GUIDANCE

Multi-factor authentication for online services

Advice for organisations on implementing multi-factor authentication (or two-factor authentication) to protect against password guessing and theft on online services.

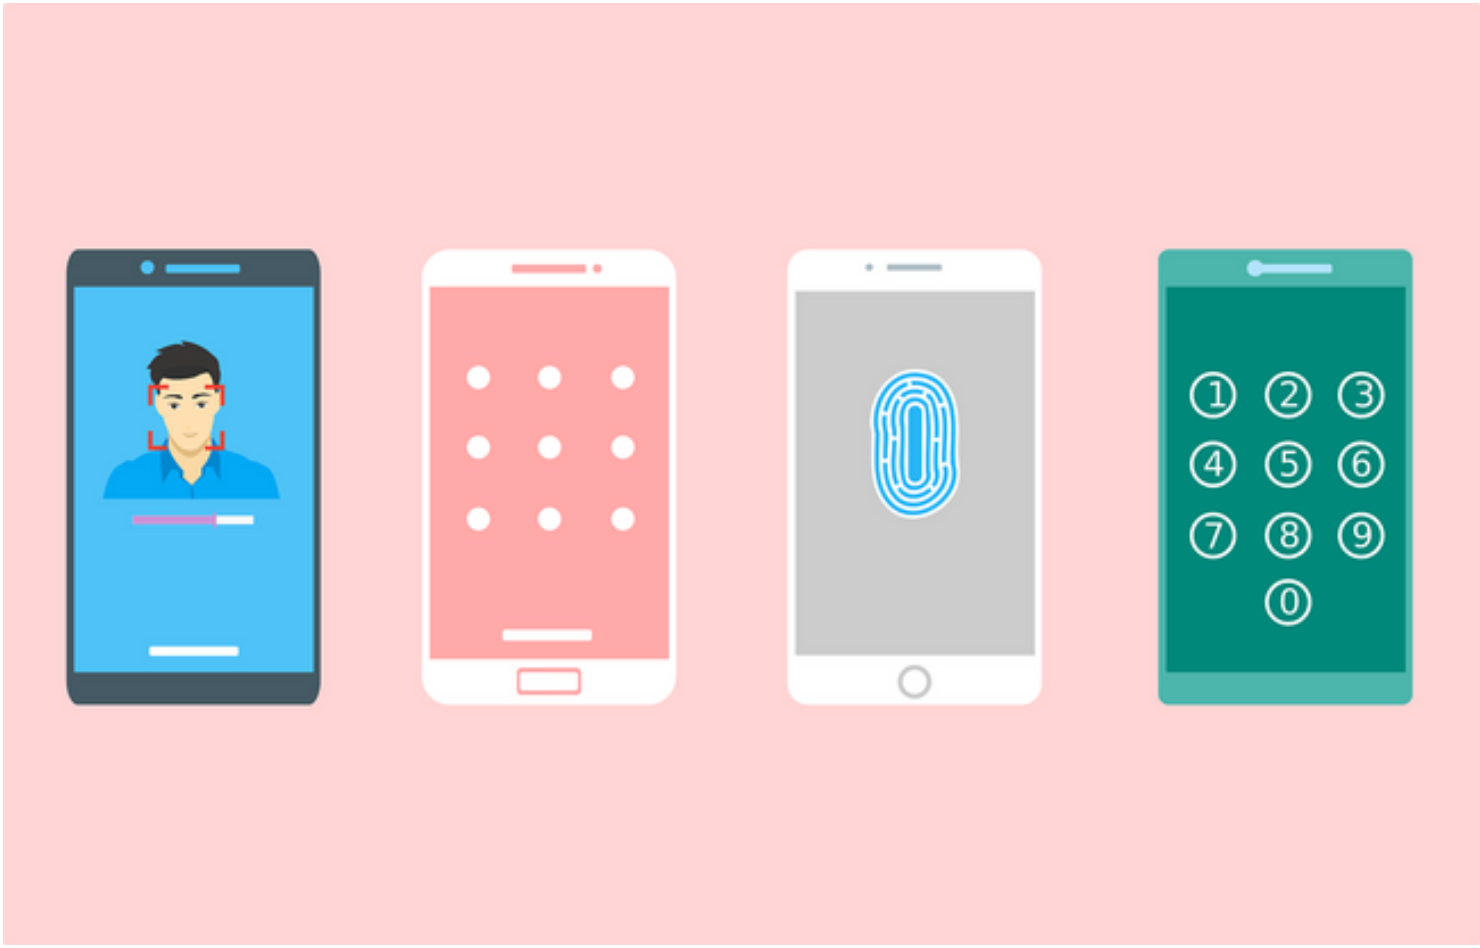

This guidance describes how to use multi-factor authentication (MFA) to mitigate against password guessing and theft, including brute force attacks. MFA can also be called two-step verification or 2-factor authentication (2FA).

- This guidance is primarily for senior decision makers in larger organisations, and administrators responsible for configuring access to online and enterprise services that requires users to authenticate to the service.
- Information about implementing 2FA on personal accounts is covered in separate guidance.
- This guidance does not cover authentication to devices, which can be found in the End User Devices Guidance.

Introduction

Enterprise services have historically used a password to authenticate a user to services hosted on the enterprise's internal network. Enterprises are now using more services that are directly connected to the Internet, to enable collaboration, to make remote working easier, or to benefit from shared services. The services themselves are commonly hosted in the cloud, or may allow remote connectivity to on-premise services that would previously have only been accessible from an organisation's internal network. In both cases, authentication becomes the main decision point for whether a user or attacker can gain access to a service.

In recent attacks, adversaries have logged in to services using guessed or stolen passwords, thus granting that attacker the same permissions (sometimes privileged access) as the legitimate user. This is prevalent when there is no way to differentiate between the actual user and an attacker pretending to be them. These attacks work in a number of ways including:

- Trying passwords from leaked datasets in case a user has reused a password on other services.
- Socially engineering account details from users using techniques such as phishing.
- Attempting to log into multiple user accounts using one of the passwords on the commonly used password lists. This is known as password spraying and works on the assumption that some users will have chosen passwords on those lists. Even if an account automatically locks after a number of failed

attempts, by trying multiple user accounts, statistically an attacker will still be able to break into a few. It has the advantage of being less obvious to monitoring than a brute force attack ,where multiple passwords are tried against a single account.

The NCSC password guidance describes these issues in more detail and gives advice for users and administrators needing to cope with an ever-increasing number of passwords.

When to use an extra factor

As long as passwords are used for authentication, there will always be a chance that users and administrators will choose machine-guessable passwords and be susceptible to social engineering. Therefore:

- Organisations should choose Cloud and Internet-connected services that offer a form of multi-factor authentication.
- All users, including administrators, should use multi-factor authentication when using Cloud and Internet-connected services. This is particularly important when authenticating to services that hold sensitive or private data.
- Administrators should, wherever possible, be required to use multi-factor authentication.
- Organisations should consider carefully the use of services which only allow for single-factor authentication.

The service will need to check the extra factor whenever there is a need to authenticate a user. The exact implementation will vary per-service. Common implementations include:

- The user needs an extra authentication factor when they are logging on to a service using a device that they have not used before. The service tracks devices previously used and so will not need the extra factor every time. It may be necessary to opt in to the service remembering the device by selecting a 'remember my device' option.

- The user needs an extra authentication factor every time they log on to a service. This is more usual for services that have a higher impact if they're compromised, such as an email account or online banking.
- The user needs to re-verify themselves using an extra factor when performing high risk actions - such as changing a password or transferring money.
- The user is prompted to use their extra authentication factor when the authentication has been determined to be high risk, such as the connection coming from a different part of the world than is normal for that user.

Choosing extra authentication factors

Your choice of extra factor will vary depending on the service you are using, because the service might not support your first choice. Enterprises should consider a single sign-on (SSO) solution to give the best user experience. A multi-factor token will then usually only be needed to register a new device. On personal devices, or in bring your own device (BYOD) scenarios where SSO is unavailable, you should consider using the techniques that have the best user experience.

Using a managed/enterprise device as an extra factor

These techniques work by limiting access to online services to devices that are managed by your organisation (or another organisation that you trust, so long as you can verify the device is theirs). The device being used to access the online service becomes the extra factor. Online services will reject any authentication attempts that haven't come from a managed device hence preventing an external attacker from authenticating against the service from elsewhere. Common methods to implement this include:

- The managed device passes a trusted claim signed by the enterprise's identity service. This ensures that only devices that are enrolled in and managed by the enterprise can authenticate against the cloud service. This claim can be combined with a user identity to also allow seamless single sign-on.

- Managed devices may already have a hardware backed credential such as an enterprise certificate that cannot be removed from the device. The use of such a credential ensures that only devices that have been configured by the enterprise can authenticate against the cloud service. The use of a user-identifying credential can also allow seamless single sign-on.
- Organisations can configure cloud services to only accept authentication attempts from within their trusted enterprise networks. This ensures that users can only authenticate if they are either directly connected to that trusted network or have remote access to it over a VPN.

Using an app on a trusted device as an extra factor

These techniques use the knowledge that a user possesses a specific device to prove they are who they say they are. To be effective, the user will need to unlock the device or app to prove that it is theirs. The extra factor will be better protected on devices that are well-configured and regularly updated - such as described in the NCSC EUD Security Guidance - and have hardware-backed credential storage. It is best if the app is used on a corporately managed and configured device as it will likely be better defended against malware. However there is still value in implementing multi-factor authentication using an app installed on a personal device if users do not have corporately managed one.

- An authenticator app generates a single-use password that changes every minute or so. This may use a generic authenticator app that works with many services or one that is bespoke to the service. The user types the code in to the service what they are authenticating to.
- An authenticator app or a feature built into the device receives a push notification that prompts the user to confirm or deny that they are currently trying to log in to a named service.

Using a physically separate extra factor

These techniques use the knowledge that a user has a physical security token, which proves they are who they say they are. Some types will require the user to unlock them before use, others just require proof of possession. Security tokens which require the user to carry an additional item with them do not offer an ideal user experience. Examples of physical security tokens include:

- FIDO universal 2nd factor authenticators such as YubiKey. Devices contain a cryptographic key which you can use to authenticate to a service on your laptop or phone via USB, Bluetooth or NFC. A single device can often be used as an extra factor for many services.
- Smartcards have long been used to authenticate against systems that require proof of possession of an item, which is unlocked by a PIN code. Organisations that already use smartcards can use them as a second factor to authenticate to online services.
- Bespoke devices such as RSA tokens and chip-and-PIN card readers generate a single-use code each time a user logs in. These are effective although generally only work on a single service. They should therefore only be considered for specialist use to avoid the user needing to carry too many of these devices.
- Single-use authenticator codes aka backup codes are designed to be used when the user loses access to their usual second factor. They will only work once and will need to be protected prior to use – we suggest encrypting them using a password manager, or printing them and keeping them in a safe place.

Using a known or trusted account as an extra factor

These techniques send codes to a registered email address or phone number. The email authentication technique is only effective if the user has a separate password for the registered email account and (ideally) also implements multi-factor authentication. The extra factor will be better protected if the SMS or email is only accessed from devices that are well-configured and regularly updated - such as described in the NCSC EUD Security Guidance.

- The service emails a single-use code to an address registered for that user. If the user can see the email then it proves that they own that email address. Services that send a code for the user to type in are preferred over ones that send a clickable link, as it is difficult for a user to distinguish between a legitimate email and a phishing email.
- The service sends an SMS message containing a single-use code or makes a voice call in which a single-use code is read out to the phone number registered for that user. When receiving a phone call the user should be required to interact with the call to ensure that the code cannot be recorded on an unprotected voicemail service. Services often give the user a choice of whether they prefer a message or a call. SMS is not the most secure type of MFA, but still offers a huge advantage over not using any MFA. Any multi-factor authentication is better than not having it at all. However, if there are alternatives available that will work for your use case, we recommend you use these instead of SMS.

Using another piece of knowledge as an extra factor

These techniques ask the user for an extra piece of information as a second factor. These schemes are usually designed so that an attacker will need to monitor several successful authentication attempts to reveal the full second factor. However they are susceptible to the same types of attacks as passwords (such as phishing or guessing common answers), and if a fixed second factor is used infrequently (such as when logging into a new device) a

user is likely to forget their answer and will need to reset. For these reasons we do not recommend them as effective extra factors. The following examples demonstrate why:

- The user is asked to confirm one or more answers that they have already provided in response to a set of security questions. These are usually designed to be easy to remember such as place of birth or favourite holiday location. This leads to a second factor that can be easily guessed so this method should not be considered to be an extra factor and is not recommended.
- The user is asked to answer a series questions whose answers are derived from data associated with the user. A financial institute could ask questions based on recent transactions or a social network could ask for context around friendships or photographs. The answers should not be guessable or derivable by others but will be something that the user knows simply by knowing their own behaviour. This balance can be difficult to achieve and requires the service to hold sufficient data on the user to generate a sufficiently random question that can't be easily guessed by an attacker. Additionally, the question itself should not reveal anything sensitive to an attacker.
- The user is asked to provide specific characters from a memorable password or numeric code. This technique partially mitigates against credential theft through phishing attacks however is as susceptible to guessing as normal passwords. It is also easy for a user to forget and will not usually be as well protected by the services when compared with passwords as they cannot be usefully protected by hashing. Therefore, this method is not recommended as an effective extra factor.

Additional considerations

The introduction of multi-factor authentication to online services may require your IT helpdesk to offer extra services to support users. If users lose their extra

factor, they will need a way of reporting and replacing it. This could be offered directly by the service or via an enterprise portal. You will need to consider how your account reset and multi-factor token replacement processes verify that the user is who they say they are. You will need to ensure that an attacker cannot use these processes to bypass multi-factor authentication.

You will need to consider how administrators can gain access to the service if multi-factor authentication is unavailable. This could be caused by a service configuration or the loss of an authentication token. Accounts such as an emergency or 'break glass' account that use a single authentication factor should be the subject of increased protective monitoring so that its misuse can be easily detected.

Defence in depth

Authentication requests should report successful and unsuccessful authentication requests to enterprise audit and monitoring systems. This allows the monitoring system to highlight unusual activity and contribute to a log of malicious activity post-breach. Services that email the user when they log in are an effective way of detecting unauthorised authentications, but only if the user has a route to report unexpected emails.

Some online services adjust authentication requirements based on things such as the physical location associated with the connecting IP address, whether a connection comes from a higher risk network such as Tor or deviations from a user's normal usage patterns. A service may choose to force multi-factor authentication when the risk is deemed too high, or block authentication entirely. These approaches can reduce the success rate of brute force attacks however are most effective when used alongside multi-factor authentication.

PUBLISHED

14 June 2018

REVIEWED

3 December 2018

VERSION

1.0

WRITTEN FOR

Cyber security professionals