Introduction to identity and access management

This guidance provides a primer on the essential techniques, technologies and uses of access management.

Introduction

If identity and access management procedures and controls are badly designed or implemented, they can give attackers an easy way to gain access to your systems which could appear legitimate. You should consider the security of all the aspects of identity:

- ensuring that a new user is who they say they are, and the level of trust and access you give them is commensurate with their personal and professional background
- binding an identified user to an identity within your system with an appropriate method of authentication
- ensuring that the authentication method gives you confidence that when an identity is used, it is being used by the member of staff whose identity you have previously validated
- applying the principle of least privilege to limit the access or functionality that different users have

As well as this, it is important that your access management systems are robustly designed and administered.

About this guidance

This guidance provides a primer on the essential techniques, technologies and uses of identity and access management. It's aimed primarily at technical staff.

This guidance should:

- help you understand basic principles to follow when designing user access management for sensitive systems or roles
- help you understand basic architectural good practice when designing and administering access management systems
- give you further reading for some topics, if you want to find out more

---

Overview

Identity and access management refers to the collection of policies, processes and systems which support binding an individual (or in some cases a system) to a set permissions within your system.

These permissions may allow the individual to:

- perform functions (such as altering and industrial control process)
- access data (such as staff records)
- administer your system

An access management system is comprised of a number of technical components, including; directory services, authentication components and the parts of your system that consume authentication and authorisation information.

Identity and access management can be broken down into the following broad areas:

- **Policy -** your strategy governing who is authorised to access systems, data or functionality, how they can request access, when their access should be revoked and whether any particular operations should require multiple users to collaborate
- **Identity management -** how you establish the identity of a person, both at point of first contact and subsequent interactions with your systems or processes

- **Privileged user management -** the additional processes and controls you should put in place to safeguard the most sensitive operations in the system
- **Architectural design -** secure design of the computer systems that support the above areas
- **Operations and monitoring -** the supporting processes and technology to identify and enable investigation of breaches of policy or controls

---

Policy

An identity and access management policy would typically cover:

- who should have access to certain systems, data or functionality, and why
- the circumstances under which they would be granted or revoked access - typically managed with a joiners, leavers and movers process
- analysis of which actions or processes, if any, should require multiple people to perform them
- which actions should be recorded, how audit records are acquired and how they are safeguarded against tampering

Various international security standards contain detailed identity and access management policies which you can follow and be assessed against. For example, ISO27002 (specifically section 9 - Access Management) or if operating an industrial control system IEC 62443-2-1:2011, sections 4.3.3.5 - 4.3.3.7, are relevant.

In addition to the above mentioned standards, the following guidance may be useful:

- CPNI's Physical security guidance

Identity management

It is important that you appropriately identify anyone who could have access to your systems at the point of first contact, establishing their true identity and a method of future authentication. When establishing a person's identity, consider:

- the level of access they will have to your systems - the more sensitive or privileged their access, the stronger the identity verification you should perform
- the reliance you are placing on the assertions of any third parties - for example, if equipment maintenance is performed by an external contractor, you should gain sufficient confidence in the contractor's identity proofing procedures
- whether certain roles should require more stringent background checks or security clearances

After establishing an initial identity, you must bind a user to this identity using a method of authentication. The strength of the authentication you choose to use should depend on the sensitivity or privilege of a user's access.

When designing an authentication approach, it is important to consider both the physical environment which the user will perform their authentication from, and the trustworthiness of the devices they will use to perform that authentication. If the users can authenticate from an untrusted location (e.g. a cafe) or from a less trusted device (their home PC), this should factor in your authentication requirements.

Different methods of authentication have different strengths and limitations. Below are some of the basic properties of some common authentication methods:

- **passwords** – the most basic form of credentials, used to authenticate a user based on *something they know* - however if the password becomes known or is guessed by an attacker, then they can impersonate the legitimate user
- **two-factor authentication** – the "second factor" in this context should add *something you have* to the *something you know* - the second factor could be a physical token or your phone (such as through an authentication app or an SMS message). Adding a second factor considerably increases

the difficulty for an attacker to subvert authentication checks, though some actors have shown they are capable of subverting some two-factor techniques, such as SMS messages

- **hardware-backed certificate authentication** – a cryptographic key and certificate stored in a hardware chip on a user's device (such as a Trusted Platform Module) or on a user's smart card - they can provide high confidence in the identity of a user or device if supported by other physical and technical controls. However, to gain this confidence, the end user device upon which the token is utilised must be trusted
- **biometric** – the use of an aspect of the user's physiology to identify them, such as a finger print or facial recognition.These can be useful for reducing the password burden on users. Whilst they cannot be "shoulder surfed" by a nearby attacker, some forms are easier to spoof than others

We recommend you review and implement where appropriate:

- CPNI's guidance on Personnel and people security and Pre-employment screening
- the NCSC's Password guidance

Privileged user management

Anyone with access that enables them to affect change which would be felt beyond their immediate job role could be considered a privileged user. For example:

- administering systems or networks which are critical to a business (e.g. database admins, system admins)
- accessing systems used to perform a critical function (e.g. industrial control systems, approving financial payments)
- developers who are able to commit changes to code repositories upon which your business relies (e.g. software whose sales provide a material line

of profit, code which could affect control of an essential service, or software used by customers)

The most important aspect of privileged user management is ensuring that actions performed by privileged user accounts are indeed the actions of your privileged users, as opposed to a malicious third party. As well as strongly authenticating privileged users, there are other actions you can take to reduce the risk of privileged user accounts being misused:

- **issue separate user accounts and credentials to users who have a need to perform both privileged and typical day-to-day functions -** privileged accounts should not be used for reading email or browsing the web unless protections have been deployed to ensure such actions are performed in a less privileged user context
- **avoid users perfoming privileged actions from untrusted devices -** unless you are confident in how an end user device has been secured, it is prudent to assume it has been compromised. The risk of allowing an untrusted device to connect to your systems increases with the level of functionality or information exposed to those devices, so whilst such access may be appropriate for some operations it may not be acceptable for all
- **when working across network boundaries or zones, prefer to "browse down" from the more trusted environment to the less trusted environment rather than "browse up" -** *browse up* occurs when a privileged function (such as operating an industrial control system or gaining console access to a production database server) is performed from a less trusted system or network. Given that we should assume the less trusted system is *compromised*, the integrity of the privileged action is undermined. To maintain the integrity of the more trusted environment a *browse down* approach should be used to perform functions in less trusted environments from within the more trusted one

Some of the above techniques can have a detrimental effect on usability, and may be unpopular with privileged users. However, they can make a significant difference to the security of computer systems that employ them, and with thoughtful design, the impact on usability can be reduced.

As well as authenticating privileged users, you should ensure that their actions are appropriately controlled. This could be achieved by:

- Requiring independent confirmation and approval of actions. For example, code updates need to be reviewed by a team leader before acceptance. Or, when an administrator accesses certain critical servers, an alert is raised that is followed up for confirmation of actions.
- Requiring support tickets to be issued before certain actions can be undertaken. For example, administrators are only able to log onto servers if another team has raised a ticket requesting a fix against a server.
- Monitoring privileged user actions. This is not just logging, but defining rules that could detect suspicious activity and actively reviewing events. For example, highlight when certain commands are run by administrators, or when code changes are made at odd times, or in an unusual quantity.

Consider reading the following guidance to learn more:

- Microsoft's recommend Tiered Administration Model is a good example of the browse down model described above
- NCSC's blog on Protecting your management interfaces provides a more detailed rationale for some of the recommendations above
- NCSC's Design Principles may be useful when considering privileged access
- NCSC's Secure developement and deployment guidance provides information on how to reduce risks around developing and deploying code

---

Architecture design

When building an identity and access management system, it's important to remember that the identity system itself will be of interest to attackers, since if subverted it can provide access to your systems.

We recommend that:

- If presenting services to untrusted networks, such as the Internet, you should isolate the external-facing access management components from the rest of your systems
- If using Microsoft's Active Directory for identity and access management, we recommend the use of Microsoft's Administrative Tier Model
- If using a Single Sign On (SSO) or federated access management approach, especially over untrusted networks, validate that the identity assertion you receive has come from a trusted source. For example, in the scenario where your system accepts a SAML identity assertion from a SSO service, check the SAML assertion is signed by the Certificate Authority you expect

---

Operational technology (OT)

In organisations that manage Operational Technology (OT) - such as industrial control systems in the energy sector - we recommend that:

- OT is not administered from an enterprise IT estate. As discussed earlier, this is considered "browse up," and means that the integrity of your OT systems would be only as good as that of your IT estate.

- OT systems should not rely solely on systems in a lower trust domain for authentication and authorisation. For example, if access to OT systems was controlled by a corporate user directory service within your enterprise IT domain, a compromise of your IT estate would result in a compromise of your OT estate. The higher number of users and higher risk activities usually performed on IT systems typically means they should be assumed to be compromised when considering the security of OT.

- When data needs to be transferred between IT and OT systems, it is advisable to 'push' data from the OT estate to the IT estate, as opposed to

allowing systems in the IT estate to reach into OT systems. For example, push from data historians to enterprise IT repositories rather than enabling two-way communications

- Where system to system communications are needed across IT to OT boundaries, ensure they are in an inspectable format, and are monitored or (preferably) validated at the boundary. For example, if a system on the IT estate needs to make an HTTP GET request to a system on the OT estate, a web application firewall could be deployed at the boundary to only allow access to the specific API endpoint and provide validation of any query parameters

---

Operations and monitoring

Given the high value to an attacker of compromising your identity and access management systems they should be given priority for security maintenance. This means, among other things, prompt application of security patches across your estate (or otherwise mitigating security issues), practicing good user and privileged user management, and applying appropriate protective monitoring.

Additionally, we recommend:

- designing your access control systems to allow for easy monitoring of account usage and accesses
- being able to tie all user actions in the system to the user that performed them (e.g. in a web service this might mean ensuring that access tokens are linked to all API calls performed)

The following guides provide further information on these topics:

- Vulnerability management
- Protective monitoring

- Operational technologies

PUBLISHED

22 January 2018

REVIEWED

22 January 2018

VERSION

1.0

WRITTEN FOR

Large organisations

Cyber security professionals

Public sector