CERT-EU Security Guidance 22-001

Cybersecurity mitigation measures against critical threats

CERT-EU Team
ver. **1.0**
09-03-2022

Contents

History:

- *09/03/2022 - v1.0 - Initial publication.*

1 Introduction

On February 14, ENISA and CERT-EU made a joint publication strongly encouraging all EU-based organisations to implement a set of cybersecurity best practices.

Building on this joint publication, CERT-EU is making available the following specific implementation recommendations. By applying these systematically, organisations can boost their cybersecurity defence and resilience. This would allow them to:

1. Improve their cybersecurity posture to fend off a wide range of attacks and limit the number of cybersecurity incidents.
2. Detect and react to cyber operations that may be carried off by sophisticated threat actors.

If you have suggestions that could help improve this document, please contact us at services@cert.europa.eu. We always appreciate constructive feedback.

2 Scope and audience

This document addresses mitigations against a wide range of threats, including critical ones.

The audience of this document is mainly security officers (e.g. CISOs). It is also aimed at decision makers (both in IT and general management) and entities that support organisational risk management.

3 Cybersecurity posture improvement

3.1 Strategic aspects

1. Build your cybersecurity strategy centred on your key assets.
2. Put people at the centre of your cybersecurity strategies and create a culture where protecting your organisation from threats is everyone's responsibility, from the board of directors to the frontline personnel.
3. Create an effective, lean and efficient governance structure. Oversight from your governing body and strong leadership from the top management are essential to successfully manage your cybersecurity risk. Identify your key controls, measure their coverage and effectiveness and report them to your governing body and top management.

3.2 Operational aspects

1. Follow the security best practices proposed by vendors to harden their products and **give priority to the secure management of high-privileged accounts and key assets**. Whenever possible, benchmark your systems against well-established frameworks, such as the Center for Internet Security's CIS Benchmarks or The Profile from the Cyber Risk Institute (CRI).
2. Follow best practices for identity and access management.
3. Leverage penetration testing and conduct Red team exercises to identify and quickly fix weaknesses on your publicly exposed systems and high-value assets.

4 Prevention and resilience

1. Ensure all services which can be remotely accessed require multi-factor authentication (MFA). Applications that are publicly accessible without MFA are highly susceptible to brute-force attacks, password spraying or remote access with stolen credentials. These applications include, but are not limited to, RDP, VPN services, external facing corporate portals (extranets) and email access (e.g. OWA or Exchange Online). If possible, avoid using SMS and voice calls to provide the second factor and consider deploying phishing resistant tokens such as authenticator apps or FIDO2 (Fast IDentity Online) security keys.

2. Threat actors are known to also target credentials used by staff members in their private life such as passwords for accessing personal emails or social media accounts. Staff members should never reuse the same password for corporate services. Moreover, they should be strongly encouraged to use MFA in their private usage whenever an application supports it (e.g. LinkedIn, social media platforms). In addition, users can be advised to check on Have I Been Pwned if their personal email addresses are present in any known data breach and to change immediately any compromised password on the relevant websites and/or applications. The use of password managers should be encouraged to generate and/or securely store complex, solid passwords to authenticate to corporate applications as well as private ones.

3. Ensure that software is up-to-date, prioritising updates which address known exploited vulnerabilities on internet-facing assets: web servers, email gateways, virtualisation solutions, reverse proxies, routers, load balancers, VPN products, etc. Reengineer your vulnerability management processes in order to deploy high and critical severity patches as quickly as possible. Encourage also your users to regularly patch their personal systems at home (e.g. computers, smartphones, tablets, connected devices like TV sets, videogame consoles, home routers, etc.).

4. Pay special attention to secure your cloud environments before moving critical assets there. Use the strong security controls that are available on cloud platforms and properly segment cloud system management from on-premise system management to ensure that threat actors cannot jump from one environment to the other.

5. Review your backup strategy and use the so-called 3-2-1 rule approach which states that organisations should keep three complete copies of their data, two of which are locally stored but on different types of media, and at least one copy stored off-site. Your backup strategy should be fully aligned with your business needs by setting explicit recovery time (RTOs) and recovery point objectives (RPOs). Ensure access to backups is controlled, limited and logged. Ensure also that your restoration procedures are well documented and tested at least once a year. It is strongly recommended to increase the frequency of backups for critical data. Users should be trained to save data only on data centre storage devices or, if applicable, on the corporate cloud storage and not on their workstations.

6. Change all default credentials and disable protocols that do not support MFA or use weak authentication (e.g. cleartext passwords or outdated and vulnerable authentication or encryption protocols).

7. Employ appropriate network segmentation and restrictions to limit access and utilise additional attributes (such as device information, environment and access paths) when making access decisions.

8. Conduct regular training to ensure that IT and system administrators have a solid understanding of your organisation's security policy and associated procedures. Vigilantly monitoring the misuse of sysadmin tools can help you prevent attackers from breaching

your network and moving laterally.

9. Create a resilient email security environment by enabling antispam filtering and stripping email attachments from active content, adding a secure email gateway configured to automatically follow field-tested policies and playbooks designed to prevent malicious emails from reaching their recipients.
10. Leverage end-user awareness and phishing exercises. Organise regular cyber awareness events to train your users on common phishing techniques (e.g. identifying spoofed/suspicious messages, macros in documents, compressed files) and the effects of phishing attacks. Provide tailored training to senior management and their assistants as well as to IT system administrators.
11. Protect your web assets from denial-of-service attacks by using a CDN (Content Delivery Network) which will expand your web assets' footprint across multiple servers or use the native high-availability features of cloud platforms. Automate the disaster recovery runbooks for on-premise systems and ensure that you can move workloads to the disaster recovery site with a single click if possible.
12. Block or severely limit egress internet access for servers or other devices that are seldom rebooted, as they are coveted by threat actors for establishing backdoors and creating persistent beacons to Command and Control (C2) infrastructure.
13. Use a security scorecard to track the progress of your key controls and mitigation activities.
14. Work with your procurement and legal teams to add clear cybersecurity clauses in contracts with suppliers before establishing them. Require your suppliers (e.g. Managed Service Providers) and, if applicable, their subcontractors to comply with a minimum set of field-tested cybersecurity standards and regularly review their compliance with those standards. Consider seeking damages or terminating contracts in case of non-compliance.
15. **Last but not least**, apply the more specific technical guidance on hardening against sophisticated attacks in Annex A.

5 Detection

1. Implement robust log collection. You may refer to ENISA's guidance on proactive detection (pp. 12 - 18) which comprehensively covers sources of telemetry.
2. Use carefully curated cyber threat intelligence to proactively search your logs for possible signs of compromise.
3. Detect traces of compromise in your network through well-conceived, regular threat hunting based, for example, on the MITRE ATT&CK® framework.
4. Monitor the activities of the machines in your network with appropriate tools like Endpoint Detection and Response (EDR) and User and Entity Behaviour Analytics (UEBA) since a substantial part of network traffic is encrypted nowadays. This applies both to servers and endpoint systems and it is crucial to ensure that your EDR is set to alert mode if monitoring or reporting is disabled, or if communication is lost with a host agent for more than a reasonable amount of time.
5. Train your users to immediately report any suspicious activity to their local security team.

6 Incident Response

1. Make sure you have the procedures to reach out and swiftly communicate with your national or governmental CSIRT as well as provide access to forensics evidence when asked. Ensure these procedures are known to all the local incident handlers and that they have been tested in advance.
2. At the minimum, you should keep an up-to-date list of the key contacts of your strategic suppliers and service providers.
3. Avoid common mistakes in incident handling such as:

- ignoring a security event without assessing what triggered it and the potential impact,
- preemptively blocking or probing infrastructure used by threat actors (pinging, using `nslookup`, browsing, etc.),
- mitigating the affected systems before responders can collect and/or recover evidence,
- ignoring telemetry sources, such as network, system and access logs,
- fixing the symptoms, ignoring the root causes and doing partial containment and recovery,
- forgoing keeping a detailed record of actions taken and the event timeline.

4. Incident response requires communication amongst several internal stakeholders and it is strongly recommended to have clear, concise communication guidelines prepared and tested in advance.
5. If personal data is in the scope of the incident, seek the advice of your Data Protection Officer and/or your legal team.

7 References

In addition to the guidance documents and recommendations provided by your national or governmental CSIRT, please refer to the online resources below for additional advice.

7.1 ENISA guidance

1. Guidance on Secure Backups, 1 September 2021.
2. Proactive Detection, section on monitoring, pp. 12 - 18, 26 May 2020.
3. Proactive detection – Measures and information sources, 26 May 2020.
4. How to set up CSIRT and SOC, 10 December 2020.
5. Standards and tools for exchange and processing of actionable information, 19 January 2015.
6. Good Practice Guide for Incident Management, 20 December 2010.

7.2 Additional resources

- Framework for Improving Critical Infrastructure Cybersecurity, Version 1.1, NIST, 16 April 2018.
- Joint Cybersecurity Advisory: Technical Approaches to Uncovering and Remediating Malicious Activity, Canadian Center for Cybersecurity, 1 September 2020.
- Turning on Two-Factor Authentication – Microsoft Accounts, Cyber.gov.au, 1 July 2020.
- Protecting Microsoft 365 from on-premises attacks, Argonsys, 18 December 2020.
- Implement Cybersecurity Measures Now to Protect Against Potential Critical Threats, CISA Insights, 18 January 2022.
- The Five Anchors of Cyber Resilience: Why some enterprises are hacked into bankruptcy, while others easily bounce back, Phillimon Zongo, ISBN-10 0648007847.

8 Credits

We would like to warmly thank our colleagues from ENISA for the useful feedback and suggestions they have provided to improve this guidance in the framework of our structured cooperation.

We would also like to thank the European Food Safety Authority, the European Union Agency for the Operational Management of Large-Scale IT Systems in the Area of Freedom, Security and Justice (eu-LISA), the Translation Centre For the Bodies of the EU (CDT), CERT-SE, GOV-CERT.LU and Freddy Dezeure for their suggestions and contributions.

9 Annex A: Hardening recommendations

To better prevent and detect sophisticated cyber attacks, please follow the hardening recommendations below.

9.1 Protect external facing assets

- Patching as quickly as possible external facing assets (e.g. web, email, DNS and application servers) should be a key priority since the exploitation of vulnerabilities have been the initial entry vector behind countless significant incidents in the EU. For example, last year, threat actors achieved initial access by taking advantage of unpatched vulnerabilities in widely used software such as Microsoft Exchange, F5, VMware, Pulse Connect Secure, Cisco routers and Oracle Weblogic.
- Good practices involve automated patching of all non-critical systems and controlled updates on critical assets. The virtualisation solutions that are in place in the majority of the networks help ensure quick rollbacks by creating snapshots of critical servers before patching.

9.2 Enforce MFA

- Applications that are publicly accessible without using MFA are highly susceptible to brute-force attacks, password spraying or remote access with stolen credentials.
- Beside the technical aspects, **it is of paramount importance to educate users on the risks of misuse or abuse of MFA** and to ensure that they will not confirm a push notification on their mobile devices if they haven't tried to login.
- Use Conditional Access and AD Identity protection whenever possible.

9.3 Prevent lateral movement in your on-premise environment

- Deploy a solid endpoint firewall policy to restrict communications between endpoints to the permitted flows only. For any specific applications that may require inbound connectivity to endpoints, the local firewall policy should be configured with specific IP address exceptions for systems that are authorised to initiate inbound connections to such devices. Typical cases that require exceptions are for instance Skype for Business and Cisco Jabber.
- At a minimum, the following common ports and protocols, leveraged for lateral movement, should be blocked in Windows environments between workstation-to-workstation and workstations to non-domain controllers and non-file servers:
  - SMB (TCP/445, TCP/135, TCP/139)
  - Remote Desktop Protocol (TCP/3389)
  - Windows Remote Management (WinRM) / Remote PowerShell (TCP/80, TCP/5985, TCP/5986)

  - Windows Management Instrumentation (WMI) (dynamic port range assigned through Distributed Component Object Model or DCOM).
- Create different network segments, by implementing VLANs or a similar mechanism, and create controlled traffic flows between the segments using filtering functions that protect against unwanted traffic flowing freely in the network.

9.4 Protect your M365 from on-premise compromise

9.4.1 Fully isolate your M365 administrator accounts

M365 administrator accounts should be:

1. Mastered in Azure AD.
2. Authenticated with MFA.
3. Secured by Azure AD conditional access.
4. Accessed only from Azure managed workstations.

These are restricted usage accounts. There should be no on-premise accounts with administrative privileges in Microsoft 365. Moreover, system administrators must never reuse their normal account passwords for their high-privileged accounts.

9.4.2 Manage devices from M365

Use Azure AD Join and cloud-based mobile device management (MDM) to eliminate dependencies on your on-premise device management infrastructure, which may adversely affect device and security controls.

9.4.3 Ensure no on-premise account has elevated privileges on M365

No on-premise account should have elevated privileges on M365.

Accounts accessing on-premise applications that require NTLM, LDAP, or Kerberos authentication need an account in the organisation's on-premise identity infrastructure. Ensure that these accounts, including service accounts, are not included in privileged cloud roles or groups and that changes to these accounts cannot impact the integrity of your cloud environment. Privileged on-premise software must not be capable of impacting M365 privileged accounts or roles.

9.4.4 Use Azure AD for cloud authentication

Use Azure AD for cloud authentication to eliminate dependencies on your on-premise credentials. Always use strong authentication, such as Windows Hello, FIDO, Microsoft Authenticator or Azure AD MFA.

9.5 Protect accounts and avoid credential exposure

Threat actors will look for privileged accounts as part of their reconnaissance efforts. Once identified, they will attempt to obtain credentials for these accounts in order to achieve lateral movement, persistence, and act on their objectives.

9.5.1 Restrict logons to privileged accounts

Privileged and service account credentials are commonly used for lateral movement and persistence. Limit the number of high-privileged accounts and groups to the minimum. The use of shared or generic accounts should be discouraged.

In addition, accounts that have privileged access throughout an environment should not be used on standard workstations and laptops, but rather from designated systems (e.g. privileged access workstations (PAWs)) that reside in restricted and protected VLANs.

9.5.2 Identify and review non-computer accounts configured with a SPN

Accounts with service principal names (SPNs) are commonly targeted by threat actors for privilege escalation. Using Kerberos, any domain user can request a Kerberos service ticket (Ticket Granting Service or TGS) from a domain controller for any account configured with an SPN.

Non-computer accounts are likely configured with guessable (non-random) passwords. Regardless of the domain function level or the host's Windows version, SPNs that are registered under a non-computer account will use the legacy RC4-HMAC encryption suite rather than Advanced Encryption Standard (AES). The key used for encryption and decryption of the RC4-HMAC encryption type represents an unsalted NTLM hash version of the account's password, which could be derived if the ticket is cracked.

9.5.3 Use managed or group managed service accounts

Organisations with static service accounts should review the feasibility of migrating the service accounts to managed service accounts (MSAs) or group managed service accounts (gMSAs).

MSAs were first introduced with the Windows Server 2008 R2 Active Directory schema (domain-functional level). They provide automatic password management (30-day rotation) for dedicated service accounts that are associated with running services on specific endpoints.

It is also crucial that service or unmonitored accounts do not have domain admin rights.

9.5.4 Protect credentials when using RDP

Restricted Admin Mode for RDP can be enabled for all end-user systems assigned to personnel that perform Remote Desktop connections to servers or workstations with administrative credentials. This feature can limit the in-memory exposure of administrative credentials on a destination endpoint when accessed using RDP.

To leverage restricted admin RDP, you can invoke the following command:

```
mstsc[.]exe /RestrictedAdmin
```

9.6 Enable controlled folder access to fend off ransomware attacks

Controlled folder access can help protect data from being encrypted by ransomware. Beginning with Windows 10 version 1709+ and Windows Server 2019+, controlled folder access was introduced within Windows Defender Antivirus (as part of Windows Defender Exploit Guard).

Once controlled folder access is enabled, applications and executable files are assessed by Windows Defender Antivirus, which then determines if an application is malicious or safe. If an application is determined to be malicious or suspicious, it will be blocked from making changes to any files in a protected folder.

Once enabled, controlled folder access will apply to a number of system folders and default locations, including:

- `C:\Users\<username>\Documents`
- `C:\Users\Public\Documents`
- `C:\Users\<username>\Pictures`

- `C:\Users\Public\Pictures`
- `C:\Users\<username>\Videos`
- `C:\Users\Public\Videos`
- `C:\Users\<username>\Music`
- `C:\Users\Public\Music`
- `C:\Users\<username>\Desktop`
- `C:\Users\Public\Desktop`
- `C:\Users\<username>\Favorites`

Note: For controlled folder access to fully function, Windows Defender's Real Time Protection setting must be enabled.

9.7 Mitigations for techniques used by sophisticated threat actors

9.7.1 Whitelist applications

Only permitted software may be executed. To this endeavour, system administrators should implement restrictions on which software is allowed to run on any particular system. This is normally achieved using dedicated solutions such as Microsoft Software Restriction Policies (SRP) or AppLocker.

Safe defaults allow applications to run only from specific directories or when they are digitally signed by trusted authors. In general, an executable should not be allowed to run unless an exception is granted.

9.7.2 Disable active content

While disabling VBA macros in Office documents is an efficient control for blocking specific malware families, it is often difficult to disable macro support entirely due to specific business needs.

Nevertheless, you should assess the use of macros in your environment, and if possible disable or at least limit their availability by implementing restrictions such as blocking macros in documents originating from the internet, allowing only digitally signed macros or making macro support available only to specific users or groups.

Moreover, check if your email gateway or proxy server can sanitise transiting documents by removing active content (e.g. scripts, macros in Office documents and JavaScript in PDFs). If these features are available, you should enable them.

9.7.3 Restrict RDP access

Block Remote Desktop Protocol (RDP) connections originating from untrusted external addresses unless an exception exists. Review these exceptions on a regular basis.

9.7.4 Periodically invalidate krbtgt credentials

A compromised *krbtgt* account may grant attackers privileged persistence in an Active Directory infrastructure. For this reason, it is considered a good practice to periodically invalidate its credentials by performing a password change.

For technical reasons, **the password change needs to be performed twice** to properly clear the password history. Please refer to AD Forest Recovery - Resetting the krbtgt password.

9.7.5 Protect AD FS private keys

As reported by CISA and other organisations, the threat actor behind the SolarWinds Orion campaign obtained PKI keys, certificate files and the private encryption key from an Active Directory Federation Services (AD FS) container to decrypt corresponding SAML signing certificates.

According to MITRE ATT&CK, mitigations for this technique include the following:

1. Ensure only authorised keys are allowed access to critical resources and audit access lists regularly.
2. When possible, store keys on separate cryptographic hardware instead of on the local system.
3. Use strong passphrases for private keys to make cracking difficult.
4. Ensure permissions are properly set on folders containing sensitive private keys to prevent unintended access.

For a more comprehensive guide for securing AD FS, please refer to Microsoft's Best practices for securing Active Directory Federation Services.

If you suspect SAML abuse (e.g. through a Golden SAML attack), apply Microsoft's Advice for incident responders on recovery from systemic identity compromises to remediate this type of activity and uphold administrative control of your environment. Consider also an emergency rotation of your AD FS certificates.

9.7.6 Prevent web shell deployment

Ensure that all your web servers are securely configured. All unnecessary services and ports should be disabled or blocked, and all necessary ones restricted where feasible. In addition, conduct regular system and application vulnerability scans. While this method does not protect against zero-day vulnerabilities, it might in any case highlight possible areas of concern.

Last but not least, deploy a Web Application Firewall (WAF).

9.7.7 Block anonymisation services

Deny all inbound connections from known anonymisation services, such as commercial Virtual Private Networks (VPN) and The Onion Router (Tor), when such access is not required.

9.7.8 Deploy honeypots and honeytokens

Deploy honeypot devices and honeytokens or set up at least one SQL server and a file server as honeypots. After obtaining the initial foothold, attackers might scan the target network looking for content of interest.

9.7.9 Detect post-compromise activity in Microsoft cloud environments

Please refer to CISA alert AA21-008A for tools and procedures to detect post-compromise activity in Azure.

10 Annex B: Examples of techniques, tactics and procedures used by sophisticated threat actors

Preparation

- Collects public or sensitive documents for use as credible lures to enable malware delivery,

- impersonates government officials in spear-phishing campaigns,
- uses a bait website to download an archive containing likely stolen documents that have been weaponised with malware.

Spear-phishing

- Conducts spear-phishing operations, sometimes including zero-day exploits,
- carries out spear-phishing campaigns using the messaging services of social networks (e.g. LinkedIn Messaging),
- conducts spear-phishing operations to deliver weaponised documents,
- conducts high-volume spear-phishing campaigns containing self-extracting (SFX) archives which launch a payload (mimicking official government documents).

Credential harvesting infrastructure

- Creates websites spoofing email account login pages for targeted organisations,
- uses spoofed domains that mimic those of target organisations to collect credentials,
- collects credentials via a spoofed login page on an infrastructure masquerading as legitimate services.

Weaponisation

- Creates documents to deliver malware through the exploitation of client software,
- uses the Dynamic Data Exchange mechanism in Microsoft Office documents instead of macros,
- uses `mshta.exe` to open malicious URLs,
- uses `.docx` files weaponised with obfuscated VBS,
- uses custom .NET loaders for PowerShell Empire.

Direct exploitation

- Exploits the web systems of targets,
- exploits zero-day flaws in web browsers,
- Uses `sqlmap` to probe for vulnerabilities in SQL servers.

Supply chain attacks

- Breaches intermediary parties in order to have a better chance of compromising the real target by using trusted relationships with the initial victim,
- implants malicious code within legitimate infrastructure or tools to be delivered by the legitimate vendor or repository (e.g. SolarWinds Orion),
- compromises Microsoft customer-service tools and then using information from that to breach customers,
- compromises CSPs and uses the privileged access and credentials belonging to these providers to compromise downstream customers,
- infects intermediary parties in order to have a better chance of infecting the real target by using trusted relationships with the initial victim,
- deploys malware via trojanised software executables,
- delivers destructive payloads through supply chain attacks using compromised software update mechanisms.

Strategic web compromise

- Uses strategic web compromise,
- Uses payloads often disguised as software updates.

Compromise of air-gapped networks

- Compromises air-gapped networks by using, for example, Usbstealer, a USB jumping malware.

TLP Definition

TLP	Disclosure	Message
RED	Not for disclosure, restricted to participants only.	Recipients may not share TLP:RED information with any parties outside of the specific exchange, meeting, or conversation in which it was originally disclosed.
AMBER	Limited disclosure, restricted to participants' organisations.	Recipients may share TLP:AMBER information only with members of their own organisation.
GREEN	Limited disclosure, restricted to the community.	Subject to standard copyright rules, TLP:GREEN information may be distributed with peers and partner organisations within their sector or community, but not via publicly accessible channels.
WHITE	Disclosure is not limited.	TLP:WHITE information may be distributed freely.