Preventing Lateral Movement

Guidance for preventing lateral movement in enterprise networks.

This guidance explains how system owners can prevent and detect lateral movement within their enterprise networks. It will help you to:

- improve the chances of spotting an intruder
- increase the difficulty for an attacker to reach their goal once inside your network

Implementing the recommended security controls outlined below - including monitoring to detect the early stages of lateral movement - can reduce the potential for serious damage.

Platform-specific guidance

- The steps below can be applied in networks regardless of the platforms in use. However, for guidance on specific platforms see our Mobile Device Security Collection.
- To learn more about lateral movement in an enterprise environment (in this case using Windows infrastructure), please refer to the following whitepaper: CERT-EU Security Whitepaper 17-002: Detecting Lateral Movements in Windows Infrastructure (PDF).

What is lateral movement?

After an attacker has gained an initial foothold in a network, they will typically look to broaden and cement that foothold whilst gaining further access to valuable data or systems. This activity is known as lateral movement.

Following the initial compromise of a host, the first step in lateral movement is to perform internal reconnaissance of the network. This gives the attacker an idea of their location within the network, and its overall structure. To solidify their presence and maintain persistence, the attacker will usually try to compromise additional hosts and escalate their privileges, ultimately gaining control of their target (such as a domain controller, a critical system, or sensitive data).

Any credentials that the attacker collects will give them (what appears to be) legitimate access to more hosts and servers. Once the goal has been reached, data can be exfiltrated, or systems and devices sabotaged.

Why protect against lateral movement?

The security controls advised in the NCSC's Mitigating Malware guidance can reduce the risk of an initial attack succeeding. **However, you should assume that an attacker with sufficient time and resources will eventually be successful.** It's therefore important to:

- detect breaches as soon as possible
- implement internal security controls to reduce the damage done by an attacker post-breach

Networks with strong boundary protection but no internal security give attackers free rein to traverse the network once they have gained access. The chances of achieving their goals will increase the longer that they're able to maintain a foothold.

Protecting your organisation

Applying the following protections will buy time and make it easier to detect attempts at lateral movement.

1. Protect credentials

All credentials on a network, especially those of administrator accounts, should be adequately protected to prevent attackers using them to gain access to devices and systems.

A common type of attack involves stealing a security token to gain access to another device or server. 'Pass the hash' is an example of this, where a stolen hash is used to authenticate the attacker. Passwords should not be stored in plain text by users or systems, and password hashes should be protected to prevent attackers easily accessing them.

Credentials that are used to authenticate to a device (as well as credentials used to authenticate to services) will both need to be protected by the device. Devices that support hardware-backed credential storage will better protect these credentials. Work credentials should not be entered into any other device except for those approved for work use, as these devices may not adequately protect the credentials.

In summary:

- Do **not** store passwords in plain text, and ensure password hashes are stored in protected areas.
- Use devices with hardware-backed credential storage where possible.
- Only use work credentials on devices and services that have been approved for work use.

2. Deploy good authentication practices

Authentication should be easy for the user, but make it difficult for an attacker to gain access. Follow the NCSC password guidance to ensure your policy follows best practice. For example, do not re-use passwords for different systems, and consider the use of password managers in your organisation. This will limit the number of users storing credentials in plain text.

Logon restrictions (such as password lockout and throttling) reduce the chances of an attacker authenticating with a host if the credentials have not already been acquired. Ensure that a single account cannot grant access to all devices and components across an enterprise, particularly if those accounts are privileged.

Multi-factor authentication (MFA) should be used for internet-facing services to combat brute forcing and password guessing attacks. MFA can also be used as a physically separate factor on high-privilege devices that malware cannot use remotely.

Single sign-on (SSO) can be used to limit the number of passwords in use and reduce the chance of them being stolen. We also encourage the use of alternative technical authentication methods such as biometrics, single-use sign-in links (magic links), smartcards, and hardware-backed PINs.

In summary:

- Follow the NCSC password guidance, and do not re-use passwords for different systems.
- Consider using passwords managers in your organisation.
- Enable logon restrictions/throttling.
- Use multi-factor authentication for internet-facing services and high-risk accounts.
- Where possible use alternative authentication methods to passwords.

3. Protect high privilege accounts

Local and domain administrative accounts - with access to most systems and data - are powerful tools in a network. Their use should be tightly controlled and locked down.

Administrators should use separate accounts; one for day-to-day business use (such as web browsing and emails), and a privileged administrator account that should only be used on separate admin devices. This reduces the risk of an infected device being used for admin purposes.

Administrator accounts should be prevented from browsing the web and accessing emails, and only be used when a task requires elevated permissions.

In summary:

- Administrators should use a normal account for normal user activities, and a separate administrator account for administrator activities only.
- Use separate devices for normal accounts and administrator accounts if possible. If not, consider using the 'browse down' approach.
- Lock down administrator accounts to prevent high risk actions such as browsing the web and accessing emails.

4. Apply the principle of least privilege

The principle of 'least privilege' (where accounts and users have the minimum amount of access needed to perform their role) should be implemented wherever possible. A tiering model for administrative accounts ensures they only have access to the specific administrative capabilities needed, rather than all of them. Using various tiers of administrative accounts limits the number of very high privileged accounts in use, and reduces the access an attacker gains if a lower privilege administrator account is compromised.

Accounts with full privilege across an enterprise (such as a domain admin, global admin, or cloud admin account) should **not** normally be used. Whilst they are required for some tasks (such as initially building a network, performing upgrades, creating new privileged accounts, or disaster recovery), lower tier administrative accounts should be used for most other tasks.

Using time-based privileged access can help reduce the impact of a leaked admin credential, especially as it will be audited every time the user requests or receives it. Identifying high-risk devices, services and users can help in planning granted privileges, ensuring that those with the highest risk have the lowest privileges.

In summary:

- Use a tiering model for administrative accounts so that they do not have any unnecessary access or privileges.

- Only use accounts with full privileges across an enterprise when absolutely necessary.
- Consider the use of time-based privileges to further restrict their use.
- Identify high risk devices, services and users to minimise their accesses.

5. Lock down devices

Any device or system that is part of your network (even those not directly connected to the internet) can be targeted in the lateral movement stage of an attack. All devices should be kept up to date, with the latest patches deployed as soon as possible. Automated updates can also be used to simplify this process, although it is important to ensure that redundant pairs of devices update at different times to maintain redundancy.

Endpoints should be configured securely by following the NCSC mobile device guidance. If possible, applications should be allow listed so that only approved applications can run. This can also be done by using an architecture that only allows applications to be installed and run if they originate from a trusted source.

In addition to firewalls on the network boundary, **local** firewalls on hosts should be enabled to restrict unnecessary inbound and outbound traffic. By default, firewalls should block all inbound connections (such as SMB) and only allow those that you need. The list of approved connections should be regularly reviewed to remove any that are no longer needed.

Secure boot mechanisms should be enabled where possible, to ensure the integrity of the boot process on devices, and increase the difficulty for an attacker to gain persistence on a device.

Finally, follow our Macro Security guidance to reduce the risk from malicious Macros.

In summary:

- Apply patches to all devices as soon as they are released, and use automated updates where possible.
- Use allow listing to control and restrict the use of applications.

- Follow the Macro Security guidance.
- Enable local firewalls on hosts.
- Use secure boot mechanisms if available.
- Follow the NCSC mobile device guidance.

6. Segregate networks as sets

Network segmentation (or segregation), involves splitting up a network into various network segments. This greatly increases the difficulty for an attacker to reach their goal once in the network, as their point of entry may not have any means of reaching the target data or system.

Systems and data that do not need to communicate or interact with each other should be separated into different network segments, and only allow users to access a segment where needed. As stated in our guidance on network security, "Segregate networks as sets: identify, group and isolate critical business systems and apply appropriate network security controls to them."

These security controls should ensure that all data and connections originating from within the network boundary are **not** automatically trusted. The ISO 27001 and 27002 standards provide some insight into network segmentation and best practice for implementing this in a network.

In summary

- Segregate networks as sets: identify, group and isolate critical business systems and apply appropriate network security controls to them.

7. Monitor networks

It is vital to monitor your network for any security events that may be of interest. As new vulnerabilities are constantly discovered, determined attackers will eventually gain access, regardless of how well your network is protected. Once this happens, monitoring the network is the only way you can identify a breach, then react. The Network Monitoring section from our '10 Steps to Cyber Security' provides a starting point, and our Security operations centre (SOC) buyers guide provides more details on the process of monitoring, and what to look for.

The basis of monitoring is the recording and storing of logs for potentially interesting security events. Systems can then analyse these logs and look for suspicious behaviour that may signal an attacker has compromised your network, and alert those responsible. You should enable any logging and security auditing features in the systems and technologies that your network uses (such as those on firewalls and other network architecture), and also logging within operating systems.

Knowing the location of the high value assets within a network allows you to provide more detailed and sensitive alerting. High value assets can include important services and servers (such as a domain controller) in the network, in addition to various users and accounts. Some notable users include:

- privileged users (due to the access they have)
- directorate accounts (due to the information they may contain)
- social media accounts (due to the potential for reputational damage if compromised)

You should be familiar with your network as a whole, including its structure and how it's used. Maintain an audit of all devices that can connect to the network, and update it regularly to help identify illegitimate use. Unusual activity can be present on the network protocol layer, but also in application-specific circumstances, such as credential usage and authentication events.

Attackers will try to blend in with your usual network traffic using legitimate tools and systems to move laterally, meaning it is often overlooked by typical AV software and is much harder to spot. Being aware of the common tools and processes that an attacker could utilise will greatly increase your chances of identifying them.

The biggest challenge in network monitoring is identifying genuine security incidents, rather than false positives which are common in the large volume of ‘noise’ present in a network. Understanding your network and the typical behaviour of its users can help alleviate the problem of false positives, as you become more adept at spotting unusual activity. By segmenting a network, you have the opportunity to focus monitoring on the focal points of traffic flow that are created between segments.

In summary:

- Follow our network monitoring and Security operations centre (SOC) buyers guide publications.
- Enable any logging and auditing features on your systems, and use them to detect unusual activity.
- Keep an audit or record of all devices that can connect to your network, and understand the high value assets.
- Understand and become familiar with your network, and how it's typically used.

8. Consider using honeypots

Honeypots are systems set up for the sole purpose of being attacked.

Production honeypots, set up internally in a network as a decoy for real systems, can be valuable tools in detecting an intrusion into your network. As honeypots are not legitimate systems on the network (and contain no real data or services) unexpected connections can be assumed to be hostile activity (because genuine users have no need to access the honeypot). If you detect interactions with the honeypot it should be immediately investigated. Production honeypots should be used to complement network monitoring and other intrusion detection techniques.

Research honeypots do not benefit the network directly, but are set up to gather information about the latest techniques that attackers use.

Using honeypots does introduce some risks. Research honeypots are inherently risky, as they encourage attackers to interact with them. Production honeypots are less risky, but still introduce some risk to an organisation, depending on their level of complexity. For example, they could be exploited and used as a platform to launch attacks on legitimate systems in the network during lateral movement.

For these reasons, honeypots should only be used if you have assessed the impact of incorrect implementation, and have the relevant expertise in your organisation to do so.

In summary:

- Consider the use of a production honeypot in your organisation, provided you have the expertise to do so and understand the risks involved.

PUBLISHED

8 February 2018

REVIEWED

10 March 2021

VERSION

1.0

WRITTEN FOR

Large organisations

Public sector

Cyber security professionals