88702 SIC 68320

1. **Attempt *any three* of the following (15)**

a. What are the importance of information protection? Explain with example.

Information is an important asset. The more information you have at your command, the better you can adapt to the world around you. In business, information is often one of the most important assets a company possesses. . Organizations typically choose to deploy more resources to control information that has higher sensitivity Organizations classify information in different ways in order to differently manage aspects of its handling, such as labeling (whether headers, footers, and watermarks specify how it should be handled), distribution (who gets to see it), duplication (how copies are made and handled), release (how it is provided to outsiders), storage (where it is kept), encryption (if required), disposal (whether it is shredded or strongly wiped), and methods of transmission (such as e-mail, fax, print, and mail).

Companies may have confidential information, such as research and development plans, manufacturing processes, strategic corporate information, product roadmaps, process descriptions, customer lists and contact information, financial forecasts, and earnings announcements, that is intended for internal use on a need-to-know basis. Loss or theft of confidential information could violate the privacy of individuals, reduce the company's competitive advantage, or cause damage to the company.

Specialized information or secret information may include trade secrets, such as formulas, production details, and other intellectual property, proprietary methodologies and practices that describe how services are provided, research plans, electronic codes, passwords, and encryption keys. If disclosed, this type of information may severely damage the company's competitive advantage. It is usually restricted to only a few people or departments within a company and is rarely disclosed outside the company.

Egg on Their Faces: A Case Study Egghead Software was a well-known software retailer who discovered in 2000 that Internet attackers might have stolen as many as 3.7 million credit card numbers from its web site, housed offsite at an e-commerce service provider that lacked good security. This information quickly made the news, and as a result, Egghead's corporate identity was more than just tarnished—it was destroyed. Customers fled in droves. The media coverage ruined the company's reputation. Egghead's stock price dropped dramatically, along with its sales. Cost-cutting measures, including layoffs, followed. The chain reaction finally concluded with Egghead's bankruptcy and subsequent acquisition by Amazon.com.

b. Explain about various components which are used to build a security program.

There are many components that go into the building of a security program:

- **Authority:** The security program must include the right level of responsibility and authorization to be effective. Usually, the security organization is responsible for information protection, risk management, monitoring, and response. It might also be responsible for enforcement, such as reprimanding or even terminating employees or contract workers, but more commonly that authority is vested in the Human Resources department. A resourcing plan is an ongoing strategy for providing the headcount needed to operate the security function. Insourcing, outsourcing, offshoring, and the like are factored into a resourcing plan, which describes how employees, contractors,

consultants, service providers, and temporary workers will be leveraged to fuel the progress of security implementations, operations, and improvement.

- **Framework:** A security framework provides a defensible approach to building the program. The security policy provides a framework for the security effort. The policy describes the intent of executive management with respect to what must be done to comply with the business requirements. The policy drives all aspects of technical implementations, as well as policies and procedures.
- **Assessmen**t: Assessing what needs to be protected, why, and how leads to a strategy for improving the security posture. **A risk analysis** provides a perspective on current risks to the organization's assets. This analysis is used to prioritize work efforts and budget allocation, so that the greater risks can receive a greater share of attention and resources. A risk analysis results in a well-defined set of risks that the organization is concerned about. **A gap analysis** compares the desired state of the security program with the actual current state and identifies the differences. **Remediation planning** takes into account the risks, gaps, and other objectives of the security program, and puts them together into a prioritized set of steps to move the security program from where it is today to where it needs to be at a future point.

- **Planning** : Planning produces priorities and timelines for security initiatives. • Action The actions of the security team produce the desired results based on the plans. **A roadmap** is a plan of action for how to implement the security remediation plans. It describes when, where, and what is planned. The roadmap is useful for managers who need the information to plan activities and to target specific implementation dates and the order of actions. It is also useful for implementers who will be responsible for putting everything together.

The security architecture documents how security technologies are implemented, at a relatively high level.

The project plans detail the activities of the individual contributors to the various security implementations.

Action :The actions of the security team produce the desired results based on the plans. Procedures describe how processes are performed by people on an ongoing basis to produce the desired outcomes of the security program in a repeatable, reliable fashion. Maintenance and support are part of maintaining the ongoing operations of the security program and its associated technologies, as part of a normal lifecycle of planning, updating, reviewing, and improving.

The actions that should be taken when a security event occurs are defined in the incident response plan.

- **Maintenance** : The end stage of the parts of the security program that have reached maturity is to maintain them.

Policy enforcement is necessary to ensure that the intentions of management are carried out by the various people responsible for the behavior and actions defined in the security policies.

Security awareness programs are used to educate employees, business partners, and other stakeholders about what behaviors are expected of them, what actions they should take under various circumstances to comply with security policies, and what consequences may ensue if they don't follow the rules..

c. What are the three recognized variants of malicious mobile code. Explain it.

There are three generally recognized variants of malicious mobile code: viruses, worms, and Trojans.

The lifecycle of malicious mobile code looks like this:

1. Find 2. Exploit 3. Infect 4. Repeat

   A virus is a self-replicating program that uses other host files or code to replicate. Most viruses infect files so that every time the host file is executed, the virus is executed too. A virus infection is simply another way of saying the virus made a copy of itself (replicated) and placed its code in the host in such a way that it will always be executed when the host is executed. Viruses can infect program files, boot sectors, hard drive partition tables, data files, memory, macro routines, and scripting files.

   Anatomy of a Virus: The damage routine of a virus (or really of any malware program) is called the payload. The vast majority of malicious program files do not carry a destructive payload beyond the requisite replication. This means they aren't intentionally designed by their creators to cause damage. However, their very nature requires that they modify other files and processes without appropriate authorization, and most end up causing program crashes of one type or another. Error-checking routines aren't high on the priority list for most attackers.

   A computer worm uses its own coding to replicate, although it may rely on the existence of other related code to do so. The key to a worm is that it does not directly modify other host code to replicate. A worm may travel the Internet trying one or more exploits to compromise a computer, and if successful, it then writes itself to the computer and begins replicating again. An example of an Internet worm is Bugbear. Bugbear was released in June 2003, arriving as a file attachment in a bogus e-mail. In unpatched Outlook Express systems, it can execute while the user is simply previewing the message. In most cases, it requires that the end user execute the file attachment. Once launched, it infects the PC, harvests e-mail addresses from the user's e-mail system, and sends itself out to new recipients. It adds itself into the Windows startup group so it gets executed each time Windows starts. Bugbear looks for and attempts to gain access to weakly password-protected network shares and terminates antivirus programs. It also drops off and activates a keylogging program, which records users' keystrokes in an attempt to capture passwords. The captured keystrokes, and any cached dial-up passwords that are found, are then e-mailed to one of ten predefined e-mail addresses. Lastly, Bugbear

opens up a back door service on port 1080 to allow attackers to manipulate and delete files. Bugbear was one of the most successful worms of 2003.

Trojans :Trojan horse programs, or Trojans, work by posing as legitimate programs that are activated by an unsuspecting user. After execution, the Trojan may attempt to continue to pose as the other legitimate program (such as a screensaver) while doing its malicious actions in the background. Many people are infected by Trojans for months and years without realizing it. If the Trojan simply starts its malicious actions and doesn't pretend to be a legitimate program, it's called a direct-action Trojan. Direct-action Trojans don't spread well because the victims notice the compromise and are unlikely, or unable, to spread the program to other unsuspecting users. An example of a direct-action Trojan is JS.ExitW. It can be downloaded and activated when unsuspecting users browse malicious web sites. In one case, this Trojan posed as a collection of Justin Timberlake pictures and turned up in a search using Google. The link, instead of leading to the pictures, downloaded and installed the JS.ExitW Trojan. When activated, JS.ExitW installs itself in the Windows startup folder as an HTML application (.hta) that shuts down Windows. Because it is in the startup folder, this has the consequence of putting infected PCs in a never-ending loop of starts and shutdowns. Luckily, this Trojan does no real damage. Unfortunately, many Trojans aren't so harmless.

d. Write a short note on Network-Layer Attack.

Network-layer attacks attempt to compromise network devices and protocol stacks. Network-layer attacks include packet-sniffing and protocol-anomaly exploits.

Packet Sniffing . Encryption is used to prevent packet-sniffing (also known as packet capturing or protocol analyzing) attacks. Sniffing occurs when an unauthorized third party captures network packets destined for computers other than their own. Packet sniffing allows the attacker to look at transmitted content and may reveal passwords and confidential data. In order to use sniffing software, an attacker must have a promiscuous network card and specialized packet driver software, must be connected to the network segment they want to sniff, and must use sniffer software. By default, a network interface card (NIC) in a computer will usually drop any traffic not destined for it. By putting the NIC in promiscuous mode, it will read any packet going by it on the network wire. Note that in order for a sniffer to capture traffic, it must physically be able to capture it. On switched networks, where each network drop is its own collision domain, packet sniffing by intruders can be more difficult, but not impossible. Packet-sniffing attacks are more common in areas where many computer hosts share the same collision domain (such as a wireless segment or local LAN shared over an Ethernet hub) or over the Internet where the attacker might insert a sniffer in between source and destination traffic. For example, on a LAN, a less privileged user may sniff traffic originating from an administrative account, hoping to get the password. There are several open source sniffing tools, including tcpdump (or WinDump, the Windows version) and the easier-to-use Ethereal (www.ethereal.com).

Protocol-Anomaly Attacks Most network protocols were not created with security in mind. A rogue attacker can create malformed network packets that do not follow the

intended format and purpose of the protocol, with the result that the attacker is able to either compromise a remote host or network, or compromise a confidential network data stream. Network-layer attacks are most often used to get past firewalls and to cause DoS attacks. DoS attacks are common against big e-commerce sites. In one type of DoS attack, the attacker machines send massive amounts of TCP SYN packets. This is the first of three packets sent during a normal TCP handshake used to begin a communication session. The victim machine responds with the expected ACK/SYN packet, which is normal, and then awaits an answering ACK from the originator. However, the ACK packet never comes, leaving the TCP connection in an open state, waiting for an extended period of time. When sent millions of these packets, the attacked operating system is overtaxed with open connections all in a waiting state. Often the victim machine has to reboot to clear all the open connections. If they do reboot without doing something to stop the DoS attack, it just happens again and again. Often the originating address of the malicious ACK packets is faked, so there is no way to simply block the originating IP address. This is just one type of DoS attack, and there are dozens of ways to cause them.

Network-layer attacks usually require that the attacker create malformed traffic, which can be created by tools called packet injectors or traffic generators. Packet injectors are used by legitimate sources to test the throughput of network devices or to test the security defenses of firewalls and IDSs. There are dozens of commercial and open source packet generators that allow a fair amount of flexibility in generating TCP/IP traffic, permitting different protocols (TCP, UDP, and ICMP), packet sizes, payload contents, packet flow rates, flag settings, and customized header options. Attackers can even manually create the malformed traffic as a text file and then send it using a traffic replay tool.

e. Explain the two most common approaches of security.

There are two approaches can take to preserve the confidentiality, integrity, availability, and authenticity of electronic and physical assets such as the data on your network:

- Build a defensive perimeter around those assets and trust everyone who has access inside
- Use many different types and levels of security controls in a layered defense-indepth approach

The concepts of the lollipop and the onion is used to depict the two most common approaches to security.

The Lollipop Model The most common form of defense, known as perimeter security, involves building a virtual (or physical) wall around objects of value. Perimeter security is like a lollipop with a hard, crunchy shell on the outside and a soft, chewy center on the inside, as illustrated in Figure . Consider the example of a house—it has walls, doors, and windows to protect what's inside (a perimeter). But does that make it impenetrable? No, because a determined attacker can find a way in—either by breaking through the perimeter, or exploiting some weakness in it, or convincing someone inside to let them in. By comparison, in network security, a firewall is like the house—it is a perimeter that can't keep out all attackers. Yet the firewall is the most common choice

for controlling outside access to the internal network, creating a virtual perimeter around the internal network (which is usually left wide open). This often creates a false sense of security, because attackers can break through, exploit vulnerabilities, or compromise the network from the inside.

Figure The lollypop model of defense

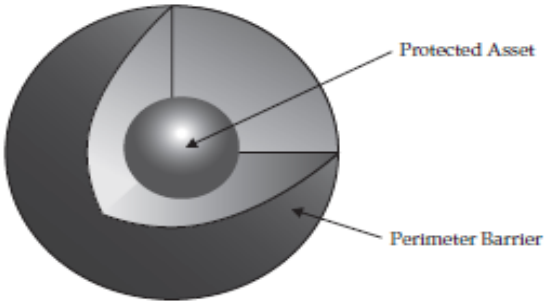

The Onion Model A better approach is the onion model of security. It is a layered strategy, often referred to as defense in depth. This model addresses the contingency of a perimeter security breach occurring. It includes the strong wall of the lollipop but goes beyond the idea of a simple barrier, as depicted in Figure A layered security architecture, like an onion, must be peeled away by the attacker, layer by layer, with plenty of crying.

Figure onion model of defense

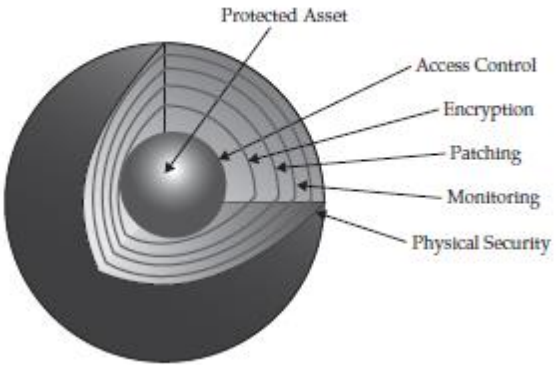

Consider what happens when an invader picks the front door lock or breaks a window to gain entry to a house. The homeowner may hide cash in a drawer and may store valuable jewels in a safe. These protective mechanisms address the contingency that the perimeter security fails. They also address the prospect of an inside job. The same principles apply to network security. The onion model addresses these contingencies.

A firewall alone provides only one layer of protection against threats originating from the Internet, and it does not address internal security needs. With only one layer of protection, which is common on networks connected to the Internet, all a determined individual has to do is successfully attack that one system to gain full access to everything on the network. A layered security architecture provides multiple levels of protection against internal and external threats. The more layers of controls that exist, the better the protection against a failure of any one of those layers.

f. Explain the best practices for network defence.

There are many countermeasures you can implement to minimize the risk of a successful attack, such as securing the physical environment, hardening the operating systems, keeping patches updated, using an antivirus scanner, using a firewall, securing network share permissions, using encryptions, securing applications, backing up the system, creating a computer security defense plan, and implementing ARP poisoning defences

Secure the Physical Environment

A basic part of any computer security plan is the physical aspect. Of course, mission-critical servers should be protected behind a locked door, but regular PCs need physical protection too. Depending on your environment, PCs and laptops might need to be physically secured to their desks. There are several different kinds of lockdown devices, from thin lanyards of rubber-coated wire to hardened metal jackets custom-made to surround a PC. If anyone leaves their laptop on their desk overnight, it should be secured. There are also other steps that need to be taken on every PC in your environment. **Password Protect Booting**

Consider requiring a boot-up password before the operating system will load. This can usually be set in the CMOS/BIOS and is called a user or boot password. This is especially important for portable computers, such as laptops and tablets and smartphones. Smallform-factor PCs are the most likely candidates to be stolen. Since most portable devices often contain personal or confidential information, password-protecting the boot sequence might keep a nontechnical thief from easily seeing the data on the hard drive or storage RAM. If a boot-up password is reset on a tablet or smartphone, often it requires that the data be erased too, so confidentiality and privacy are assured.

Password Protect CMOS

The CMOS/BIOS settings of a computer contain many potential security settings, such as boot order, remote wake-up, and antivirus boot-sector protection. It is important to ensure that unauthorized users do not have access to the CMOS/BIOS settings. Most CMOS/ BIOSs allow you to set up a password to prevent unauthorized changes. The password should not be the same as other administrative passwords, but for simplicity's sake, a common password can be used for all machines.

Disable Booting from USB and CD

Disabling booting from USB storage devices and optical drives will prevent boot viruses from those devices and stop attackers from bypassing operating system security by loading a different operating system on the computer.

Harden the Operating System

To reduce the attack surface of the operating system by removing unnecessary software, disabling unneeded services, and locking down access:

1. Reduce the attack surface of systems by turning off unneeded services.

2. Install secure software.

3. Configure software settings securely.

4. Patch systems regularly and quickly.

5. Segment the network into zones of trust and place systems into those zones based on their communication needs and Internet exposure.

6. Strengthen authentication processes.

7. Limit the number (and privileges) of administrators

Keep Patches Updated

An attacker's best friend is an unpatched system. In most cases, the vulnerabilities used are widely known, and the affected vendors have already released patches for system administrators to apply.

Use an Antivirus Scanner (with Real-Time Scanning)

In today's world, an antivirus (AV) scanner is essential. It should be deployed on your desktop, with forced, automatic updates, and it should be enabled for real-time protection.

Use Firewall Software

Almost as important as an AV scanner is the firewall. Firewalls have come a long way since their days of simple port filtering

Secure Network Share Permissions

One of the most common ways an attacker or worm breaks into a system is through a network share (such as NetBIOS or SMB) with no password or a weak password. Folders and files accessed remotely over the network should have discretionary ACLs (DACLs) applied using the principle of least privilege and should have complex passwords.

Use Encryption

Most computer systems have many encryption opportunities. Use them. Linux and Unix administrators should be using SSH instead of Telnet or FTP to manage their computers. The latter utilities work in plaintext over the network, whereas SSH is encrypted.

Secure Applications

Managing your applications and their security should be a top priority of any administrator. Applications can be managed by configuring application security, installing

applications to nonstandard directories and ports, locking down applications, securing P2P services, and making sure your application programmers code securely.

Back Up the System

With the notable exception of stolen confidential information, the most common symptom of damage from malware is modified, corrupted, or deleted files. Worms and viruses often delete files, format hard drives, or intentionally corrupt data. Even malware that does nothing intentionally wrong to a system's files is maliciously modifying a system just by being present. Security experts cannot always repair the damage and put the system back to the way it was prior to the exploit. This means it's important to keep regular, tested backups of your system. The backup should include all your data files at a minimum, and a complete system backup ensures a quicker recovery in the event of a catastrophic exploit event.

Implement ARP Poisoning :

ARP poisoning attacks are one of the most common and effective threats against network infrastructures (especially wireless networks). They are a form of man-in-the-middle (MITM) attack that allows an attacker to intercept and modify network traffic, invisibly. Thus, these attacks merit their own special countermeasures. There are a few ways an organization can defend against an ARP poisoning attack. Defenses include implementing static ARP tables, configuring port rate limiting, or using DHCP snooping with dynamic ARP inspection (DAI). The most effective defense is a combination of the latter two methods.

Create a Computer Security Defense Plan

Steps to creating a Computer Security Defense plan:

1. Inventory the assets you have to protect.

2. Decide the value of each asset and its chance of being exploited in order to come up with a quantifiable exposure risk.

3. Using the steps outlined in this chapter (and summarized next), develop a plan to tighten the security on your protected assets. Assets with the highest exposure risk should be given the most protection, but make sure all assets get some baseline level of security. 4. Develop and document security baseline tools and methods.

5. Use vulnerability testing tools to confirm assets have been appropriately configured.

6. Do periodic testing to make sure security settings stay implemented. 7. Change and update the plan as dictated by new security events and risks.

2. Attempt *any three* of the following (15)

a. Define authentication. Explain two parts of authentication.

Authentication is the process by which people prove they are who they say they are. It's composed of two parts: a public statement of identity (usually in the form of a username) combined with a private response to a challenge (such as a password).

Usernames and Passwords

In the familiar method of password authentication, a challenge is issued by a computer, and the party wishing to be identified provides a response. If the response can be validated, the user is said to be authenticated, and the user is allowed to access the system. Otherwise, the user is prevented from accessing the system.

Other password-based systems, including Kerberos, are more complex, but they all rely on a simple fallacy: they trust that anyone who knows a particular user's password is that user. Many password authentication systems exist. The following types of systems are commonly used today:

- Local storage and comparison
- Central storage and comparison
- Challenge and response
- Kerberos
- One-time password (OTP)

Local Storage and Comparison

Early computer systems did not require passwords. Whoever physically possessed the system could use it. As systems developed, a requirement to restrict access to the privileged few was recognized, and a system of user identification was developed. User passwords were entered in simple machine-resident databases by administrators and were provided to users. Often, passwords were stored in the database in plaintext format (unencrypted), because protecting them wasn't really a high priority. Anyone who was able to open and read the file could determine what anyone else's password was. The security of the database relied on controlling access to the file, and on the good intentions of all the administrators and users. Administrators were in charge of changing passwords, communicating changes to the users, and recovering passwords for users who couldn't remember them.

Central Storage and Comparison

When passwords are encrypted, authentication processes change. Instead of doing a simple comparison, the system must first take the user-entered, plaintext password and encrypt it using the same algorithm used for its storage in the password file. Next, the newly encrypted password is compared to the stored encrypted password. If they match, the user is authenticated. This is how many operating systems and applications work today.

Sometimes the password entered by the user is encrypted, passed over the network in this state, and then compared by the remote server to its stored encrypted password. This is the ideal situation. Unfortunately, some network applications transmit passwords in cleartext—telnet, FTP, rlogin, and many others do so by default. Even systems with secure local, or even centralized, network logon systems may use these and other applications which then transmit passwords in cleartext. If attackers can capture this data in flight, they can use it to log in as that user. In addition to these network applications, early remote authentication algorithms (used to log in via dial-up

connections), such as Password Authentication Protocol (PAP), also transmit cleartext passwords from client to server.

Kerberos

Kerberos is a network authentication system based on the use of tickets. In the Kerberos standard (RFC 1510), passwords are key to the system, but in some systems certificates may be used instead. Kerberos is a complex protocol developed at the Massachusetts Institute of Technology to provide authentication in a hostile network. Its developers, unlike those of some other network authentication systems, assumed that malicious individuals as well as curious users would have access to the network

One-Time Password Systems

Two problems plague passwords. First, they are (in most cases) created by people. Thus, people need to be taught how to construct strong passwords, and most people aren't taught (or don't care enough to follow what they're taught). These strong passwords must also be remembered and not written down, which means, in most cases, that long passwords cannot be required. Second, passwords do become known by people other than the individual they belong to. People do write passwords down and often leave them where others can find them. People commonly share passwords despite all your warnings and threats. Passwords are subject to a number of different attacks. They can be captured and cracked, or used in a replay attack in which the passwords are intercepted and later used to repeat authentication.

One solution to this type of attack is to use an algorithm that requires the password to be different every time it is used. In systems other than computers, this has been accomplished with the use of a one-time pad. When two people need to send encrypted messages, if they each have a copy of the one-time pad, each can use the day's password, or some other method for determining which password to use. The advantage, of course, to such a system is that even if a key is cracked or deduced, it is only good for the current message. The next message uses a different key.

on authentication systems. The Kerberos authentication process follows these steps,

1. A user enters their password.
2. Data about the client and possibly an authenticator is sent to the server.

The authenticator is the result of using the password (which may be hashed or otherwise manipulated) to encrypt a timestamp (the clock time on the client computer). This authenticator and a plaintext copy of the timestamp accompany a login request, which is sent to the Kerberos authentication server (AS)—this is the KRB_AS_REQ message. This is known as pre-authentication and may not be part of all Kerberos implementations.

b. Explain the authorization systems.

The counterpart to authentication is authorization. Authentication establishes who the user is; authorization specifies what that user can do. Typically thought of as a way of establishing access to resources, such as files and printers, authorization also ddresses the suite of privileges that a user may have on the system or on the network. In its ultimate use, authorization even specifies whether the user can access the system at all. There are a variety of types of authorization systems, including user rights, role-based authorization, access control lists, and rule-based authorization.

User Rights

Privileges or user rights are different from permissions. User rights provide the authorization to do things that affect the entire system. The ability to create groups, assign users to groups, log in to a system, and many more user rights can be assigned. Other user rights are implicit and are rights that are granted to default groups—groups that are created by the operating system instead of by administrators. These rights cannot be removed.

Role-Based Authorization (RBAC)

Each job within a company has a role to play. Each employee requires privileges (the right to do something) and permissions (the right to access particular resources and do specified things with them) if they are to do their job. Early designers of computer systems recognized that the needs of possible users of system. would vary, and that not all users should be given the right to administer the system.

Access Control Lists (ACLs)

Attendance at some social events is limited to invitees only. To ensure that only invited guests are welcomed to the party, a list of authorized individuals may be provided to those who permit the guests in. If you arrive, the name you provide is checked against this list, and entry is granted or denied. Authentication, in the form of a photo identification check, may or may not play a part here, but this is a good, simple example of the use of an access control list (ACL).

Information systems may also use ACLs to determine whether the requested service or resource is authorized. Access to files on a server is often controlled by information that is maintained on each file. Likewise, the ability for different types of communication to pass a network device can be controlled by ACLs.

Rule-Based Authorization

Rule-based authorization requires the development of rules that stipulate what a specific user can do on a system. These rules might provide information such as "User Alice can access resource Z but cannot access resource D." More complex rules specify combinations, such as "User Bob can read file P only if he is sitting at the console in the data center." In a small system, rule-based authorization may not be too difficult to maintain, but in larger systems and networks, it is excruciatingly tedious and difficult to administer.

c. Briefly explain public key Cryptography.

Public Key Cryptography

One of the method for exchanging a session key is to use public key cryptography. This algorithm is asymmetric—it uses a set of related keys. If one key is used to encrypt the message, the other is used to decrypt it, and vice versa. This means that if each party holds one of the keys, a session key can be securely exchanged. In the typical arrangement, each party has their own set of these asymmetric keys. One of the key pairs is known as the private key and the other as the public key. Public keys are exchanged and private keys are kept secret. Even if a public key becomes, well, public, it does not compromise the system. It's meant to be shared openly.

Key Exchange

Public/private key pairs can be used to exchange session keys. To do so, each party that needs to exchange keys generates a key pair. The public keys are either exchanged among the parties or kept in a database. The private keys are kept secret. When it is necessary to exchange a key, one party can encrypt it using the public key of the other. The encrypted key is then transmitted to the other party. Since only the intended recipient holds the private key that is related to the public key used to encrypt the session key, only that party can decrypt the session key. The confidentiality of the session key is assured, and it can then be used to encrypt communications between the two parties. The steps are outlined here and are illustrated in Figure

Figure using public key cryptography for key exchange.

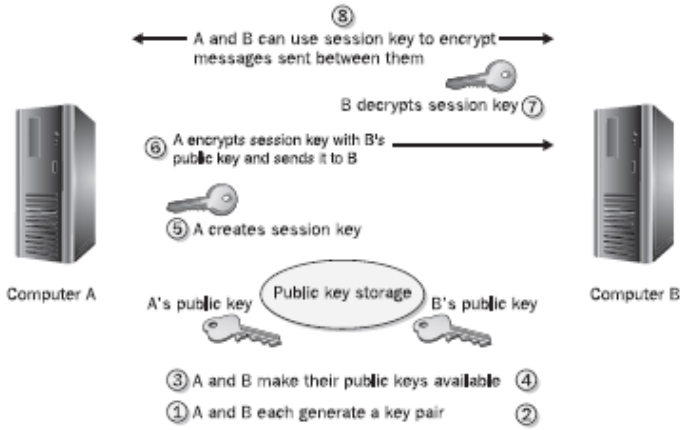

d. What are the three primary categories of storage infrastructure in modern storage security.

Modern storage environments can be considered as separate IT infrastructures of their own. Many organizations are now dividing their IT organizations along the lines of networks, servers, and storage—acknowledging that storage merits a place alongside these long-venerated institutions.

Storage infrastructure can often be found on a dedicated LAN, with servers, arrays, and NAS appliances, with specialized operating systems to support the storage. Storage can also be located in multiple sites, including geographically diverse regional distributions, and even third-party and Internet locations. In securing these components, you must take into account three primary categories:

- Storage networks
- Arrays
- Servers

Storage Networks

Separation of duties should be applied within the storage infrastructure. Since all storage devices are connected physically, either over a network or through a storage connection protocol, separating access to the physical servers prevents a storage administrator from connecting a rogue server into the environment and then provisioning it access to restricted logical unit numbers (LUNs). A LUN is the mechanism an array uses to present its storage to a host operating system. Likewise,

while someone may connect a server to the environment and configure it, methods of protecting the LUNs are applied so that the server cannot gain access to restricted LUNs.

Port Zoning :The most notable characteristic of port zoning is that the accessibility of the host to the LUNs is defined by the switch port. The advantage to zoning in this manner is that an intruder cannot connect a host to the switch, enable spoofing of a good WWN, and access LUNs of another host. Since the protection is enforced on the port interface, the intruder would need to disconnect the good host interface and connect the intruding host into the defined port. All this would need to be done without any alerts being flagged by the host operating system, which is practically impossible.

WWN Zoning The alternative to port zoning, in which the zones are created relative to the ports the servers are connected to on the switch, is WWN zoning, which defines the individual zone based on the WWN ID of the host bus adapter (HBA). The WWN is very much like the MAC address of a network card. It is a 16-digit hexadecimal number that uniquely identifies the HBA within the SAN fabric. These numbers are assigned in much the same way as MAC addresses are assigned to OEM manufacturers, with the first eight digits assigned to specific manufacturers and the rest of the numbers assigned by the manufacturers.

Arrays

Another area of risk is the storage array itself. When LUNs are created, it is necessary for the array to provide a screen to prevent the data that resides on the array from being accessed by other hosts that are able to connect to the array. Storage arrays are therefore equipped with a mechanism that provides protection known as LUN masking. This allows multiple hosts to communicate with the array and only access LUNs that are assigned through the application that provides the LUN-masking protection. Consider the differences in protection between zoning and LUN masking.

Figure security area of arrays

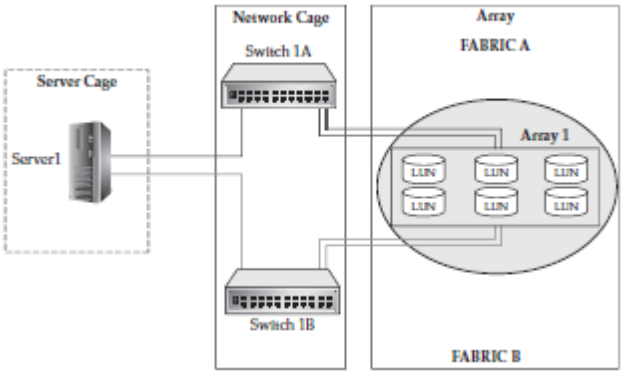

Servers

. Storage administrators often have limited control over what can or cannot be done on the host, as this administration is handled by the systems administrators. However, in many organizations, the systems administrator is also the storage administrator, which means that person has full access to both the storage and the systems that use it. As long as the data “rests” on the server, the potential to access that data exists. Many options are available to protect that data while it is at rest on the server. The concern of the

storage administrator is what happens if someone is able to access the data either locally or remotely. In the worst-case scenario, an attacker may obtain access to the server and escalate his authority to attempt to read the data. In order to keep the data secure in this scenario, it is necessary to implement data encryption. Therefore, when securing data, a comprehensive solution is necessary. The operating system must be secured and patched, file permissions must be planned and applied to reduce access as much as possible, and monitoring needs to be performed. Finally, confidential data should also be encrypted to protect it from unwanted access.

Figure security area of servers

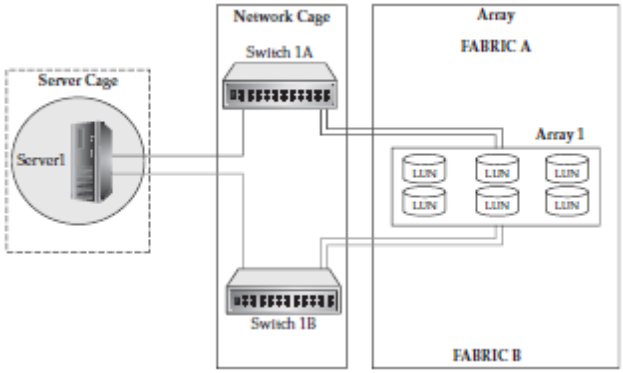

e. Write a short note on integrity risks.

Integrity Risks

Integrity risks affect both the validity of information and the assurance that the information is correct. Some government regulations are particularly concerned with ensuring that data is accurate. If information can be changed without warning, authorization, or an audit trail, its integrity cannot be guaranteed.

Malfunctions

Computer and storage failures that corrupt data damage the integrity of that data.

- Defense Make sure the storage infrastructure you select has appropriate RAID redundancy built in and that archives of important data are part of the service. • Detection Employ integrity verification software that uses checksums or other means of data verification.
- Deterrence Due to the nature of data, because there is no human element involved, there isn't much that can be done.
- Residual risk: Technology failures that damage data may result in operational or compliance risk (especially relating to Sarbanes-Oxley requirements for publicly traded companies to ensure the integrity of their financial data).

Data Deletion and Data Loss

Data can be accidentally or intentionally destroyed due to computer system failures or mishandling. Such data may include financial, organizational, personal, and audit trail information.

- Defense: Ensure that your critical data is redundantly stored and housed in more than one location.
- Detection: Maintain and review audit logs of data deletion.

- Deterrence :Maintain educational and awareness programs for individuals who access and manage data. Ensure that data owners are assigned that have authority and control over data and responsibility for its loss.
- Residual risk: Once critical data is gone, if it can't be restored, it is gone forever.

Data Corruption and Data Tampering

Changes to data caused by malfunction in computer or storage systems, or by malicious individuals or malware, can damage the integrity of that data. Integrity can also be damaged by people who modify data with intent to defraud.

- Defense :Utilize version control software to maintain archive copies of important data before it is modified. Ensure that all data is protected by antivirus software. Maintain role-based access control over all data based on least privilege principles, pursuant to job function and need to know.
- Detection: Use integrity-checking software to monitor and report alterations to key data

. • Deterrence :Maintain educational and awareness programs for individuals who access and manage data. Ensure that data owners are assigned that have authority and control over data and responsibility for its loss.

- Residual risk :Corrupted or damaged data can cause significant issues because valid, reliable data is the cornerstone of any computing system.

Accidental Modification

Perhaps the most common cause of data integrity loss, accidental modification occurs either when a user intentionally makes changes to data but makes the changes to the wrong data or when a user inputs data incorrectly.

- Defense Utilize version control software to maintain archive copies of important data before it is modified. Maintain role-based access control over all data based on least privilege principles, pursuant to job function and need to know.
- Detection Use integrity-checking software to monitor and report alterations to key data.
- Deterrence Maintain educational and awareness programs for individuals who access and manage data. Ensure that data owners are assigned that have authority and control over data and responsibility for its loss.
- Residual risk Corrupted or damaged data can cause significant issues because valid, reliable data is the cornerstone of any computing system.

f. Explain Database-Level Security.

Databases are commonly used to host many different databases and applications, and users should have different types of permissions based on their job functions. Once a user has been allowed to connect to a server (through the use of a server login), the user will be given only the permissions that are granted to that login. This process of determining permissions is generally known as authorization.

The first type of database-level security is generally used to determine to which database(s) a user has access. Database administrators can specify whether or not certain databases can be accessed by a user login. For example, one login may be

granted permissions to access only the Human Resources database and not any system databases or databases used by other applications.

Database Administration Security

One important task related to working with a relational database is maintenance of the server itself. Important tasks include creating databases, removing unneeded databases, managing disk space allocation, monitoring performance, and performing backup and recovery operations. Database platforms allow the default systems administrator account to delegate permissions to other users, allowing them to perform these important operations. As an example, Microsoft's SQL Server platform provides built-in server-level roles, including Database Creators, Disk Administrators, Server Administrators, Security Administrators, and many others.

Database Roles and Permissions

Having a valid server login only allows a user the permission to connect to a server. In order to actually access a database, the user's login must be authorized to use it.

The general process begins with specifying to which database(s) a login may connect. Then, permissions must be assigned within the database. The details here do vary between types of relational database platforms, but the overall concepts are the same. Generally, database administrators will create "groups" or "roles," and each of these will contain users. Specific permissions (which we'll look at in the next section) are assigned to the roles. This process is quite similar to the best practices that are suggested for most modern network operating systems. Additionally, some relational database platforms allow groups to be nested, thereby allowing you to create a hierarchy of permissions.

Object-Level Security

Relational databases support many different types of objects. Tables, however, are the fundamental unit of data storage. Each table is generally designed to refer to some type of entity (such as an Employee, a Customer, or an Order). Columns within these tables store details about each of these items (FirstName or CustomerNumber are common examples). Permissions are granted to execute one or more of the most commonly used SQL commands. These commands are

- SELECT : Retrieves information from databases. SELECT statements can obtain and combine data from many different tables, and can also be used for performing complex aggregate calculations.
- INSERT :Adds a new row to a table.
- UPDATE: Changes the values in an existing row or rows.
- DELETE: Deletes rows from a table.

The ANSI Standard SQL language provides for the ability to use three commands for administering permissions to tables and other database objects:

- GRANT :Specifies that a particular user or role will have access to perform a specific action.
- REVOKE: Removes any current permissions settings for the specified users or roles.
- DENY :Prevents a user or role from performing a specific action.

Other Database Objects for Security

Views

Perhaps the most commonly used method of controlling data access is views. A view is a logical relational database object that actually refers to one or more underlying database tables. Views are generally defined simply as the result of a SELECT query. This query, in turn, can pull information from many different tables and can also perform common calculations on the data.

Stored Procedures

Database logic can become significantly complex, and common operations often must be performed by many different users. Thankfully, databases offer developers the ability to create and reuse SQL code through the use of objects called stored procedures. Stored procedures can be used to perform any function that is possible through the use of standard SQL commands. Additionally, they can take arguments (much like functions and subroutines in other programming languages), making them very flexible.

Triggers

Triggers are designed to automatically be "fired" whenever specification actions take place within a database. For example, you might create a trigger on the SalesOrder table that will automatically create a corresponding row in the Invoice table. Or, you might create a trigger that performs complex data validation

3. Attempt *any three* of the following (15)

a. Explain the Cisco Hierarchical Internetworking model.

The legacy Cisco Hierarchical Internetworking model, which most network engineers are intimately familiar with, is a common design implemented in large-scale networks today, although many new types of purposed designs have been developed that support emerging technologies like class fabrics, lossless Ethernet, layer two bridging with trill or IEEE 802.1aq, and other data center–centric technologies.

The Cisco three-tier model is derived from the Public Switched Telephone Network (PSTN) model, which is in use for much of the world's telephone infrastructure. The Cisco Hierarchical Internetworking model, three main layers commonly referred to as the core, distribution, and access layers:

- Core layer : Forms the network backbone and is focused on moving data as fast as possible between distribution layers. Because performance is the core layer's primary focus, it should not be used to perform CPU-intensive operations such as filtering, compressing, encrypting, or translating network addresses for traffic.
- Distribution layer Sits between the core and the access layer. This layer is used to aggregate access-layer traffic for transmission into and out of the core.
- Access layer Composed of the user networking connections.

Filtering, compressing, encrypting, and address-translating operations should be performed at the access and distribution layers.

The Cisco model is highly scalable. As the network grows, additional distribution and access layers can be added seamlessly. As the need for faster connections and more bandwidth arises, the core and distribution equipment can be upgraded as required. This model also assists corporations in achieving higher levels of availability by allowing

for the implementation of redundant hardware at the distribution and core layers. And because the network is highly segmented, a single network failure at the access or distribution layers does not affect the entire network.

Although the Cisco three-tier model is perhaps the most commonly known and referenced model for designing LAN environments, it has its limitations and is rapidly being supplanted by newer models aimed at addressing the specific needs of highly virtualized data centers, the specific needs of different industry verticals, and the specific needs of cloud computing and multitenancy environments.

Figure Cisco Hierarchical Internetworking model

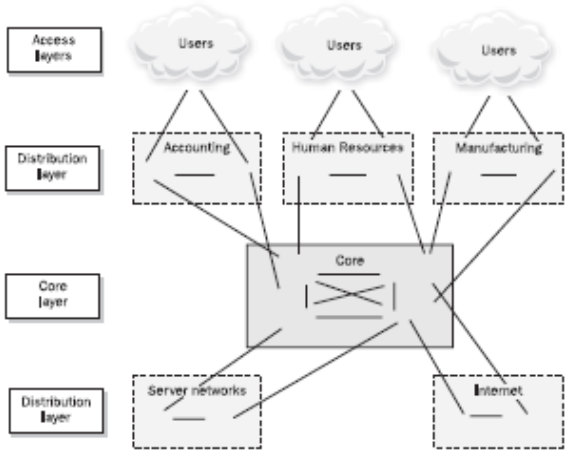

b. Briefly explain network availability and security.

Availability

Network availability requires that systems are appropriately resilient and available to users on a timely basis (meaning, when users require them). The opposite of availability is denial of service, which is when users cannot access the resources they need on a timely basis. Denial of service can be intentional (for example, the act of malicious individuals) or accidental (such as when hardware or software fails). Unavailable systems cost corporations real dollars in lost revenue and employee productivity, and they can hurt organizations in intangible ways through lost onsumer confidence and negative publicity. Business availability needs have driven some organizations to construct duplicate data centers that perform real-time irroring of systems and data to provide failover and reduce the risk of a natural disaster or terrorist attack destroying their only data center.

Depending on the specific business and risk factors, redundancy often increases both cost and complexity. Determining the right level of availability and redundancy is an important design element, which is best influenced by a balance between business requirements and resource availability.

The best practice for ensuring availability is to avoid single points of failure within the architecture. This can require redundant and/or failover capabilities at the hardware, network, and application functions. A fully redundant solution can be extremely expensive to deploy and maintain, because as the number of failover mechanisms

increases, system complexity increases, which alone can raise support costs and complicate troubleshooting.

Implementing a redundant firewall or router solution is only one step in achieving a full high-availability network architecture. For example, a high-availability firewall solution provides no value when both firewalls are plugged into the same switch. The switch becomes a single point of failure, and any interruption in its normal operation would take both firewalls off the network, negating any benefit of the firewall failover mechanism. The same holds true of a router—if there is only a single router between the firewalls and the rest of the network, the failure of that router would also cause an outage.

Figure full high availability network design

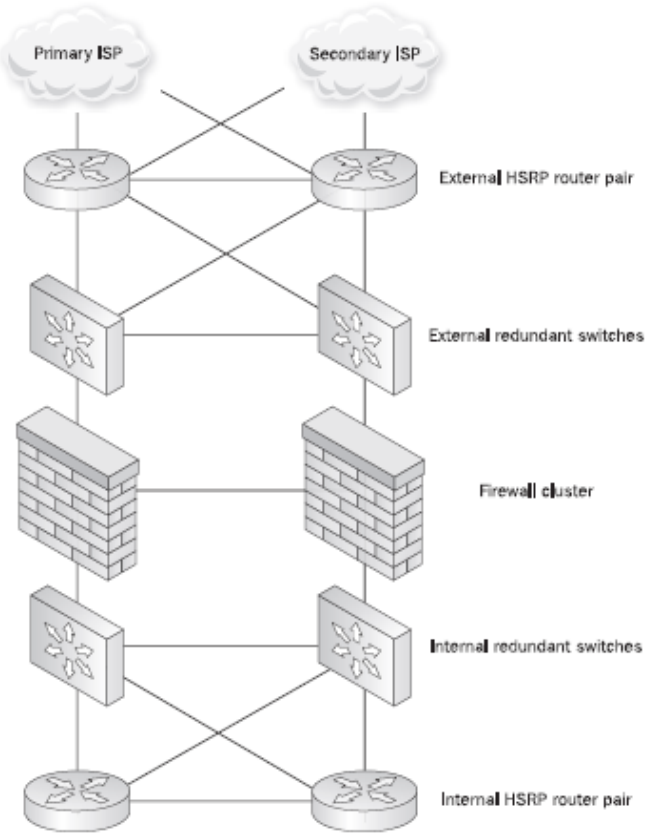

Security

Each element on a network performs different functions and contains data of differing security requirements. Some devices contain highly sensitive information that could damage an organization if disseminated to unauthorized individuals, such as payroll records, internal memorandums, customer lists, and even internal job-costing documents. Other devices have more exposure due to their location on the network. For example, internal file servers will be protected differently than publicly available web servers.

When designing and implementing security in network and system architectures, it is helpful to identify critical security controls and understand the consequences of a failure in those controls. For example, firewalls protect hosts by limiting what services users can connect to on a given system. Firewalls can allow different sets of users selective access to different

services, such as allowing system administrators to access administrative services while preventing non-administrative users from accessing those same services. This provides an additional level of control over that provided by the administrative mechanisms themselves. By denying a non-administrative user the ability to connect to the administrative service, that user is prevented from mounting an attack directly on that service without first circumventing the firewall.

However, simply restricting users to specific services may be insufficient to achieve the desired level of security. For example, it is necessary to allow traffic through the firewall to connect to various authorized services. In order for an organization to send and receive e-mail, firewalls must be configured to permit e-mail traffic. As Chapter 15 will discuss, firewalls have limited capability in preventing attacks directed at authorized applications, so overall network security is dependent on the proper and secure operation of those applications.

Flaws, such as a buffer overflows, can allow an attacker to turn a vulnerable server into a conduit through the firewall. Once through the firewall, the attacker can mount attacks against infrastructure behind the protection of the firewall. If the server is on the internal network, the entire network could be attacked without the protection provided by the firewall, but if the server is on a separate firewalled segment instead of the internal network, only the hosts on the same subnet could be directly attacked. Because all traffic exiting that subnet still must pass back through the firewall, it can still be relied upon to protect any additional communications from this compromised subnet to any other internal subnets. In addition to the best practice of segmenting the traffic, using the advanced inspection

c. Write a short note on hubs and switches.

Hubs

Hubs were dumb devices used to solve the most basic connectivity issue: how to connect more than two devices together. They transmitted packets between devices connected to them, and they functioned by retransmitting each and every packet received on one port out through all of its other ports without storing or remembering any information about the hosts connected to them. This created scalability problems for legacy half-duplex Ethernet networks, because as the number of connected devices and volume of network communications increased, collisions became more frequent, degrading performance.

A collision occurs when two devices transmit a packet onto the network at almost the exact same moment, causing them to overlap and thus mangling them. When this happens, each device must detect the collision and then retransmit their packet in its entirety. As more devices are attached to the same hub, and more hubs are interconnected, the chance that two nodes transmit at the same time increases, and collisions became more frequent. In addition, as the size of the network increases, the distance and time a packet is in transit over the network also increases, making collisions even more likely. Thus, it is necessary to keep the size of such networks very small to achieve acceptable levels of performance.

Although most modern "hubs" offer 100-Mbps full-duplex or gigabit connectivity (there are no half-duplex connections in gigabit networks—the Gigabit Ethernet standard is always full duplex) to address the collision issue, and actually do perform some type of switching, the basic behavior of a hub still cannot address the scaling problem of a single broadcast domain. For that reason, hubs are rarely if ever seen anymore in enterprise network environments. Thus, we'll say little more about them.

Switches

Switches are the evolved descendents of the network hub. From a network operation perspective, switches are layer two devices and routers are layer three devices (referring to their level of operation in the OSI stack), though as technology advances, switches are being built with capabilities at all seven layers of the OSI model, such as the UTM functions mentioned earlier.

Switches were developed to overcome the historical performance shortcomings of hubs. Switches are more intelligent devices that learn the various MAC addresses of connected devices and transmit packets only to the devices they are specifically addressed to. Since each packet is not rebroadcast to every connected device, the likelihood that two packets will collide is significantly reduced. In addition, switches provide a security benefit by reducing the ability to monitor or "sniff" another workstation's traffic. With a hub, every workstation would see all traffic on that hub; with a switch, every workstation sees only its own traffic.

A switched network cannot absolutely eliminate the ability to sniff traffic. An attacker can trick a local network segment into sending it another device's traffic with an attack known as ARP poisoning. ARP poisoning works by forging replies to ARP broadcasts.

For example, suppose malicious workstation Attacker wishes to monitor the traffic of workstation Victim, another host on the local switched network segment. To accomplish this, Attacker would broadcast an ARP packet onto the network containing Victim's IP address but Attacker's MAC address. Any workstation that receives this broadcast would update its ARP tables and thereafter would send all of Victim's traffic to Attacker. This ARP packet is commonly called a gratuitous ARP and is used to announce a new workstation attaching to the network. To avoid alerting Victim that something is wrong, Attacker would immediately forward any packets received for Victim to Victim. Otherwise Victim would soon wonder why network communications weren't working. The most severe form of this attack is where the Victim is the local router interface. In this situation, Attacker would receive and monitor all traffic .

d. Explain the features of firewall.

Today's firewalls are expected to do much more than simply block traffic based on the outward appearance of the traffic (such as the TCP or UDP port). As applications have become increasingly complex and adaptive, the firewall has become more sophisticated in an attempt to control those applications. You should expect at least the following capabilities from your firewall.

Application Awareness

The firewall must be able to process and interpret traffic at least from OSI layers three through seven. At layer three, it should be able to filter by IP address; at layer four by port; at layer five by network sessions; at layer six by data type, and, most significantly, at layer seven to properly manage the communications between applications.

Accurate Application

Fingerprinting The firewall should be able to correctly identify applications, not just based on their outward appearance, but by the internal contents of their network communications as well. Correct application identification is necessary to ensure that all applications are properly covered by the firewall policy configuration.

Granular Application

Control In addition to allowing or denying the communication among applications, the firewall also needs to be able to identify and characterize the features of applications so they can be managed appropriately. File transfer, desktop sharing, voice and video, and in-application games are examples of potentially unwanted features that the firewall should be able to control.

Bandwidth Management (QoS)

The Quality of Service (QoS) of preferred applications, which might include Voice over IP (VoIP) for example, can be managed through the firewall based on real-time network bandwidth availability. If a sporting event is broadcast live via streaming video on a popular web site, your firewall should be able to proactively limit or block access so all those people who want to watch it don't bring down your network. The firewall should integrate with other network devices to ensure the highest possible availability for the most critical services.

e. Explain the five different types of wireless attacks.

Since Wi-Fi primarily operates at layer two in the OSI stack, most of the attacks against it occur at layer two. But wireless attacks, such as jamming, can also occur at layer one. In this section, there are five types of wireless attacks.

Wired Side Leakage

Network attacks—whether on the wired or wireless network—typically begin with some form of reconnaissance. On wireless networks, reconnaissance involves promiscuously listening for wireless packets using a wireless sniffer so the attacker can begin to develop a footprint of the wireless network. We will ideally focus on layer two packets, whereby we are not connected to an access point. If the attacker were connected to an access point, then he or she could sniff layer three and above.

Broadcast and multicast traffic run rampant on most wired networks, using the protocols such as NetBIOS, OSPF, and HSRP among others that were designed to be chatty about their topology information because they were envisioned to be used only on protected internal networks. What many administrators don't realize is that when they connect wireless to their wired networks this broadcast and multicast traffic can leak into the wireless airspace if not properly segmented and firewalled. Most access points and wireless switches allow this traffic to leak into the airspace without being blocked.

Figure netwrk device traffic can leak on to the wireless air space.

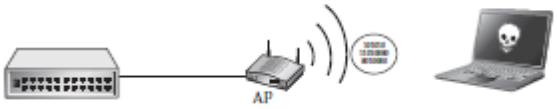

Rogue Access Points

The most common type of rogue access point involves a user who brings a consumer-grade access point like a Linksys router into the office. Many organizations attempt to detect rogue APs through wireless assessments. It is important to note that although you may detect access points in your vicinity, it is equally important to validate if they are connected to your physical network. The definition of a rogue AP is an unsanctioned wireless access point connected to your physical network. Any other visible AP that's not yours is simply a neighbouring access point.

What are the countermeasures against the possible abuse of wireless LAN.

Vetting out the potential rogue APs requires some prior knowledge of the legitimate wireless environment and sanctioned access points. This approach for detecting rogue APs involves determining the anomalous access points in the environment and, therefore, is really a best effort approach.

Misconfigured Access Points

Enterprise wireless LAN deployments can be riddled with misconfigurations. Human error coupled with different administrators installing the access points and switches can lead to a variety of misconfigurations. For example, an unsaved configuration change can allow a device to return to its factory default setting if, say, the device reboots during a power outage. And numerous other misconfigurations can lead to a plethora of vulnerabilities. Therefore, these devices must be monitored for configurations that are in line with your policies. Some of this monitoring can be done on the wired side with WLAN management products. Additionally, mature wireless IPS products can also monitor for misconfigured access points if you predefine a policy within the wireless IPS to monitor for devices not compliant with policy. Modern systems have different considerations—the controller-based approach largely prevents this issue, but some organizations, especially smaller ones, will still face this type of problem. Human error on the controller side poses a larger and more significant risk—all the access points will have a problem or configuration vulnerability, not just one.

Wireless Phishing

Since organizations are becoming more disciplined with fortifying their wireless networks, trends indicate that wireless users have become the low-hanging fruit. Enforcing secure Wi-Fi usage when it concerns human behavior is difficult. The average wireless user is simply not familiar with the threats imposed by connecting to an open Wi-Fi network at a local coffee shop or airport.

In addition, users may unknowingly connect to a wireless network that they believe is the legitimate access point but that has, in fact, been set up as a honeypot or open network specifically to attract unsuspecting victims. For example, they may have a network at home called "Linksys." As a result, their laptop may automatically connect to any other network known as "Linksys." This built-in behavior can lead to an

accidental association to a malicious wireless network, more commonly referred to as wireless phishing.

Once an attacker gains access to the user's laptop, not only could the attacker pilfer information such as sensitive files, but the attacker could also harvest wireless network credentials for the user's corporate network. This attack may be far easier to perform than attacking the enterprise network directly. If an attacker can obtain the credentials from a wireless user, he or she can then use those credentials to access the corporate enterprise wireless network, bypassing any encryption or safety mechanisms employed to prevent more sophisticated attacks.

Client Isolation

Users are typically the easiest target for attackers, especially when it comes to Wi-Fi. When users are associated to an access point, they can see others attempting to connect to the access point. Ideally, most users connect to the access point to obtain Internet access or access to the corporate network, but they can also fall victim to a malicious user of that same wireless network.

In addition to eavesdropping, a malicious user can also directly target other users as long as they're associated to the same access point. Specifically, once a user authenticates and associates to the access point, he or she obtains an IP address and, therefore, layer three access. Much like a wired network, the malicious wireless user is now on the same network as the other users of that access point, making them direct targets for attack.

f. What are the countermeasures against the possible abuse of wireless LAN.

These countermeasures include

- Secure replacements for WEP
- user authentication
- Intrusion detection and anomaly tracking on wireless LANs

Wireless Security Standards The IEEE "i" task group developed a unified wireless security standard, parts of which have been implemented by many wireless equipment and software manufacturers in order to mitigate known 802.11 security problems. Originally known as 802.11i, this standard is now widely known as WPA2, which stands for Wi-Fi Protected Access version 2. WPA2 replaced WPA, which was a hybrid of the old, insecure WEP standard that was backward compatible for existing wireless infrastructures. WPA used RC4 encryption, which is weaker than the AES encryption used in WPA2. WPA2 is the current, best security sol solution for wireless networks and is expected to remain so for the foreseeable future.

Temporal Key Integrity Protocol andCounter Mode with CBC-MAC Protocol

The WPA2 architecture can be split on two "layers:" encryption protocols and 802.11x port-based access control protocols. The Temporal Key Integrity Protocol (TKIP) and the Counter Mode with CBC-MAC Protocol (CCMP) are WPA2 encryption protocols on 802.11 LANs. TKIP encrypts each data packet with a unique encryption key. To increase key strength, TKIP includes four additional algorithms:

• A cryptographic message integrity check to protect packets
• An initialization-vector (IV) sequencing mechanism that includes hashing
• A per-packet key-mixing function to increase cryptographic strength

• A rekeying mechanism to provide key generation every 10,000 packets
802.1x-Based Authentication and EAP Methods
The 802.1x standard was originally designed to implement layer two user authentication on wired networks. On wireless networks, 802.1x can also be used for the dynamic distribution of WEP keys. Because wireless LANs have no physical ports, an association between the wireless client and the access point is assumed to be a network access port. In terms of 802.1x, the wireless client is defined as a *supplicant* (or *peer*), and the access point, as an *authenticator* (similar to an Ethernet switch on wired LANs).

Wireless Intrusion Detection and Prevention
The preceding points notwithstanding, intrusion detection on wireless networks should always cover the data-link layer. Many applications claim to be wireless IDS systems but detect new MAC addresses on a LAN only as long as these addresses are not permitted by an ACL. Such functionality is implemented in the firmware of some access points as well. Of course, anyone able to bypass MAC-based ACL will Bypass MAC-based "IDS." A true wireless IDS is a dedicated 802.11 (or 802.15) protocol analyzer supplied with an attack signature database or knowledge base and inference engine, as well as an appropriate report and alarm interface. Some suspicious events to look for on a wireless LAN include
• Probe requests (a good indication of someone using active scanning mode)
• Beacon frames from unsolicited access points or ad hoc wireless clients
• Floods of disassociate/deauthenticate frames (man-in-the-middle attack?)
Associated but not authenticated hosts (attempts to guess the shared key?)
• Frequent reassociation frames on networks without enabled roaming, and frequent packet retransmits ("hidden node," bad link, or possible DoS attack?)
• Multiple incorrect SSIDs on closed networks (SSID brute-forcing?)
• Suspicious SSIDs such as "AirJack" (or plain old "31337")
• Frames with unsolicited and duplicated MAC addresses
• Randomly changing MAC addresses (attackers using Wellenreiter or FakeAP)
• Frames transmitted on other 802.11 channels within the five-channel range, or frames with different SSIDs transmitted on the same channel (misconfigured and probably unsolicited host, interference, DoS?)
• Hosts not using implemented cryptographic solutions (should not be there)
• Multiple EAP authentication requests and responses (brute-forcing EAP-LEAP?)
• Malformed and oversized EAP frames and various EAP frame floods (802.1x DoS attack?)
• 802.11 frame sequence numbers that don't match the established sequence cycle (man-in-the-middle attacks, MAC spoofing on LAN?)
• ARP spoofing and other attacks originating from wireless LANs

4. Attempt *any three* of the following (15)

a. Explain Intrusion Defence System types and detection models.

Depending on what assets you want to protect, an IDS(Intrusion Defence System) can protect a host or a network. All IDSs follow one of two intrusion detection models—anomaly (also called profile, behavior, heuristic, or statistical) detection or signature (knowledge-based) detection— although some systems use parts of both when it's advantageous. Both anomaly and signature detection work by monitoring a wide population of events and triggering based on predefined behaviours.

Anomaly-Detection (AD) Model

Anomaly detection (AD) was proposed in 1985 by noted security laureate Dr. Dorothy E. Denning, and it works by establishing accepted baselines and noting exceptional differences. Some IDS vendors refer to AD systems as behavior-based since they look for deviating behaviors. If an IDS looks only at network packet headers for differences, it is called protocol anomaly detection.

Several IDSs have anomaly-based detection engines. Several massively distributed AD systems monitor the overall health of the Internet, and a handful of high-risk Internet threats have been minimized over the last few years because unusual activity was noticed by a large number of correlated AD systems.

The goal of AD is to be able to detect a wide range of malicious intrusions, including those for which no previous detection signature exists. By learning known good behaviors during a period of "profiling," in which an AD system identifies and stores all the normal activities that occur on a system or network, it can alert to everything else that doesn't fit the normal profile. Anomaly detection is statistical in nature and works on the concept of measuring the number of events happening in a given time interval for a monitored metric.

Following are some events AD systems can monitor and trigger alerts from:

- Unusual user account activity
- Excessive file and object access
- High CPU utilization
- Inappropriate protocol use
- Unusual workstation login location
- Unusual login frequency
- High number of concurrent logins
- High number of sessions
- Any code manipulation
- Unexpected privileged use or escalation attempts
- Unusual content

AD Advantages

AD systems are great at detecting a sudden high value for some metric. For example, when the SQL Slammer worm ate up all available CPU cycles and bandwidth on affected servers and networks within seconds of infection, you can bet AD systems went off. They did not need to wait until an antivirus vendor released an updated signature.

AD Disadvantages

Because AD systems base their detection on deviation from what's normal, they tend to work well in static environments, such as on servers that do the same thing day in and day out, or on networks where traffic patterns are consistent throughout the day. On more dynamic systems and networks that, therefore, have a wider range of normal behaviors, false positives can occur when the AD triggers on something that wasn't captured during the profiling period.

Signature-Detection Model

Signature-detection or misuse IDSs are the most popular type of IDS, and they work by using databases of known bad behaviors and patterns. This is nearly the exact opposite

of AD systems. When you think of a signature-detection IDS, think of it as an antivirus scanner for network traffic. Signature-detection engines can query any portion of a network packet or look for a specific series of data bytes. The defined patterns of code are called signatures, and often they are included as part of a governing rule when used within an IDS.

Signatures are byte sequences that are unique to a particular malady. A byte signature may contain a sample of virus code, a malicious combination of keystrokes used in a buffer overflow, or text that indicates the attacker is looking for the presence of a particular file in a particular directory. For performance reasons, the signature must be crafted so it is the shortest possible sequence of bytes needed to detect its related threat reliably. It must be highly accurate in detecting the threat and not cause false positives. Signatures and rules can be collected together into larger sets called signature databases or rule sets.

Signature-Detection Rules

Rules are the heart of any signature-detection engine. A rule usually contains the following information as a bare minimum:

- Unique signature byte sequence
- Protocol to examine (such as TCP, UDP, ICMP)
- IP port requested
- IP addresses to inspect (destination and source)
- Action to take if a threat is detected (such as allow, deny, alert, log, disconnect)

Advantages of Signature Detection

Signature-detection IDSs are proficient at recognizing known threats. Once a good signature is created, signature detection IDSs are great at finding patterns, and because they are popular, a signature to catch a new popular attack usually exists within hours.

Disadvantages of Signature Detection

Although signature-detection IDS are the most popular type of IDS, they have several disadvantages as compared to an AD IDS.

Cannot Recognize Unknown Attacks

Performance Suffers as Signatures or Rules Grow

b. Write a short note on Security Information and Event Management .

Multiple security systems can report to a centralized Security Information and Event Management (SIEM) system, bringing together logs and alerts from several disparate sources.

SIEM platforms take the log files, find commonalities (such as attack types and threat origination), and summarize the results for a particular time period. For example, all logs and alerts from all IDSs, perimeter firewalls, personal firewalls, antivirus scanners, and operating systems can be tied together. Events from all logs are then gathered, analyzed, and reported on from one location. SIEMs offer the ultimate in event correlation, giving you one place to get a quick snapshot of your system's security or to get trend information. SIEMs can also coordinate signature and product updates.

SIEMs have a huge advantage over individual IDS systems because they have the capability to collect and analyze many different sources of information to determine

what's really happening. As a result, the SIEM can significantly reduce false positives by verifying information based on other data. That data comes from many sources, including workstations, servers, computing infrastructure, databases, applications, network devices, and security systems. Because all those sources generate a vast amount of real-time data, SIEM products need to be fast and effective, with a significant amount of storage and computing power.

A SIEM is one of the most important tools used by security operations and monitoring staff, because it provides one-stop visibility into many different areas of the information processing environment and attacks against those areas.

SIEM can do

Data Aggregation

SIEMs collect information from every available source that is relevant to a security event. These sources take the form of alerts, real-time data, logs, and supporting data. Together, these provide the correlation engine of the SIEM with information it can use to make decisions about what to bring to the security administrator's attention. Consider the following examples of specific data sources consumed by a SIEM.

Alerts

When is an alert real, and when is it a false positive? This is the key question associated with an IDS, and a source of frustration for security administrators in charge of tuning IDSs. This is where a SIEM enters the picture. The SIEM's key function is to validate security alerts using many different sources of data to reduce false positives, so only the most reliable alerts get sent on to the security administrator. Thus, the alerts from all IDS sources as well as all other security monitoring systems should be given only to the SIEM, so it can decide which ones to pass along.

Real-Time Data

Real-time data such as network flow data (for instance, Cisco's NetFlow and similar traffic monitoring protocols from other vendors) gives the SIEM additional information to correlate. Streaming this data into the SIEM provides important information about normal and abnormal traffic patterns that can be used in conjunction with alerts to determine whether an attack is in progress. For example, an unusually high amount of SMTP traffic that accompanies several malware alerts may result in a high confidence alert that an e-mail.

Logs

Logs are different from events, in that they are a normal part of system activity and usually meant for debugging purposes. Logs can be an important additional data source for a SIEM, however. Logs contain valuable information about what's happening on a system, and they can give the SIEM a deeper view into what's happening.

Supporting Data

You can enhance the quality of a SIEM's correlation even more by providing the SIEM with supporting data that has been previously collected. Data can be imported into the SIEM, and it will use that data to make comparative determinations.

Analysis

A SIEM takes all the data given to it and makes decisions, so the security administrator can focus on the most important alerts. For this reason, event correlation is a SIEM's

most important feature. The correlation engine of every SIEM product is its most distinguishing feature.

c. What are Voice Over IP components. Explain it.

Call Control

The call control element (the "brains" of the operation) of a VoIP system can be either a purposed appliance, a piece of software that runs on a common or specialized server operating system, or a piece of network hardware embedded or integrated into another networking component such as a switch blade or software module (soft switch).

In the enterprise, the original IP phone systems were traditional digital time-division multiplexing (TDM) systems with an IP-enabled component, designed like digital systems. They eventually evolved into full IP-based systems (IPPBX). They have now evolved far beyond the early designs that mimicked the "old thinking" of voice networks by leveraging the tools and resiliency available in IP networking, high-availability server architecture, and virtualization..

Primarily responsible for call setup and teardown, signaling, device software serving, and feature configuration, call control is one of the easier pieces of the voice infrastructure to protect. This does not mean that security for this component should be taken lightly.

Voice and Media Gateways and Gatekeepers

The voice (or media) gateway is the pathway to the outside world. This component is what allows termination to a PSTN, transcoding between TDM and IP networks, media termination, and other types of analog/digital/IP interface required in today's multimediarich IP infrastructures. Gateways are configured to use dial peers (defined as "addressable endpoints") to originate and receive calls. Some gateways are directly managed by the call control elements via a control protocol (MGCP or H.248), whereas others operate in a more independent, stand-alone capacity (H.323 or SIP). Voice gateways can also run soft switches and perform primary (or survivable) call processing or "all-in-one" functions, an approach commonly used in the SMB space.

MCUs(Multi Conference Unit)

Conferencing and collaboration is used extensively within and across all enterprises as part of the fundamental communications capability that connects all users to each other. At the heart of this technology is the conference bridge, or multi-conference unit (MCU), a multiport bridging system for audio, video, and multimedia collaboration. The trend between internally hosted MCUs and provider-hosted MCUs has been stuck in the yoyo of corporate decision making, with each specific situation warranting one direction or the other based on cost to own, cost to operate, features, and security. Special attention should be paid to MCU functionality, whether they are hosted on premise or externally, in order to make sure they are secure.

Hardware Endpoints

Endpoint compromises today are frequently targeted at mobile devices, and much of the attention in the industry right now is focused on how to secure the mobile environment. The hardware phone or video codec, sitting quietly idle in the office but running 24/7, may, however, become an important tool for advanced corporate espionage, eavesdropping, or denial of service attacks. Modern VoIP phones have a fair

bit of intelligence built into them and offer a previously unavailable avenue—some phones have a built-in layer two switch and are capable of executing XML scripts or Java code locally. Video codecs run all kinds of custom code required for video conferencing and content sharing and are sometimes directly exposed to the Internet. None of these devices have particularly robust mechanisms for authenticating to their control components, unless a diligent administrator goes out of his or her way to enable them. Generally, these local capabilities are used to make the devices more interactive and functional, but they can be exploited in a variety of ways.

Software Endpoints

Enterprise desktop strategy focuses on convergence and extending simple, useful technologies to end users. This focus is intended to increase overall productivity and collaboration. One component of this strategy is the soft phone or voice and video-enabled chat client. This is a piece of software that runs on a PC or mobile device and acts like a hardware endpoint by registering to the call control element(s) as a device.

We are installing soft client on our mobile devise because of two reasons :one is cost and second by running the soft client, you can extend your enterprise features to the mobile user, including functionality not typically available on mobile devices such as consolidated extension-based or URI dialing. Some enterprises are even using direct inward system access (DISA) features or forking in order to make the mobile device itself an augmentation of the desk phone, creating a Single Number Reach (SNR) environment and automatically employing intelligent features like tail-end hop-off without direct user invocation

Voicemail Systems

The major component of a VoIP-based telephony system is the voicemail system. Auto attendants, direct inward system access (DISA) features used for manual call forwarding, automatic call forwarding, and other voicemail features are a “standard” component of enterprise life, which nearly everyone has come to expect and rely on. Unfortunately, they have historically been one of the easiest systems to abuse for three main reasons:

- Access to mailboxes is typically numeric-only, and people find long strings of numbers difficult to remember. Easy (and often default) passwords are commonplace. War dialers can be set up to target these systems and record successful logins for attackers to return to later. Anyone who has ever built a voicemail system knows the practice of initially setting everyone’s default password to their extension, or perhaps the last four digits of their direct inward dialing (DID) phone number, or some other easy-tofigure-out formula. This is a good opportunity to stretch your creative brain muscle and come up with something better.
- Since voicemail systems have never really been considered a “key” component of an enterprise infrastructure, much less attention has been paid to securing these systems than to, say, the enterprise ERP or financial systems. Keep in mind, access to this type of functionality in the wrong hands can cause permanent damage to an organization in financial (and worse) ways.

- More often than not system-level access to and from the outside world is not carefully controlled or audited, as some of a voicemail system's convenience "features" need outside access in order to work properly.

d. Write a short note on Private Bank Exchange(PBX).

A Private Branch Exchange (PBX) is a computer-based switch that can be thought of as a local phone company. Following are some common PBX features:

- Multiple extensions
- Voicemail
- Call forwarding
- Fax management
- Remote control (for support)

Hacking a PBX

Attackers hack PBXs for several reasons:

- To gain confidential information (espionage)
- To place outgoing calls that are charged to the organization's account (and thus free to the attacker)
- To cause damages by crashing the PBX

Common attacks:

Administrative Ports and Remote Access

Administrative ports are needed to control and diagnose the PBX. In addition, vendors often require remote access via a modem to be able to support and upgrade the PBX. This port is the number one hacker entry point. An attacker can connect to the PBX via the modem; or if the administrative port is shared with a voice port, the attacker can access the port from outside the PBX by calling and manipulating the PBX to reach the administrative port. Just as with administrative privileges for computers, when attackers have remote administrative privileges, "they own the box" and can use it to make international calls or shut down the PBX.

Voicemail

An attacker can gain information from voicemail or even make long-distance phone calls using a "through-dial" service. (After a user has been authenticated by the PBX, that user is allowed to make calls to numbers outside the PBX.) An attacker can discover a voicemail password by running an automated process that "guesses" easy passwords such as "1111," "1234," and so on.

Denial of Service

A PBX can be brought down in a few ways:

- PBXs store their voicemail data on a hard drive. An attacker can leave a long message, full of random noises, in order to make compression less effective—whereby a PBX might have to store more data than it anticipated. This can result in a crash.
- An attacker can embed codes inside a message. (For example, an attacker might embed the code for message rewinding. Then, while the user listens to the message, the PBX will decode the embedded command and rewind the message in an endless loop.)

Securing a PBX

Here is a checklist for securing a PBX:

- Connect administrative ports only when necessary.
- Protect remote access with a third-party device or a dial-back.
- Review the password strength of your users' passwords.
- Allow passwords to be different lengths, and require the # symbol to indicate the end of a password, rather than revealing the length of the password.
- Disable all through-dialing features.
- If you require dial through, limit it to a set of predefined needed numbers.
- Block all international calls, or limit the number of users who can initiate them.
- Block international calls to places such as the Caribbean that fraudsters tend to call.
- Train your help desk staff to identify attempted PBX hacks, such as excessive hangups, wrong number calls, and locked-out mailboxes.
- Make sure your PBX model is immune to common DoS attacks.

e. Explain different classic security model.

Classic Security Models

The most famous security models are Bell-LaPadula, Biba, and ClarkWilson. These three models are often mentioned in computing textbooks, and they form the foundation of most current operating system models. But practically speaking, most of them are little used in the real world, functioning only as security references. Those designing operating system security models have the liberty of picking and choosing from the best of what the famous models have, without being encumbered by their myriad details.

Bell-LaPadula

The Bell-LaPadula model was one of the first attempts to formalize an information security model. The Bell-LaPadula model was designed to prevent users and processes from reading above their security level. This is used within a data classification system—so a given classification cannot read data associated with a higher classification—as it focuses on sensitivity of data according to classification levels.

In addition, this model prevents objects and processes with any given classification from writing data associated with a lower classification. This aspect of the model caused a lot of consternation in the security space. Most operating systems assumed that the need to write below one's classification level is a necessary function. But the military influence on which Bell-LaPadula was created mandated that this be taken into consideration.

Biba

Biba is often known as a reversed version of Bell-LaPadula, as it focuses on integrity labels, rather than sensitivity and data classification. (Bell-LaPadula was designed to keep secrets, not to protect data integrity.)

Biba covers integrity levels, which are analogous to sensitivity levels in Bell-LaPadula, and the integrity levels cover inappropriate modification of data. Biba attempts to preserve the first goal of integrity, namely to prevent unauthorized users from modifying data.

Clark-Wilson

Clark-Wilson attempts to define a security model based on accepted business practices for transaction processing. Much more real-world-oriented than the other models described, it articulates the concept of well-formed transactions that

- Perform steps in order
- Perform exactly the steps listed
- Authenticate the individuals who perform the steps

f. Write a short note on trust worthy computing.

The four goals of the Trustworthy Computing initiative are

- **Security** As a customer, you can expect to withstand attack. In addition, you can expect the data is protected to prevent availability problems and corruption.
- **Privacy** You have the ability to control information about yourself and maintain privacy of data sent across the network.
- **Reliability** When you need your system or data, they are available.
- **Business integrit**y The vendor of a product acts in a timely and responsible manner, releasing security updates when a vulnerability is found.

To track and assure its progress in complying with the Trustworthy Computing initiative, Microsoft created a framework to explain its objectives: that its products be secure by design, secure by default, and secure in deployment, and that it provide communications .

Secure by design simply means that all vulnerabilities are resolved prior to shipping the product. Secure by design requires three steps.

1. Build a secure architecture. This is imperative. Software needs to be designed with security in mind first and then features.
2. Add security features. Feature sets need to be added to deal with new security vulnerabilities.
3. Reduce the number of vulnerabilities in new and existing code. The internal process at Microsoft was revamped to make developers more conscious of security issues while designing and developing software.

5. Attempt *<u>any three</u>* of the following (15)

a. Define virtual machine. How hypervisor responsible for managing all guest OS installations on a VM server.

In computing, a virtual machine is an emulation of a computer system. Virtual machines are based on computer architectures and provide functionality of a physical computer. Their implementations may involve specialized hardware, software, or a combination.

In addition to securing the VMs themselves, additional steps are needed to secure the virtual environment as a whole. The risks associated with VMs are a superset of those associated with physical servers along with a new set of risks based on the controllability of the individual virtual machines through a centralized management platform (sometimes referred to as a hypervisor or virtual machine monitor). National

Institute of Standards and Technology, or NIST, has published an excellent set of security practices for VMs in Special Publication 800-125.

The hypervisor is responsible for managing all guest OS installations on a VM server, and the service console provides a centralized location for managing all the servers in a virtual environment. As a result, a compromise of the hypervisor or service console has the potential to inflict significant damage as this would effectively allow all security controls on the virtual servers to be bypassed.

Hypervisor and service console servers need to be properly patched and secured, as well as logically separated through the use of isolated networks with strict access controls. The administration interfaces should reside on a network separate from the virtual machines themselves, one that is inaccessible from all VMs and other application servers on the network. Firewalls should be used to block access attempts from the virtual machines to the management consoles. This setup prevents attacks and malware on VMs from reaching the service consoles and affecting other VMs.

Because the hypervisor has so much power, and consequent damage and abuse potential, its administrative access should be strictly controlled.

Administrative access to the hypervisor is like having administrative access to all the VMs it controls. Any supervisory account for the hypervisor needs to be controlled in the same way you would protect privileged accounts for server and network administrator use. As with those other privileged accounts, consider using alternatives to passwords. A password associated with an administrative account for the hypervisor has the potential to be shared, or written down, despite your policies, threats, and warnings. The password may also be intercepted in various ways, such as by keyloggers or network sniffers. Password secrecy can never be guaranteed. Multifactor authentication—using tokens (portable digital one-time password generators), biometrics, and smart cards—is a better choice for hypervisor access. Limit physical access to the hardware as well. Despite any technical defenses that are in place, an attacker with physical access to the machine hardware is going to have an easier time getting into the system.

Limiting the number of administrators and their privileges is another practice that can reduce the risks of hypervisor attacks via administrator accounts. Hypervisor administrators should not use the same privileged accounts they also use to manage VMs and other systems, owing to the greater damage potential of hypervisors. Finally, someone other than the administrator, preferably someone with a security or audit function, should perform a periodic review of administrator activities. This check helps ensure that administrators haven't intentionally or inadvertently reduced system security level, altered the VMs, or cloned images inappropriately.

b. What is cloud computing. Explain the types of cloud services.

Cloud computing provides a way to increase capacity or add capabilities on the fly without investing in new infrastructure, training new personnel, or licensing new software. It encompasses any subscription-based or pay-per-use service that, in real time over the Internet, extends existing IT capabilities.

Types of Cloud Services

The term "cloud" is thrown around a lot these days, and it's used pretty loosely. Everybody wants to get in on the cloud phenomenon, so there are many types of services that get branded as cloud services. The following are the most common types of services with which we find the term "cloud" associated.

- **Infrastructure-as-a-Service (IaaS**) This type of service allows consumers to provision processing, storage, and networking resources, allowing them to deploy and run their own operating systems or applications in their own cloud environment.
- **Software-as-a-Service (SaaS)** This type of cloud computing delivers a single application through the browser to customers using a multitenant architecture.
- **Utility computing** Companies that offer storage and virtual servers that IT can access on demand. Early enterprise adopters mainly use utility computing for supplemental, non-mission-critical needs, but it is envisaged that one day it may replace parts of the data center.
- **Platform-as-a-Service (PaaS)** This form of cloud computing delivers development environments as a service. You build your own applications that run on the provider's infrastructure and are delivered to your users via the Internet from the provider's servers.
- **Web services** in the Cloud Web service providers offer APIs that enable developers to exploit functionality over the Internet, rather than delivering full-blown applications.
- **Managed service providers** (MSP) One of the oldest forms of cloud computing, a managed service is basically an application exposed to IT rather than to end users. Examples include virus scanning services, e-mail spam filtering services, application monitoring services, and managed security services.
- **Service commerce platforms** Similar to an automated service bureau and most common in trading environments, a service commerce platform is a service hub that users interact with, such as an expense management system, to order travel or secretarial services from a common platform that then coordinates the service delivery and pricing within the specifications set by the user.
- **Internet integration** The integration of cloud-based services mainly serving SaaS providers using in-the-cloud integration technology.

c. Explain the application security practices and decisions that appear in most secure development lifecycle.

Security Training

Typically, a security training program for development teams includes technical securityawareness training for everyone and role-specific training for most individuals. Role-specific training goes into more detail about the security activities a particular individual participates in, and the technologies in use (for developers).

Secure Development Infrastructure

At the beginning of a new project, source code repositories, file shares, and build servers must be configured for team members' exclusive access, bug tracking software must be configured to disclose security bugs only according to organization policies,

project contacts must be registered in case any application security issues occur, and licenses for secure development tools must be acquired.

Security Requirements

Security requirements may include access control matrices, security objectives (which specify actions attackers with specific privileges should not be able to perform), abuse cases, references to policies and standards, logging requirements, security bug bars, assignment of a security risk or impact level, and low-level security requirements such as key sizes or how specific error conditions should be handled.

Secure Design

Secure design activities usually revolve around secure design principles and patterns. They also frequently include adding information about security properties and responsibilities to design documents.

Threat Modeling

Threat modeling is a technique for reviewing the security properties of a design and identifying potential issues and fixes. Architects can perform it as a secure design activity, or independent design reviewers can perform it to verify architects' work. There is a variety of threat modeling methodologies to choose from.

Secure Coding

Secure coding includes using safe or approved versions of functions and libraries, eliminating unused code, following policies, handling data safely, managing resources correctly, handling events safely, and using security technology correctly.. Security

Code Review

To find security issues by inspecting application code, development teams may use static analysis tools, manual code review, or a combination. Static analysis tools are very effective at finding some kinds of mechanical security issues but are usually ineffective at finding algorithmic issues like incorrect enforcement of business logic. Static analysis tools usually require tuning to avoid high numbers of false positives. Manual code review by someone other than the code author is more effective at finding issues that involve code semantics, but requires training and experience. Manual code review is also time-consuming and may miss mechanical issues that require tracing large numbers of lines of code or remembering many details.

Security Testing

To find security issues by running application code, developers and independent testers perform repeatable security testing, such as fuzzing and regression tests for past security issues, and exploratory security testing, such as penetration testing.

Security Documentation

When an application will be operated by someone other than the development team, the operator needs to understand what security the application needs the deployment environment to provide, what settings can affect security, and how to handle any error messages that have security impact. The operator also needs to know if a release fixes any vulnerabilities in previous releases.

Secure Release Management

When an application will be shipped, it should be built on a limited-access build server and packaged and distributed in such a way that the recipients can verify it is

unchanged. Depending on the target platform, this may mean code signing or distributing signed checksums with the binaries.

Dependency Patch Monitoring

Any application that includes third-party code should monitor that external dependency for known security issues and updates, and issue a patch to update the application when any are discovered.

Product Security Incident Response

Like operational security incident response , product security incident response includes contacting people who should help respond, verifying and diagnosing the issue, figuring out and implementing a fix, and possibly managing public relations. It does not usually include forensics.

Decisions to Proceed

Any decision to ship an application or continue its development should take security into account. At ship time, the relevant question is whether the application can be reasonably expected to meet its security objectives. Frequently, this means that security validation activities have occurred and no critical or high-severity security issues remain open. Decisions to continue development should include some indicator of expected security risk, so that business stakeholders can draw conclusions regarding the expected business risk.

d. Explain the reasons for remote administration security. What are advantages of web remote administration.

Remote administration is needed for various reasons:

- Relocated servers An administrator needs an interface to administer any relocated web servers (computers that belong to an organization but that are physically located at the ISP).
- Outsourced services Managing security products requires knowledge that some organizations don't possess, so they often outsource their entire security management to a firm specializing in that area. In order to save costs, that firm needs to manage all the security products through the Internet.
- Physical distance An administrator may need to manage a large number of computers in the organization. Some organizations span several buildings (or cities), and physically attending the computers can be a tedious and time-consuming task. Additionally, physical access may be limited to the actual data centers.

There are some advantages of remote web administration:

- Quick development time Developing a web interface is faster than developing a GUI client, in terms of development, debugging, and deployment.
- OS support A web interface can be accessed from all the major OSs by using a browser (unless the developers used an OS-specific solution, like ActiveX, which only runs on Windows).
- Accessibility A web interface can be accessed from any location on the Internet. An administrator can administrate even if he's not in the office.
- User learning curve An administrator knows how to use a browser, so the learning curve for the administrator will be shorter.

e. Explain the security considerations for choosing a secure site location.

There are many security considerations for choosing a secure site location, a few of which are

- Accessibility •
  To the site
- From the site (in the event of evacuation)
- Lighting
- Proximity to other buildings
- Proximity to law enforcement and emergency response
- RF and wireless transmission interception
- Utilities reliability
- For a data center, the loss of power may be overcome through the use of generators, but if the water supply is cut off, the AC units will be unable to cool the servers
- Construction and excavation (past and present)

Accessibility

Accessibility of the site is typically the first consideration, and with good reason. If a site is located too remotely to be practical, usability and commutability are affected. However, by the same token, if the site is accessible easily to you, it probably is to others also. Conversely, you must consider potential evacuation. For example, bomb threats, fires, terrorist attacks, anthrax mailings, and SARS are potential catalysts for evacuation.

Lighting

Proper lighting, especially for organizations with 24×7 operations, should be evaluated and taken into consideration. Threats to employee safety, as well as the potential for break-ins, are more common under poor lighting conditions. Establish from the outset as many physical barriers between your business environment and undesirable people and circumstances as practical. Mirrored windows or windows with highly reflective coatings should face north-south rather than east-west to avoid casting sun glare into trafficked areas. Lighting should be positioned in such a way that it never blinds those leaving the building at night.

Proximity to Other Buildings

Know who your neighbors are. For instance, sharing a building with a branch of law enforcement would be considered less of a risk than sharing a building with "XYZ Computer Ch40s Klub." The closer the proximity to other buildings and companies, the higher the probability is for a physical security incident to occur. Also consider the fact that whatever problems an adjacent or connected building might have could potentially become your problem as well.

Proximity to Law Enforcement and Emergency Response

Another consideration is the location's relative proximity to law enforcement and/or emergency response units. If the area has a history of crime, but you've chosen the site anyway, consider the possibility that the incident may not get a response within a framework that you consider ideal. Similarly, if an emergency service unit were to be called to respond to an incident at this location, consider what the impact would be for any delay and if this latency in response would be justified.

RF and Wireless Transmission Interception

As wireless networking becomes more prevalent, especially in metropolitan areas, wireless hacking and hijacking become more of a threat. Other "airborne" protocols that should be taken into consideration include radio frequency devices, cordless phones, cell phones, PIMs, and mobile e-mail devices. Test drive for existing protocols with scanners, and avoid heavily trafficked frequency ranges wherever possible. Using encryption for sensitive traffic is indispensible.

Utilities Reliability

Office buildings provide work space for employees who need to be productive and reliable in their work. Power outages can seriously interfere with productivity, as can phone service and network outages. Some of these things can be compensated for, but some can't.

Construction and Excavation

Construction and excavation can take your entire network and communications infrastructure down with one fell swoop of a backhoe's bucket. Take a look at past construction activities in the area, and the impact (if any) that they had on the immediate vicinity. Town or city records will usually provide the information you need regarding any construction/excavation/ demolition, both past and present. Make it a point to ask people in the vicinity about power/ telecom outages.

f. Explain the different factors for securing the assets with physical security devices.

Locks

Locks aren't just for doors anymore. Anything of value that is capable of "growing legs and wandering away" should have a lock or be secured in a location that has a lock. Your physical security vulnerability assessment probably came across a few unsecured laptops, smartphones, tablets, MP3 players, jewelry, keys, and other assorted items. Lock up the device or valuable and make it a point to educate the asset owner on the importance of securing the item.

Doors and File Cabinets

Check for locked doors where applicable; you'll be surprised at the results. Make sure the lock on the door functions correctly and can withstand sufficient force. A broken or nonfunctioning lock is only slightly better than no lock at all. File cabinets containing sensitive information or valuable equipment should be kept locked when not in use. The keys to these should also be kept out of common reach.

Laptops

Laptops at the office, when not in transport, should be physically locked to the desk or in the docking station. Cable locks are a relatively small price to pay to ensure the laptop (and confidential information) doesn't fall into the wrong hands. Laptop theft is at an all-time high; most disappear right under the nose of the owner. One second it's here, the next it's gone. All personnel should be instructed to be especially wary when traveling with a laptop.

Data Centers, Wiring Closets, Network Rooms

All of these areas should have common access controls, as they all perform a similar function. Make sure these rooms are kept locked. If automatic entry-tracking mechanisms are not in use, ensure an access log is kept.

Entry Controls

Entry controls have their own security considerations that will undoubtedly vary with your security plan and business needs. When looking at the various options, you must first consider the site in which the entry controls will be deployed. Some of the most common types of deployment scenarios are for an existing structure with a single tenant, for a suite in a multitenant building, for a campus group of buildings with specific public entrances, and for a high-rise building.

Building Access Control Systems

For existing structures, there may be equipment already in place that can be reused. Multitenant buildings typically have access control systems that control entrance into the building or entrance to a special parking lot that is common to the entire building. If you plan to implement an access control system that is not compatible with an existing system, multiple access cards may be necessary. Many of the access control systems can support many of the card technologies, and there are even cards that support multiple types of technology and can work on several different incompatible systems.

Mantraps

A mantrap is an area designed to allow only one authorized individual entrance at any given time. These are typically used as an antitailgating mechanism—to prevent an unauthorized person from closely following an authorized person through an open door, for example— and are most commonly used in high-security areas, cash handling areas, and data centers.

Building and Employee Ids

Typically, one of the first things any organization does after hiring new employees is to provide them with ID badges. Building and/or employee identification should be displayed at all times, and anyone who lacks a visible ID should be challenged. Far too often, an individual becomes friendly with the security guard and, eventually, the guard just waves them through without showing valid identification.

Biometrics

Biometric devices have come a long way in the past several years and continue to gain traction both in the entry control market and the network authentication market. A biometric device is classified as any device that uses distinctive personally identifiable characteristics or unique physical traits to positively identify an individual.

Security Guards

The best deterrent seems to be security guards. But guards are not there merely as a deterrent. Here's what the New York State Department of Labor says a security guard's responsibilities include: "A security guard is employed by an organization, company, or agency to patrol, guard, monitor, preserve, protect, support, and maintain the security and safety of personnel and property.

==========