

Information Commissioners Guidance

Principle (a) Lawfulness, Fairness and Transparency

At a glance:

- You must identify valid grounds under the GDPR (known as a 'lawful basis') for collecting and using personal data.
- You must ensure that you do not do anything with the data in breach of any other laws.
- You must use personal data in a way that is fair. This means you must not process the data in a way that is unduly detrimental, unexpected or misleading to the individuals concerned.
- You must be clear, open and honest with people from the start about how you will use their personal data.

Checklist:

Lawfulness

- We have identified an appropriate lawful basis (or bases) for our processing.
- If we are processing special category data or criminal offence data, we have identified a condition for processing this type of data.
- We don't do anything generally unlawful with personal data.

Fairness

- We have considered how the processing may affect the individuals concerned and can justify any adverse impact.
- We only handle people's data in ways they would reasonably expect, or we can explain why any unexpected processing is justified.
- We do not deceive or mislead people when we collect their personal data.

Transparency

- We are open and honest, and comply with the transparency obligations of the right to be informed.

Principle (b) Purpose limitation

At a glance:

- You must be clear about what your purposes for processing are from the start.
- You need to record your purposes as part of your documentation obligations and specify them in your privacy information for individuals.
- You can only use the personal data for a new purpose if either this is compatible with your original purpose, you get consent, or you have a clear obligation or function set out in law.

Checklist:

- We have clearly identified our purpose or purposes for processing.
- We have documented those purposes.
- We include details of our purposes in our privacy information for individuals.
- We regularly review our processing and, where necessary, update our documentation and our privacy information for individuals.
- If we plan to use personal data for a new purpose other than a legal obligation or function set out in law, we check that this is compatible with our original purpose or we get specific consent for the new purpose.

In brief:

- What is the purpose limitation principle?
- Why do we need to specify our purposes?
- How do we specify our purposes?
- Once we collect data for a specified purpose, can we use it for other purposes?
- What is a 'compatible' purpose?

What is the purpose limitation principle?

Article 5(1)(b) says:

“1. Personal data shall be:

(b) collected for specified, explicit and legitimate purposes and not further processed in a manner that is incompatible with those purposes; further processing for archiving purposes in the public interest, scientific or historical research purposes or statistical purposes shall, in accordance with Article 89(1), not be considered to be incompatible with the initial purposes.”

In practice, this means that you must:

- be clear from the outset why you are collecting personal data and what you intend to do with it;
- comply with your documentation obligations to specify your purposes;
- comply with your transparency obligations to inform individuals about your purposes; and
- ensure that if you plan to use or disclose personal data for any purpose that is additional to or different from the originally specified purpose, the new use is fair, lawful and transparent.

Why do we need to specify our purposes?

This requirement aims to ensure that you are clear and open about your reasons for obtaining personal data, and that what you do with the data is in line with the reasonable expectations of the individuals concerned.

Specifying your purposes from the outset helps you to be accountable for your processing, and helps you avoid ‘function creep’. It also helps individuals understand how you use their data, make decisions about whether they are happy to share their details, and assert their rights over data where appropriate. It is fundamental to building public trust in how you use personal data.

There are clear links with other principles – in particular, the fairness, lawfulness and transparency principle. Being clear about why you are processing personal data will help you to ensure your processing is fair, lawful and transparent. And if you use data for unfair, unlawful or ‘invisible’ reasons, it’s likely to be a breach of both principles.

Specifying your purposes is necessary to comply with your accountability obligations.

How do we specify our purposes?

If you comply with your documentation and transparency obligations, you are likely to comply with the requirement to specify your purposes without doing anything more:

- You need to specify your purpose or purposes for processing personal data within the documentation you are required to keep as part of your records of processing (documentation) obligations under Article 30.
- You also need to specify your purposes in your privacy information for individuals.

However, you should also remember that whatever you document, and whatever you tell people, this cannot make fundamentally unfair processing fair and lawful.

If you are a small organisation and you are exempt from some documentation requirements, you may not need to formally document all of your purposes to comply with the purpose limitation principle. Listing your purposes in the privacy information you provide to individuals will be enough. However, it is still good practice to document all of your purposes. For more information, read our documentation guidance.

If you have not provided privacy information because you are only using personal data for an obvious purpose that individuals already know about, the “specified purpose” should be taken to be the obvious purpose.

You should regularly review your processing, documentation and privacy information to check that your purposes have not evolved over time beyond those you originally specified (‘function creep’).

Once we collect personal data for a specified purpose, can we use it for other purposes?

The GDPR does not ban this altogether, but there are restrictions. In essence, if your purposes change over time or you want to use data for a new purpose which you did not originally anticipate, you can only go ahead if:

- the new purpose is compatible with the original purpose;
- you get the individual’s specific consent for the new purpose; or

- you can point to a clear legal provision requiring or allowing the new processing in the public interest – for example, a new function for a public authority.

If your new purpose is compatible, your processing will be lawful and you don't need a new lawful basis for the further processing. However, you should remember that if you originally collected the data on the basis of consent, you usually need to get fresh consent to ensure your new processing is fair and lawful. See our lawful basis guidance for more information.

You also need to make sure that you update your privacy information to ensure that your processing is still transparent.

What is a 'compatible' purpose?

The GDPR specifically says that the following purposes should be considered to be compatible purposes:

- archiving purposes in the public interest;
- scientific or historical research purposes; and
- statistical purposes.

Otherwise, the GDPR says that to decide whether a new purpose is compatible (or as the GDPR says, "not incompatible") with your original purpose you should take into account:

- any link between your original purpose and the new purpose;
- the context in which you originally collected the personal data – in particular, your relationship with the individual and what they would reasonably expect;
- the nature of the personal data – eg is it particularly sensitive;
- the possible consequences for individuals of the new processing; and
- whether there are appropriate safeguards - eg encryption or pseudonymisation.

As a general rule, if the new purpose is either very different from the original purpose, would be unexpected, or would have an unjustified impact on the individual, it is likely to be incompatible with your original purpose. In practice, you are likely to need to ask for specific consent to use or disclose data for this type of purpose.

Example

A GP discloses his patient list to his wife, who runs a travel agency, so that she can offer special holiday deals to patients needing recuperation. Disclosing the information for this purpose would be incompatible with the purposes for which it was obtained.

There are clear links here with the lawfulness, fairness and transparency principle. In practice, if your intended processing is fair, you are unlikely to breach the purpose limitation principle on the basis of incompatibility.

Principle (c) Data minimisation

At a glance:

You must ensure the personal data you are processing is:

- adequate – sufficient to properly fulfil your stated purpose;
- relevant – has a rational link to that purpose; and
- limited to what is necessary – you do not hold more than you need for that purpose.

Checklist:

- We only collect personal data we actually need for our specified purposes.
- We have sufficient personal data to properly fulfil those purposes.
- We periodically review the data we hold, and delete anything we don't need.

In brief:

- What is the data minimisation principle?
- How do we decide what is adequate, relevant and limited?
- When could we be processing too much personal data?
- When could we be processing inadequate personal data?
- What about the adequacy and relevance of opinions?

What is the data minimisation principle?

Article 5(1)(c) says:

“1. Personal data shall be:

(c) adequate, relevant and limited to what is necessary in relation to the purposes for which they are processed (data minimisation)”

So you should identify the minimum amount of personal data you need to fulfil your purpose. You should hold that much information, but no more.

This is the first of three principles about data standards, along with accuracy and storage limitation.

The accountability principle means that you need to be able to demonstrate that you have appropriate processes to ensure that you only collect and hold the personal data you need.

Also bear in mind that the GDPR says individuals have the right to complete any incomplete data which is inadequate for your purpose, under the right to rectification. They also have right to get you to delete any data that is not necessary for your purpose, under the right to erasure (right to be forgotten).

How do we decide what is adequate, relevant and limited?

The GDPR does not define these terms. Clearly, though, this will depend on your specified purpose for collecting and using the personal data. It may also differ from one individual to another.

So, to assess whether you are holding the right amount of personal data, you must first be clear about why you need it.

For special category data or criminal offence data, it is particularly important to make sure you collect and retain only the minimum amount of information.

You may need to consider this separately for each individual, or for each group of individuals sharing relevant characteristics. You should in particular consider any specific factors that an individual brings to your attention – for example, as part of an objection, request for rectification of incomplete data, or request for erasure of unnecessary data.

You should periodically review your processing to check that the personal data you hold is still relevant and adequate for your purposes, and delete anything you no longer need. This is closely linked with the storage limitation principle.

When could we be processing too much personal data?

You should not have more personal data than you need to achieve your purpose. Nor should the data include irrelevant details.

Example

A debt collection agency is engaged to find a particular debtor. It collects information on several people with a similar name to the debtor. During the enquiry some of these people are discounted. The agency should delete most of their personal data, keeping only the minimum data needed to form a basic record of a person they have removed from their search. It is appropriate to keep this small amount of information so that these people are not contacted again about debts which do not belong to them.

If you need to process particular information about certain individuals only, you should collect it just for those individuals – the information is likely to be excessive and irrelevant in relation to other people.

Example

A recruitment agency places workers in a variety of jobs. It sends applicants a general questionnaire, which includes specific questions about health conditions that are only relevant to particular manual occupations. It would be irrelevant and excessive to obtain such information from an individual who was applying for an office job.

You must not collect personal data on the off-chance that it might be useful in the future. However, you may be able to hold information for a foreseeable event that may never occur if you can justify it.

Example

An employer holds details of the blood groups of some of its employees. These employees do hazardous work and the information is needed in case of accident. The employer has in place safety procedures to help prevent accidents so it may be that this data is never needed, but it still needs to hold this information in case of emergency.

If the employer holds the blood groups of the rest of the workforce, though, such information is likely to be irrelevant and excessive as they do not engage in the same hazardous work.

If you are holding more data than is actually necessary for your purpose, this is likely to be unlawful (as most of the lawful bases have a necessity element) as well as a breach of the data minimisation principle. Individuals will also have the right to erasure.

When could we be processing inadequate personal data?

If the processing you carry out is not helping you to achieve your purpose then the personal data you have is probably inadequate. You should not process personal data if it is insufficient for its intended purpose.

In some circumstances you may need to collect more personal data than you had originally anticipated using, so that you have enough information for the purpose in question.

Example

A group of individuals set up a club. At the outset the club has only a handful of members, who all know each other, and the club's activities are administered using only basic information about the members' names and email addresses. The club proves to be very popular and its membership grows rapidly. It becomes necessary to collect additional information about members so that the club can identify them properly, and so that it can keep track of their membership status, subscription payments etc.

Data may also be inadequate if you are making decisions about someone based on an incomplete understanding of the facts. In particular, if an individual asks you to supplement incomplete data under their right to rectification, this could indicate that the data might be inadequate for your purpose.

Obviously it makes no business sense to have inadequate personal data – but you must be careful not to go too far the other way and collect more than you need.

What about the adequacy and relevance of opinions?

A record of an opinion is not necessarily inadequate or irrelevant personal data just because the individual disagrees with it or thinks it has not taken account of information they think is important.

However, in order to be adequate, your records should make clear that it is opinion rather than fact. The record of the opinion (or of the context it is held in) should also contain enough information to enable a reader to interpret it correctly. For example, it should state the date and the author's name and position.

If an opinion is likely to be controversial or very sensitive, or if it will have a significant impact when used or disclosed, it is even more important to state the circumstances or the evidence it is based on. If a record contains an opinion that summarises more detailed records held elsewhere, you should make this clear.

Example

A GP's record may hold only a letter from a consultant and it will be the hospital file that contains greater detail. In this case, the record of the consultant's opinion should contain enough information to enable detailed records to be traced.

Principle (d) Accuracy

At a glance:

- You should take all reasonable steps to ensure the personal data you hold is not incorrect or misleading as to any matter of fact.
- You may need to keep the personal data updated, although this will depend on what you are using it for.
- If you discover that personal data is incorrect or misleading, you must take reasonable steps to correct or erase it as soon as possible.
- You must carefully consider any challenges to the accuracy of personal data.

Checklist

- We ensure the accuracy of any personal data we create.
- We have appropriate processes in place to check the accuracy of the data we collect, and we record the source of that data.

- We have a process in place to identify when we need to keep the data updated to properly fulfil our purpose, and we update it as necessary.
- If we need to keep a record of a mistake, we clearly identify it as a mistake.
- Our records clearly identify any matters of opinion, and where appropriate whose opinion it is and any relevant changes to the underlying facts.
- We comply with the individual's right to rectification and carefully consider any challenges to the accuracy of the personal data.
- As a matter of good practice, we keep a note of any challenges to the accuracy of the personal data.

In brief:

- What is the accuracy principle?
- When is personal data 'accurate' or 'inaccurate'?
- What about records of mistakes?
- What about accuracy of opinions?
- Does personal data always have to be up to date?
- What steps do we need to take to ensure accuracy?
- What should we do if an individual challenges the accuracy of their personal data?

What is the accuracy principle?

Article 5(1)(d) says:

"1. Personal data shall be:

(d) accurate and, where necessary, kept up to date; every reasonable step must be taken to ensure that personal data that are inaccurate, having regard to the purposes for which they are processed, are erased or rectified without delay ('accuracy')"

This is the second of three principles about data standards, along with data minimisation and storage limitation.

There are clear links here to the right to rectification, which gives individuals the right to have inaccurate personal data corrected.

In practice, this means that you must:

- take reasonable steps to ensure the accuracy of any personal data;
- ensure that the source and status of personal data is clear;
- carefully consider any challenges to the accuracy of information; and
- consider whether it is necessary to periodically update the information.

When is personal data ‘accurate’ or ‘inaccurate’?

The GDPR does not define the word ‘accurate’. However, the Data Protection Act 2018 does say that ‘inaccurate’ means “incorrect or misleading as to any matter of fact”. It will usually be obvious whether personal data is accurate.

You must always be clear about what you intend the record of the personal data to show. What you use it for may affect whether it is accurate or not. For example, just because personal data has changed doesn’t mean that a historical record is inaccurate – but you must be clear that it is a historical record.

Example

If an individual moves house from London to Manchester a record saying that they currently live in London will obviously be inaccurate. However a record saying that the individual once lived in London remains accurate, even though they no longer live there.

Example

The Postcode Address File (PAF) contains UK property postal addresses. It is structured to reflect the way the Royal Mail delivers post. So it is common for someone to have a postal address linked to a town in one county (eg Stoke-on-Trent in Staffordshire) even if they actually live in another county (eg Cheshire) and pay council tax to that council. The PAF file is not intended to accurately reflect county boundaries.

What about records of mistakes?

There is often confusion about whether it is appropriate to keep records of things that happened which should not have happened. Individuals understandably do not want their

records to be tarnished by, for example, a penalty or other charge that was later cancelled or refunded.

However, you may legitimately need your records to accurately reflect the order of events – in this example, that a charge was imposed, but later cancelled or refunded. Keeping a record of the mistake and its correction might also be in the individual's best interests.

Example

A misdiagnosis of a medical condition continues to be held as part of a patient's medical records even after the diagnosis is corrected, because it is relevant for the purpose of explaining treatment given to the patient, or for other health problems.

It is acceptable to keep records of mistakes, provided those records are not misleading about the facts. You may need to add a note to make clear that a mistake was made.

Example

An individual finds that, because of an error, their account with their existing energy supplier has been closed and an account opened with a new supplier. Understandably aggrieved, they believe the original account should be reinstated and no record kept of the unauthorised transfer. Although this reaction is understandable, if their existing supplier did close their account, and another supplier opened a new account, then records reflecting what actually happened will be accurate. In such cases it makes sense to ensure that the record clearly shows that an error occurred.

Example

An individual is dismissed for alleged misconduct. An Employment Tribunal finds that the dismissal was unfair and the individual is reinstated. The individual demands that the employer deletes all references to misconduct. However, the record of the dismissal is accurate. The Tribunal's decision was that the employee should not have been dismissed on those grounds. The employer should ensure its records reflect this.

What about accuracy of opinions?

A record of an opinion is not necessarily inaccurate personal data just because the individual disagrees with it, or it is later proved to be wrong. Opinions are, by their very nature, subjective and not intended to record matters of fact.

However, in order to be accurate, your records must make clear that it is an opinion, and, where appropriate, whose opinion it is. If it becomes clear that an opinion was based on inaccurate data, you should also record this fact in order to ensure your records are not misleading.

Example

An area of particular sensitivity is medical opinion, where doctors routinely record their opinions about possible diagnoses. It is often impossible to conclude with certainty, perhaps until time has passed or tests have been done, whether a patient is suffering from a particular condition. An initial diagnosis (which is an informed opinion) may prove to be incorrect after more extensive examination or further tests. However, if the patient's records reflect the doctor's diagnosis at the time, the records are not inaccurate, because they accurately reflect that doctor's opinion at a particular time. Moreover, the record of the doctor's initial diagnosis may help those treating the patient later, and in data protection terms is required in order to comply with the 'adequacy' element of the data minimisation principle.

If an individual challenges the accuracy of an opinion, it is good practice to add a note recording the challenge and the reasons behind it.

How much weight is actually placed on an opinion is likely to depend on the experience and reliability of the person whose opinion it is, and what they base their opinion on. An opinion formed during a brief meeting will probably be given less weight than one derived from considerable dealings with the individual. However, this is not really an issue of accuracy. Instead, you need to consider whether the personal data is "adequate" for your purposes, in line with the data minimisation principle.

Note that some records that may appear to be opinions do not contain an opinion at all. For example, many financial institutions use credit scores to help them decide whether to provide credit. A credit score is a number that summarises the historical credit information on a credit report and provides a numerical predictor of the risk involved in granting an individual credit. Credit scores are based on a statistical analysis of individuals' personal data, rather than on a subjective opinion about their creditworthiness. However, you must ensure the accuracy (and adequacy) of the underlying data.

Does personal data always have to be up to date?

This depends on what you use the information for. If you use the information for a purpose that relies on it remaining current, you should keep it up to date. For example, you should update your employee payroll records when there is a pay rise. Similarly, you should update your records for customers' changes of address so that goods are delivered to the correct location.

In other cases, it will be equally obvious that you do not need to update information.

Example

An individual places a one-off order with an organisation. The organisation will probably have good reason to retain a record of the order for a certain period for accounting reasons and because of possible complaints. However, this does not mean that it has to regularly check that the customer is still living at the same address.

You do not need to update personal data if this would defeat the purpose of the processing. For example, if you hold personal data only for statistical, historical or other research reasons, updating the data might defeat that purpose.

In some cases it is reasonable to rely on the individual to tell you when their personal data has changed, such as when they change address or other contact details. It may be sensible to periodically ask individuals to update their own details, but you do not need to take extreme measures to ensure your records are up to date, unless there is a corresponding privacy risk which justifies this.

Example

An organisation keeps addresses and contact details of previous customers for marketing purposes. It does not have to use data matching or tracing services to ensure its records are up to date – and it may actually be difficult to show that the processing involved in data matching or tracing for these purposes is fair, lawful and transparent.

However, if an individual informs the organisation of a new address, it should update its records. And if a mailing is returned with the message 'not at this address' marked on the envelope – or any other information comes to light which suggests the address is no longer accurate – the organisation should update its records to indicate that the address is no longer current.

What steps do we need to take to ensure accuracy?

Where you use your own resources to compile personal data about an individual, then you must make sure the information is correct. You should take particular care if the information could have serious implications for the individual. If, for example, you give an employee a pay increase on the basis of an annual increment and a performance bonus, then there is no excuse for getting the new salary figure wrong in your payroll records.

We recognise that it may be impractical to check the accuracy of personal data someone else provides. In order to ensure that your records are not inaccurate or misleading in this case, you must:

- accurately record the information provided;
- accurately record the source of the information;
- take reasonable steps in the circumstances to ensure the accuracy of the information; and
- carefully consider any challenges to the accuracy of the information.

What is a 'reasonable step' will depend on the circumstances and, in particular, the nature of the personal data and what you will use it for. The more important it is that the personal data is accurate, the greater the effort you should put into ensuring its accuracy. So if you are using the data to make decisions that may significantly affect the individual concerned or others, you need to put more effort into ensuring accuracy. This may mean you have to get independent confirmation that the data is accurate. For example, employers may need to check the precise details of job applicants' education, qualifications and work experience if it is essential for that particular role, when they would need to obtain authoritative verification.

Example

An organisation recruiting a driver will want proof that the individuals they interview are entitled to drive the type of vehicle involved. The fact that an applicant states in his work history that he worked as a Father Christmas in a department store 20 years ago does not need to be checked for this particular job.

If your information source is someone you know to be reliable, or is a well-known organisation, it is usually reasonable to assume that they have given you accurate information. However, in some circumstances you need to double-check – for example if

inaccurate information could have serious consequences, or if common sense suggests there may be a mistake.

Example

A business that is closing down recommends a member of staff to another organisation. Assuming the two employers know each other, it may be reasonable for the organisation to which the recommendation is made to accept assurances about the individual's work experience at face value. However, if a particular skill or qualification is needed for the new job role, the organisation needs to make appropriate checks.

Example

An individual sends an email to her mobile phone company requesting that it changes its records about her willingness to receive marketing material. The company amends its records accordingly without making any checks. However, when the customer emails again asking the company to send her bills to a new address, they carry out additional security checks before making the requested change.

Even if you originally took all reasonable steps to ensure the accuracy of the data, if you later get any new information which suggests it may be wrong or misleading, you should reconsider whether it is accurate and take steps to erase, update or correct it in light of that new information as soon as possible.

What should we do if an individual challenges the accuracy of their personal data?

If this happens, you should consider whether the information is accurate and, if it is not, you should delete or correct it.

Remember that individuals have the absolute right to have incorrect personal data rectified – see the right to rectification for more information.

Individuals don't have the right to erasure just because data is inaccurate. However, the accuracy principle requires you to take all reasonable steps to erase or rectify inaccurate data without delay, and it may be reasonable to erase the data in some cases. If an individual asks you to delete inaccurate data it is therefore good practice to consider this request.

Principle (e) Storage limitation

At a glance:

- You must not keep personal data for longer than you need it.
- You need to think about – and be able to justify – how long you keep personal data. This will depend on your purposes for holding the data.
- You need a policy setting standard retention periods wherever possible, to comply with documentation requirements.
- You should also periodically review the data you hold, and erase or anonymise it when you no longer need it.
- You must carefully consider any challenges to your retention of data. Individuals have a right to erasure if you no longer need the data.
- You can keep personal data for longer if you are only keeping it for public interest archiving, scientific or historical research, or statistical purposes.

Checklist

- We know what personal data we hold and why we need it.
- We carefully consider and can justify how long we keep personal data.
- We have a policy with standard retention periods where possible, in line with documentation obligations.
- We regularly review our information and erase or anonymise personal data when we no longer need it.
- We have appropriate processes in place to comply with individuals' requests for erasure under 'the right to be forgotten'.
- We clearly identify any personal data that we need to keep for public interest archiving, scientific or historical research, or statistical purposes.

In brief:

- What is the storage limitation principle?
- Why is storage limitation important?
- Do we need a retention policy?
- How should we set retention periods?
- When should we review our retention?
- What should we do with personal data that we no longer need?
- How long can we keep personal data for archiving, research or statistical purposes?
- How does this apply to data sharing?

What is the storage limitation principle?

Article 5(1)(e) says:

“1. Personal data shall be:

(e) kept in a form which permits identification of data subjects for no longer than is necessary for the purposes for which the personal data are processed; personal data may be stored for longer periods insofar as the personal data will be processed solely for archiving purposes in the public interest, scientific or historical research purposes or statistical purposes in accordance with Article 89(1) subject to implementation of the appropriate technical and organisational measures required by this Regulation in order to safeguard the rights and freedoms of the data subject (‘storage limitation’)”

So, even if you collect and use personal data fairly and lawfully, you cannot keep it for longer than you actually need it.

There are close links here with the data minimisation and accuracy principles.

The GDPR does not set specific time limits for different types of data. This is up to you, and will depend on how long you need the data for your specified purposes.

Why is storage limitation important?

Ensuring that you erase or anonymise personal data when you no longer need it will reduce the risk that it becomes irrelevant, excessive, inaccurate or out of date. Apart from helping

you to comply with the data minimisation and accuracy principles, this also reduces the risk that you will use such data in error – to the detriment of all concerned.

Personal data held for too long will, by definition, be unnecessary. You are unlikely to have a lawful basis for retention.

From a more practical perspective, it is inefficient to hold more personal data than you need, and there may be unnecessary costs associated with storage and security.

Remember that you must also respond to subject access requests for any personal data you hold. This may be more difficult if you are holding old data for longer than you need.

Good practice around storage limitation - with clear policies on retention periods and erasure - is also likely to reduce the burden of dealing with queries about retention and individual requests for erasure.

Do we need a retention policy?

Retention policies or retention schedules list the types of record or information you hold, what you use it for, and how long you intend to keep it. They help you establish and document standard retention periods for different categories of personal data.

A retention schedule may form part of a broader 'information asset register' (IAR), or your general processing documentation.

To comply with documentation requirements, you need to establish and document standard retention periods for different categories of information you hold wherever possible. It is also advisable to have a system for ensuring that your organisation keeps to these retention periods in practice, and for reviewing retention at appropriate intervals. Your policy must also be flexible enough to allow for early deletion if appropriate. For example, if you are not actually using a record, you should reconsider whether you need to retain it.

If you are a small organisation undertaking occasional low-risk processing, you may not need a documented retention policy.

However, if you don't have a retention policy (or if it doesn't cover all of the personal data you hold), you must still regularly review the data you hold, and delete or anonymise anything you no longer need.

How should we set retention periods?

The GDPR does not dictate how long you should keep personal data. It is up to you to justify this, based on your purposes for processing. You are in the best position to judge how long you need it.

You must also be able to justify why you need to keep personal data in a form that permits identification of individuals. If you do not need to identify individuals, you should anonymise the data so that identification is no longer possible.

For example:

- You should consider your stated purposes for processing the personal data. You can keep it as long as one of those purposes still applies, but you should not keep data indefinitely 'just in case', or if there is only a small possibility that you will use it.

Example

A bank holds personal data about its customers. This includes details of each customer's address, date of birth and mother's maiden name. The bank uses this information as part of its security procedures. It is appropriate for the bank to retain this data for as long as the customer has an account with the bank. Even after the account has been closed, the bank may need to continue holding some of this information for legal or operational reasons for a further set time.

Example

A bank may need to retain images from a CCTV system installed to prevent fraud at an ATM machine for several weeks, since a suspicious transaction may not come to light until the victim gets their bank statement. In contrast, a pub may only need to retain images from their CCTV system for a short period because incidents will come to light very quickly. However, if a crime is reported to the police, the pub will need to retain images until the police have time to collect them.

Example

A tracing agency holds personal data about a debtor so that it can find that individual on behalf of a creditor. Once it has found the individual and reported to the creditor, there may

be no need to retain the information about the debtor – the agency should remove it from their systems unless there are good reasons for keeping it. Such reasons could include if the agency has also been asked to collect the debt, or because the agency is authorised to use the information to trace debtors on behalf of other creditors.

- You should consider whether you need to keep a record of a relationship with the individual once that relationship ends. You may not need to delete all personal data when the relationship ends. You may need to keep some information so that you can confirm that the relationship existed – and that it has ended – as well as some of its details.

Example

A business may need to keep some personal data about a previous customer so that they can deal with any complaints the customer might make about the services they provided.

Example

An employer should review the personal data it holds about an employee when they leave the organisation's employment. It will need to retain enough data to enable the organisation to deal with, for example, providing references or pension arrangements. However, it should delete personal data that it is unlikely to need again from its records – such as the employee's emergency contact details, previous addresses, or death-in-service beneficiary details.

Example

A business receives a notice from a former customer requiring it to stop processing the customer's personal data for direct marketing. It is appropriate for the business to retain enough information about the former customer for it to stop including that person in future direct marketing activities.

- You should consider whether you need to keep information to defend possible future legal claims. However, you could still delete information that could not possibly be relevant to such a claim. Unless there is some other reason for keeping it, personal data should be deleted when such a claim could no longer arise.

Example

An employer receives several applications for a job vacancy. Unless there is a clear business reason for doing so, the employer should not keep recruitment records for unsuccessful applicants beyond the statutory period in which a claim arising from the recruitment process may be brought.

- You should consider any legal or regulatory requirements. There are various legal requirements and professional guidelines about keeping certain kinds of records – such as information needed for income tax and audit purposes, or information on aspects of health and safety. If you keep personal data to comply with a requirement like this, you will not be considered to have kept the information for longer than necessary.
- You should consider any relevant industry standards or guidelines. For example, we have agreed that credit reference agencies are permitted to keep consumer credit data for six years. Industry guidelines are a good starting point for standard retention periods and are likely to take a considered approach. However, they do not guarantee compliance. You must still be able to explain why those periods are justified, and keep them under review.

You must remember to take a proportionate approach, balancing your needs with the impact of retention on individuals' privacy. Don't forget that your retention of the data must also always be fair and lawful.

When should we review our retention?

You should review whether you still need personal data at the end of any standard retention period, and erase or anonymise it unless there is a clear justification for keeping it for longer. Automated systems can flag records for review, or delete information after a pre-determined period. This is particularly useful if you hold many records of the same type.

It is also good practice to review your retention of personal data at regular intervals before this, especially if the standard retention period is lengthy or there is potential for a significant impact on individuals.

If you don't have a set retention period for the personal data, you must regularly review whether you still need it.

However, there is no firm rule about how regular these reviews must be. Your resources may be a relevant factor here, along with the privacy risk to individuals. The important thing to remember is that you must be able to justify your retention and how often you review it.

You must also review whether you still need personal data if the individual asks you to. Individuals have the absolute right to erasure of personal data that you no longer need for your specified purposes.

What should we do with personal data that we no longer need?

You can either erase (delete) it or anonymise it.

You need to remember that there is a significant difference between permanently deleting personal data and taking it offline. If personal data is stored offline, this should reduce its availability and the risk of misuse or mistake. However, you are still processing personal data. You should only store it offline (rather than delete it) if you can still justify holding it. You must be prepared to respond to subject access requests for personal data stored offline, and you must still comply with all the other principles and rights.

The word 'deletion' can mean different things in relation to electronic data, and we recognise it is not always possible to delete or erase all traces of the data. The key issue is to ensure you put the data beyond use. If it is appropriate to delete personal data from a live system, you should also delete it from any back-up of the information on that system.

Alternatively, you can anonymise the data so that it is no longer "in a form which permits identification of data subjects".

Personal data that has been pseudonymised – eg key-coded – will usually still permit identification. Pseudonymisation can be a useful tool for compliance with other principles such as data minimisation and security, but the storage limitation principle still applies.

How long can we keep personal data for archiving, research or statistical purposes?

You can keep personal data indefinitely if you are holding it only for:

- archiving purposes in the public interest;
- scientific or historical research purposes; or
- statistical purposes.

Although the general rule is that you cannot hold personal data indefinitely 'just in case' it might be useful in future, there is an inbuilt exception if you are keeping it for these archiving, research or statistical purposes.

You must have appropriate safeguards in place to protect individuals. For example, pseudonymisation may be appropriate in some cases.

This must be your only purpose. If you justify indefinite retention on this basis, you cannot later use that data for another purpose - in particular for any decisions affecting particular individuals. This does not prevent other organisations from accessing public archives, but they must ensure their own collection and use of the personal data complies with the principles.

How does this apply to data sharing?

If you share personal data with other organisations, you should agree between you what happens once you no longer need to share the data. In some cases, it may be best to return the shared data to the organisation that supplied it without keeping a copy. In other cases, all of the organisations involved should delete their copies of the personal data.

Example

Personal data about the customers of Company A is shared with Company B, which is negotiating to buy Company A's business. The companies arrange for Company B to keep the information confidential and use it only in connection with the proposed transaction. The sale does not go ahead, and Company B returns the customer information to Company A without keeping a copy.

The organisations involved in an information-sharing initiative may each need to set their own retention periods, because some may have good reasons to retain personal data for longer than others. However, if you all only hold the data for the purposes of the data-sharing initiative and it is no longer needed for that initiative, then all organisations with copies of the information should delete it.

Principle (f) Integrity and confidentiality

At a glance:

- A key principle of the GDPR is that you process personal data securely by means of 'appropriate technical and organisational measures' – this is the 'security principle'.

- Doing this requires you to consider things like risk analysis, organisational policies, and physical and technical measures.
- You also have to take into account additional requirements about the security of your processing – and these also apply to data processors.
- You can consider the state of the art and costs of implementation when deciding what measures to take – but they must be appropriate both to your circumstances and the risk your processing poses.
- Where appropriate, you should look to use measures such as pseudonymisation and encryption.
- Your measures must ensure the ‘confidentiality, integrity and availability’ of your systems and services and the personal data you process within them.
- The measures must also enable you to restore access and availability to personal data in a timely manner in the event of a physical or technical incident.
- You also need to ensure that you have appropriate processes in place to test the effectiveness of your measures, and undertake any required improvements.

Checklists

We undertake an analysis of the risks presented by our processing and use this to assess the appropriate level of security we need to put in place.

When deciding what measures to implement, we take account of the state of the art and costs of implementation.

We have an information security policy (or equivalent) and take steps to make sure the policy is implemented.

Where necessary, we have additional policies and ensure that controls are in place to enforce them.

We make sure that we regularly review our information security policies and measures and, where necessary, improve them.

We have put in place basic technical controls such as those specified by established frameworks like Cyber Essentials.

- We understand that we may also need to put other technical measures in place depending on our circumstances and the type of personal data we process.
- We use encryption and/or pseudonymisation where it is appropriate to do so.
- We understand the requirements of confidentiality, integrity and availability for the personal data we process.
- We make sure that we can restore access to personal data in the event of any incidents, such as by establishing an appropriate backup process.
- We conduct regular testing and reviews of our measures to ensure they remain effective, and act on the results of those tests where they highlight areas for improvement.
- Where appropriate, we implement measures that adhere to an approved code of conduct or certification mechanism.
- We ensure that any data processor we use also implements appropriate technical and organisational measures.

In brief:

- What does the GDPR say about security?
- Why should we worry about information security?
- What do we need to protect with our security measures?
- What level of security is required?
- What organisational measures do we need to consider?
- What technical measures do we need to consider?
- What if we operate in a sector that has its own security requirements?
- What do we do when a data processor is involved?
- Should we use pseudonymisation and encryption?
- What are 'confidentiality, integrity, availability' and 'resilience'?
- What are the requirements for restoring availability and access to personal data?
- Are we required to ensure our security measures are effective?

- What about codes of conduct and certification?
- What about our staff?

What does the GDPR say about security?

Article 5(1)(f) of the GDPR concerns the 'integrity and confidentiality' of personal data. It says that personal data shall be:

'Processed in a manner that ensures appropriate security of the personal data, including protection against unauthorised or unlawful processing and against accidental loss, destruction or damage, using appropriate technical or organisational measures'

You can refer to this as the GDPR's 'security principle'. It concerns the broad concept of **information security**.

This means that you must have appropriate security to prevent the personal data you hold being accidentally or deliberately compromised. You should remember that while information security is sometimes considered as cybersecurity (the protection of your networks and information systems from attack), it also covers other things like physical and organisational security measures.

You need to consider the security principle alongside Article 32 of the GDPR, which provides more specifics on the security of your processing. Article 32(1) states:

'Taking into account the state of the art, the costs of implementation and the nature, scope, context and purposes of processing as well as the risk of varying likelihood and severity for the rights and freedoms of natural persons, the controller and the processor shall implement appropriate technical and organisational measures to ensure a level of security appropriate to the risk'

Why should we worry about information security?

Poor information security leaves your systems and services at risk and may cause real harm and distress to individuals – lives may even be endangered in some extreme cases.

Some examples of the harm caused by the loss or abuse of personal data include:

- identity fraud;
- fake credit card transactions;

- targeting of individuals by fraudsters, potentially made more convincing by compromised personal data;
- witnesses put at risk of physical harm or intimidation;
- offenders at risk from vigilantes;
- exposure of the addresses of service personnel, police and prison officers, and those at risk of domestic violence;
- fake applications for tax credits; and
- mortgage fraud.

Although these consequences do not always happen, you should recognise that individuals are still entitled to be protected from less serious kinds of harm, for example embarrassment or inconvenience.

Information security is important, not only because it is itself a legal requirement, but also because it can support good data governance and help you demonstrate your compliance with other aspects of the GDPR.

The ICO is also required to consider the technical and organisational measures you had in place when considering an administrative fine.

What do our security measures need to protect?

The security principle goes beyond the way you store or transmit information. Every aspect of your processing of personal data is covered, not just cybersecurity. This means the security measures you put in place should seek to ensure that:

- the data can be accessed, altered, disclosed or deleted only by those you have authorised to do so (and that those people only act within the scope of the authority you give them);
- the data you hold is accurate and complete in relation to why you are processing it; and
- the data remains accessible and usable, ie, if personal data is accidentally lost, altered or destroyed, you should be able to recover it and therefore prevent any damage or distress to the individuals concerned.

These are known as 'confidentiality, integrity and availability' and under the GDPR, they form part of your obligations.

What level of security is required?

The GDPR does not define the security measures that you should have in place. It requires you to have a level of security that is 'appropriate' to the risks presented by your processing. You need to consider this in relation to the state of the art and costs of implementation, as well as the nature, scope, context and purpose of your processing.

This reflects both the GDPR's risk-based approach, and that there is no 'one size fits all' solution to information security. It means that what's 'appropriate' for you will depend on your own circumstances, the processing you're doing, and the risks it presents to your organisation.

So, before deciding what measures are appropriate, you need to assess your information risk. You should review the personal data you hold and the way you use it in order to assess how valuable, sensitive or confidential it is – as well as the damage or distress that may be caused if the data was compromised. You should also take account of factors such as:

- the nature and extent of your organisation's premises and computer systems;
- the number of staff you have and the extent of their access to personal data; and
- any personal data held or used by a data processor acting on your behalf.

We cannot provide a complete guide to all aspects of security in all circumstances for all organisations, but this guidance is intended to identify the main points for you to consider.

What organisational measures do we need to consider?

Carrying out an information risk assessment is one example of an organisational measure, but you will need to take other measures as well. You should aim to build a culture of security awareness within your organisation. You should identify a person with day-to-day responsibility for information security within your organisation and make sure this person has the appropriate resources and authority to do their job effectively.

Example

The Chief Executive of a medium-sized organisation asks the Director of Resources to ensure that appropriate security measures are in place, and that regular reports are made to the board.

The Resources Department takes responsibility for designing and implementing the organisation's security policy, writing procedures for staff to follow, organising staff training, checking whether security measures are actually being adhered to and investigating security incidents.

Clear accountability for security will ensure that you do not overlook these issues, and that your overall security posture does not become flawed or out of date.

Although an information security policy is an example of an appropriate organisational measure, you may not need a 'formal' policy document or an associated set of policies in specific areas. It depends on your size and the amount and nature of the personal data you process, and the way you use that data. However, having a policy does enable you to demonstrate how you are taking steps to comply with the security principle.

Whether or not you have such a policy, you still need to consider security and other related matters such as:

- co-ordination between key people in your organisation (eg the security manager will need to know about commissioning and disposing of any IT equipment);
- access to premises or equipment given to anyone outside your organisation (eg for computer maintenance) and the additional security considerations this will generate;
- business continuity arrangements that identify how you will protect and recover any personal data you hold; and
- periodic checks to ensure that your security measures remain appropriate and up to date.

What technical measures do we need to consider?

Technical measures are sometimes thought of as the protection of personal data held in computers and networks. Whilst these are of obvious importance, many security incidents can be due to the theft or loss of equipment, the abandonment of old computers or hard-

copy records being lost, stolen or incorrectly disposed of. Technical measures therefore include both physical and computer or IT security.

When considering physical security, you should look at factors such as:

- the quality of doors and locks, and the protection of your premises by such means as alarms, security lighting or CCTV;
- how you control access to your premises, and how visitors are supervised;
- how you dispose of any paper and electronic waste; and
- how you keep IT equipment, particularly mobile devices, secure.

In the IT context, technical measures may sometimes be referred to as 'cybersecurity'. This is a complex technical area that is constantly evolving, with new threats and vulnerabilities always emerging. It may therefore be sensible to assume that your systems are vulnerable and take steps to protect them.

When considering cybersecurity, you should look at factors such as:

- system security – the security of your network and information systems, including those which process personal data;
- data security – the security of the data you hold within your systems, eg ensuring appropriate access controls are in place and that data is held securely;
- online security – eg the security of your website and any other online service or application that you use; and
- device security – including policies on Bring-your-own-Device (BYOD) if you offer it.

Depending on the sophistication of your systems, your usage requirements and the technical expertise of your staff, you may need to obtain specialist information security advice that goes beyond the scope of this guidance. However, it's also the case that you may not need a great deal of time and resources to secure your systems and the personal data they process.

Whatever you do, you should remember the following:

- your cybersecurity measures need to be appropriate to the size and use of your network and information systems;

- you should take into account the state of technological development, but you are also able to consider the costs of implementation;
- your security must be appropriate to your business practices. For example, if you offer staff the ability to work from home, you need to put measures in place to ensure that this does not compromise your security; and
- your measures must be appropriate to the nature of the personal data you hold and the harm that might result from any compromise.

A good starting point is to make sure that you're in line with the requirements of Cyber Essentials – a government scheme that includes a set of basic technical controls you can put in place relatively easily.

You should however be aware that you may have to go beyond these requirements, depending on your processing activities. Cyber Essentials is only intended to provide a 'base' set of controls, and won't address the circumstances of every organisation or the risks posed by every processing operation.

What if we operate in a sector that has its own security requirements?

Some industries have specific security requirements or require you to adhere to certain frameworks or standards. These may be set collectively, for example by industry bodies or trade associations, or could be set by other regulators. If you operate in these sectors, you need to be aware of their requirements, particularly if specific technical measures are specified.

Although following these requirements will not necessarily equate to compliance with the GDPR's security principle, the ICO will nevertheless consider these carefully in any considerations of regulatory action. It can be the case that they specify certain measures that you should have, and that those measures contribute to your overall security posture.

Example

If you are processing payment card data, you are obliged to comply with the Payment Card Industry Data Security Standard. The PCI-DSS outlines a number of specific technical and organisational measures that the payment card industry considers applicable whenever such data is being processed.

Although compliance with the PCI-DSS is not necessarily equivalent to compliance with the GDPR's security principle, if you process card data and suffer a personal data breach, the ICO will consider the extent to which you have put in place measures that PCI-DSS requires particularly if the breach related to a lack of a particular control or process mandated by the standard.

What do we do when a data processor is involved?

If one or more organisations process personal data on your behalf, then these are data processors under the GDPR. This can have the potential to cause security problems – as a data controller you are responsible for ensuring compliance with the GDPR and this includes what the processor does with the data. However, in addition to this, the GDPR's security requirements also apply to any processor you use.

This means that:

- you must choose a data processor that provides sufficient guarantees about its security measures;
- your written contract must stipulate that the processor takes all measures required under Article 32 – basically, the contract has to require the processor to undertake the same security measures that you would have to take if you were doing the processing yourself; and
- you should ensure that your contract includes a requirement that the processor makes available all information necessary to demonstrate compliance. This may include allowing for you to audit and inspect the processor, either yourself or an authorised third party.

At the same time, your processor can assist you in ensuring compliance with your security obligations. For example, if you lack the resource or technical expertise to implement certain measures, engaging a processor that has these resources can assist you in making sure personal data is processed securely, provided that your contractual arrangements are appropriate.

Should we use pseudonymisation and encryption?

Pseudonymisation and encryption are specified in the GDPR as two examples of measures that may be appropriate for you to implement. This does not mean that you are obliged to

use these measures. It depends on the nature, scope, context and purposes of your processing, and the risks posed to individuals.

However, there are a wide range of solutions that allow you to implement both without great cost or difficulty. For example, for a number of years the ICO has considered encryption to be an appropriate technical measure given its widespread availability and relatively low cost of implementation. This position has not altered due to the GDPR — if you are storing personal data, or transmitting it over the internet, we recommend that you use encryption and have a suitable policy in place, taking account of the residual risks involved.

When considering what to put in place, you should undertake a risk analysis and document your findings.

What are ‘confidentiality, integrity, availability’ and ‘resilience’?

Collectively known as the ‘CIA triad’, confidentiality, integrity and availability are the three key elements of information security. If any of the three elements is compromised, then there can be serious consequences, both for you as a data controller, and for the individuals whose data you process.

The information security measures you implement should seek to guarantee all three both for the systems themselves and any data they process.

The CIA triad has existed for a number of years and its concepts are well-known to security professionals.

You are also required to have the ability to ensure the ‘resilience’ of your processing systems and services. Resilience refers to:

- whether your systems can continue operating under adverse conditions, such as those that may result from a physical or technical incident; and
- your ability to restore them to an effective state.

This refers to things like business continuity plans, disaster recovery, and cyber resilience

What are the requirements for restoring availability and access to personal data?

You must have the ability to restore the availability and access to personal data in the event of a physical or technical incident in a ‘timely manner’.

The GDPR does not define what a 'timely manner' should be. This therefore depends on:

- who you are;
- what systems you have; and
- the risk that may be posed to individuals if the personal data you process is unavailable for a period of time.

The key point is that you have taken this into account during your information risk assessment and selection of security measures. For example, by ensuring that you have an appropriate backup process in place you will have some level of assurance that if your systems do suffer a physical or technical incident you can restore them, and therefore the personal data they hold, as soon as reasonably possible.

Example

An organisation takes regular backups of its systems and the personal data held within them. It follows the well-known '3-2-1' backup strategy: three copies, with two stored on different devices and one stored off-site.

The organisation is targeted by a ransomware attack that results in the data being encrypted. This means that it is no longer able to access the personal data it holds.

Depending on the nature of the organisation and the data it processes, this lack of availability can have significant consequences on individuals – and would therefore be a personal data breach under the GDPR.

The ransomware has spread throughout the organisation's systems, meaning that two of the backups are also unavailable. However, the third backup, being stored off-site, allows the organisation to restore its systems in a timely manner. There may still be a loss of personal data depending on when the off-site backup was taken, but having the ability to restore the systems means that whilst there will be some disruption to the service, the organisation are nevertheless able to comply with this requirement of the GDPR.