.

Interactive mode is started by simply typing `nslookup` and pressing ENTER. Your default server name will display, along with its IP address, and a caret (>) will await entry of your next command. In this scenario, we have entered interactive mode and set the type to MX, which we all know means “Please provide me with all the mail exchange servers you know about.”

15. DNS Poisoning

Joe accesses the company website, `www.anybusi.com`, from his home computer and is presented with a defaced site containing disturbing images. He calls the IT department to report the website hack and is told they do not see any problem with the site—no files have been changed, and when accessed from their terminals (inside the company), the site appears normally. Joe connects over VPN into the company website and notices the site appears normally. Which of the following might explain the issue?

DNS poisoning makes the most sense here. In many cases (such as mine right here in my own work-from-home office), a VPN connection back to the company forces you to use the company DNS instead of your local resolution. In this example, Joe’s connection from home uses a different DNS server for lookups than that of the business network. It is entirely possible someone has changed the cache entries in his local server to point to a different IP than the one hosting the real website—one that the hackers have set up to provide the defaced version. The fact the web files haven’t changed, and it seems to be displaying just fine from inside the network also bears this out. If it turns out Joe’s DNS modification is the only one in place, there is a strong likelihood that Joe is being specifically targeted for exploitation—something Joe should take very seriously. Lastly, the `HOSTS` and `LMHOSTS` files can also play a big role in this kind of scenario—however, if an attacker already has that kind of access to Joe’s computer, he has bigger problems than the corporate website.

16. SOA

One way to mitigate against DNS poisoning is to restrict or limit the amount of time records can stay in cache before they are updated. Which DNS record type allows you to set this restriction?

The SOA record holds all sorts of information, and when it comes to DNS poisoning, the TTL is of primary interest. The shorter the TTL, the less time records are held in cache. While it won’t prevent DNS poisoning altogether, it can limit the problems a successful cache poisoning attack cause.

17. Active Directory–integrated DNS server.

If you have a Windows Active Directory (AD) network, having AD-integrated DNS servers has some great advantages. For example (and directly from Microsoft, I might add), “with directory-integrated storage, dynamic updates to DNS are conducted based upon a multimaster update model. In this model, any authoritative DNS server, such as a domain controller running a DNS server, is designated as a primary source for the zone. Because the master copy of the zone is maintained in the Active Directory database, which is fully replicated to all domain controllers, the zone can be updated by the DNS servers operating at any domain controller for the domain.” Zones are also replicated and synchronized to new domain controllers automatically whenever a new one is added to an Active Directory domain, and directory replication is faster and more efficient than standard DNS replication.

But having an Active Directory server facing externally is a horrible idea.

18. EDGAR Database

The EDGAR Database—<https://www.sec.gov/edgar.shtml>

—holds various competitive intelligence information on businesses and is an old favorite of EC-Council. Per the website, “All companies, foreign and domestic, are required to file registration statements, periodic reports, and other forms electronically through EDGAR. Anyone can access and download this information for free. Here you’ll find links to a complete list of filings available through EDGAR and instructions for searching the EDGAR database.” Finally, one more note on EDGAR and the SEC: They have purview only over publicly traded companies. Privately held companies are not regulated or obligated to put information in EDGAR. Additionally, even publicly traded companies might not provide information about privately owned subsidiaries, so be careful and diligent.

- **SpiderFoot** is a free, open source, domain Footprinting tool. According to the site, “it will scrape the websites on that domain, as well as search Google, Netcraft, Whois and DNS to build up information.”
- **Sam Spade** is a DNS Footprinting tool.
- **pipl.com** is a site used for “people search.” For footprinting, pipl.com can use so-called “deep web searching” for loads of information you can use. The following is from the site: “Also known as ‘invisible web,’ the term ‘deep web’ refers to a vast repository of underlying content, such as documents in online databases that general-purpose web crawlers cannot reach. The deep web content is estimated at 500 times that of the surface web yet has remained mostly untapped due to the limitations of traditional search engines.”

19. Traceroute

What method does traceroute use to map routes traveled by a packet?

By manipulating the Time-To-Live (TTL) parameter

Traceroute (at least on Windows machines) tracks a packet across the Internet by incrementing the TTL on each packet it sends by one after each hop is hit and returns, ensuring the response comes back explicitly from that hop and returns its name and IP address. This provides route path and transit times. It accomplishes this by using ICMP ECHO packets to report information on each “hop” (router) from the source to destination. As an aside, Linux machines use a series of UDP packets by default to carry out the same function in traceroute.

20. Improving DNS Security

Brad is auditing an organization and is asked to provide suggestions on improving DNS security. Which of the following would be valid options to recommend? (Choose all that apply.)

- Implementing a split-horizon operation and
- Restricting zone transfers

Split-horizon DNS (also known as split-view or split DNS) is a method of providing different answers to DNS queries based on the source address of the DNS request. It can be accomplished with hardware or software solutions and provides one more step of separation between you and the bad guys.

Restricting zone transfers to only those systems you desire to have them is always a good idea. If you leave it open for anyone to grab, you are just asking for trouble. DNSSEC should also be included, but is not an option listed.

21. A zone file consists of which records? (Choose all that apply.)

☒ **A.** PTR

☒ **B.** MX

☐ **C.** SN

☒ **D.** SOA

☐ **E.** DNS

☒ **F.** A

☐ **G.** AX

☑ **A, B, D, F.** A zone file contains a list of all the resource records in the namespace zone. Valid resource records are as follows:

SRV	Service This record defines the hostname and port number of servers providing specific services, such as a Directory Services server.
SOA	Start of Authority This record identifies the primary name server for the zone. The SOA record contains the hostname of the server responsible for all DNS records within the namespace, as well as the basic properties of the domain.
PTR	Pointer This record maps an IP address to a hostname (providing for reverse DNS lookups). You don't absolutely need a PTR record for every entry in your DNS namespace, but PTR records are usually associated with e-mail server records.
NS	Name Server This record defines the name servers within your namespace. These servers are the ones that respond to your client's requests for name resolution.
MX	Mail Exchange This record identifies your e-mail servers within your domain.
CNAME	Canonical Name This record provides for domain name aliases within your zone. For example, you may have an FTP server and a web service running on the same IP address. CNAME records could be used to list both within DNS for you.
A	Address This record maps an IP address to a hostname and is used most often for DNS lookups.

22. OSRFramework

Within the OSRFramework, which tool verifies if a username/profile exists in up to 306 different platforms?

- usufy.py

The OSRFramework (<https://github.com/i3visio/osrframework>) is an open-source research framework in Python that helps you in the task of user profiling by making use of different open source intelligence (OSINT) tools.

The framework design itself is reminiscent of the Metasploit framework. It also has a web-based GUI that does the work for you if you like to work without the command line. In other words, it is a set of libraries used to perform OSINT tasks, helping you gather more, and more accurate, data using multiple applications in one easy-to-use package.

Usufy.py is but one of the tools in the framework, and it verifies if a username/profile exists in up to 306 different platforms.

Domainfy.py this tool verifies the existence of a given domain (per the site, in up to 1567 different registries).

Mailfy.py this tool checks if a username (e-mail) has been registered in e-mail providers.

Searchfy.py this tool looks for profiles using full names and other info in up to seven platforms. As an aside, ECC words this differently by saying the tool queries the OSRFramework platform itself.

23. Google operators

A colleague enters the following into a Google search string:

```
intitle:intranet inurl:intranet intext:"finance"
```

Which of the following statements is most correct concerning this attempt?

The search engine will respond with only those pages having the word intranet in the title and URL and with finance in the text.

This is a great Google hack that is listed on several websites providing Google hacking examples. Think about what you are looking for here—an internal page (intranet in title and URL) possibly containing finance data. Don't you think that would be valuable? This example shows the beauty of combining Google hacks to really burrow down to what you want to grab. Granted, an intranet being available from the Internet, indexed by Google and open enough for you to touch it, is unlikely, but these are questions concerning syntax, not reality.

24. Archive.org

The Internet Archive (<http://archive.org>) is a nonprofit organization “dedicated to build an Internet library. Its purposes include offering permanent access for researchers, historians, scholars, people with disabilities, and the general public to historical collections that exist in digital format.” The good-old Wayback Machine has been used for a long time to pull up old copies of websites, for good and maybe not-so-good purposes. Archive.org includes “snapshots of the World Wide Web,” which are archived copies of pages taken at various points in time dating back to 1996. As an additional note, Archive.org is only going to pull and store pages that were linked, shared, or commonly available, so do not assume every page ever put up by anyone anywhere will always be available.

25. Computer Security Incident Response Team (CSIRT)?

Which of the following is a primary service of the U.S. Computer Security Incident Response Team (CSIRT)?

- CSIRT provides an incident response service to enable a reliable and trusted single point of contact for reporting computer security incidents worldwide.

Per its website www.csirt.org, the Computer Security Incident Response Team (CSIRT) “provides 24x7 Computer Security Incident Response Services to any user, company, government agency or organization. CSIRT provides a reliable and trusted single point of contact for reporting computer security incidents worldwide. CSIRT provides the means for reporting incidents and for disseminating important incident-related information.” A privately held company that started in 2001, CSIRT seeks “to raise awareness among its customers of computer security issues and provides information for secure protection of critical computing infrastructure and equipment against potential organized computer attacks.”

26. Regional Registries

- **APNIC** handles the Asia and Pacific realms.
- **RIPE** handles Europe, Middle East, and parts of Central Asia/Northern Africa. If you're wondering, the name is French and stands for Réseaux IP Européens.

- **ARIN** service region includes Canada, many Caribbean and North Atlantic islands, and the United States. Caribbean islands falling under ARIN include Puerto Rico, the Bahamas, Antigua, American and British Virgin Islands, Turks and Caicos Islands, and the Cayman Islands (among others).
- **LACNIC** handles Latin America and parts of the Caribbean. It stands for Latin America and Caribbean Network Information Center. LACNIC coverage includes most of South America, Guatemala, French Guiana, Dominican Republic, and Cuba (among others). Exam takers most often get this one and ARIN confused.

References

<https://learning.oreilly.com/library/view/ceh-certified-ethical/9781260455090/ch1.xhtml#ch1>