Computer Network Security

Lesson 11

Intrusion Detection Systems

Lecturer: Dr Msagha J Mbogholi, PhD

Flashback from Lesson 10

- Digital forensics includes the forensics of computers, media, networks, mobile devices, and so on.
- Email and web forensics is about using email details (obvious and hidden) to find information that can be used for investigative purposes.
- Data forensics purposes to find hidden data in media using forensic tools; this data can then be used as evidence in court.
- Mobile forensics is the use of mobile information from both user or/and provider in order to extract information that can be used in an investigation.
- Network forensics uses hosts and agents on the network to get information on the goings-on in the network. Network forensic tools can then string together events to give them meaning.
- A good software forensic tool should preserve the integrity of the data it is investigating.

Content

- Introduction to IDS
- IDS Types
- Methods of Detection
- IDS Setup

Part 1

Introduction to IDS

Introduction

- In previous lessons we have learnt about the different types of attacks that a network is exposed to. We even went further to suggest a few ways in which a network can be protected, for example by use of antivirus.
- In our previous lesson (lesson 10) intrusion detection systems (IDS) were mentioned; as the name implies these are systems that are designed to detect network intrusions.
- A key question the learner may be asking is what's the point of detecting intrusions and then leaving it at that? Why use all time and resources to detect intrusions instead of just dealing with them?
- Or better still why not just prevent these intrusions from taking place in the first instance? Of course this is more like taking a proactive approach to security, right?

Introduction (cont'd)

- The reasons for having an IDS can be seen from the following perspectives:
  - First if an intrusion is detected in good time then steps can be taken to prevent any damage from occurring; in any case if the intrusion is already taking place then further damage can be prevented.
  - Secondly just as a dog is to a burglar (they don't like to invade places where dogs are guarding, who would anyway?) a good IDS serves well to discourage intruders.
  - Thirdly information gathered from the IDS will help in development of better IDS and prevention measures.

Introduction (cont'd)

- The intruders can be categorized in three groups:
  - Masquerader – normally an external user who attempts to access the system by pretending to be a legitimate user (they steal a legitimate user's credentials by other means)
  - Misfeasor – this is someone who engages in improper conduct, either by accessing applications they're not supposed to, or by abusing their system given privileges.
  - Clandestine user – this is a user who wrongfully gets administrative/supervisory rights in the system and uses them to do wrong so that they escape audit records capturing their actions.

Introduction (cont'd)

- Let us generalize the types of attacks the IDS is meant to detect. In previous lessons we learnt of the specific types that can be categorized into one or more of the following groups of attacks:
  - Group 1 – destroy information maliciously on the network
  - Group 2 – seize confidential information
  - Group 3 – limit or prevent legitimate users from gaining access to resources on the network
- With this in mind let us move on to discussing the different types of IDS

Part 2

IDS Types

Introduction

- Intrusion detection has been around historically for a while if you think about it.
- You probably apply this subconsciously in your day to day life without knowing; your mind makes a mental note of where you left your bag. When you come back and find it moved that's intrusion detection! I am sure you can think of many other instances of such nature.
- In networking IDS are categorized into 2 groups:
  - Host based IDS
  - Network based IDS

Host-based IDS (HIDS)

- As the name implies a HIDS is one which resides on and monitors a host.
- The HIDS is a software (not hardware) and it monitors events on the host device.
- For example in Windows based systems the HIDS will monitor the system logs, namely system log, application log, and security events log.
- The HIDS will compare the events in the different logs to the rules that have been placed in it; if it finds log entries that match the rules it will then sound off an alarm (warning). For example if a file has been altered or deleted the HIDS will sound an alarm.
- HIDS have their own log files which protect them in the event an attacker modifies the logs of the host device.

HIDS

- For more effectiveness certain HIDS are loaded on the operating system (OS) of the respective host.
- The reason for this is that most attackers target the vulnerabilities of the OS; by placing the HIDS here an attack can be detected very quickly and easily.
- Nonetheless the tradeoff is that should the attacker be able to overcome the IDS then they will gain access to the host and they can disable the host itself; imagine if this were the DNS or email server!
- A HIDS can also be centrally managed via a console. In this scenario the manager HIDS will have all the managed HIDS send their information to a central console and it is from here that alerts are generated.

HIDS Scenario

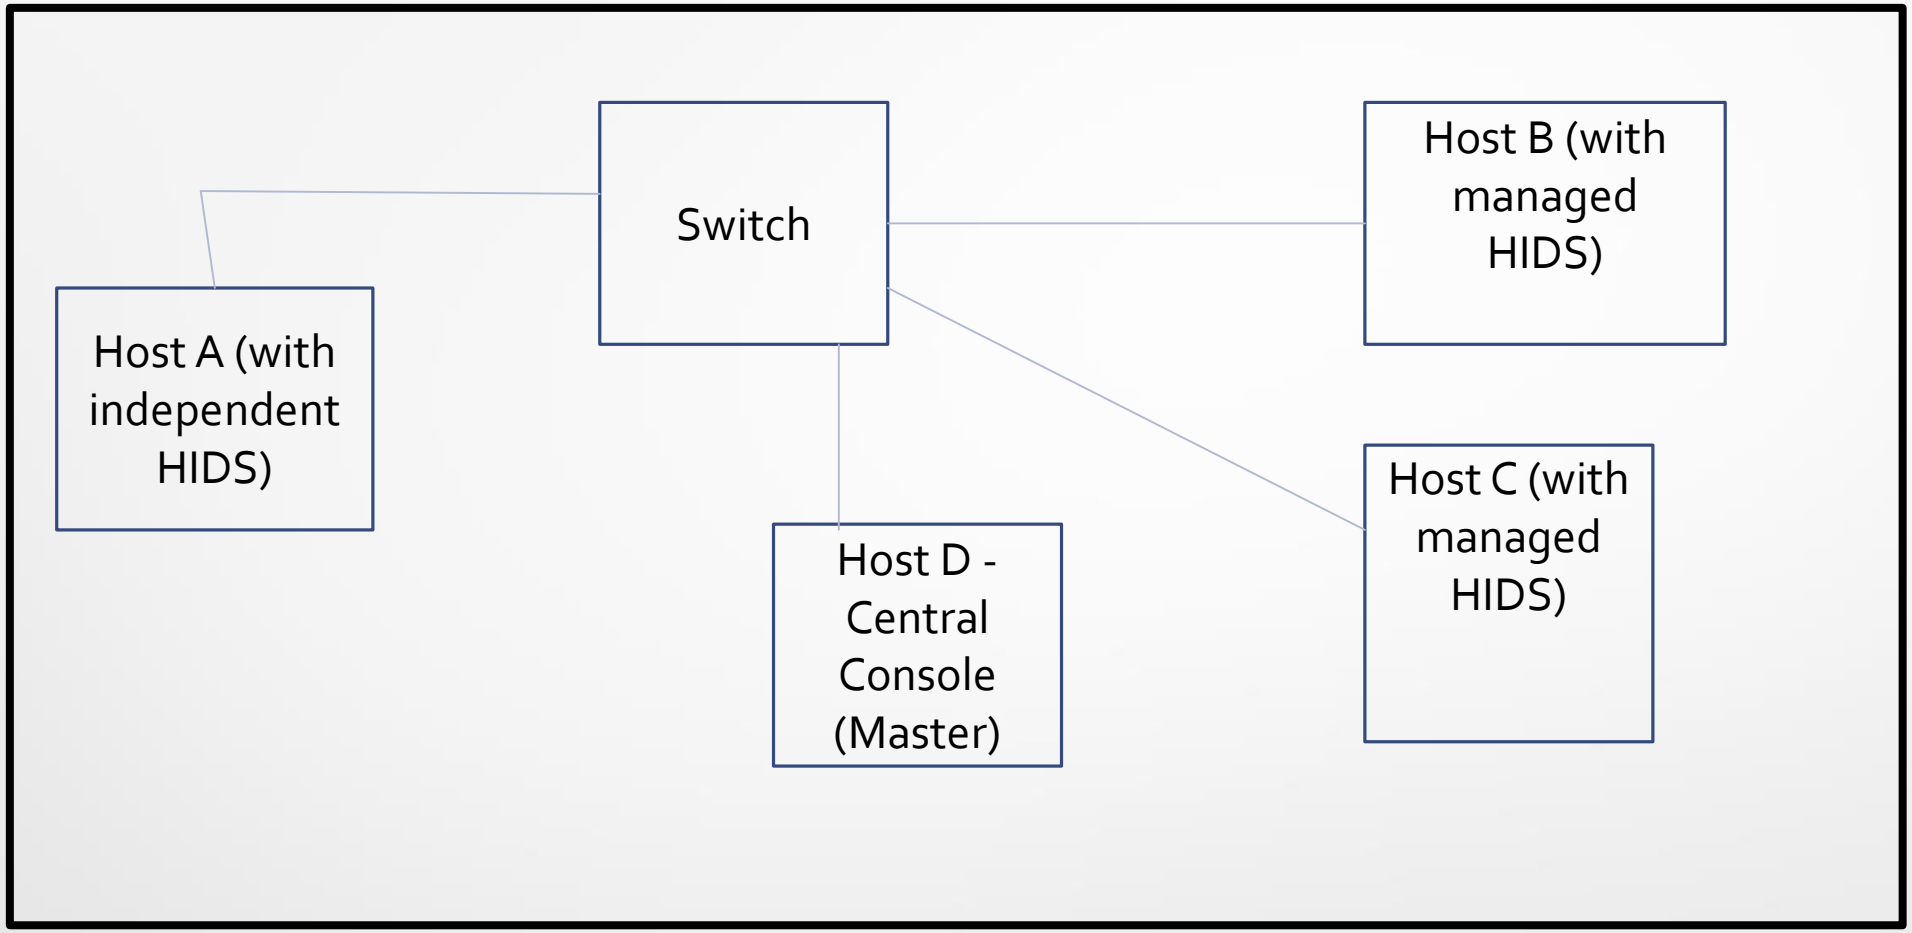

Fig 1. A network with a standalone HIDS and managed HIDS

HIDS

- Fig 1 shows a dual scenario where there is a standalone HIDS and a managed one.
- All 4 hosts are connected to each other in a star topology in a given network.
- Host A has a HIDS installed on it and functions independently; should something be out of normal based on the configured rules it will sound off an alarm which is then sent to the administrator.
- Hosts B and C are referred to as managed hosts. They are managed from the central (master) console located at host D. The difference here is that all managed hosts will be monitored from the central console as opposed to host A. The master console thus monitors all information provided by the managed hosts and alerts the network manager when there's a possible intrusion.
- The difference between the two should now be clear.

HIDS (cont'd)

There are several advantages to using HIDS:

- They are very efficient at detecting local events on hosts in a networked environment.
- Since they reside on the host they are equipped at processing encrypted information (since it will already have been decrypted)
- It is protocol free (only works at host level)
- By examining the system logs HIDS are more adept at detecting attacks.
- The HIDS is also capable of identifying misfeasor unauthorized access attempts.

HIDS (cont'd)

- HIDS also present the following disadvantages:
- It requires more time and effort to install and manage a HIDS on a network; this is because it has to be configured on each individual host.
- As described earlier attacks can be focused on either the host OS or the IDS itself, presenting two vulnerable centers of attack.
- The HIDS scope is limited to the host in which it is installed; it therefore can't detect an attack on multiple devices on the network. Consequently it also can't detect attacks on devices such as routers, switches or bridges on the network.
- It is vulnerable to DoS attacks
- IDS require a lot of disk space due to their audit logs (which are separate from host system logs).
- Installing a HIDS has performance implications on the host.
- The HIDS is limited to alerts on pre-configured rules; it can't alert on attacks that are not configured in the rules.

Network – Based IDS (NIDS)

- The NIDS is a software running on a dedicated device.
- The concept is rather simple; it monitors network traffic looking for signatures that match what is in its rules.
- If it finds something that is in its rules it will trigger an alarm.
- Sounds quite simple doesn't it? Well not quite.
- Let us get into some details.

NIDS

- In setting up the NIDS there are 3 approaches that can be deployed.
- In the first approach the NIDS is configured on a dedicated device with a network interface card (NIC).
- This NIC is then set up in promiscuous mode. This means that it will monitor ALL traffic passing through it. In so doing it can analyze the traffic as it looks for any signs of attack. The NIDS are mostly signature based meaning they look for signatures that have been configured in their rules; this also presents a downside obviously since any signature that hasn't been configured in the NIDS will pass through unnoticed. Consequently some NIDS allow for customization, meaning that an experienced administrator can configure additional signatures into the NIDS.
- Fig.2 demonstrates this setup.

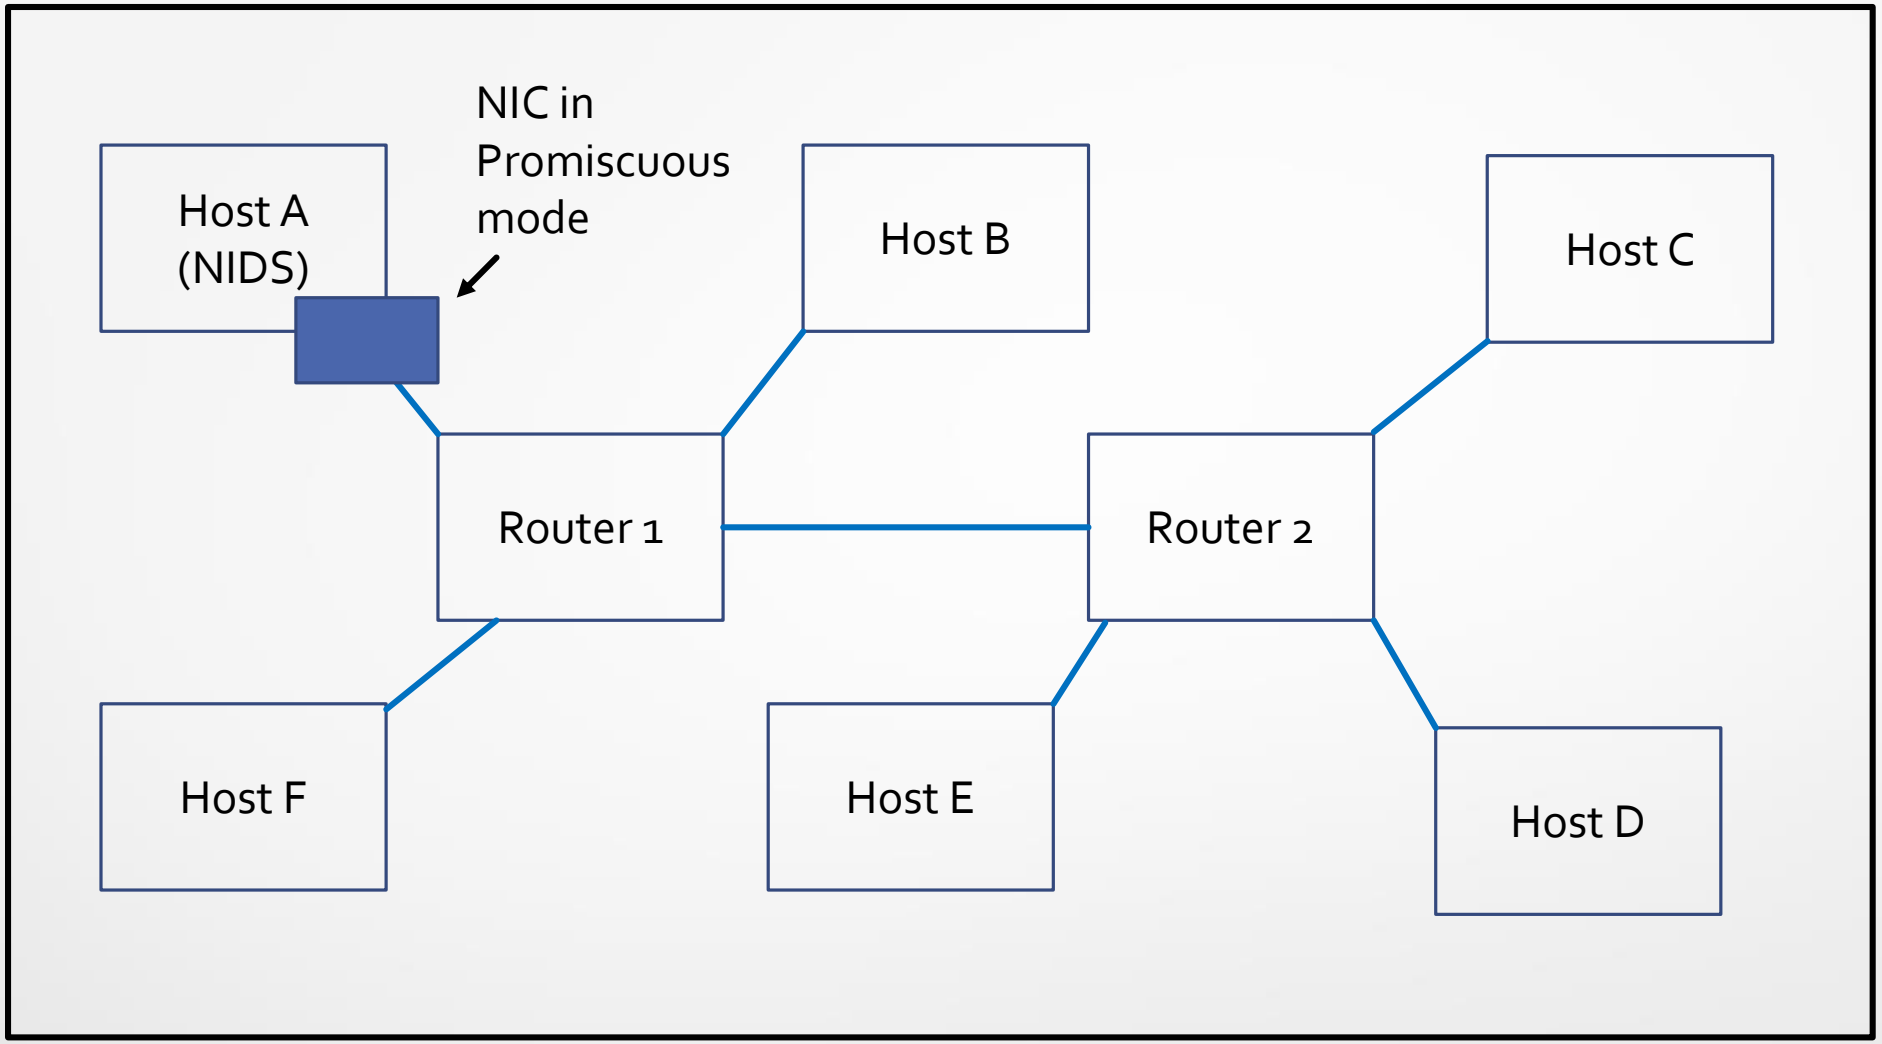

Fig 2. NIDS setup on a LAN with NIC in promiscuous mode

NIDS (cont'd)

- Fig 2 demonstrates a LAN setup in a given organization. The setup is a hierarchical model with router 1 and router 2 providing distribution at the distribution layer.
- There are 6 hosts configured in this setup.
- The NIDS is configured in host A. The NIC at A is then set in promiscuous mode so that it can monitor all the traffic flowing in the network.
- In this setup even traffic received from the core router will also be monitored once it gets to the distribution layer, meaning that no traffic will go unmonitored at any given time.
- The traffic is compared with the signatures in the NIDS and when a match is found the alarm is sounded by the NIDS

NIDS (cont'd)

- The second approach is to configure 2 NICs on the host. In this case the NIC that is set in promiscuous mode is configured to operate in stealth mode.
- This means that it has no IP address configured on it. Configuring it this way means that it will be transparent on the network being monitored and thus not respond to any communication (including pings).
- The second NIC card will be configured on the NIDS and will send alarms to a different network from the one being monitored.
- Fig 3 demonstrates this setup.

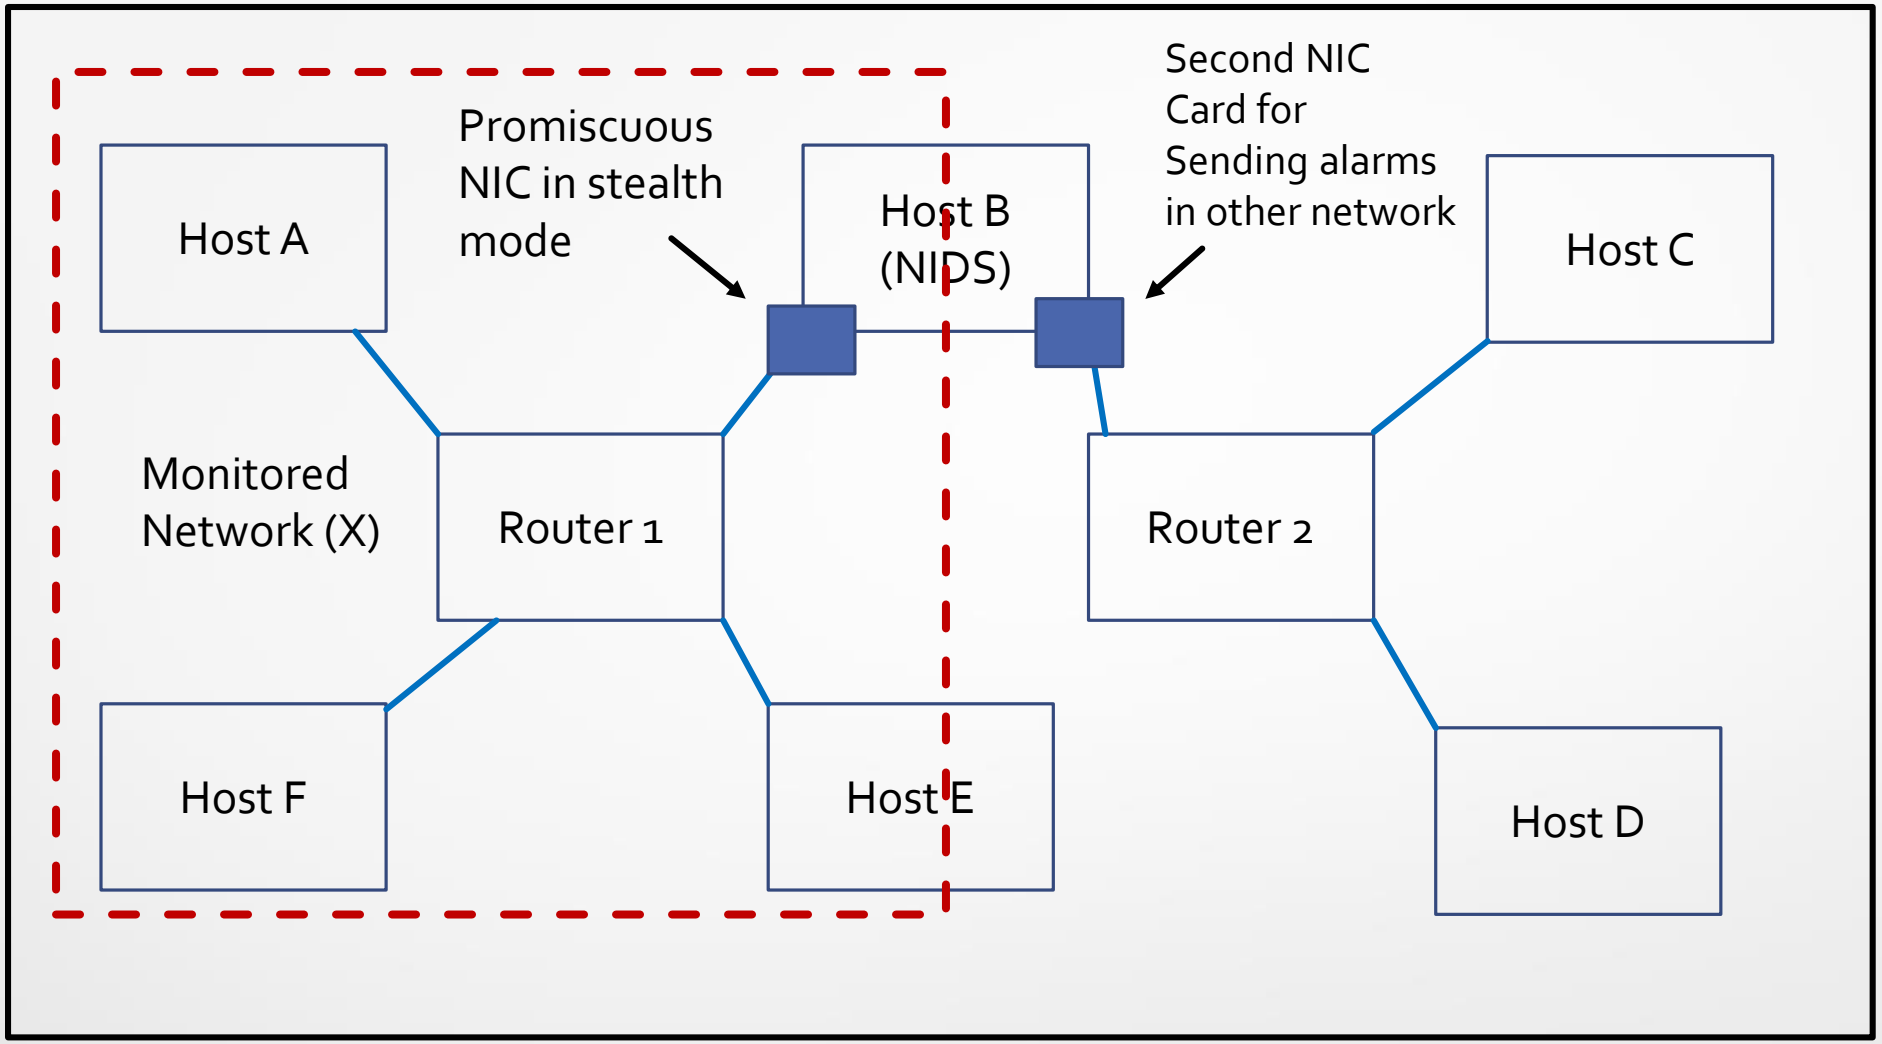

Fig 3. NIDS setup using 2 Network Interface Cards.

NIDS

- Fig 3 demonstrates this setup. In this case the NIDS has been configured in host B.
- Host B has one NIC card configured in stealth promiscuous mode. This card monitors the traffic in network segment X. Since it is in stealth mode it is transparent to devices in X and will not respond to any communication (including pings).
- If a packet matches information found in the rules the NIDS will send the alarm to the unmonitored segment where the administrator is.
- Therefore X is unaware of the unmonitored network and vice versa. This is a good approach security wise.

NIDS

- In both approaches (fig 1 and fig 2) the NIDS has been centralized and configured for a particular network.
- A NIDS can also be configured for a distributed network architecture. These architectures used the distributed computing model. In a distributed computing model the nodes at different sites share workload and data using complex algorithms.
- Fig 3 shows how this is set up

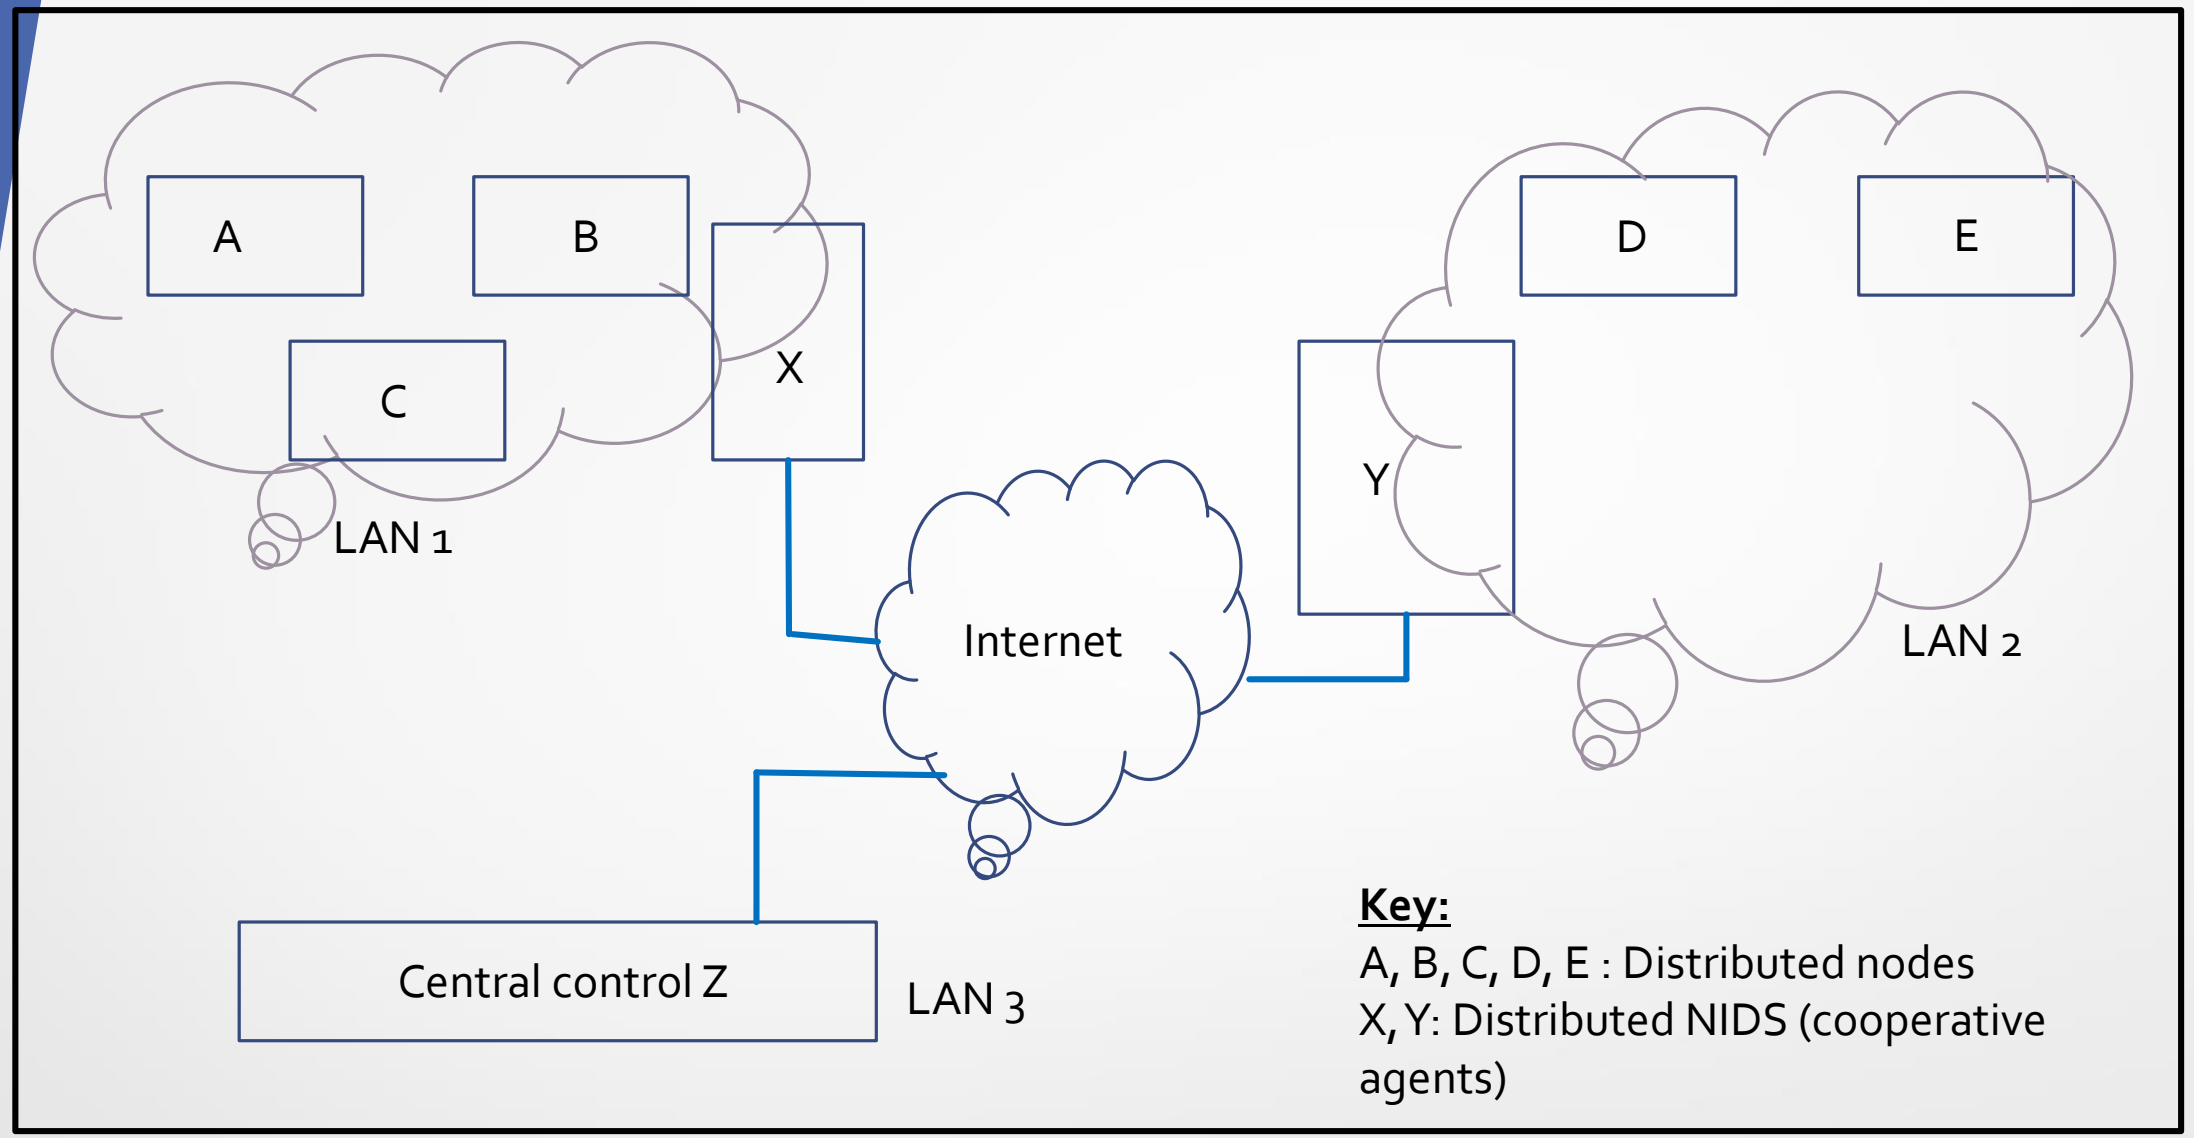

Fig 3. NIDS setup in distributed network

NIDS

- Fig 3 shows an NIDS setup in a distributed network.
- A, B, C, D, and E are nodes in the network. The first 3 nodes are in LAN 1, while the remaining 2 are in LAN 2.
- The NIDS is distributed using cooperative agents. The cooperative agents are in X and Y.
- The cooperative agents work together to monitor the whole distributed network as one. There is a central console residing in LAN3 which is part of the distributed network.
- All alerts / alarms are sent to the central console by the cooperative agents.
- With this kind of setup the distributed network is protected better as the cooperative agents work together to monitor the whole network traffic.

NIDS

There are some advantages associated with the use of NIDS:

- Due to the transparent configuration of the NIDS an attacker will not be aware that they are being monitored.
- The NIDS can be used in a network to monitor a large number of potential target hosts.
- Due to the promiscuous setting of the NIC card the NIDS is able to capture all packets on the network thus avert attack on a target host.

NIDS

The disadvantages of NIDS include:

- It is only able to alert regarding rules that have been preconfigured in it; any attack that is not configured will go unnoticed and potentially succeed.
- It can’t examine encrypted packets, thus all encrypted traffic will pass through.
- Special settings are required for all switched networks.

NIDS or HIDS?

- I am sure this question is running through your mind.
- Which is better to use between the two? The truth is there is no one size fits all!
- The advantages and disadvantages of both have been listed.
- An organization will have to choose which to use based on their individual requirements such as budget (in a LAN more HIDS will be required vs one single NIDS) and also the level of threats they perceive they are exposed to.

Part 3

Methods of Detection

Introduction

- So far we have discussed the workings of IDS by referring to rules.
- However there are different "rules" that are used by different IDS.
- These can be categorized into four categories based on how the monitoring is done.
- The four categories are:
  - Anomaly-based
  - Signature – based
  - Behavior – based
  - Heuristic

Audit Records

- The use of audit records has been described earlier in this lesson. The audit logs capture the activity of users and this can be used by the IDS.
- There are two approaches to the use of audit records:
- First is making use of already in-built accounting features in the system under investigation. All multiuser OS collects information regarding the activity of individual users. Even without a lot of operational knowledge you can do this in Windows. In Windows 7 you just go to start and key in "logs" in the text box. A whole list will be presented to you and pick any folder to examine logs. Event logs is a good place to start the examination. Audit logs will be found under Windows logs in the event viewer. If using a newer version of Windows say 10 then the approach would be the same. These type of records are referred to as native audit records since they are found already inbuilt in the OS.

Audit Records

- The second approach is to collect information that is of specific interest to the IDS from the audit records. The disadvantage of this is that the host will be running two accounting packages. On the other hand the advantage is that this information can be collected and used on different devices, incorporating both portability and platform independence. These type of records are referred to as detection – specific audit records. The name is self- explanatory.
- Stallings and Brown (2015) refers to Denning (1987) for the fields to be used in detection – specific audit records. These are: Subject, action, object, exception-condition, resource-usage, and time-stamp.

Anomaly Detection

- This monitoring method is based on statistical monitoring. There are two parameters that are monitored.
- In the first approach a baseline is created after observing behavior that is considered normal for some time, i.e. the normal activities of the system. (this creates a profile which is the parameter to be monitored)
- This baseline is then used to compare operations of the system; if a deviation is noticed then an alarm is triggered.
- The distinct advantage of this is that alarms are triggered without having to understand the cause of the deviation; on the downside this can lead to quite a few false positives. This is because what is defined as "normal" will not always be the case.

Anomaly Detection

- Another area of concern is the time it takes to create a baseline; an attack can occur while the baseline is still being created.
- The second approach is to monitor established thresholds (also known as threshold analysis). An event is monitored based on the number of times it occurs; if it surpasses a particular number (established threshold parameter) an alarm is triggered.
- Audit records are normally used as the input for determination of thresholds and baseline parameters.

Signature Detection

- This method works a lot like antivirus systems.
- There is a database of signatures from which the IDS compares activities.
- If it finds activities that are similar to a signature in the database it triggers an alarm.
- Just like with antiviruses the signature database needs to be updated frequently in order to detect new signatures.
- Further this provides a challenge in terms of overheads and possibly disk space with time.
- Another foreseeable challenge is in the rigid nature of signatures; a slight variation of the signature will make the threat pass undetected resulting in an attack.

Rule –based Detection

- This method is also known as behavior – based detection.
- This approach is similar to the statistical anomaly approach; the key difference is that there is no baseline that is created.
- Audit records are examined in order to identify activities and generate rules based on these activities.
- The IDS then observes the happenings in the system and when it observes a deviation from the rules it generates an alert.
- This method therefore just compares the rules based on past behavior and alerts when a deviation is detected.
- This method can more quickly stop new attacks as it just monitors behavior; no reading of signatures and no baseline establishment.

Heuristic Detection

- It is an intelligent form of detection.
- It makes use of an algorithm to detect if a threat exists.
- It is more efficient than other forms of IDS and may find threats that the other three have missed.

Part 4

Setting up the IDS

Steps in the Setup

- Setting up the IDS will require the development of an IDS policy. The policy will determine the parameters within which the IDS will operate in the organizational environment. This involves determining 5 key parameters:
- The objective/purpose/goal of the IDS
- What the IDS will monitor (and consequently what it won't)
- Response actions to be taken when attacks are detected.
- Determination of thresholds (for use in determining what's normal and what's not)
- Implementation of the Policy.

Tools

- Some examples of IDS available in the market today include:
  - Snort (free)
  - Stealthwatch
  - Solarwinds security event manager
  - OSSEC (free)
  - Zeek
  - Security Onion

Summary

- The two IDS types are host-based (HIDS) and network-based (NIDS)
- IDS are categorized according the method of detection used. There are four categories – anomaly based, signature based, rules based and heuristic.
- An IDS policy needs to be setup by an organization that will capture the parameters associated with the use of the IDS.

References

- Ciampa, M. D. (2012). *Security+ guide to network security fundamentals* (4th ed.). Course Technology, Cengage Learning.
- Cole, E., Krutz, R. L., & Conley, J. W. (2005). *Network security bible*. Wiley Pub.
- Denning, D. (1987) "An Intrusion-Detection Model." *IEEE Transactions on Software Engineering*
- Kahate, A. (2013). *Cryptography and network security*. McGraw Hill Education.
- Maiwald, E. (2001). *Network Security: A Beginner's Guide* (1st ed.). Osborne_McGraw Hill.
- Stallings, W., & Brown, L. (2015). *Computer security: Principles and practice*. Pearson.