Council of the
European Union

Brussels, 13 April 2015
(OR. en)

7867/15

CSCI 17
CSC 81
CYBER 27

NOTE

From:	General Secretariat of the Council
To:	Delegations
Subject:	Information Assurance Security Guidelines on Intrusion Detection and Prevention Systems in CIS

Delegations will find attached the Information Assurance Guidelines on Intrusion Detection and Prevention Systems in CIS as approved by the Council Security Committee on 10 April 2015.

This page intentionally left blank

Information Assurance Security Guidelines on

Intrusion Detection and Prevention Systems in CIS

IASG 4-02

TABLE OF CONTENTS

I. PURPOSE AND SCOPE

1. These guidelines, agreed by the Council Security Committee in accordance with Article 6(2) of the Council Security Rules [1] (hereinafter 'CSR'), are designed to support implementation of the CSR.

2. These guidelines describe minimum standards to be observed for the purpose of selection, planning and deployment of intrusion detection and intrusion prevention systems in communication and information systems (CIS) and interconnections between them and/or other systems. Network defence requires prophylactic measures to reduce the likelihood of compromise of a CIS and to minimise the impact of such an event. Network defence also requires monitoring to detect compromise and record evidence as well as reactive measures to stop attacks and restore normal service. Intrusion detection and prevention systems support network defence by enabling both prophylactic and reactive measures to be taken by network defence Management (NDM).

3. The Council and the General Secretariat of the Council (GSC) will apply these security guidelines in their structures and CIS.

4. Member States should use security guidelines as a benchmark when EU classified information is handled in national structures, including in national CIS.

5. EU agencies and bodies established under Title V, Chapter 2, of the TEU, Europol and Eurojust should use these security guidelines as a reference for implementing security rules in their own structures.

6. In this document, intrusion detection means the process of recording events in a CIS and examining them, in order to detect violations of security policy and/or incidents such as the unauthorised logical access to CIS or the information handled by CIS. Intrusion detection systems (IDS) are sophisticated electronic surveillance systems which automate this process.

[1] Council Decision 2013/488/EU, OJ L274 of 23.09.2013, p.1

Intruder alarms (burglar alarms) are out of scope of this document. Intrusion prevention systems (IPS) have all functionality provided by IDS and in addition have the capability to take action in an attempt to block detected intrusions without human intervention at the time of the event. Throughout this document IDS and IPS are considered together unless otherwise stated, using the term IDPS for brevity. In this document the term CIS can mean a single CIS or a group of CIS supported by the NDM.

7. Although a web application firewall (WAF) could be considered a special type of IDPS, its use is not covered in this document. The use of a WAF to detect known and unknown types of attacks against web applications is, however, a way to provide additional protection and provide defence in depth of web-based CIS.

8. As compromise of CIS could affect the security or essential interests of the EU, its Member States and partners, it is important to ensure that IDPS are deployed in CIS to a degree commensurate with the level of risk and the threat scenario established by a risk assessment process. The use of IDPS should be co-ordinated by the NDM of the CIS.

II. IDPS DESCRIPTION AND USES

9. IDPS as a rule feed the central monitoring and reaction system which produces alerts of potential security incidents which the IA Operational Authority (IA OA), NDM and its Incident Response Team (IRT) handle in accordance with the Information Assurance Guidelines on CIS Security Incident Handling IASG 4-03.

10. IDPS is to be used primarily to collect events in a CIS, to analyse them to determine whether they constitute a threat to the security of the CIS and to take action if this is the case. Preventive action usually means that the IDPS modifies the security features of the CIS, for example by:

    (a) tearing down the connection that is being used for the attack causing the events;

    (b) changing access control lists to block access to the victim resource or service from the attack source or to block all access to the victim resource or service;

(c) applying patches to the victim resource to eliminate vulnerabilities or weaknesses which enable the attack to succeed;

(d) changing firewall rules;

or

(e) filtering the attack data and changing it (e.g. by acting as an application “proxy”).

11. If no prevention action is activated, IDPS is typically used to report the suspicious events to security support personnel for further investigation and manual corrective action if required.

12. IDPS consist of

(a) a front end[2] to collect event information:
sensors which provide information about events to the back end;

and

(b) a back end:

an analysis component (software and/or hardware) which records, correlates and archives events to decide whether the event observed is benign or malicious and, if malicious, attempts to take action to prevent the attack or policy violation affecting the CIS. The back-end can be a standalone IDPS component, a component of a central monitoring and reaction system, or a combination of both. A back end typically consists of :

i. a console, nowadays usually via an encrypted access to the web interface, to monitor events, generate reports and modify the IDPS settings;

ii. a computer (server or servers) which performs collection, correlation and analysis of the detected events and takes action as foreseen by the IDPS rule set;

and

iii. a database to store the events detected and the actions taken in response.
Action can be indirect:- e.g. sending an alert to NDM personnel, or direct:- e.g. modifying firewall rules, intercepting network traffic, changing security settings of components of the CIS, etc. in real time.

[2] Note that the user interface for IDPS is the back-end, not the front end which consists of listeners to collect information.

13. Events can be any kind of unusual behaviour of an accidental or malicious nature. A recorded event can be both success or failure - e.g. successful or failed attempts to access a service offered by the CIS. IDPS cannot a-priori know what is "usual" or "normal" behaviour so that in the initial phase of deployment, experts thoroughly knowledgeable of the monitored CIS must be involved to advise whether events being reported by the IDPS are "expected behaviour" or not, and whether all events which NDM considers to be potential security incidents are being captured by the IDPS. This effort, referred to commonly as "tuning" the IDPS, is not a one-off but an ongoing exercise as described in sections IV, V and VI of this document.

14. Events can be, for example

    (a) disk crash, failure of a network card, lack of disc space, processor overload, process termination;

    (b) a much higher frequency of connections using a specific network protocol (TCP/UDP/ICMP);

    (c) recognition of the "signature" of a known attack method or of a known malware item;

    (d) a host which usually never uses a certain protocol suddenly making intensive use of it;

    (e) hosts which do not usually communicate suddenly exchanging information;

    (f) violations of acceptable use policy of the organisation such as downloading files to removable media which are not “approved”;
    or

    (g) transmission of classified or sensitive information to unauthorised recipients.

15. IDPS can and should also be used for:

    (a) recording the type and frequency of successful and unsuccessful attacks on a CIS, in order to fine tune network defence measures and report to management on the threat scenario which the CIS faces;

(b) acting as a second line of defence by duplicating and supplementing boundary protection measures such as filtering and firewall rules;

(c) deterring potential attackers, for example by displaying a warning banner indicating that monitoring and detection of unauthorised use is in place on all methods[3] which can be used to interactively access the CIS wherever technically feasible.

16. It should be noted that IDPS systems require operation of its components in ways which deviate from or violate standards and security regulations. Deployment and use of IDPS should therefore be authorised by the NDM and the CIS business owner.

17. Most IDPS can be set to a monitoring-only mode which collect events and record what action would have been taken in response to the events detected. Sensors are of various types and several such types of IDPS should be used in order to have an effective service. Annex I describes characteristics of the sensor types currently in common use:

(a) Network IDPS

sensors which monitor traffic, usually at important sections of the communication lines of a CIS and analyse network traffic to identify suspicious events. Such sensors are technically modified to be able to listen to all traffic in their network segment not only to traffic directed to their MAC (hardware) address. Further the devices are set up in a way that they are "invisible" to other devices and tools in the network. Network IDPS sensors are therefore often attached via a network TAP which splits traffic at critical points in the network and sends a replica of this traffic to the sensors. Instead of a network TAP, a port on a network switch can act as connector if it is configured to act as a mirroring or SPAN port which can listen to all traffic passing through a port or V-LAN of a switch and send it to the IDPS sensor. While as reported in publicly available documentation, TAP connection has advantages over the use of SPAN ports, a higher

[3] for example unencrypted or encrypted "terminal", web, file transfer using Telnet, FTP, HTTP, SCP, SSL, SSH, etc…

number of sensors are needed so that a decision based on a risk assessment must be taken on which kind of Network IDPS connections to use at which points in a specific network. If network traffic being monitored must pass through the sensor, it is termed an inline sensor. Such sensors are usually preferred as the sensor can itself perform the prevention task, rather than the off-line sensors which instruct another network device to do it.

(b) Host IDPS

sensors which usually require an agent – a software or hardware component which resides on the host device and analyses various operating system, layered product or application activity, searching for unusual or forbidden behaviour of the monitored objects. Such IDPS functionality is often provided as an add-on to malware protection solutions “anti-virus”, or other security tools, for various operating systems, applications or layered products. The sensors and agents are usually very specific to the type of object for which they are made.

(c) Network Behaviour Analysis (NBA) IDPS

sensors which examine network traffic to identify threats that generate unusual traffic flows. This kind of IDPS usually analyses flow information provided by boundary protection such as routers firewalls or other network devices, for example NetFlow, sFlow, or IPFIX. The connection of NBA sensors is usually identical to that of network IDPS sensors – network TAP or spanning port of a switch.

(d) Wireless IDPS

sensors which can be incorporated into existing wireless network devices – both wireless clients and access points or repeaters, or can be a separate set of wireless devices which listen to the wireless traffic in the area covered by the wireless network of the CIS being monitored.

18. Honeypots can be considered to be special IDPS sensors and can be important additions to IDPS. They are not usually part of the core IDPS system but the information they collect complements it. Honeypots should not form part of any production CIS. They should be servers or workstations, occasionally network devices, which are never accessed by authorised clients as they offer no services to them. Any access to a honeypot is therefore unauthorised.

19. Honeypots can be of two kinds:

    (a) One method is to set up very slow hardware purposely loaded with many vulnerable software packages configured to offer a large attack surface and little security. This enables Network defence Management to detect attacks being used in practice and potentially also identify their source.

    (b) Alternatively the honeypot can be an exact replica of the production system but loaded with no live or sensitive information, some of which can be made attractive to potential attackers by adding protective markings or "interesting" words, in order to enable NDM to identify which attacks can be effective against the production system and what attackers are interested in doing.

20. Honeypots must on the other hand be isolated logically, for example by forcing all devices in the production CIS to drop[4] or refuse all connections from the honeypot. The activity of attackers is usually slowed down by the poor processing power of the honeypot and its slow network connections, enabling the security team to examine the mechanism being used in the attack and to prophylactically protect the production CIS from it. With a honeypot there is also a higher probability of being able to identify the source and thus the perpetrator of the attack, be it a faulty device or a malicious attacker.

[4] Dropping a connection means a target ignoring a network (protocol) connection request without sending any reply to the source. The attacker thus has no idea whether the target exists at all, whereas if a connection is refused, the attacker knows that there is a device at the target network address as it receives a connection refusal message.

III. IDPS DESIGN

21. The legal framework around the use of monitoring systems must be determined and taken into consideration when deploying IDPS. Consultation with affected users, data protection, legal experts and management is strongly recommended to explain the purpose and nature of the monitoring and ensure correct notification of the parties affected. In certain legal systems, the use of monitoring tools may need to be advertised to clients e.g. by a 'legal' or 'warning' banner which is displayed before attempted login, on all methods used for interactively accessing the resources of the CIS being monitored, wherever technically feasible. Such a legal banner as well as the use of IDPS must be carefully drafted ensuring that it conforms to the legal system(s) applicable to the CIS and its clients. If system monitoring by IDPS or other methods violates legal or regulatory conditions for monitoring user activity, e.g. data protection law, failure to advertise the use of monitoring tools correctly could result in the collected information being considered illegal snooping and being rejected in court, even possibly leading to prosecution of the organisation performing the monitoring.

22. Before embarking on an IDPS deployment, the CIS business owner, SAA and IAOA must estimate the threat scenario to which the CIS is to be exposed. They must then work with NDM to define which events are likely to be security-relevant in order to be able to specify the functional capabilities of the IDPS to be implemented – what kind of events, how many in a specific period of time, scalability, response time, reporting capabilities, etc.

23. Running IDPS is an ongoing task requiring in depth knowledge of the security features of the CIS components, the IDPS components themselves, and security techniques, all of which change over time due to security or functional updates and patches. Not only must suitably trained experts be engaged to set up and run the IDPS, but also continuous professional education must be required of the experts supporting the IDPS, both in new attack and defence techniques and in security features of CIS components, such training being supported by the NDM.

24. Since IDPS sensors can cause an increase on the load on the components of the monitored CIS, on the components of the IDPS, as well as on the support team which has to manage and tune the IDPS, the number of sensors needs to be kept to the minimum necessary but still sufficient to monitor network activity. It is also strongly discouraged to deploy all planned sensors at once.

25. IDPS sensors console, analysis, and database servers must communicate with one another and with the security support team members. When designing the IDPS, the level of sensitivity and/or classification of this IDPS traffic must be set and documented.

26. As a rule IDPS components must themselves be protected against attack and all IDPS traffic protected against casual eavesdropping by encrypting it and/or confining it to a secure management network, separate from the CIS being monitored. Unless the IDPS traffic is transmitted inside a secured area or contiguous secured areas as defined by the Council Security Rules, encryption products approved to the classification level of the IDPS traffic must be used when IDPS traffic is sent over untrusted communication networks.

27. The type of IDPS sensors in the CIS must be carefully planned. As a rule, various types of sensor will be needed to be able to detect unexpected events of various kinds. Depending on the expected threat scenario and the type of event which is considered to be most relevant as an indication of a security incident, sensors of the types network, host, network behaviour analysis and, if applicable wireless must be chosen. Annex II lists events which should be collected from the various types of sensor.

28. The type of IDPS sensor must also be carefully chosen. the risk of using "in-line" sensors where the IDPS function is embedded in CIS components, which reduces their performance, should be balanced against that of using "off-line" sensors which may not be able to block attacks as quickly as in-line sensors, or which may react to events in ways that are different from those of the CIS components.

29. As IDPS can on the other hand introduce new risks and as its compromise can result in a dramatic lowering of the security of the monitored CIS, the IDPS itself must be protected by network defence measures. Further, as overloaded IDPS sensors could be configured to either block all traffic or pass all traffic, the capability of the IDPS must be tailored to the expected volume of traffic. This must be done not only when it is initially set up but also continually during its operation, adjusting the nature and capacity of the components to the observed type and volume of traffic handled.

30. The placement of sensors in the CIS must also be carefully planned. For example by considering that

(a) wireless IDPS sensors, when deployed, must cover at least the same volume as used by the CIS for wireless connectivity; on the other hand it is sometimes desirable to extend the wireless IDPS coverage beyond that used by the monitored CIS to detect attempted attacks or the presence of neighbouring rogue wireless access points or scanners;

(b) as a rule, it is strongly discouraged to place network-IDPS and Network Behaviour Analysis sensors outside the protected CIS as they would be overwhelmed by the "background noise" of the untrusted networks; some network-based sensors should be placed directly behind the external Boundary Protection Services which filter the connections between the monitored CIS and all other partly trusted or untrusted networks; often, additional network-based sensors are placed at strategic points in the internal network, for example "network-near" authentication servers or devices handling highly sensitive or critical information;

(c) host-based IDPS should be initially placed only on critical devices or devices which themselves, or because of the information they handle, are considered to be particularly sensitive to inadvertent or malicious disruption;

(d) deployment of agents and solutions for malware protection, a form of host-based IDPS, should take into account the need for ensuring that alerts are produced when individual components of the CIS are unreachable, have obsolete malware protection engines or obsolete sets of malware signatures. Such events must be alerted to the support team so that action can be taken to correct the situation.

31. IDPS rules must be carefully designed. It is quite common for the built-in rule-set of IDPS toolkits to block desired functionality while failing to detect unusual behaviour which could be a security incident. The rule set defines the way the IDPS recognises and reacts to an event. Rules should not be confused with "signatures" of malware or known attacks, although some rules might rely on attack detection using attack or malware signatures, checksums of illicit images, etc. The rule set of any such toolkit used must be therefore be checked and modified by NDM personnel on an ongoing basis. The active rule set must be documented as a function of time to enable correlation of the performance of the IDPS to changes in its configuration.

32. The process for generating rules must be defined and the persons authorised to do this identified. It is also important that decisions as to what sensors to deploy, where and the reasons for doing so are documented for further reference.

33. Based on the expected type and frequency of events related to feared threats exploiting vulnerabilities of the CIS components, alerting and protection rules must be activated or deactivated as considered appropriate. Exchange of information with other trusted security teams, regular review of publicly available security information, etc. are strongly recommended to obtain information on the types and probability of attacks being detected on CIS of a similar type and thus be able to make an informed choice of which IDPS rules to activate in the initial phase.

34. Sensors typically report very large numbers of suspicious events until experts in intrusion detection have adjusted (tuned) their configuration to give the correct balance between false positives – events suspected to be violations of policy but which are normal behaviour – and correct detection of real incidents. It is recommended to tune the IDPS so that it collects the maximum number of real events, at the expense of reporting a number of false positives.

35. Prior to deploying IDPS, features of the CIS which define IDPS design must be documented, recording at least:

(a) where sensitive, important or interesting information is stored, who accesses it and how it is protected;

(b) what kind of information will most likely be needed in a security incident;

(c) which devices are most likely sources of security-relevant information;

(d) which parts of the CIS would cause most damage and disruption if compromised; and

(e) how many personnel resources are trained and available for support of IDPS.

IV. IDPS TESTING AND INITIAL TUNING

36. It is strongly discouraged to use sets of real events from a production CIS for testing purposes of IDPS.

37. IDPS must be first deployed in a test environment – an isolated replica of the production CIS – to enable the members of the security support team to familiarise themselves with the output of the IDPS and to tune it before deployment on the production CIS, as well as to determine whether the IDPS affects the functioning of the monitored CIS.

38. Further, it is strongly recommended to initially deploy very few sensors, as otherwise the security team will be overwhelmed with a multitude of real and false positive events.

39. Performance and load testing are also necessary to enable the size of the final IDPS components such as database services, analysis servers, and sensor appliances or workstations to be correctly estimated prior to scale-up and roll out to the CIS to be monitored. It is again strongly discouraged to use real live data from “production” systems of any kind for testing. The SAA in consultation with the NDM and IAOA must decide on the methodology to be used for testing performance and suitability for the CIS to be monitored as, while there are no generally accepted and universally applicable standards or methods for IDPS testing, various approaches are available.

40. As described in paragraph 28, any sensors deployed on the active servers, network devices or workstation components of the CIS can cause high processing loads and sluggish behaviour of the CIS, a decision must be taken whether to choose independent IDPS appliances rather than deploy IDPS sensors on the components of the CIS itself. The use of custom IDPS appliances (hardware and firmware) is not only more robust but also relieves the monitored systems from the work of detecting events and communicating them to the IDPS back-end systems.

41. During initial testing, it is strongly recommended to deactivate the attack prevention functionality of IDPS sensors until the security team gains experience in which events are high priority, which are less important, and which must be ignored. Once this has been established, prevention functionality, in parallel to alerting of the security support team, can be gradually enabled on IDPS sensors in the test phase.

42. During testing it must be determined whether to purchase an integrated IDPS solution consisting of all types of sensor required, the IDPS console, analysis and action server, as well as database servers from a single vendor, or to use SIEM or GRC tools to integrate IDPS components from different vendors in order to remain vendor-independent and/or to be able to choose the best offering for each kind of IDPS required.

43. Testing must also determine the performance of IDPS components under high load, which might cause many kinds of sensor to fail or pass all observed events without processing them. Such load testing must be performed to determine whether changes are needed to the planned IDPS deployment, such as the adding IDPS load balancers and/or choosing sensors which perform well under high load, for example by analysing and logging the first few of a flood of identical events and ignoring the rest.

44. In addition, as described in paragraph 28 and 40, a decision must be taken during the testing phase whether to use the network components of the CIS themselves for IDPS connectivity or to deploy an independent, physically or logically (e.g. V-LAN) separated security-management network for IDPS purposes, with access to the security-management network restricted to trained and trusted members of Network Defence Management (NDM).

45. Alerts must be sent to the NDM independently of whether the IDPS takes preventive action to block the attack or not. Prior to or at least at the start of testing, a decision must be taken on which alerts shall be sent to NDM and by which method, depending on the risk they present to the CIS. Alerts can be notified for example by flashing a warning on a console of tools[5] monitoring the CIS, sending an instant voice or data message, or by email.

46. The initial configuration and the IDPS configurations modified as a result of testing must be documented and the changes justified. A "first final" description of the IDPS must be documented in as much detail as possible at the end of testing and used as a basis for further IDPS development as well as for accreditation of the CIS monitored.

V. IDPS ROLL OUT TO PRODUCTION

47. Once initial tuning has been performed and tested in an isolated test CIS, IDPS can be incrementally rolled out on the CIS to be monitored.

48. If the IDPS devices used for testing are redeployed on the production CIS, all "event" and "action" test records must be wiped from the IDPS database prior to redeployment, while retaining the configuration of the IDPS components and duplicating them in the production CIS.

49. To ensure minimum disruption of service and maximum benefit from the use of IDPS when it is released to production, sensors must initially be deployed only on a small number of CIS components which are

    (a) important or critical to the business of the organisation using them;
    (b) well known to the IAOA and the NDM; and
    (c) physically easily reachable.

50. The IDPS rules and sensors usually require renewed tuning when they are moved from the test environment to the production CIS to adapt them to the live situation, both just after the move and as an ongoing exercise.

[5] Common system and network monitoring products e.g.: (open source) NAGIOS, Zabbix, (closed source) HP- Openview, CA-Unicenter TNG, IBM-Tivoli, etc.

51. The configuration of the IDPS must be regularly monitored and reviewed. This must be done as, no matter how thorough the testing, the production CIS may show unforeseen events or sequences of events which could result in service interruptions unless properly processed.

52. It is therefore strongly recommended to disable all “prevention” functionality of any IDPS in the initial deployment stages on the production CIS in order to allow the NDM to tune the IDPS rule set to the real situation.

53. This is particularly important for “Network Behaviour Analysis” (NBA) sensors, which must learn what “normal” behaviour is in order to be able to detect abnormal behaviour. After any changes are implemented in the CIS, be it in volume or type of processing performed, NBA sensors must be tuned again to adjust them to the new situation.

54. It is strongly recommended to maintain a record of all configurations of the IDS as a function of time. Adding new malware definitions or adding new tests to a running IDPS can cause a drastic change in the number and nature of events detected by the IDPS. This does not necessarily mean that the monitored CIS have suddenly become insecure, just that more events are being reported. This permits correlation of changes in output to IDPS configuration changes, especially if IDPS results are used for generating security metrics for reporting to higher management and security decision-making.

55. When the NDM is satisfied with the ratio of false positives to real alerts, the prevention features of the IDPS can be gradually activated. Change control must be implemented to alert affected support teams and users to the potential of service interruptions following such activation.

56. Experience shows that it is advisable to analyse IDPS data on a dedicated computer, not on that which is recording and correlating events as the processing of analytical tasks often causes overload resulting in late entry of events and missing of security incidents.

VI. IDPS MAINTENANCE AND SCALE UP

57. The IDPS should be run for several months to gain experience before performing modifications or upgrades. This enables further fine tuning of the rule sets and sensor locations as experience is gained.

58. Once the IDPS is considered stable, the NDM should consider whether to increase or decrease the number of sensors and types of events being monitored. If IDPS rules never report events, they can usually be safely deactivated after assessing the risk, e.g. the rule would not be deactivated if the event detected is considered to be of critical importance. On the other hand, it may be necessary to change rules to detect or prevent events such as:
    (a) sending sensitive information outside the monitored CIS (data loss prevention function)
    (b) new attack techniques
    (c) exploitation of new vulnerabilities reported by vendors, security expert organisations, etc.
    (d) unauthorised attachment or modification of devices, wireless access points, etc.

59. Any changes to the IDPS should be covered by change orders in order to alert support staff and the NDM of the CIS.

60. All IDPS settings, the location of the sensors and the response rules must be revised at regular intervals, at least every 12 months, as well as after any security incidents affecting the monitored CIS.

    Following major changes to the IDPS, the need to reaccredit the monitored CIS must be determined by the SAA in co-operation with the NDM and IAOA and the CIS must be submitted to renewed accreditation if considered necessary.

VII. CONTINUOUS IMPROVEMENT AND SECURITY METRICS

61. IDPS output should not be used without comment to generate security metrics as an increase in the number and type of events intercepted, which indicate that the IDPS is actually performing its function, may have the effect of disconcerting recipients of such reports.

62. Judiciously processed, IDPS output can on the other hand provide useful information to the NDM as well as higher management as to the effectiveness and efficiency of the IDPS and other security measures protecting the CIS being monitored.

63. IDPS provides a picture of the real threat scenario to the CIS covered by IDPS. It should report how many events of a certain type are being detected. This should be used by technical security personnel to refine the network defence measures of the monitored CIS.

64. The IDPS security support team should also check whether any of the events feared during the IDPS design phase are not generating any reports, in which case the cause should be determined and appropriate corrective action taken, e.g. by moving sensors to better detect the events, or by disabling the rules if it turns out that the threat is not real.

VIII. GLOSSARY

GRC	Governance Risk management and Compliance:- toolkits which not only have SIEM-type functionality but can also perform or trigger other tasks for establishing compliance to regulatory or legal requirements.
IPFIX	Internet Protocol Flow Information Export (IPFIX) is an IETF protocol, as well as the name of the working group defining the protocol for network traffic measurement - It is a common, universal standard of export for Internet Protocol flow information from routers, probes and other devices that are used by mediation systems, accounting/billing systems and network management systems to facilitate services such as measurement, accounting and billing. The IPFIX standard defines how IP flow information is to be formatted and transferred from an exporter to a collector. At time of writing IPFIX is defined by RFC7011, and RFC7012, RFC7013, RFC7014, RFC7015 http://datatracker.ietf.org/wg/ipfix/
Netflow	This is the precursor of IPFIX: Netflow is a feature that was introduced on Cisco routers that give the ability to collect IP network traffic as it enters or exits an interface. Devices that support NetFlow can collect IP traffic statistics on all interfaces where NetFlow is enabled, and later export those statistics as NetFlow records, toward at least one NetFlow collector - typically a server that does the actual traffic analysis. http://www.cisco.com/c/en/us/products/ios-nx-os-software/ios-netflow/index.html
sFlow	sFlow® is an industry standard technology for monitoring high speed switched networks. It gives complete visibility into the use of networks enabling performance optimization, accounting/billing for usage, and defense against security threats. www.sFlow.org drives the widespread adoption of sFlow by end users, network equipment and software vendors.
SIEM	Security Incident and Event Management: toolkits which can consolidate and correlate input from different sources of security information, using predefined rule sets in order to detect unusual events, perform prepared responses to such events and/or inform network defence Management about them.
SPAN (port)	Switched Port ANalyser , also called RSPAN - Remote Switched Port ANalyser or Mirror Port - a port on a network switch which is set up to copy the network traffic on another port or an entire VLAN on the same switch.
TAP	Test Access Point - a network ID(P)S connection which splits network traffic sending a copy for analysis while leaving the normal network traffic unaffected

ANNEX I - SENSOR TYPES: STRENGTHS AND WEAKNESSES[6]

I NETWORK IDPS

I.1 Strengths

(a) identification of devices. Lists of IP addresses of source devices can identify which hosts are active in the network.

(b) identification of the type of device. The use of fingerprinting techniques on the network traffic can identify the type of device and the version of its operating system.

(c) identification of network protocols (ICMP, TCP, UDP ports, etc.) used by each host;

(d) identification of services (applications) on hosts by "fingerprinting" can identify which layered products are being used to offer services such as email, web services, database engines, enterprise resource planning, document management software, etc. as well as the version of the product being used.

(e) network mapping. The network paths between "active" devices on the network can often be determined.

(f) recording of entire network conversations. network sensors can often perform packet capture and reconstruct the entire exchange of information during a specific network connection.

[6] For further details see NIST Computer Security Division (CSD) DRAFT Special Publication 800-94 Revision 1, Guide to Intrusion Detection and Prevention Systems (IDPS) (July 2012) or its successor documents.

I.2 Weaknesses

(a) it is not possible to detect of attacks using encrypted network traffic: encrypted traffic (HTTPS, SSL, SSH and other forms of encryption) cannot be analysed by typical network sensors;

(b) it is often not possible to detect attacks against the IDPS itself ;

(c) sensor may not perform under high load: large numbers of connections or other high network traffic situations may overload network sensors;

(d) it is not possible to detect hosts with no network address;

(e) differences in the handling of network traffic: sensors might interpret network traffic in a different way than the target hosts in the monitored network(s) so may miss attacks due to such misinterpretation;

(f) it only detects attacks at the point of the network to which it is attached (via a SPAN port or TAP)

II HOST SENSORS

II.1 Strengths

The data collection capabilities of host IDPS sensors vary widely and the type of monitoring performed depends on the specific agent which is used to monitor activity of the host. An agent is a piece of software specially tailored to permanently reside in memory and monitor certain features of the host. Some or all of these tasks listed below could be performed by host IDPS agents.

(a) addition or modification of user accounts and/or of account rights and privileges;

(b) analysis of log files of the operating system and/or of layered products;

(c) integrity of processes needed by layered products such as malware protection, email packages, database engines and web server software;

(d) detection of malware or other potentially undesirable programs;

(e) network traffic analysis (on the host's network cards);

(f) changes in network configuration;

(g) addition of storage devices – fixed or removable;

(h) other changes in host configuration – physical and custom settings

(i) file integrity checking;

(j) unauthorised (or unusual) access to critical files and folders, e.g. to install malware disguised as print drivers;

II.2 Weaknesses

(a) conflict of the IDPS agent with other system monitoring agents or boundary protection devices such as local firewalls, VPN client or server software;

(b) slower processing – network and file access being usually the most affected;

(c) requirement to reboot the host after installing or updating the agents;

(d) late alerting of events and hence late response as the agent usually sends its information for analysis at regular intervals, not continuously.

III NETWORK BEHAVIOUR ANALYSIS SENSORS

III.1 Strengths

NBA sensors can detect events listed above for network IDPS sensors and in addition:

(a) which protocols are being used for communication between hosts;

(b) which hosts are communicating with which others;

(c) the volume of such traffic as a function of time, often displayed graphically;

(d) changes in any of the above: unusual network behaviour, unauthorised connections between devices, etc. even if it is encrypted to hide attacks from signature-based detection;

III.2 Weaknesses

(a) delay in reporting attacks – NBA sensors must first learn what is normal then after observing changes in behaviour, will report the change to central monitoring;

(b) as a result of the above, rapidly evolving attacks can be detected too late to prevent service disruptions in the network

(c) stealth attacks which do not significantly affect network behaviour in a short time go undetected.

IV WIRELESS SENSORS

IV.1 Strengths

(a) Detection of unauthorised wireless devices. All wireless sensors typically detect active wireless devices by their hardware (MAC) address as well as other techniques such as fingerprinting to identify the make, type and operating system version of the device.

(b) Identification of wireless networks. Networks can be identified by their names - their SSIDs (Service Set IDentifier). They can often also determine the type of wireless communication being used (802.11 a,b,g,n, ac), whether encryption is enabled, which channels are being used, as well as various other characteristics of the various WLANs detected – signal strength etc. The relative strengths of signals from a rogue access point can sometimes therefore be used to calculate the approximate location of the unauthorised device and triangulation can be used if directional wireless sensor devices are available.

IV.2 Weaknesses

(a) Wireless sensors cannot detect passive eavesdropping. While the correct use of encryption at base station and client level will prevent casual eavesdropping on wireless network traffic, the use of no or weak encryption methods such as WEP can permit unauthorised interception of wireless communications.

(b) Channel limitations. As wireless traffic is carried over different channels in two different frequency bands[7], if the sensor can only monitor one channel at a time it may miss attacks being carried out over other channels. More modern sensors therefore have several powerful radios and several powerful antennae and can scan multiple channels simultaneously.

(c) Range limit. An attack from a weak transmitter in the vicinity of the monitored network devices may be missed by sensors due to the signal being too weak to enable adequate monitoring.

The above two weaknesses can be mitigated by deploying multiple sensors having overlapping detection ranges and having the different sensors monitoring different sets of channels. Special tools are also available to optimise the placement of wireless sensors.

[7] 14 or 21 channels + special channels in some regions, 2.4 and 5 GHz

ANNEX II - EVENTS TO MONITOR

Below are lists of typical features of events which should be collected by the IDPS where possible. The lists below are additive - i.e. any information in any list should be collected for all events if possible, the analysis function of the IDPS being subsequently used to filter and report the events collected.

I. General Event Information

(a) Name of sensor

(b) Sensor feature generating the event (functionality/process ID)

(c) Date of event

(d) Time of event

(e) Date when logged

(f) Time when logged

II. Network Event Information

(a) Source device (host) name

(b) Source hardware (MAC) address

(c) Source network (ip) address

(d) Source domain

(e) Source account (name or user id)

(f) Source account privileges

(g) Target device (host) name (dns name, netbios name etc.)

(h) Target MAC address

(i) Target ip address

(j) Target domain

(k) Target account (name or user id)

(l) Target account privileges

(m) Protocol (tcp/icmp/udp/etc.)

(n) Target network service (arp, icmp-echo, tcp-telnet, tcp-http, etc.)

(o) Source (tcp) port number (0-65535)

(p) Target (tcp) port number (0-65535)

(q) Communication direction (source-target/target-source)

(r) ‘Flow' details (duration, volume, frequency, etc.)

III. File Information

(a) File location (path)

(b) File creation date/time

(c) File modification date/time

(d) Old file name

(e) Old file extension

(f) Old file size

(g) Old file acl (permissions, owner)

(h) Old file hash (specify hashing algorithm)

(i) New file name

(j) New file extension

(k) New file size

(l) New file acl

(m) New file hash

IV. Email Information

(a) Mail server type

(b) Mail server device name

(c) Mail server name (when several mail services are running on one device)

(d) Source mail address

(e) Target mail address

(f) Mail relay(s) name(s)

V. Web Access Information

(a) Web server type (Apache, IIS, etc.)

(b) Web site type (asp, php, etc.)

(c) Web server name

(d) Accessed resource (database record/file) location (path)

(e) Source browser

(f) Source browser version

(g) Request method (get, post, etc.)

(h) Type of attack (classified as per WASC TC2 or later versions)

VI. Wireless Information

(a) MAC address of source

(b) MAC address of target

(c) Channel used

(d) SSID of target

(e) Target device type

(f) Version of 802.11 (a, b, g, n, ac)

VII. Custom Information

Several Host IDPS monitoring specific applications generate event information which is more detailed and specific to the product used and its configuration, for example the presence of unauthorised or dangerous software, access to blocked web sites or services, etc.
It would exceed the limited scope of this Annex to go into detail for such events but allowance must be made for them by the IDPS either by ensuring that the IDPS can handle the specific format of such events or using what is often termed a "connector" to generate standard format event information from all sensor output.