

Next steps

This document has dealt only with the logs that can be recorded by individual computers and other systems. A great deal of useful information and early warnings can also be obtained by looking at computers and networks in combination. Any organisation that is concerned to protect its own systems and reputation should also be developing systems to monitor these systems. Two examples are given below of what can, and should, be done.

Network Flows

Networks can be characterised very effectively by knowing what flows of traffic are taking place along them. Flows can be classified by their source and destination IP addresses, or groups of addresses, and ports, along with the volume of traffic making up the flow. For example it would be quite normal in an open-access workstation room to see large numbers of packets coming from the HTTP ports of external machines into the workstation room – web browsing is a legitimate use for these types of systems – but much less normal to see traffic flowing out from the HTTP port on the workstations. The latter situation may indicate that someone has set up a web server on a workstation, either with or without authority, or that systems have been compromised and are being controlled by a remote intruder. Thus simple flow information can be effective in detecting both security problems and breaches of policy.

Network flows are often the only way to trace denial of service attacks, since these commonly use forged addresses to conceal their origin. If addresses are forged then an attack can be hard to trace at the IP level; instead information taken from routers and switches will be needed to determine which ports or interfaces are carrying the traffic. Such information can rarely be gathered from a single central point but needs effective, secure reporting and management to be set up on the network devices when they are deployed.

Intrusion Detection

Network flow monitoring examines where packets are going to and from; network Intrusion Detection Systems (IDS) examine the content of packets or search for patterns in time. For example an IDS might be configured to check HTTP packets for the commands used by well-known attack tools, or could detect that a particular IP address had sent packets to all the addresses in a range. IDSes can be very effective at warning of known problems, though they are less good at identifying new, suspicious activity. Some IDS packages can take action to respond to a detected threat, either by blocking the hostile traffic or by targeting a response at the apparent source. However these options, sometimes known as Intrusion Prevention Systems (IPS), run the risk of denying service unnecessarily (any system, computer or human, will occasionally make mistakes) or of taking reprisals against an innocent party. In the worst case an IPS may even assist an attacker if it can be made to react in a way that harms the organisation. For example if an attacker can persuade an IPS that it is being attacked by a vital internal or external system, such as a DNS resolver, then the IPS may effectively cut off the organisation's access to that system or the Internet.

Source URL: <https://community.jisc.ac.uk/library/janet-services-documentation/next-steps>