

Port and address scanning

Information for a Janet organization on scanning activity that may affect their network.

Within an organization's own network scanning activity may be a legitimate form of audit or of information gathering; but it is almost never acceptable otherwise without the express permission of the managers of the target network.

- If you observe scanning of your own network from outside you are under attack, even though you may not have been singled out as the only target.
- If one of your own systems is scanning addresses anywhere else in Janet or in the Internet, that too is a form of network abuse, and once you notice (or Janet CSIRT tells you) you must ensure that it stops as soon as possible.

Patterns of scanning

Address range scans

The most common pattern is to attempt to connect to the same (TCP or UDP) port or the same small selection of ports at some or all of the IP addresses in some range.
One port, many addresses.

Port scans

In another pattern of scanning, the scanning system attempts to connect to many or all ports on one system at a time.
One address, many ports.

Sequence and speed

A very patient scanner might make connection attempts at a very low rate (one or two a day) from a range of different source addresses, and going through your address range in an apparently random order. To distinguish such a stealth scanning exercise from a collection of accidents with no bearing on the target network will inevitably take some time and is normally impracticable; it would be necessary to find some additional information linking the various packets.

Currently most address scans are much less discreet; a single source address sends packets to the whole of a /24 range in a few minutes or even a few seconds. NetFlow records, or relatively simple selection and sorting of log data, will identify the pattern in the same timescale (although still after the event). Addresses scanned may be in numerical order, and may start with the same /24 or larger range as that in which the source address falls.

It may be that this hasty pattern fits better the usual nature of the scanning system. It is most

likely to be compromised, with the scanning being done by a worm or virus program on behalf of a remote unauthorised user who may only have the system available for a short time before its legitimate owner disconnects it or turns it off. The perpetrator has no need to conceal the identity of the computer which is the direct source of the scanning.

Purpose and impact

The abuse which gives rise to scanning is normally an attempt to take control of as many target systems as possible and to install malicious software which at the very least gives unintended access to the intruder or the intruding malware. It may result in the unauthorised transfer of sensitive data, inclusion in a botnet where hundreds of compromised targets are used for some nefarious purpose by a single controller, activity by the target system unintended by its owner, new scanning by the target system, and possibly damage to it.

Notionally the process has several phases:

- identifying a potential target,
- detecting that it is running and connected to the Internet,
- finding whether some particular service is available,
- finding whether the service is vulnerable to the particular method of entry which the scanning computer is attempting,
- and finally sending the data which will cause a target to cooperate in the abuse.

All except the last of these can be regarded as information gathering.

Most commonly all the phases are applied to each target at one time, so that within a few seconds the legitimate user has lost control of their machine. It may also happen that information is gathered in one pass and further processed to identify vulnerable systems or to prepare attacks specific to each of a smaller number of targets which will be revisited later and may then be taken over very quickly.

Response

Scanning directed at your network

If you depend heavily on Network Address Translation (NAT) and use only a few public addresses, the segment of an address range scan which directly affects you will be over before you or any automated system can spot what is going on.

If you have a large external address space and a system to alert you to attempted intrusions, and it has spotted a scanning pattern:

- it may be worth blocking traffic to and from the source address until you believe the scan will have moved out of your address space;
- you should look at any log or flow records you have, try to identify which of your addresses responded in some way to the scan (packets returned, larger flows) and consider checking them;
- you may report the event to Janet CSIRT.

Scanning originating in your network

If you become aware (perhaps because Janet CSIRT have alerted you) that one or more of your computers are scanning addresses outside your network, you **MUST** stop them as soon as practicable:

- identify the machines concerned using address translation logs (NAT gateway) and login records if necessary;
- remove them from the network if possible;
- make them safe before reconnecting them.

This is the same procedure as for a worm or virus in your network (described in

?Dealing with worms or viruses?), and scanning is almost always associated with those. An alternative short-term response is to block the destination port at your router or firewall.

Reducing the risks from scanning

Scanning directed at your network

A ?default deny? policy for inbound traffic at your router or firewall will ensure that only the addresses and computers you actively wish to be reachable from outside your network are candidates for an outside scan.

If you cannot apply such a policy you may still be able to block certain ports, and you **SHOULD** block some permanently; see ?Blocking LAN service ports?.

Each computer in your network should only be running the services, and listening on the ports, it needs to for its intended purpose. Although difficult to enforce in general, this is normally straightforward for server systems. You may be able to find which ports are listening with *netstat* on the machine itself, or with *nmap* or *mscan* from elsewhere in your network.

Make sure the operating system and anti-virus products on all your computers are promptly updated.

Scanning originating in your network

The comments above about a ?default deny? policy and ?Blocking LAN service ports? apply equally to outbound traffic.

Most scanning is due to compromised or infected computers attempting to compromise or infect further ones, and operating system and anti-virus updates will help to prevent that initial compromise.

You should ensure that local policy provides for your NAT gateways to log their address translations, for the logs to be retained for some period (a week should be adequate for the purpose of tracing the source of a scan) and then to be destroyed.

Related pages

- [Blocking LAN service ports](#) ^[1]
 - [Dealing with worms or viruses](#) ^[2]
-

Source URL: <https://community.jisc.ac.uk/library/janet-services-documentation/port-and-address-scanning>

Links

[1] <http://community.ja.net/library/janet-services-documentation/blocking-lan-service-ports>

[2] <http://community.ja.net/library/janet-services-documentation/dealing-worms-or-viruses>