System Intrusion Detection and Prevention 13

13.1 Definition

The psychology and politics of ownership have historically dictated that individuals and groups tend to protect valuable resources. This grew out of the fact that once a resource has been judged to have value, no matter how much protection given to it, there is always a potential that the security provided for the resource will at some point fail. This notion has driven the concept of system security and defined the disciplines of computer and computer network security. Computer network security is made up of three principles: prevention, detection, and response. Although these three are fundamental ingredients of security, most resources have been devoted to detection and prevention because if we are able to detect all security threats and prevent them, then there is no need for response.

Intrusion detection is a technique of detecting unauthorized access to a computer system or a computer network. An intrusion into a system is an attempt by an outsider to the system to illegally gain access to the system. Intrusion prevention, on the other hand, is the art of preventing an unauthorized access of a system's resources. The two processes are related in a sense that while intrusion detection passively detects system intrusions, intrusion prevention actively filters network traffic to prevent intrusion attempts. For the rest of the chapter, let us focus on these two processes.

13.2 Intrusion Detection

The notion of intrusion detection in computer networks is a new phenomenon born, according to many, from a 1980 James Anderson's paper, "Computer Security Threat Monitoring and Surveillance." In that paper, Anderson noted that computer audit trails contained vital information that could be valuable in tracking misuse and understanding user behavior. The paper, therefore, introduced the concept of

J.M. Kizza, *Guide to Computer Network Security*, Computer Communications and Networks, DOI 10.1007/978-3-319-55606-2_13

"detecting" misuse and specific user events and has prompted the development of intrusion detection systems.

An *intrusion* is a deliberate unauthorized attempt, successful or not, to break into, access, manipulate, or misuse some valuable property and where the misuse may result into or render the property unreliable or unusable. The person who intrudes is an *intruder*.

Aurobindo Sundaram [1] divides intrusions into six types as follows:

- Attempted break-ins, which are detected by atypical behavior profiles or violations of security constraints. An intrusion detection system for this type is called anomaly-based IDS.
- Masquerade attacks, which are detected by atypical behavior profiles or violations of security constraints. These intrusions are also detected using anomaly-based IDS.
- Penetrations of the security control system, which are detected by monitoring for specific patterns of activity.
- Leakage, which is detected by atypical use of system resources.
- Denial of service, which is detected by atypical use of system resources.
- Malicious use, which is detected by atypical behavior profiles, violations of security constraints, or use of special privileges.

13.2.1 The System Intrusion Process

The intrusion process into a system includes a number of stages that start with the identification of the target, followed by reconnaissance that produces as much information about the target as possible. After enough information is collected about the target and weak points are mapped, the next job is to gain access into the system and finally the actual use of the resources of the system. Let us look at each one of these stages.

13.2.1.1 Reconnaissance

Reconnaissance is the process of gathering information about the target system and the details of its workings and weak points. Hackers rarely attack an organization network before they have gathered enough information about the targeted network. They gather information about the type of information used in the network, where it is stored, how it is stored, and the weak entry points to that information. They do the reconnaissance through system scanning for vulnerabilities.

Although vulnerability assessment is not intrusion, it is part of the intrusion process in that it proceeds the intrusion itself. Vulnerability assessment is an automated process in which a scanning program sends network traffic to all computers or selected computers in the network and expects receiving return traffic that will indicate whether those computers have known vulnerabilities. These vulnerabilities may include weaknesses in operating systems, application software, and protocols.

Through the years and as technology improved, vulnerability assessment itself has gone through several generations, including using code or script downloaded from the Internet or freely distributed that was compiled and executed for specific hardware or platforms.

Once they have identified the target system's vulnerability, then they just go in for a kill.

13.2.1.2 Physical Intrusion

Besides scanning the network for information that will eventually enable intruders to illegally enter an organization network, intruders also can enter an organization network masquerading as legitimate users. They do this through a number of ways ranging from acquiring special administrative privileges to low-privilege user accounts on the system. If the system doesn't have the latest security patches, it may not be difficult for the hacker to acquire these privileges. The intruder can also acquire remote access privileges.

13.2.1.3 Denial of Service

Denial-of-service (DoS) attacks are where the intruder attempts to crash a service (or the machine), overload network links, overload the CPU, or fill up the disk. The intruder is not trying to gain information, but to simply act as a vandal to prevent you from making use of your machine.

Common Denial-of-Service Attacks

- Ping of death sends an invalid fragment, which starts before the end of packet, but extends past the end of the packet.
- SYN flood sends a TCP SYN packet, which starts connections, very fast, leaving the victim waiting to complete a huge number of connections, causing it to run out of resources, and dropping legitimate connections.
- Land/Latierra sends a forged SYN packet with identical source/destination address/port so that the system goes into an infinite loop trying to complete the TCP connection.
- WinNuke sends an OOB/URG data on a TCP connection to port 139 (NetBIOS Session/SMB), which causes the Windows system to hang.

13.2.2 The Dangers of System Intrusions

The dangers of system intrusion manifests are many including the following:

- Loss of personal data that may be stored on a computer: Personal data loss means a lot and means different things to different people depending on the intrinsic value attached to the actual data lost or accessed. Most alarming in personal data loss is that the way digital information is lost is not the same as the loss of physical data. In physical data loss, you know that if it gets stolen, then somebody has it, so you may take precautions. For example, you may report to

the police and call the credit card issuers. However, this is not the same with digital loss because in digital loss, you may even never know that your data was lost. The intruders may break into the system and copy your data and you never know. The damage, therefore, from digital personal data loss may be far greater.

- Compromised privacy: These days more and more people are keeping a lot more of their personal data online either through the use of credit or debit cards; in addition, most of the information about an individual is stored online by companies and government organizations. When a system storing this kind of data is compromised, a lot of individual data gets compromised. This is because a lot of personal data is kept on individuals by organizations. For example, a mortgage company can keep information on your financial credit rating, social security number, bank account numbers, and a lot more. Once such an organization's network is compromised, there is much information on individuals that is compromised, and the privacy of those individuals is compromised as well.
- Legal liability: If your organization network has personal information of the customer and it gets broken into, thus compromising personal information that you stored, you are potentially liable for damages caused by a hacker either breaking into your network or using your computers to break into other systems. For example, if a hacker does two- or three-level hacking using your network or a computer on your network, you can be held liable. A two-level hacking involves a hacker breaking into your network and using it to launch an attack on another network.

13.3 Intrusion Detection Systems (IDSs)

An *intrusion detection system* (*IDS*) is a system used to detect unauthorized intrusions into computer systems and networks. Intrusion detection as a technology is not new; it has been used for generations to defend valuable resources. Kings, emperors, and nobles who had wealth used it in rather an interesting way. They built castles and palaces on tops of mountains and sharp cliffs with observation towers to provide them with a clear overview of the lands below where they could detect any attempted intrusion ahead of time and defend themselves. Empires and kingdoms grew and collapsed based on how well intrusions from the enemies surrounding them could be detected. In fact, according to the Greek legend of the Trojan Horse, the people of Crete were defeated by the Greeks because the Greeks managed to penetrate the heavily guarded gates of the city walls.

Through the years, intrusion detection has been used by individuals and companies in a variety of ways including erecting ways and fences around valuable resources with sentry boxes to watch the activities surrounding the premises of the resource. Individuals have used dogs, flood lights, electronic fences, closed-circuit television, and other watchful gadgets to be able to detect intrusions.

As technology has developed, a whole new industry based on intrusion detection has sprung up. Security firms are cropping up everywhere to offer individual and

property security—to be a watchful eye so that the property owner can sleep or take a vacation in peace. These new systems have been made to configure changes, compare user actions against known attack scenarios, and be able to predict changes in activities that indicate and can lead to suspicious activities.

In Sect. 13.2, we outlined six subdivisions of system intrusions. These six can now be put into three models of intrusion detection mechanisms: *anomaly-based* detection, *signature-based* detection, and *hybrid* detection. In anomaly-based detection, also known as behavior-based detection, the focus is to detect the behavior that is not normal or behavior that is not consistent with normal behavior. Theoretically, this type of detection requires a list of what is normal behavior. In most environments this is not possible, however. In real-life models, the list is determined from either historical or empirical data. However, neither historical nor empirical data represent all possible acceptable behavior. So a list has got to be continuously updated as new behavior patterns not on the list appear and are classified as acceptable or normal behavior. The danger with this model is to have unacceptable behavior included within the training data and later be accepted as normal behavior. Behavior-based intrusion detections, therefore, are also considered as rule-based detection because they use rules, usually developed by experts, to be able to determine unacceptable behavior.

In signature-based detection, also known as misuse-based detection, the focus is on the signature of known activities. This model also requires a list of all known unacceptable actions or misuse signatures. Since there are an infinite number of things that can be classified as misuse, it is not possible to put all these on the list and still keep it manageable. So only a limited number of things must be on the list. To do this and therefore be able to manage the list, we categorize the list into three broad activities:

- Unauthorized access
- Unauthorized modification
- Denial of service

Using these classifications, it is then possible to have a controlled list of misuse whose signatures can be determined. The problem with this model, though, is that it can detect only previously known attacks.

Because of the difficulties with both the anomaly-based and signature-based detections, a hybrid model is being developed. Much research is now focusing on this hybrid model [1].

13.3.1 Anomaly Detection

Anomaly based systems are “learning” systems in a sense that they work by continuously creating “norms” of activities. These norms are then later used to detect anomalies that might indicate an intrusion. Anomaly detection compares

observed activity against expected normal usage profiles "leaned." The profiles may be developed for users, groups of users, applications, or system resource usage.

In anomaly detection, it is assumed that all intrusive activities are necessarily anomalous. This happens in real life too, where most "bad" activities are anomalous, and we can, therefore, be able to character profile the "bad elements" in society. The anomaly detection concept, therefore, will create, for every guarded system, a corresponding database of "normal" profiles. Any activity on the system is checked against these profiles and is deemed acceptable or not based on the presence of such activity in the profile database.

Typical areas of interest are threshold monitoring, user work profiling, group work profiling, resource profiling, executable profiling, static work profiling, adaptive work profiling, and adaptive rule-based profiling.

Anonymous behaviors are detected when the identification engine takes observed activities and compares them to the rule-based profiles for significant deviations. The profiles are commonly for individual users, groups of users, system resource usages, and a collection of others as discussed below [2]:

- Individual profile: A collection of common activities a user is expected to do and with little deviation from the expected norm. This may cover specific user events such as the time being longer than usual usage, recent changes in user work patterns, and significant or irregular user requests.
- Group profile: This is a profile that covers a group of users with a common work pattern, resource requests and usage, and historic activities. It is expected that each individual user in the group follows the group activity patterns.
- Resource profile: This includes the monitoring of the use patterns of the system resources such as applications, accounts, storage media, protocols, communications ports, and a list of many others the system manager may wish to include. It is expected, depending on the rule-based profile, that common uses will not deviate significantly from these rules.
- Other profiles: These include executable profiles that monitor how executable programs use the system resources. This, for example, may be used to monitor strange deviations of an executable program if it has an embedded Trojan worm or a trapdoor virus. In addition to executable profiles, there are also the following profiles: work profile which includes monitoring the ports, static profile whose job is to monitor other profiles periodically updating them so that those profiles cannot slowly expand to sneak in intruder behavior, and a variation of the work profile called the adaptive profile which monitors work profiles, automatically updating them to reflect recent upsurges in usage. Finally, there is also the adoptive rule-based profile which monitors historic usage patterns of all other profiles and uses them to make updates to the rule base [3].

Besides being embarrassing and time-consuming, the concept also has other problems. As pointed out by Sundaram [1], if we consider that the set of intrusive activities only intersects the set of anomalous activities instead of being exactly the same, then two problems arise:

- Anomalous activities that are not intrusive are classified as intrusive.
- Intrusive activities that are not anomalous result in false negatives, that is, events are not flagged intrusive, though they actually are.

Anomaly detection systems are also computationally expensive because of the overhead of keeping track of, and possibly updating, several system profile metrics.

13.3.2 Misuse Detection

Unlike anomaly detection where we labeled every intrusive activity anomalous, the misuse detection concept assumes that each intrusive activity is representable by a unique pattern or a *signature* so that slight variations of the same activity produce a new signature and therefore can also be detected. Misuse detection systems are, therefore, commonly known as *signature systems*. They work by looking for a specific signature on a system. Identification engines perform well by monitoring these patterns of known misuse of system resources. These patterns, once observed, are compared to those in the rule base that describe "bad" or "undesirable" usage of resources. To achieve this, a knowledge database and a rule engine must be developed to work together. Misuse pattern analysis is best done by expert systems, model-based reasoning, or neural networks.

Two major problems arise out of this concept:

- The system cannot detect unknown attacks with unmapped and unarchived signatures.
- The system cannot predict new attacks and will, therefore, be responding after an attack has occurred. This means that the system will never detect a new attack.

In a computer network environment, intrusion detection is based on the fact that software used in all cyber attacks often leave a *characteristic signature*. This signature is used by the detection system, and the information gathered is used to determine the nature of the attack. At each different level of network investigative work, there is a different technique of network traffic information gathering, analysis, and reporting. Intrusion detection operates on already gathered and processed network traffic data. It is usually taken that the anomalies noticed from the analysis of this data would lead to distinguishing between an intruder and a legitimate user of the network. The anomalies resulting from the traffic analyses are actually large and noticeable deviations from historical patterns of usage. Identification systems are supposed to identify three categories of users: legitimate users, legitimate users performing unauthorized activities, and of course intruders who have illegally acquired the required identification and authentication.

13.4 Types of Intrusion Detection Systems

Intrusion detection systems are also classified based on their monitoring scope. There are those that monitor only a small area and those that can monitor a wide area. Those that monitor a wide area are known as network-based intrusion detection, and those that have a limited scope are known as host-based detections.

13.4.1 Network-Based Intrusion Detection Systems (NIDSs)

Network-based intrusion detection systems have the whole network as the monitoring scope. They monitor the traffic on the network to detect intrusions. They are responsible for detecting anomalous, inappropriate, or other data that may be considered unauthorized and harmful occurring on a network. There are striking differences between NIDS and firewalls. Recall from Chap. 11 that firewalls are configured to allow or deny access to a particular service or host based on a set of rules. Only when the traffic matches an acceptable pattern is it permitted to proceed regardless of what the packet contains. An NIDS also captures and inspects every packet that is destined to the network regardless of whether it is permitted or not. If the packet signature based on the contents of the packet is not among the acceptable signatures, then an alert is generated.

There are several ways an NIDS may be run. It can either be run as an independent standalone machine where it promiscuously watches over all network traffic or it can just monitor itself as the target machine to watch over its own traffic. For example, in this mode, it can watch itself to see if somebody is attempting a SYN flood or a TCP port scan.

While NIDSs can be very effective in capturing all incoming network traffic, it is possible that an attacker can evade this detection by exploiting ambiguities in the traffic stream as seen by the NIDS. Mark Handley, Vern Paxson, and Christian Kreibich list the sources of these exploitable ambiguities as follows [4]:

- Many NIDSs do not have complete analysis capabilities to analyze a full range of behavior that can be exposed by the user and allowed by a particular protocol. The attacker can also evade the NIDS: even if the NIDS does perform analysis for the protocol.
- Since NIDSs are far removed from individual hosts, they do not have full knowledge of each host's protocol implementation. This knowledge is essential for the NIDS to be able to determine how the host may treat a given sequence of packets if different implementations interpret the same stream of packets in different ways.
- Again, since NIDSs do not have a full picture of the network topology between the NIDS and the hosts, the NIDS may be unable to determine whether a given packet will even be seen by the hosts.

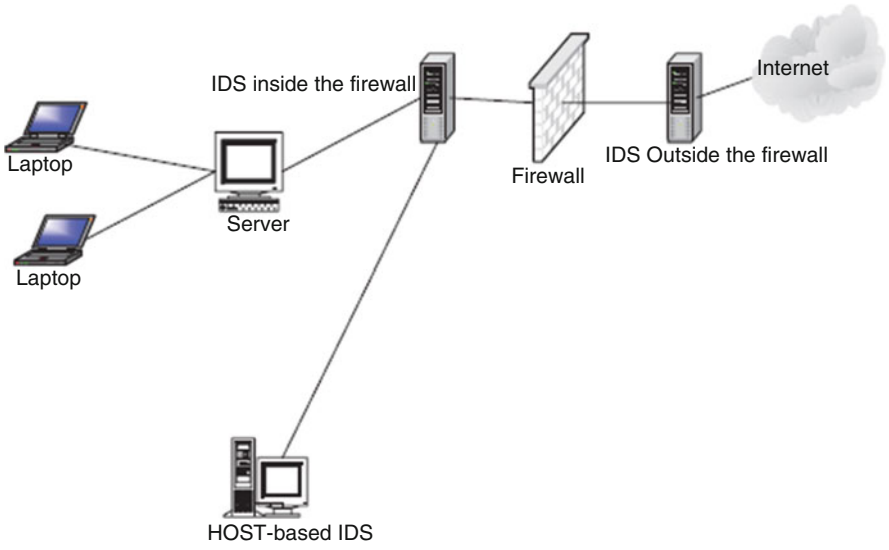

Fig. 13.1 The architecture of a network-based intrusion detection system

13.4.1.1 Architecture of a Network-Based Intrusion Detection

An intrusion detection system consists of several parts that must work together to produce an alert. The functioning of these parts may be either sequential or sometimes parallel [5, 6]. The parts are shown in Fig. 13.1.

Network Tap/Load Balancer

The network tap, or the load balancer as it is also known, gathers data from the network and distributes it to all network sensors. It can be a software agent that runs from the sensor or hardware, such as a router. The load balancer or tap is an important component of the intrusion detection system because all traffic into the network goes through it, and it also prevents packet loss in high-bandwidth networks. Certain types of taps have limitations in selected environments such as switched networks. In networks where there are no load balancers, sensors must be placed in such a way that they are responsible for traffic entering the network in their respective subnetwork.

Network Sensor/Monitoring

The network sensor or monitor is a computer program that runs on dedicated machines or network devices on mission critical segments. In networks with a load balancer, the sensors receive traffic from the balancer. In other networks without a load balancer, the sensors receive live traffic from the network and separate it between suspicious and normal traffic. A sensor can be implemented as an agent on a mission critical destination machine in a network. They are either anomaly based or signature based. Promiscuous mode sensors, which are sensors

that detect anything that seems like a possible attempt at intrusion, run on dedicated machines.

Analyzer

The analyzer determines the threat level based on the nature and threat of the suspicious traffic. It receives data from the sensors. The traffic is then classified as either safe or an attack. Several layers of monitoring may be done where the primary layer determines the threat severity; secondary layers then determine the scope, intent, and frequency of the threat.

Alert Notifier

It contacts the security officer responsible for handling incidents whenever a threat is severe enough according to the organization's security policy. Standard capabilities include on-screen alerts, audible alerts, paging, and e-mail. Most systems also provide SNMP so that an administrator can be notified. Frequent alerts for seemingly trivial threats must be avoided because they result in a high rate of false positives. It must also be noted that not reporting frequently enough because the sensors are set in such a way that they ignore a number of threats, many of them being real, result in false negatives which results in the intrusion detection system providing misleading sense of security.

Because the performance of the intrusion detection system depends on the balancing of both false positives and false negatives, it is important to use intrusion detection systems that are adjustable and can, therefore, offer balancing capabilities.

Command Console/Manager

The role of the command console or manager is to act as the central command authority for controlling the entire system. It can be used to manage threats by routing incoming network data to either a firewall or to the load balancer or straight to routers. It can be accessed remotely so the system may be controlled from any location. It is typically a dedicated machine with a set of tools for setting policy and processing collected alarms. On the console, there is an assessment manager, a target manager, and an alert manager. The console has its own detection engine and database of detected alerts, for scheduled operations and data mining.

Response Subsystem

The response subsystem provides the capabilities to take action based on threats to the target systems. These responses can be automatically generated or initiated by the system operator. Common responses include reconfiguring a router or a firewall and shutting down a connection.

Database

The database is the knowledge repository for all that the intrusion detection system has observed. This can include both behavioral and misuse statistics. These statistics are necessary to model historical behavior patterns that can be useful during

damage assessment or other investigative tasks. Useful information need not necessarily be indicative of misuse. The behavioral statistics help in developing the patterns for the individual, and the misuse statistics aid in detecting attempts at intrusion.

13.4.1.2 Placement of IDS Sensors

The position to place network IDS sensors actually depends on several factors, including the topology of the internal network to be protected, the kind of security policy the organization is following, and the types of security practices in effect. For example, you want to place sensors in places where intrusions are most likely to pass. These are the network "weak" points. However, it is normal practice to place IDS sensors in the following areas [6]:

- Inside the DMZ: We saw in Chap. 11 that the DMZ is perhaps the most ideal place to put any detection system because almost all attacks enter the protected internal network through the DMZ. IDS sensors are, therefore, commonly placed outside of the organization's network's first firewall in the DMZ. The IDS sensors in the DMZ can be enhanced by putting them into zoned areas. Another good location for IDS sensors is inside each firewall. This approach gives the sensors more protection, making them less vulnerable to coordinated attacks. In cases where the network perimeter does not use a DMZ, the ideal locations then may include any entry/exit points such as on both sides of the firewall, dial-up servers, and on links to any collaborative networks. These links tend to be low bandwidth (T1 speeds) and are usually the entry point of an external attack.
- Between the firewall and the Internet: This is a frequent area of unauthorized activity. This position allows the NIDS to "see" all Internet traffic as it comes into the network. This location, however, needs a good appliance and sensors that can withstand the high volume of traffic.
- Behind the network front firewall: This is a good position; however, most of the bad network traffic has already been stopped by the firewall. It handles all the bad traffic that manages to get through the firewall.
- Inside the network: Commonly placed in strategic points and used to "see" segments of the network. Network segments like these are usually the suspected weak areas of the network. The problem with this approach, however, is that the sensors may not be able to cover all the targets it is supposed to. Also it may cause the degradation of the network performance.

Figure 13.2 shows the various places where ID sensors can be deployed.

13.4.1.3 Advantages of Network-Based Intrusion Detection Systems

Although both NIDSs and HIDSs (Sect. 13.4.2) have different focuses, areas of deployment, and deployment requirements, using NIDS has the following advantages [7]:

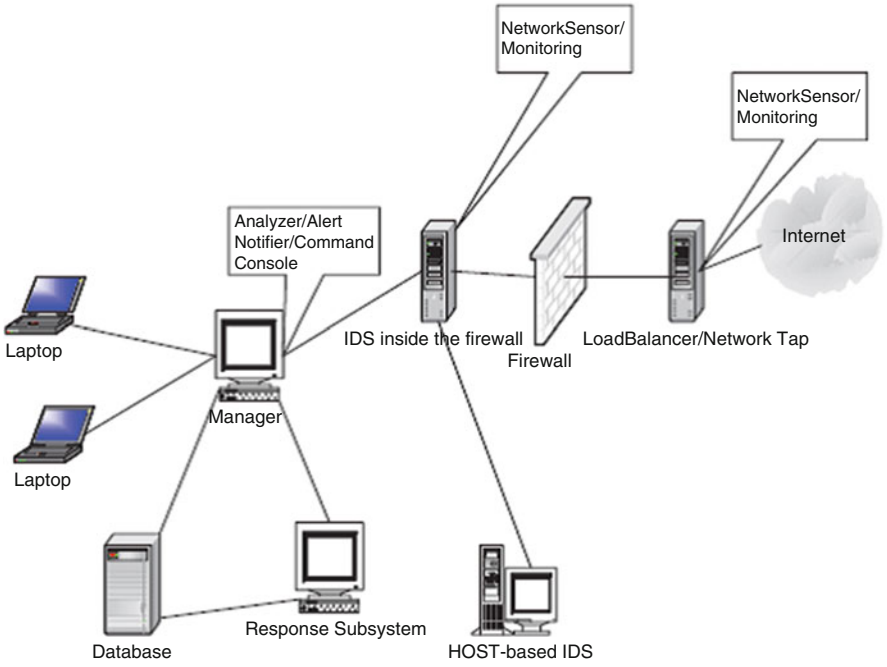

Fig. 13.2 The various places of placing the IDS sensors

- Ability to detect attacks that a host-based system would miss: Because NIDSs are on dedicated machines that are routinely protected, it is more difficult for an **attacker** to remove the evidence than it is with HIDSs which are near or at the attacker's desk. Also, since NIDSs use live network traffic and it is this traffic that is captured by NIDSs when there is an attack, this also makes it difficult for an attacker to remove evidence.
- Difficulty to remove evidence: Because NIDSs are on dedicated machines that are routinely protected, it is more difficult for an attack to remove the evidence than it is with HIDSs which are near or at the attacker's desk. Also, since NIDSs use live network traffic and it is this traffic that is captured by NIDSs when there is an attack, this also makes it difficult for an attacker to remove evidence.
- Real-time detection and response: Because the NIDSs are at the most opportune and strategic entry points in the network, they are able to detect foreign intrusions into the network in real time and report as quickly as possible to the administrator for a quick and appropriate response. Real-time notification, which many NIDSs have now, allows for a quick and appropriate response and can even let the administrators allow the intruder more time as they do more and targeted surveillance.
- Ability to detect unsuccessful attacks and malicious intent: Because the HIDSs are inside the protected internal network, they never come into contact with many types of attack since such attacks are many times stopped by the outside firewall. NIDSs, especially those in the DMZ, come across these attacks (those

that escape the first firewall) that are later rejected by the inner firewall and those targeting the DMZ services that have been let in by the outer firewall. Besides showing these attacks, NIDSs can also record the frequency of these attacks.

13.4.1.4 Disadvantages of NIDS

Although NIDSs are very well suited to monitor all the networks coming into the network, they have limitations [8]:

- Blind spots: Deployed at the borders of an organization network, NIDS are blind to the whole inside network. As sensors are placed in designated spots, especially in switched networks, NIDSs have blind spots—sometimes whole network segments they cannot see.
- Encrypted data: One of the major weaknesses of NIDS is on encrypted data. They have no capabilities to decrypt encrypted data. Although they can scan unencrypted parts of the packet such as headers, they are useless to the rest of the package.

13.4.2 Host-Based Intrusion Detection Systems (HIDS)

Recent studies have shown that the problem of organization information misuse is not confined only to the "bad" outsiders, but the problem is more rampant within organizations. To tackle this problem, security experts have turned to inspection of systems within an organization network. This local inspection of systems is called *host-based intrusion detection systems* (HIDS).

Host-based intrusion detection is the technique of detecting malicious activities on a single computer. A host-based intrusion detection system is, therefore, deployed on a single target computer, and it uses software that monitors operating system-specific logs, including system, event, and security logs on Windows systems and syslog in Unix environments to monitor sudden changes in these logs. When a change is detected in any of these files, the HIDS compares the new log entry with its configured attack signatures to see if there is a match. If a match is detected, then this signals the presence of an illegitimate activity.

Although HIDSs are deployable on a single computer, they can also be put on a remote host, or they can be deployed on a segment of a network to monitor a section of the segment. The data gathered, which sometimes can be overwhelming, is then compared with the rules in the organization's security policy. The biggest problem with HIDSs is that given the amount of data logs generated, the analysis of such raw data can put significant overhead not only on the processing power needed to analyze this data but also on the security staff needed to review the data.

Host sensors can also use user-level processes to check key system files and executables to periodically calculate their checksum and report changes in the checksum.

13.4.2.1 Advantages of Host-Based Intrusion Detection Systems

HIDSs are new kids on the intrusion detection block. They came into widespread in use in the early and mid-1980s when there was a realization after studies showed that a large number of illegal and illegitimate activities in organization networks actually originated from within the employees. Over the succeeding years as technology advanced, the HIDS technology has also advanced in tandem. More and more organizations are discovering the benefits of HIDSs on their overall security. Besides being faster than their cousins the NIDSs, because they are dealing with less traffic, they offer additional advantages including the following [7]:

- Ability to verify success or failure of an attack quickly: Because they log continuing events that have actually occurred, they have information that is more accurate and less prone to false positives than their cousins, the NIDSs. This information can accurately infer whether an attack was successful or not quickly, and a response can be started early. In this role, they complement the NIDSs, not as an early warning but as a verification system.
- Low-level monitoring: Because they monitor at a local host, they are able to "see" low-level local activities such as file accesses, changes to file permissions, attempts to install new executables or attempts to access privileged services, changes to key system files and executables, and attempts to overwrite vital system files or to install Trojan horses or backdoors. These low-level activities can be detected very quickly, and the reporting is quick and timely to give the administrator time for an appropriate response. Some of these low-level attacks are so small and far less intensive such that no NIDS can detect them.
- Near real-time detection and response: HIDSs have the ability to detect minute activities at the target hosts and report them to the administrator very quickly at a rate near real time. This is possible because the operating system can recognize the event before any IDS can, in which case, an intruder can be detected and stopped before substantial damage is done.
- Ability to deal with encrypted and switched environments: Large networks are routinely switch-chopped into many but smaller network segments. Each one of these smaller networks is then tagged with a NIDS. In a heavily switched network, it can be difficult to determine where to deploy a network-based IDS to achieve sufficient network coverage. This problem can be solved by the use of traffic mirroring and administrative ports on switches, but not as effective. HIDS provides this needed greater visibility into these switched environments by residing on as many critical hosts as needed. In addition, because the operating systems see incoming traffic after encryption has already been de-encrypted, HIDSs that monitor the operating systems can deal with these encryptions better than NIDSs, which sometimes may not even deal with them at all.
- Cost-effectiveness: Because no additional hardware is needed to install HIDS, there may be great organization savings. This compares favorably with the big costs of installing NIDS which require dedicated and expensive servers. In fact,

in large networks that are switch-chopped which require a large number of NIDSs per segment, this cost can add up.

13.4.2.2 Disadvantages of HIDS

Like their cousin the NIDS, HIDSs have limitations in what they can do. These limitations include the following [8]:

- Myopic viewpoint: Since they are deployed at a host, they have a very limited view of the network.
- Since they are close to users, they are more susceptible to illegal tampering.

13.4.3 The Hybrid Intrusion Detection System

We have noted in both Sects. 13.4.1 and 13.4.2 that there was a need for both NIDS and HIDS, each patrolling its own area of the network for unwanted and illegal network traffic. We have also noted the advantages of not using one over the other but of using one to complement the other. In fact, if anything, after reading Sects. 13.4.1.3 and 13.4.2.1, one comes out with an appreciation of how complementary these two intrusion detection systems are. Both bring to the security of the network their own strengths and weaknesses that nicely complement and augment the security of the network.

However, we also know and have noted in Sect. 13.4.1.4 that NIDSs have been historically unable to work successfully in switched and encrypted networks, and as we have also noted in Sect. 13.4.2.2, both HIDS and NIDS have not been successful in high-speed networks—networks whose speeds exceed 100 Mbps. This raises the question of a hybrid system that contains all the things that each system has and those that each system misses, a system with both components. Having both components provides greater flexibility in their deployment options.

Hybrids are new and need a great deal of support to gain on their two cousins. However, their success will depend to a great extent on how well the interface receives and distributes the incidents and integrates the reporting structure between the different types of sensors in the HIDS and NIDS spheres. Also the interface should be able to smartly and intelligently gather and report data from the network or systems being monitored.

The interface is so important and critical because it receives data, collects analysis from the respective component, coordinates and correlates the interpretation of this data, and reports it. It represents a complex and unified environment for tracking, reporting, and reviewing events.

13.5 The Changing Nature of IDS Tools

Although ID systems are assumed, though wrongly, by management and many in the network community that they protect network systems from outside intruders, recent studies have shown that the majority of system intrusions actually come from insiders. So newer IDS tools are focusing on this issue. Also, since the human mind is the most complicated and unpredictable machine ever, as new IDS tools are being built to counter systems intrusion, new attack patterns are being developed to take this human behavior unpredictability into account. To keep abreast of all these changes, ID systems must be changing constantly.

As all these changes are taking place, the primary focus of ID systems has been on a network as a unit where they collect network packet data by watching network packet traffic and then analyzing it based on network protocol patterns "norms," "normal" network traffic signatures, and network traffic anomalies built in the rule base. But since networks are getting larger, traffic heavier, and local networks more splintered, it is becoming more and more difficult for the ID system to "see" all traffic on a switched network such as an Ethernet. This has led to a new approach to looking closer at the host. So in general, ID systems fall into two categories: host based and network based.

13.6 Other Types of Intrusion Detection Systems

Although NIDS and HIDS and their hybrids are the most widely used tools in network intrusion detection, there are others that are less used but more targeting and therefore more specialized. Because many of these tools are so specialized, many are still not considered as being intrusion detection systems, but rather intrusion detection add-ons or tools.

13.6.1 System Integrity Verifiers (SIVs)

System integrity verifiers (SIVs) monitor critical files in a system, such as system files, to find whether an intruder has changed them. They can also detect other system components' data; for example, they detect when a normal user somehow acquires root/administrator level privileges. In addition, they also monitor system registries in order to find well-known signatures [9].

13.6.2 Log File Monitors (LFM)

Log file monitors (LFMs) first create a record of log files generated by network services. Then they monitor this record, just like NIDS, looking for system trends, tendencies, and patterns in the log files that would suggest that an intruder is attacking.

13.6.3 Honeypots

A *honeypot* is a system designed to look like something that an intruder can hack. They are built for many purposes, but the overriding one is to deceive attackers and learn about their tools and methods. Honeypots are also add-on/tools that are not strictly sniffer-based intrusion detection systems like HIDS and NIDS. However, they are good deception systems that protect the network in much the same way as HIDS and NIDS. Since the goal for a honeypot is to deceive intruders and learn from them without compromising the security of the network, then it is important to find a strategic place for the honeypot.

To many, the best location to achieve this goal is in the DMZ for those networks with DMZs or behind the network firewall if the private network does not have a DMZ. The firewall location is ideal because of the following [5]:

- Most firewalls log all traffic going through it; hence, this becomes a good way to track all activities of the intruders. By reviewing the firewall logs, we can determine how the intruders are probing the honeypot and what they are looking for.
- Most firewalls have some alerting capability, which means that with a few additions to the firewall rule base, we can get timely alerts. Since the honeypot is built in such a way that no one is supposed to connect to it, any packets sent to it are most likely from intruders probing the system. And if there is any outgoing traffic coming from the honeypot, then the honeypot is most likely compromised.
- The firewall can control incoming and outgoing traffic. This means that the intruders can find, probe, and exploit our honeypot, but they cannot compromise other systems.

So any firewall dedicated as a honeypot can do as long as it can control and log traffic going through it. If no firewall is used, then dedicate any machine either in the DMZ or behind a firewall for the purpose of logging all attempted accesses. Figure 13.3 shows the positioning of a honeypot.

Honeypots come in a variety of capabilities from the simplest monitoring one to two intruder activities to the most powerful monitoring many intruder activities. The simplest honeypot is a port monitor which is a simple socket-based program that opens up a listening port. The program can listen to any designed port. For example, *NukeNabbe*, for Windows, listens on ports typically scanned for by hackers. It then alerts the administrator whenever such designated ports are being scanned. The second type of honeypot is the deception system, which instead of listening quietly on a port, interacts with the intruder, responding to him or her as if it were a real server with that port number. Most deception systems implement only as much of the protocol machine as necessary to trap 90% of the attacks against the protocol [9]. The next type of honeypot is the multiprotocol deception system which offers most of the commonly hacked protocols in a single toolkit. Finally, there is a full system that goes beyond what the deception systems do to incorporate

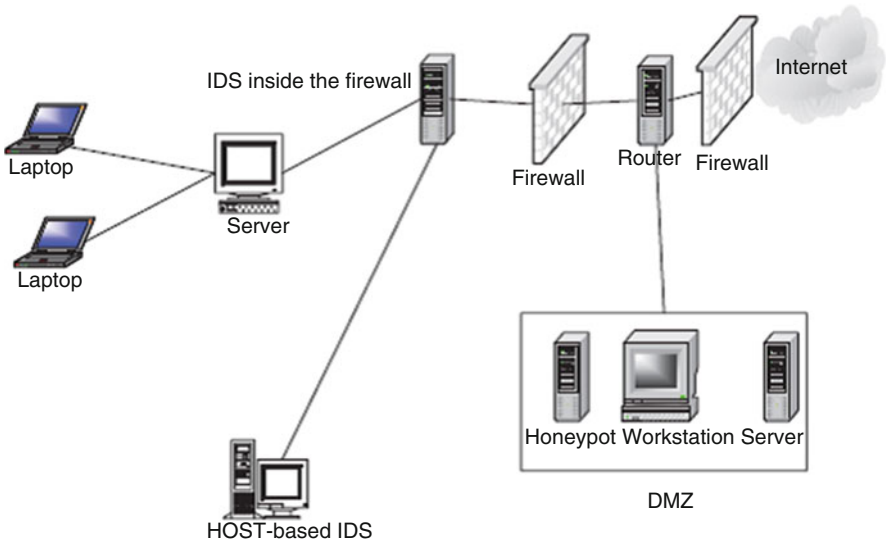

Fig. 13.3 The positioning of a honeypot

the ability to alert the system administrator on any exceptional condition. Other more complex honeypots combine a full system with NIDSs to supplement the internal logging [9].

13.6.3.1 Advantages of Honeypots

Perhaps one would wonder why a system administrator would go through the pain of setting up, maintaining, and daily responding to honeypots. There are advantages to having honeypots on a network. They include the following [9]:

- Since NIDSs have difficulties distinguishing between hostile and nonhostile activities, honeypots are more suited to digging out hostile intrusions because isolated honeypots should not normally be accessed. So if they are accessed at all, such accesses are unwanted intrusions and they should be reported.
- A honeypot can attract would-be hackers into the trap by providing a banner that looks like a system that can easily be hacked.

13.7 Response to System Intrusion

A good intrusion detection system alert should produce a corresponding response. The type of response is relative to the type of attack. Some attacks do not require responses; others require a precautionary response. Yet others need a rapid and forceful response. For the most part, a good response must consist of preplanned defensive measures that include an incident response team and ways to collect IDS logs for future use and for evidence when needed.

13.7.1 Incident Response Team

An *incident response team* (IRT) is a primary and centralized group of dedicated people charged with the responsibility of being the first contact team whenever an incidence occurs. According to Keao [6], an IRT must have the following responsibilities:

- Keeping up-to-date with the latest threats and incidents
- Being the main point of contact for incident reporting
- Notifying others whenever an incident occurs
- Assessing the damage and impact of every incident
- Finding out how to avoid exploitation of the same vulnerability
- Recovering from the incident

In handling an incident, the team must carefully do the following:

- Prioritize the actions based on the organization's security policy but taking into account the following order:
  - Human life and people's safety
  - Most sensitive or classified data
  - Costly data and files
  - Preventing damage to systems
  - Minimizing the destruction to systems
- Assess incident damage: This is through doing a thorough check on all the following: system log statistics, infrastructure and operating system checksum, system configuration changes, changes in classified and sensitive data, traffic logs, and password files.
- Alert and report the incident to relevant parties: These may include law enforcement agencies, incident reporting centers, company executives, employees, and sometimes the public.
- Recovering from incident: This involves making a postmortem analysis of all that went on. This postmortem report should include steps to take in case of similar incidents in the future.

13.7.2 IDS Logs as Evidence

First and foremost, IDS logs can be kept as a way to protect the organization in case of legal proceedings. Some people tend to view IDS as a form of wiretap. If sensors to monitor the internal network are to be deployed, verify that there is a published policy explicitly stating that the use of the network is consent to monitoring.

13.8 Challenges to Intrusion Detection Systems

While IDS technology has come a long way and there is an exciting future for it as the marriage between it and artificial intelligence takes hold, it faces many challenges. Although there are IDS challenges in many areas, more serious challenges are faced in deploying IDSs in switched environments.

13.8.1 Deploying IDS in Switched Environments

There is a particularly hard challenge faced by organizations trying to deploy IDS in their networks. Network-based IDS sensors must be deployed in areas where they can "see" network traffic packets. However, in switched networks, this is not possible because by their very nature, sensors in switched networks are shielded from most of the network traffic. Sensors are allowed to "see" traffic only from specified components of the network.

One way to handle this situation has traditionally been to attach a network sensor to a mirror port on the switch. But port mirroring, in addition to putting an overhead on the port, gets unworkable when there is an increase in traffic on that port because overloading one port with traffic from other ports may cause the port to bulk and miss some traffic.

Several solutions have been used recently including the following [8]:

- Tapping: This involves deploying a line of passive taps that administrators can tap into to listen in on Ethernet connections; by sending "copies" of the frames to a second switch with dedicated IDS sensor, overloading a port can be avoided.
- By using standard Cisco access control lists (ACL) in a Cisco appliance that includes a Cisco Secure IDS, one can tag certain frames for inspection.

Among other issues still limiting IDS technology are [2]:

- False alarms: Though the tools have come a long way and are slowly gaining acceptance as they gain widespread use, they still produce a significant number of both false positives and negatives.
- The technology is not yet ready to handle a large-scale attack. Because of its very nature, it has to literally scan every packet, every contact point, and every traffic pattern in the network. For larger networks and in a large-scale attack, it is not possible that the technology can be relied on to keep working with acceptable quality and grace.
- Unless there is a breakthrough today, the technology in its current state cannot handle very fast and large quantities of traffic efficiently.
- Probably the biggest challenge is the IDS's perceived and sometimes exaggerated capabilities. The technology, while good, is not the cure of all computer network ills that it is pumped up to be. It is just like any other good security tool.

13.9 Implementing an Intrusion Detection System

An effective IDS does not stand alone. It must be supported by a number of other systems. Among the things to consider, in addition to the IDS, in setting up a good IDS for the company network are the following [9]:

- Operating systems: A good operating system that has logging and auditing features. Most of the modern operating systems including Windows, Unix, and other variants of Unix have these features. These features can be used to monitor security critical resources.
- Services: All applications on servers such as Web servers, e-mail servers, and databases should include logging/auditing features as well.
- Firewalls: As we discussed in Chap. 11, a good firewall should have some network intrusion detection capabilities. Set those features.
- Network management platform: Whenever network management services such as OpenView are used, make sure that they do have tools to help in setting up alerts on suspicious activity.

13.10 Intrusion Prevention Systems (IPSs)

Although IDS have been one of the cornerstones of network security, they have covered only one component of the total network security picture. They have been and they are a passive component which only detects and reports without preventing. A promising new model of intrusion is developing and picking up momentum. It is the *intrusion prevention system* (IPS), which according to Andrew Yee [10] is to prevent attacks. Like their counterparts, the IDS, IPS fall into two categories: network based and host based.

13.10.1 Network-Based Intrusion Prevention Systems (NIPSs)

Because NIDSs are passively detecting intrusions into the network without preventing them from entering the networks, many organizations in recent times have been bundling up IDS and firewalls to create a model that can detect and then prevent.

The bundle works as follows. The IDS fronts the network with a firewall behind it. On the detection of an attack, the IDS then goes into the prevention mode by altering the firewall access control rules on the firewall. The action may result in the attack being blocked based on all the access control regimes administered by the firewall. The IDS can also affect prevention through the TCP resets; TCP utilizes the RST (reset) bit in the TCP header for resetting a TCP connection, usually sent as a response request to a nonexistent connection [10]. But this kind of bundling is both expensive and complex, especially to an untrained security team. The model suffers from *latency*—the time it takes for the IDS to either modify the firewall rules

or issue a TCP reset command. This period of time is critical in the success of an attack.

To respond to this need, a new technology, the IPS, is making its way into the network security arena to address this latency issue. It does this by both the intrusion detection system in-line with the firewall. Like in NIDS, NIPS architecture varies from product to product, but there is a basic underlying structure to all. These include traffic normalizer, system service scanner, detection engine, and traffic shaper [10].

13.10.1.1 Traffic Normalizer

The normalizer is in the line of network traffic to intercept traffic, resolving the traffic that has abnormalities before it sends it on. As it normalizes traffic, it may come to a point where it will discard the packet that does not conform to the set security policy criteria like if the packet has a bad checksum. It also does further activities of the firewall, thus blocking traffic based on the criteria that would normally be put in a firewall. The normalizer also may hold packet fragments and reassemble them into a packet based on its knowledge of the target system. The knowledge of the target system is provided from a reference table built by the system service scanner.

13.10.1.2 The Detection Engine

The detection engine handles all pattern matching that is not handled by the normalizer. These are patterns that are not based on protocol states.

13.10.1.3 Traffic Shaper

Before traffic leaves the NIPS, it must go through the traffic shaper for classification and flow management. The shaper classifies traffic protocol, although this may change in the future to include classification based on user and applications.

13.10.1.4 NIPS Benefits

In his extensive and thorough article "Network Intrusions: From Detection to Prevention," Andre Lee gives the following NIPS benefits:

- Zero-latency prevention: Without the NIDS and firewall bundle, NIPSs reduce this latency drastically by providing the notification within one hardwired circuitry instead of two.
- Effective network hygiene: Since many attacks are recycle attacks whose signatures are known, NIPS remove these packets quickly, although it does not do much effective anomaly analysis that is done by the NIDS.
- Simplified management: Because the would-be bundle of a NIDS and firewall are all packaged into one hardware, it reduces storage space and of course overall management.

  Although it has all these advantages, NIPSs suffer from a number of problems including the following [10]:

- Production readiness: This occurs because the technology is new and has not gotten the field testing it needs to prove effectiveness in every test.
- High availability: This occurs because it is in-line, and on the first contact with network traffic, it may not be able to withstand high traffic availability and tolerance needed by all first and head-on network devices.
- Detection effectiveness: It has not yet been tested for effectiveness of detection, and it does not ever stop everything, falling short like NIDS.

13.10.2 Host-Based Intrusion Prevention Systems (HIPSs)

Like its cousin, the NIDSs, NIPSs also have corresponding HIPS based on one host. Most HIPSs work by *sandboxing*, a process of restricting the definition of acceptable behavior rules used on HIPSs. HIPS prevention occurs at the agent residing at the host. The agent intercepts system calls or system messages by utilizing dynamic linked libraries (dll) substitution. The substitution is accomplished by injecting existing system dlls with vendor stub dlls that perform the interception. So function calls made to system dlls actually perform a jump to vendor stub code where then the bad calls are processed, evaluated, and dealt with. Most vendor stubs are kernel drivers that provide system interception at the kernel level because processes system calls can be intercepted easily.

13.10.2.1 HIPS Benefits

Again like their cousins the HIDS, HIPS have benefits that include the following [10]:

- Effective context-based prevention: HIPS are the only solution to prevention of attacks that require simulation context. HIPS agents reside on the protected host; they have complete context of the environment and are therefore more capable of dealing with such attacks.
- Effective against zero-day attacks: Since HIPS use sandboxing method to deal with attacks, they can define acceptable parameter application or operating system service behavior to enable the agent to prevent any malicious attack on the host.

  Although they have good benefits, HIPS also have disadvantages based on limitations that hamper their rapid adoption. Among these limitations are [10]:
- Deployment challenge: As we discussed in the HIDS, there are difficulties in deploying the remote agents on each and every host. These hosts need updating and are susceptible to tampering.
- Difficulty of effective sandbox configuration: It can be a challenge to define effective and nonrestrictive parameters on hosts.
- Lack of effective prevention: Because of the use of sandboxing, HIPS cannot use any standard prevention like signature prevention.

Table 13.1 Most current ID systems

(1) MetaFlows Security System *Vendor*: MetaFlows *What*: Cloud-based IDS/IPS/forensics managed security service Web site: http://www.metaflows.com	(14) AT&T Network-Based Firewall v3 *Vendor*: AT&T Corporation Web site: http://www.att.com
(2) IPS 5500 Model 75EC *Vendor*: Corero Network Security *What*: A stand-alone, purpose-built IPS Web site: http://www.toplayer.com	(15) AlertLogic Threat Manager v3.5.4 *Vendor*: AlertLogic Web site: http://www.alertlogic.com
(3) Sourcefire Next-Generation IPS *Vendor*: Sourcefire *What*: A distributed appliance-based offering modeled on the Snort detection engine Web site: http://www.sourcefire.com	(16) ZyWALL USG 200 *Vendor*: ZyXEL Communications Web site: http://www.zyxel.com
(4) NitroGuard IPS 4245 *Vendor*: NitroSecurity *What*: An intelligent packet filtering system that detects sophisticated network intrusion attempts and actively records and/or stops such attempts Web site: http://www.nitrosecurity.com	(17) Proventia IPS GX6116 *Vendor*: IBM-ISS Web site: http://www.iss.net/
(5) McAfee Network Security Platform *Vendor*: McAfee *What*: Provides threat protection for demanding networks Web site: http://www.mcafee.com	(18) NitroGuard IPS *Vendor*: NitroSecurity Web site: http://www.nitrosecurity.com/
(6) CounterSnipe APS *Vendor*: CounterSnipe Technologies *What*: Provides network-based intrusion prevention security Web site: http://www.countersnipe.com	(19) IPS 5500-150E v 5.12 *Vendor*: Corero Network Security Web site: http://info.corero.com/ppc-casestudy-idc-liquid-web.html?gclid=CJir6P7a_tICFUs7gQod9PcEDw Intrusion Detection Systems Reviews
(7) Snort *Vendor*: Open source Web site: http://www.snort.org	(20) Interceptor 1000 *Vendor*: Reflex Security Web site: http://www.reflexsecurity.com
(8) Symantec Managed IDS/IPS with Sourcefire *Vendor*: Symantec Web site: http://www.symantec.com	(21) IDS/IPS *Vendor*: SecurityMetrics Web site: http://securitymetrics.com/
(9) SecureWorks Managed IDS/IPS *Vendor*: SecureWorks Web site: http://www.secureworks.com	(22) DefensePro, Version 3.10 *Vendor*: RadWare Web site: http://www.radware.com
(10) Perimeter eSecurity Firewall and Intrusion Prevention *Vendor*: Perimeter eSecurity Web site: http://www.perimeterusa.com	(23) CounterSnipe Technologies Active Protection Software 3.0 *Vendor*: CounterSnipe Technologies Web site: http://www.countersnipe.com

(continued)

Table 13.1 (continued)

(11) Network Box v3.2 *Vendor*: Network Box USA, Inc. *Price*: from $2200 first year Web site: http://www.networkboxusa.com	(24) NetIQ Security Manager (IDS group test) *Vendor*: NetIQ Web site: http://www.netiq.com/products/sm
(12) BM Managed Protection Service *Vendor*: IBM ISS Web site: http://www.ibm.com	(25) NetScreen-IDP100 (IDS group test) *Vendor*: NetScreen Technologies Inc. Web site: http://www.netscreen.com
F-Secure Protection Service for Business v4 *Vendor*: F-Secure Web site: http://www.f-secureusa.com	(26) StealthWatch (IDS group test) *Vendor*: Lancope Web site: http://www.lancope.com
(13) Clone Systems Clone Guard Managed IDS/IPS *Vendor*: Clone Systems Web site: http://www.clone-systems.com	

Source: SC Magazine: http://www.scmagazine.com/intrusion-detection-systems/products/91/

13.11 Intrusion Detection Tools

Intrusion detection tools work best when used after vulnerability scans have been performed. They then stand watch. For the most current list of IDS systems, the reader is referred to SC Magazine at http://www.scmagazine.com/intrusion-detection-systems/products/91/.

Table 13.1 displays the most recent ID systems based on CS Magazine.

All network-based intrusion detection systems and tools can provide recon (reconnaissance) probes in addition to port and host scans. As monitoring tools, they give information on:

- Hundreds of thousands of network connections
- External break-in attempts
- Internal scans
- Misuse patterns of confidential data
- Unencrypted remote logins or a Web sessions
- Unusual or potentially troublesome observed network traffic

All this information is gathered by these tools monitoring network components and services that include the following:

- Servers for:
  - Mail
  - FTP
  - Web activities
- DNS, RADIUS, and others
- TCP/IP ports

- Routers, bridges, and other WAN connection
- Drive space
- Event log entries
- File modes and existence
- File contents

Exercises

1. Are IDSs similar to firewalls?
2. Why are system intrusions dangerous?
3. Discuss the best approaches to implementing an effective IDS.
4. Can system intrusions be stopped? Support your response.
5. For a system without a DMZ, where is the best area in the network to install a honeypot?
6. Why are honeypots important to a network? Discuss the disadvantages of having a honeypot in the network.
7. Discuss three approaches of acquiring information needed to penetrate a network.
8. Discuss ways a system administrator can reduce system scanning by hackers.
9. Discuss the benefits of system scanning.
10. Discuss as many effective ways of responding to a system intrusion as possible. What are the best? Most implementable? Most cost-effective?

Advanced Exercises

1. Snort is a software-based real-time network intrusion detection system developed by Martin Roesch. It is a good IDS that can be used to notify an administrator of a potential intrusion attempt. Download and install Snort and start using it.
2. The effectiveness of an IDS varies with the tools used. Research and develop a matrix of good and effective IDS tools.
3. If possible, discuss the best ways to combine a firewall and a honeypot. Implement this combination and comment on its effectiveness.
4. Intrusion detection hybrids are getting better. Research the products on the market and comment on them as far as their interfaces are concerned.
5. Discuss how exploits can be used to penetrate a network. Research and list ten different common exploits.

References

1. Sundaram A. An introduction to intrusion detection, ACM Digital Library. http://dl.acm.org/citation.cfm?id=332161
2. Kizza JM (2002) Computer network security and cyber ethics. McFarlans Publishers, Jefferson

3. Halme LR, Bauer KR (1995) AINT misbehaving: a taxonomy of anti-intrusion techniques. Proceedings of the 18th National Information Systems Security conference.
4. Handley M, Paxson V, Kreibich C. Network intrusion detection: evasion, traffic normalization, and end-to-end protocol semantics. ACM Digital Library. http://dl.acm.org/citation.cfm?id=1267621
5. Proctor P (2001) The practical intrusion detection handbook. Prentice Hall, Upper Saddle River
6. Innella P. The evolution of intrusion detection systems. Symantec-Connect. http://www.symantec.com/connect/articles/evolution-intrusion-detection-systems
7. Mullins M. Lock IT down: implementing an intrusion detection system on your network: how to implement a network intrusion detection system. Tech Republic. http://www.techrepublic.com/article/lock-it-down-implementing-an-intrusion-detection-system-on-your-network/
8. Internet Security Systems. Network- *vs* Host-based intrusion detection: a guide to intrusion detection technology. http://www.windowsecurity.com/archived/nvh_ids.html
9. Panko RR (2004) Corporate computer and network security. Prentice Hall, Upper Saddle River
10. Yee A (2003) Network intrusions: from detection to prevention. Int J Inform Assur Prof 8(1)