

WHITE PAPER • FOR BOARDS, DATA PROTECTION OFFICERS AND DISPOSAL OPERATORS

Data Destruction Is Binary

Why certificates without warranties fail UK GDPR accountability, and what a sufficient guarantee actually looks like.

Technology neutral and supplier agnostic. No vendor is named in this paper.

Sections 1 to 10 and 12 are supplier agnostic and name no product, supplier or certification scheme. Section 11 states the publisher's own position and is separated and labelled for that reason. Nothing here constitutes legal advice. The test set out at 9.7 is intended to be applied to the publisher on the same terms as to anyone else.

Second edition

DOCUMENT VERSION

September 2026

ISSUED

Free to share

NO REGISTRATION, NO GATE

Contents

EXECUTIVE SUMMARY	Data destruction is binary	1	SECTION 7	Warranties, responsibility, and the moment certificates lost their meaning	18
	Certificates have replaced certainty with comfort	1		7.1 What a warranty actually does	18
	The scale of it is on the public record	1		7.2 Why early destruction certificates carried real weight	18
	The regulator's position removes the ambiguity	2		7.3 The choice the industry faced	18
	The hidden liability gap	2		7.4 How responsibility was removed without changing appearances	18
	The binary choice organisations now face	2		7.5 The quiet transfer of risk back to the controller	19
	Why this paper exists	3		7.6 Why this is not a criticism of individual vendors	19
				7.7 Why responsibility is now the differentiator	19
SECTION 1	The assurance gap	4	SECTION 8	Why downstream warranties do not fix upstream disclaimers	20
	1.1 The problem is systemic, not individual	4		8.1 The common response	20
	1.2 Vendor disclosure exists, but assurance is implied elsewhere	4		8.2 How the assurance chain actually works	20
	1.3 Downstream assurance and reliance	4		8.3 The problem with relying on a downstream warranty	20
	1.4 Certification is capability, not responsibility	5		8.4 Commercial risk is not regulatory accountability	20
	1.5 The regulator's position exposes the gap	5		8.5 When downstream warranties collapse under scrutiny	21
	1.6 Why this matters	5		8.6 Article 28 is not satisfied by assumption	21
				8.7 The mismatch contract wording cannot fix	21
SECTION 2	The evidence	6	SECTION 9	What a sufficient guarantee looks like	22
	2.1 The programme, and who ran it	6		9.1 The problem with vague assurance	22
	2.2 Three models, taken from the responses themselves	6		9.2 The outcome standard	22
	2.3 What organisations said they hold	7		9.3 A guarantee must survive reconciliation	22
	2.4 The two regulators who were asked the same question	7		9.4 Control and responsibility must align	23
	2.5 What the evidence does and does not show	8		9.5 A guarantee must be device specific	23
				9.6 A guarantee must be meaningful when tested	23
SECTION 3	Process evidence and outcome evidence	9		9.7 The four questions	23
	3.1 The industry has conflated activity with result	9		9.8 Why sufficient guarantees are rare	23
	3.2 What process evidence can prove	9	SECTION 10	Why certification does not equal a guarantee	25
	3.3 What process evidence cannot prove	9		10.1 What certification is for	25
	3.4 The test that cannot fail	9		10.2 What certification actually confirms	25
	3.5 Why a standards reference does not resolve it	10		10.3 Certification does not override contracts	25
	3.6 Outcome evidence requires responsibility, not just records	10		10.4 Why it is mistaken for assurance	25
	3.7 Why this matters under UK GDPR	10		10.5 The effect of encouraged certification	25
	3.8 The practical consequence	10	SECTION 11	The Data Safe Solutions position	27
SECTION 4	Method selection, human judgement and enforced correctness	11		11.1 A different starting assumption	27
	4.1 Selection is the decision point	11		11.2 The loop at Section 5, implemented	27
	4.2 Human judgement introduces unverifiable risk	11		11.3 Certificate, system and contract are aligned	27
	4.3 Documentation does not validate judgement	11		11.4 The warranty is insured	28
	4.4 Defaults and optional controls increase exposure	11		11.5 Responsibility sits where control exists	28
	4.5 What most platforms actually do	12		11.6 What this means in an audit	28
	4.6 Enforced correctness is structurally different	12		11.7 This is not a claim that others are wrong	28
	4.7 Accountability follows control	12	SECTION 12	The decision facing boards and data controllers	29
	4.8 The unavoidable conclusion	12		12.1 This is no longer a technical debate	29
SECTION 5	The complete evidence loop	13		12.2 The decision cannot be delegated away	29
	5.1 A correct method can still fail	13		12.3 The two defensible positions	29
	5.2 A status flag is not a test	13		12.4 What this costs the organisations that got it right	29
	5.3 A test requires a known state to test against	13		12.5 What regulators and auditors will ask	30
	5.4 The test must be empirical and independent	14		12.6 The final question	30
	5.5 Two distinct gates, and both must fail closed	14	SOURCES	Sources and method	31
	5.6 The evidence must be re-checkable	14		Legislation	31
	5.7 Certify only a passed outcome	14		Regulatory correspondence	31
	5.8 The loop, stated in full	15		The Freedom of Information programme	31
	5.9 Why the chain matters more than any step in it	15		Patent	31
SECTION 6	From magnetic disks to modern storage	16		Technical standards	31
	6.1 When storage was simple, destruction was simple	16		Vendor terms	31
	6.2 Why overwrite could be guaranteed on magnetic media	16		Statement on scope	32
	6.3 The transition that broke the model	16			
	6.4 There is no longer a single correct method	16			
	6.5 Why documentation cannot close the gap	17			
	6.6 Automation is no longer an optimisation	17			

Data destruction is binary

Data is either permanently and irretrievably destroyed, or it is not.

There is no partial destruction, no acceptable probability, and no best efforts outcome that satisfies accountability once data is shown to be recoverable. If personal data can be recovered now or in the future, destruction has failed.

UK GDPR does not frame destruction as an effort. It frames it as an outcome. Article 5(2) states that "the controller shall be responsible for, and be able to demonstrate compliance with paragraph 1." Being able to demonstrate compliance requires more than evidence that an activity took place. It requires confidence that the intended outcome was achieved.

This is not a philosophical position. It is the practical reality that underpins data protection law, breach liability and regulatory enforcement.

Certificates have replaced certainty with comfort

For years, organisations have relied on Certificates of Data Destruction as evidence that risk has been transferred away from them. In most cases that belief is incorrect.

Across the data destruction industry, certificates overwhelmingly confirm that a process was executed, while the underlying software licences and service contracts explicitly refuse responsibility for the outcome. Common contractual positions include:

- Software supplied "as is"
- No warranty that data will be rendered irretrievable
- No guarantee of error free or effective operation
- Explicit exclusion or limitation of liability for data loss

As a result, certificates are issued that appear authoritative, pass audits and satisfy procurement requirements, while the supplier contractually avoids liability if the data later proves recoverable.

The certificate exists. **The responsibility does not move.**

The scale of it is on the public record

This is not a theoretical concern about what could happen. Between January and September 2026 the question was put to the UK public sector under the Freedom of Information Act, and the answers are published in full.

1,036

Freedom of Information requests issued to UK public bodies

684

Classified responses forming the analysed set

365

Named the supplier's certificate as their only assurance

437

Hold no warranty of the destruction outcome

Section 2 sets this out in full, including the two environmental regulators with IT disposal oversight duties who were asked the same questions as everyone else, and answered them.

The regulator's position removes the ambiguity

Correspondence with the Information Commissioner's Office established the test. A destruction certificate may demonstrate compliance, but only on a condition, and the condition is the whole argument of this paper:

"A destruction certificate on its [sic] own may be a sufficient way to demonstrate compliance, **as long as you have ensured that the data has been destroyed effectively**. This is likely to depend on **the system being used for destruction and how the destruction can be confirmed** by the system or the organisation providing the service."

Information Commissioner's Office, correspondence of 21 January 2026, case reference IC-471076-R5B2

On certificates that carry disclaimers, the same correspondence is direct about what a controller should do:

"I can't confirm how the ICO would treat a certificate with the disclaimers you described. Data controllers must do their due diligence and ensure that whoever they use will destroy their data effectively. **If the disclaimers don't provide reassurances and confirmations that the data controller needs to confidently say that their data will be destroyed in the way that complies with UK GDPR, they should consider using an organisation that does.**"

Information Commissioner's Office, correspondence of 21 January 2026, case reference IC-471076-R5B2

Certificates that confirm process while disclaiming outcome cannot provide the reassurance a controller needs in order to confidently assert compliance.

The hidden liability gap

Where certificates are issued alongside "as is" licences, outcome disclaimers and liability exclusions, an organisation faces a structural exposure. A certificate confirms a process ran. The supplier disclaims responsibility for effectiveness. Liability for data loss is excluded or capped. Accountability remains with the data controller.

If data believed to be destroyed later resurfaces, the existence of a certificate does not protect the organisation. It may instead demonstrate reliance on an assurance that was never given. This is not a technical nuance. It is a governance and liability gap that boards, senior officers and data protection officers carry, in many cases without knowing it.

The binary choice organisations now face

Once this position is understood there are two defensible options.

- **Accept ambiguity.** Acknowledge that the certificates the organisation relies on do not warrant irretrievability, and that residual risk remains with the controller.
- **Require certainty.** Use arrangements where the outcome is warranted, the governing licence does not contradict the certificate, and the party accepting responsibility is the party that controls the result.

There is no third option that removes the accountability gap.

Why this paper exists

This paper does not criticise regulators or assign blame to individual organisations, and it does not name individual suppliers. Its purpose is to surface a structural industry problem that has been normalised over time, and to explain why process documentation cannot compensate for contractual denial of responsibility.

Where certainty is required, **ambiguity is not a defensible position.**

The assurance gap

Where responsibility is disclosed, obscured and misunderstood.

1.1 The problem is systemic, not individual

The assurance gap in data destruction is not the product of bad actors. It is the cumulative result of reasonable decisions taken independently by vendors, service providers, certification bodies and buyers, none of which is wrong on its own terms, and which together leave responsibility for the outcome sitting with the one party least able to control it.

1.2 Vendor disclosure exists, but assurance is implied elsewhere

Data destruction software vendors generally disclose the limits of their responsibility. These disclosures sit in end user licence agreements, warranty clauses and limitation of liability provisions. Common positions include software supplied as is, no warranty of effectiveness or outcome, no guarantee of error free operation, and exclusion or limitation of liability for data loss.

From a legal standpoint, these disclosures protect the vendor. From a governance standpoint a problem emerges. The certificate produced by the software is treated by customers as the primary assurance artefact, while the licence terms that deny outcome responsibility are treated as background legal boilerplate.

This creates an asymmetry. The certificate appears to confirm successful destruction. The licence quietly denies responsibility if that conclusion is wrong. The disclosure exists, but it is not surfaced at the point where assurance is relied upon.

This is not deception. **It is misalignment.**

1.3 Downstream assurance and reliance

The gap widens where IT asset disposal providers or managed service providers sit between the controller and the underlying software. In these arrangements the disposal provider may select the destruction tooling, operate the process, issue certificates or service confirmations, and provide contractual assurances to the controller.

This creates expectation reliance. A controller may reasonably assume that the provider has selected appropriate tools, understands their limitations and is standing behind the outcome. Where the underlying software disclaims irretrievability, and that limitation is not surfaced, the controller retains risk without realising it.

This is not necessarily a failure of intent. It is a failure of alignment between what is operationally delivered, what is contractually warranted, and what is assumed by the customer. Under UK GDPR accountability remains with the controller regardless of how many intermediaries sit in the chain, so the gap accumulates at controller level unless it is actively closed.

1.4 Certification is capability, not responsibility

Independent certification and standards alignment play a valuable role. They validate that a product can perform certain functions, that a process aligns with a defined standard, and that testing has been conducted under controlled conditions.

What certification does not do is warrant live operational outcomes, override vendor licence terms, transfer liability for failure, or accept responsibility for irretrievability. In practice it is often interpreted as assurance, particularly where members of a scheme are encouraged or required to use certified products. That misunderstanding is not caused by certification itself. It is caused by how certification is read downstream.

Certification validates capability. **It does not warrant results.**

1.5 The regulator's position exposes the gap

The Information Commissioner's Office does not prescribe specific tools or certificates. It places responsibility on the controller to be able to confidently assert effectiveness, and directs a controller whose supplier's disclaimers undermine that confidence to consider using an organisation whose terms do not.

That cuts across the entire chain. Where certificates confirm process, licences disclaim outcome, service providers rely on those tools, and certification validates capability but not responsibility, the reassurance required to confidently assert compliance may simply not exist anywhere in the arrangement.

1.6 Why this matters

The assurance gap only becomes visible when something goes wrong. When data believed to be destroyed resurfaces, regulators, auditors and courts do not examine documents in isolation. They ask one practical question.

Who took responsibility for ensuring the data was **irretrievable**?

Where the answer is unclear, accountability defaults to the data controller. That is the structural problem the industry has normalised, and the problem this paper exists to surface.

The evidence

What happened when the question was put to the UK public sector.

Section 1 describes how responsibility can be disclaimed at the foundation of an assurance chain while appearing intact at the top. Whether that actually happens, and at what scale, is a question of fact rather than argument. This section answers it.

2.1 The programme, and who ran it

Between January and September 2026, **1,036 requests were issued to UK public bodies under the Freedom of Information Act 2000**, asking what evidence each organisation holds that personal data on end of life storage media was rendered irretrievable, and what assurance it relies upon.

The requests were made by **Chris Littlewood in a personal capacity, as a private individual**, not by Data Safe Solutions Ltd. They were submitted through whatdotheyknow.com under his own name, where every request and every response is published in full and permanently. Data Safe Solutions analyses that public record in this paper. It did not make the requests, and it holds no information about any responding body that is not already published.

This matters to how the evidence should be read. A supplier's own survey of its market invites the obvious objection. A body of statutory responses given under FOIA by the organisations themselves, published where anyone can read them without taking a supplier's word for any of it, does not.

2.2 Three models, taken from the responses themselves

684 classified responses form the analysed set. The categorisation is taken from each organisation's own recorded answer, given under statutory duty, rather than from any inference of ours.

MODEL	BODIES	RECORDED POSITION
1	6	Software based destruction with an explicit outcome warranty. The supplier's disclaimer has been contractually overridden.
2	241	Physical destruction. Placed outside the software based assurance question by the organisation's own choice of method.
3	437	Software based destruction with no outcome warranty required or provided, no device specific verification evidence held, and reliance placed on supplier certificates issued under terms that disclaim the outcome.

The FOI question concerned software based erasure, so the 241 organisations that moved to physical destruction placed themselves outside it. That leaves **443 bodies relying on software, of which 437 hold no warranty that the final data state is irretrievable, and six do**.

437 is the floor, not the ceiling. The classification credits an organisation with a pass wherever its response records the medium as physically destroyed, without testing whether a destruction state was evidenced or warranted. Very few of the 241 could evidence either. The true number of UK public

bodies unable to demonstrate an irretrievable final data state is higher than 437, and this paper reports only what the responses prove.

2.3 What organisations said they hold

Asked to name the basis of their assurance, organisations answered as follows.

RECORDED BASIS OF ASSURANCE	BODIES
Supplier certificate only. A process certificate from the destruction supplier, with no device level outcome evidence held.	365
Process logs only. Logs recording that the process ran, without outcome verification.	25
Hybrid evidence. Several artefacts relied on together.	16
No recorded assurance at all.	7
Standard compliance only. Alignment to a standard, without device level outcome evidence.	6
Contractual obligation only. A contractual term, without outcome evidence.	4
Explicit outcome warranty. A warrant of irretrievability.	7

Bodies relying on physical destruction are recorded under that method and counted separately. Twelve responses received since the last categorisation are not yet assigned a basis.

That is the finding, and it is the argument of this paper stated as a number. The regulator said a certificate is sufficient only where the controller has ensured that destruction was effective. **365 organisations named the certificate itself as the thing that satisfies that requirement.**

The certificate cannot carry it. It is issued by a vendor who declines, in writing, **to warrant the very thing the certificate is being used to prove.**

So the assurance runs in a circle. The controller relies on the certificate, the certificate rests on the software, and the software's licence says the outcome is not warranted. Nobody in that chain has accepted responsibility for whether the data is gone, and the only party carrying a legal duty is the one at the start of it.

2.4 The two regulators who were asked the same question

Two of the responding bodies are the environmental regulators with oversight responsibility for IT disposal. Both were asked what everyone else was asked. Their answers are on the public record.

Scottish Environment Protection Agency

Asked for the recorded basis on which the Agency determines that its use of software based destruction complies with UK GDPR, its formal response F0201023 states:

"Your request has highlighted a gap in our Data Protection Accountability obligations for this processing. This will be addressed by the drafting of a Data Protection Impact Assessment (DPIA) for the specific processing."

The characterisation as a gap is the Agency's own. No prior compliance assessment had been carried out and no DPIA was held at the time of the request. The remedial DPIA was committed to in the response itself, after an external enquiry surfaced the issue rather than an internal review.

Environment Agency

Across three successive exchanges between February and May 2026, each narrowing the position further than the last, the Agency confirmed that it does not hold recorded information demonstrating that any warranty extends to the software erasure method used or its effectiveness, and identified process based artefacts as the assurance relied upon.

Neither body is the data protection regulator, and neither response is an interpretation of Article 5(2). Their value is different and, for the purposes of this paper, greater. Two organisations with statutory oversight duties in this area were asked what evidence they hold, and found they did not hold it.

2.5 What the evidence does and does not show

It does not show that data has been recovered from any device disposed of by any of these organisations, and this paper makes no such claim. It shows something narrower and harder to answer: that across 684 recorded positions, the assurance model described in Section 1 is the dominant model in the UK public sector, and that the organisations relying on it have said so themselves, in writing, under statutory duty.

The remainder of this paper explains why that model cannot deliver what is asked of it, and what would.

Process evidence and outcome evidence

Why the distinction decides everything that follows.

3.1 The industry has conflated activity with result

Across the data destruction industry, evidence of compliance is overwhelmingly framed around process execution. The common artefacts are execution logs, sanitisation reports, verification checks, sampling results and certificates of completion.

These confirm that a defined action was taken. They do not, by themselves, confirm that the intended outcome was achieved. The distinction has blurred over time, and evidence that a process ran has been allowed to stand in for evidence that data is irretrievable. That substitution is the root of the accountability problem.

3.2 What process evidence can prove

Process evidence is valuable and necessary. It can demonstrate that a tool was executed, that a method aligned to a recognised standard was selected, that a workflow completed without reported errors, that certain addressable areas were overwritten or purged, and that an activity occurred at a specific time on a specific device.

This supports audit, traceability and operational assurance. What it does not establish is that all data on the device has been rendered irretrievable.

3.3 What process evidence cannot prove

There are inherent technical limits to what logs, reports and sampling can demonstrate, particularly on modern storage. Process evidence cannot, on its own, prove irretrievability in over provisioned areas, remapped or retired blocks, controller managed regions, firmware controlled storage areas, non addressable regions, or device specific behaviours not exposed to the host.

Sampling may confirm that some blocks appear erased. It cannot confirm that all data is unrecoverable, now or in the future. This is not a failure of tooling. It is a consequence of storage architecture.

3.4 The test that cannot fail

The most important consequence of that architecture is easy to state and rarely stated.

An overwrite can only reach the addresses the host is able to point at. On flash, the drive's own controller has already moved the original data out of those addresses and into blocks the host cannot name: remapped, spare, over provisioned. The overwrite fills the addresses it was given. The verification then reads back those same addresses, finds exactly what it just wrote, and reports success. It is telling the truth about the part of the drive that no longer holds the data.

Nothing failed. No setting was wrong, no operator slipped, no supplier cut a corner. The process ran, the check passed, the certificate was correct. **And the data is still there.**

This is why a genuine, verified, certified erasure and a drive still holding recoverable data are not a contradiction. With overwrite on flash that can be the outcome, at scale, inside disposal chains that are fully accredited and doing exactly what the standard asks of them. The certificate is not a lie. It is an answer to a question that was never the right one.

3.5 Why a standards reference does not resolve it

Referencing a recognised standard does not convert process evidence into outcome evidence. A report stating that a method aligned to NIST SP 800-88 was selected confirms intent. It does not confirm that the method was effective for the specific device, firmware and storage behaviour involved.

It is worth noting what those standards actually say. NCSC guidance, NIST SP 800-88 and IEEE 2883 all direct that flash media be sanitised by cryptographic erase or by the drive's own purge command rather than by overwriting. The standards are not the problem. Naming one in a certificate, without evidence that the method it directs was the method applied, is.

Standards describe what should be done. **They do not warrant that what was done worked.**

3.6 Outcome evidence requires responsibility, not just records

Outcome evidence is not created by additional logs or reports. It is created when an organisation explicitly accepts responsibility for the result. An outcome statement means the correct method was selected for the device, the method was enforced rather than optional, unsafe fallbacks were prevented, the result was validated, and responsibility for irretrievability is accepted.

Without that acceptance of responsibility, documentation remains evidence of attempt rather than evidence of success.

3.7 Why this matters under UK GDPR

UK GDPR does not require controllers to prove perfection. It requires them to be able to demonstrate that appropriate technical and organisational measures were effective, and Article 32(1)(d) requires a process for regularly testing, assessing and evaluating that effectiveness.

Where evidence confirms only that a process was attempted, and effectiveness is disclaimed by the party that performed it, the reassurance the regulator described may not exist. In those circumstances a paper trail does not remove accountability. It documents it.

3.8 The practical consequence

When data believed to be destroyed later resurfaces, the question asked is not what process was run. It is whether the data was irretrievable, and who took responsibility for ensuring that it was. If outcome responsibility cannot be identified, accountability defaults to the controller regardless of how complete the documentation appears.

Method selection, human judgement and enforced correctness

The single point at which destruction succeeds or fails.

4.1 Selection is the decision point

Data destruction does not fail because a tool is unavailable. It fails because the wrong method is applied to the device. Effectiveness depends on selecting a method appropriate to the specific storage technology, interface, firmware behaviour and operational state of the device being processed. If the method is incorrect, the outcome is incorrect regardless of how well the process is documented.

4.2 Human judgement introduces unverifiable risk

Across the industry, method selection is determined by a person, either selected at the machine or pre-configured in advance. In each case the correctness of the selection depends on assumptions that cannot be verified at the point of execution.

The operator cannot see firmware behaviour. The operator cannot observe controller managed regions. The operator cannot confirm how the device handles remapped or retired blocks. The selection is therefore based on belief rather than certainty.

This is not a criticism of operators, and it is worth being precise about why the human step exists at all. The category was built on offering a menu of methods, which is a reasonable design only if someone can reliably pick the right one for every device in front of them. On modern storage nobody can.

4.3 Documentation does not validate judgement

Where a method is selected manually, documentation can confirm only that a choice was made. Logs may show which option was selected. Certificates may record that the selected process completed. Signatures may attest that instructions were followed. None of those artefacts validates the underlying judgement. If the wrong method was chosen, the documentation records the failure rather than preventing it.

4.4 Defaults and optional controls increase exposure

Some systems attempt to reduce reliance on judgement through defaults or recommended configurations. A default is a general assumption applied to a specific device. Optional controls rely on the operator recognising when to change them. Where fallbacks are permitted, unsafe options remain available. Where configuration is flexible, error remains possible.

From an accountability perspective, **optional correctness is indistinguishable from error.**

4.5 What most platforms actually do

It is worth being precise about the mechanism, because the language used to describe it is misleading. The leading erasure products do not interrogate an individual drive and determine the correct sanitisation method for that device. They present a list of methods, or apply a configured policy, and then check whether the drive supports the method that has already been chosen.

That is not selection. **It is compatibility checking.**

Overwrite persists as the default because it works across the widest range of devices and gives the operator the highest chance of getting the asset through the process. On magnetic media that is sound engineering. On flash it is the failure described at 3.4.

4.6 Enforced correctness is structurally different

There is a fundamental difference between guidance and enforcement. Guidance suggests what should be done. Enforcement ensures what must be done. Enforced correctness means the device is automatically identified, the appropriate method is determined by the system, unsafe or ineffective methods are not selectable, the process cannot proceed with an incorrect configuration, and failure conditions are explicit.

This removes judgement from the point of highest risk. Training reduces error frequency but does not eliminate it, and even highly skilled operators cannot compensate for information they cannot observe.

4.7 Accountability follows control

Where method selection is left to human judgement, responsibility for correctness rests with the person or organisation making that judgement. Where selection is enforced by the system, responsibility shifts to the party that designed and controls that enforcement.

That distinction becomes decisive after a failure. The investigation focuses on the decision point: who selected the method, on what basis, with what information, under what constraints. Where selection was manual or optional, accountability traces back to the controller and their chosen process. Where it was enforced and responsibility explicitly accepted, accountability is clear and bounded.

4.8 The unavoidable conclusion

Systems that rely on human judgement at the point of method selection can provide records of activity. They cannot guarantee correctness. Where outcome certainty is required, method selection must be enforced rather than advised. That is not an enhancement. It is a prerequisite.

It is not, however, sufficient on its own. Determining the correct method establishes what should destroy the data. A correct method can still fail to complete, or report completion without achieving the result, and neither of those is detected by a system that stops at selection. What outcome evidence additionally requires is set out in the next section.

The complete evidence loop

Why correct selection is necessary, and not sufficient.

Section 4 established that the destruction method must be determined from the device rather than chosen in advance. That removes the largest single source of failure. It does not, on its own, produce outcome evidence, and a paper that stopped there would be making a smaller claim than the problem requires.

5.1 A correct method can still fail

Determining the right method for a device establishes what *should* destroy the data on it. Two things can then happen that no amount of correct selection prevents.

The command may not complete. Firmware implementations vary between manufacturers, between models, and between production runs of the same model. A device may report support for an operation and then fail to carry it out.

Or the command may report completion without achieving the result. That is the more dangerous case, because everything downstream then behaves as though the work is done.

5.2 A status flag is not a test

Where a system reads the drive's own success flag and treats that as verification, it is asking the controller whether the controller did its job. The flag is a self report from the same firmware whose behaviour is the thing in question, and it confirms only that the command returned. It does not confirm that the data is gone.

This is the same category of error described at 3.4, one layer further in. There, a verification read of host addressable space confirmed what had just been written to it. Here, a status flag confirms that an operation was accepted. In both cases something true is being reported about a question that was never the right one.

Outcome evidence cannot be produced by asking the device whether it succeeded. **It has to be tested.**

5.3 A test requires a known state to test against

A result can only be tested against a baseline. If nothing is known about the device before the operation, then afterwards there is nothing to compare, and the only available check is whether recoverable data happens to be found, which proves nothing when none is found.

Establishing a baseline means writing verification data to known locations on the device before erasure and confirming that it is genuinely present. That creates a state whose disappearance is meaningful.

It also creates a mandatory precondition. If the verification data cannot be written and confirmed present, there is no valid baseline, no trustworthy test, and therefore no basis on which the device can

later be certified.

5.4 The test must be empirical and independent

After execution, the test is to attempt to recover the verification data from the same locations. That is a before and after test of the actual result on the actual device, carried out independently of anything the device says about itself.

This is the step that converts a process record into outcome evidence. Everything before it establishes that the right thing was attempted. Only this establishes that it worked.

5.5 Two distinct gates, and both must fail closed

Two different failures can occur, and conflating them is how systems certify devices they should not.

GATE	WHAT HAS HAPPENED
The command did not complete	The determined method was attempted and did not finish. Nothing has been proved or disproved about the data. The device must not be certified.
The outcome test disproves the result	The command reported completion, but verification data has survived in locations it should not have. The device reported success and the result contradicts it. The device must not be certified.

Failing closed means the device is diverted to physical destruction or to remediation. It does not mean dropping to a weaker method that will pass. A system that substitutes an overwrite when the determined method fails has not recovered the outcome; it has recovered the certificate.

Where the correct method is the one that destroys the data rather than the one chosen to pass, **there is nothing to fall back to.**

5.6 The evidence must be re-checkable

Evidence that can only be read is evidence that must be trusted. Evidence that can be re-checked is evidence that can be relied upon under challenge, which is the only condition that matters in an investigation.

That means recording the state of the tested locations before and after, the result of the outcome test, the device state, and a signature that binds the record to the physical device rather than to a report about it. A hash of the evidence combined with a value the drive itself maintains, such as its own write counter, allows a later party to establish that the record corresponds to that unit and has not been reconstructed after the fact.

5.7 Certify only a passed outcome

A certificate should exist only for a device that has passed the outcome test. Everything that failed at any gate should already have been diverted, and should never appear in a certificate on the basis of a substitute method that was applied because the correct one did not work.

Stated that way it sounds obvious. The difficulty is that it requires a system willing to produce fewer certificates than devices processed, and to record failures as failures.

5.8 The loop, stated in full

Assembled, the requirement is a closed loop rather than a step. Each stage is supplier agnostic and can be put to any supplier as a question.

STAGE	WHAT IT ESTABLISHES
1 Interrogate the individual device	What this physical unit actually is and reports at the interface, rather than what the model number or the batch implies.
2 Determine the method from that capability	The one method that will destroy the data on this device. Determined, not selected from a list.
3 Establish a known pre-state	Verification data written and confirmed present, creating a baseline the outcome can be tested against.
4 Execute the determined method	The device native operation runs, enforced rather than advised.
5 Fail closed if it does not complete	No substitution, no weaker method, no certificate.
6 Test the outcome empirically	Attempt to recover the verification data. An independent before and after test rather than a status flag.
7 Fail closed if the test disproves it	A distinct gate. The command said it worked; the evidence says otherwise, and the evidence governs.
8 Record re-checkable evidence	Before and after state, result, device state, and a signature binding the record to the physical unit.
9 Certify only a passed outcome	The certificate becomes a record of a tested result rather than of an attempt.

5.9 Why the chain matters more than any step in it

No single stage above is remarkable on its own. Device interrogation exists. Firmware level commands exist. Logging exists. What produces outcome evidence is the chain, and specifically the two gates that stop it, because they are what make a certificate mean something.

Any link removed collapses the rest. Interrogate but do not determine, and the method is still a guess. Determine but do not seed a baseline, and there is nothing to test. Test but allow a fallback when the test fails, and the certificate records the substitute. Test and gate correctly but log only a result, and the evidence cannot be re-checked when it is challenged.

Correct selection makes the right outcome **possible**. Only a tested and gated loop makes it **evidenced**.

From magnetic disks to modern storage

Why guaranteed destruction was once possible, and why it no longer is.

6.1 When storage was simple, destruction was simple

For decades enterprise storage was dominated by rotating magnetic hard disk drives. Data was stored magnetically on spinning platters, storage behaviour was largely transparent to the host, and blocks written by the operating system mapped predictably to physical locations. Reserved and remapped areas existed, but they were limited, well understood, and small enough that overwriting the addressable media was a sound basis for treating the outcome as achieved.

In that environment destruction had one solution: overwrite the entire addressable media. If every sector was overwritten correctly, all data was rendered irretrievable. There were no alternative behaviours to account for and no ambiguity about completeness. Because there was effectively one valid method, the outcome could be warranted.

6.2 Why overwrite could be guaranteed on magnetic media

On rotating drives all user data resided in host addressable sectors, writing new data physically replaced old magnetic patterns, multiple passes increased confidence without changing the underlying mechanism, and remapped areas existed but were limited and well understood.

The relationship between process and outcome was direct and observable. Run the method correctly, achieve the outcome. This is the historical context in which destruction certificates and warranties originally made sense.

6.3 The transition that broke the model

Solid state and NVMe devices introduced fundamental architectural changes. Data is managed by controllers rather than written directly to fixed locations. Wear levelling spreads writes across physical cells. Blocks are dynamically remapped and retired. Over provisioned areas exist beyond host visibility. Firmware controls placement and movement.

These changes were essential for performance and longevity. They also broke the assumption that overwriting host addressable space destroys all data.

6.4 There is no longer a single correct method

Destruction is no longer one problem with one solution. Different storage technologies require different methods, including cryptographic erase, controller based purge operations, secure erase commands, targeted firmware supported processes, and physical destruction. The correct method depends on storage type, interface, firmware behaviour, encryption state and device capabilities.

Using the wrong method may leave data recoverable even where the process completes successfully. There is no universal fallback.

6.5 Why documentation cannot close the gap

Logs and certificates can record which method was selected. They cannot prove that the method was the correct one. This creates the critical shift in risk that defines the modern problem.

On magnetic media, correct execution equalled correct outcome. On modern storage, **correct execution does not guarantee correctness of method.**

6.6 Automation is no longer an optimisation

In a single method world, automation improved efficiency. In a multi method world it is the only way to remove selection error. A process that identifies the device, determines its capabilities, enforces the correct method and prevents unsafe alternatives removes the largest source of failure. Without it, irretrievability cannot be guaranteed.

Warranties, responsibility, and the moment certificates lost their meaning

What changed, and why nothing looked different.

7.1 What a warranty actually does

A warranty is not a report and it is not a record. It is an acceptance of responsibility for an outcome. When an organisation warrants an outcome it is saying that the result is correct, and that if the result is not correct, responsibility for that failure sits with the warranting party.

It is worth being precise about what does and does not move. A warranty transfers **contractual responsibility and commercial exposure** to the party giving it. It does not transfer the controller's **statutory accountability** under UK GDPR, which cannot be assigned by contract and remains with the controller however the supply chain is arranged. What a warranty does is give the controller something to rely on when demonstrating that its measures were effective, and a party who is answerable if they were not. That is a materially different position from holding a certificate nobody stands behind, and it is not the same thing as having transferred the duty.

7.2 Why early destruction certificates carried real weight

In the era of rotating magnetic storage, certificates were meaningful because they were implicitly tied to a single reliable outcome. There was one correct method. If it completed successfully, the data was irretrievable. Because correctness was observable and method choice was not variable, vendors and service providers could reasonably stand behind the result. Certificates functioned as warranties in practice even where the word was not used.

7.3 The choice the industry faced

As storage evolved the relationship between process and outcome broke. Multiple methods became necessary, correctness became device specific, and human judgement entered the decision path. Vendors could no longer guarantee that a single method would work for all devices in all circumstances.

At that point there were two options. Refuse to issue certificates unless the outcome could be warranted. Or continue issuing certificates while removing responsibility for whether the outcome was correct. The industry overwhelmingly chose the second.

7.4 How responsibility was removed without changing appearances

The shift did not occur by changing the format of certificates. They still look authoritative, still reference standards, still confirm successful completion. Responsibility was removed elsewhere: through software supplied as is, warranties limited to software operation rather than outcome, explicit disclaimers of effectiveness or irretrievability, exclusions of liability for data loss, and liability caps tied to licence fees.

Certificates continued to exist. The responsibility they were assumed to represent did not.

7.5 The quiet transfer of risk back to the controller

By issuing certificates while disclaiming outcome, the industry shifted risk back to customers without making that shift explicit. Most organisations did not reassess their position. Certificates looked the same, audits passed, procurement was satisfied. But what was once a transfer of responsibility became a retention of liability disguised as assurance.

7.6 Why this is not a criticism of individual vendors

This paper does not argue that vendors acted improperly. The technical complexity of modern storage made unconditional warranties difficult, and removing responsibility was commercially rational. Publishing terms that disclaim an outcome is lawful and commonplace.

The problem is not the decision. It is that the meaning of certificates changed while the reliance placed upon them did not.

7.7 Why responsibility is now the differentiator

The core question is no longer which tool is used or which standard is referenced. It is who is willing to take responsibility for the outcome. Where responsibility is accepted, certificates have meaning. Where it is disclaimed, certificates are records of activity rather than protection.

Certificates did not become weaker because standards declined or tools became less capable. They became weaker because **responsibility was removed while appearances remained the same.**

Why downstream warranties do not fix upstream disclaimers

The assumption that creates false comfort.

8.1 The common response

When concerns are raised about software disclaimers, the usual answer is that the disposal provider warrants the destruction, or that the service contract makes the provider responsible. At first glance that appears to solve the problem. It does not. A downstream warranty cannot create certainty where upstream responsibility has already been denied.

8.2 How the assurance chain actually works

In a typical service model, responsibility flows through a chain. A software vendor supplies a destruction tool. The tool is governed by licence terms and disclaimers. A disposal provider operates the tool. The provider issues a certificate or service warranty. The controller relies on that assurance.

If the foundation of the chain explicitly disclaims outcome responsibility, the links above it cannot restore certainty simply by restating it.

8.3 The problem with relying on a downstream warranty

Where a provider warrants that data has been destroyed but the underlying software disclaims irretrievability, one of two things must be true. Either the provider has independent means of guaranteeing irretrievability beyond the software, or the provider is warranting something it cannot fully control.

In most cases it is the second. The provider did not design the firmware, does not control controller behaviour, and cannot override how the software operates internally. The warranty is therefore conditional even where it is not presented as such.

8.4 Commercial risk is not regulatory accountability

Downstream warranties redistribute commercial risk. They do not change regulatory accountability. A provider may agree to compensate a customer if a failure occurs. That does not alter the fact that under UK GDPR accountability for appropriate technical and organisational measures remains with the controller.

It should also be said plainly that the processor is not insulated either. The Commissioner's office has confirmed the position directly:

"Processors can be held liable for non-compliance by relevant investigative and corrective powers of a supervisory authority (such as the ICO) and may be subject to administrative fines or other penalties. However, the ICO works on a case by case basis and I can't confirm what the outcome would be in each case."

Information Commissioner's Office, correspondence of 21 January 2026, case reference IC-471076-R5B2

8.5 When downstream warranties collapse under scrutiny

If data believed to be destroyed is later recovered, several things happen quickly. The provider refers to the software. The software vendor refers to its licence terms. Those terms disclaim outcome and liability. The downstream warranty is then exposed as unenforceable, limited, or irrelevant to the regulatory assessment. The original disclaimers resurface exactly when certainty is needed most.

8.6 Article 28 is not satisfied by assumption

Article 28(1) requires a controller to use only processors providing sufficient guarantees that processing will meet the requirements of the Regulation. A guarantee resting on tools that explicitly deny responsibility for the outcome is not, in substance, a guarantee of effectiveness. It may be a promise of effort or of process. It is not a guarantee of irretrievability.

8.7 The mismatch contract wording cannot fix

Contracts can allocate liability. They cannot change technical reality. If a destruction outcome depends on correct method selection and correct execution by software that disclaims correctness, no amount of downstream wording converts that into certainty.

Responsibility must exist **where control exists**. Where it does not, warranties become statements of intent rather than assurances of outcome.

Most downstream warranties are never tested. They exist in documents, satisfy procurement and pass audits. Their weakness only becomes visible after failure, by which point it is too late to discover that the warranty relied upon was built on disclaimers beneath it.

What a sufficient guarantee looks like

The outcome standard, stated so that it can be tested.

9.1 The problem with vague assurance

In many arrangements the language used to describe assurance is broad and imprecise. Phrases such as industry standard, certified, compliant or best practice are relied upon as proxies for certainty. They are not guarantees. They describe alignment, not responsibility. Where irretrievability matters, assurance must be specific, explicit and defensible under scrutiny.

9.2 The outcome standard

What follows is not a scheme, an accreditation or a certification. Nobody administers it. It is a definition of what *destroyed* has to mean before anyone writes it on a certificate, stated in three parts so that a buyer can test any supplier against it. A supplier meets all three or none.

Proven	Warranted	Insured
A device level record that the data on that specific device was rendered irretrievable. Not a record that a process ran.	That outcome warranted in the governing licence , and not disclaimed or contradicted anywhere else in the agreement.	Cover that responds when the outcome fails , rather than a policy that excludes the one event that matters.

The same requirement, worded for a specification

For each device processed the supplier shall deliver a **device level record evidencing that the data has been rendered irretrievable**; shall **warrant that outcome** in the governing licence, with liability for it not disclaimed elsewhere in the agreement; and shall hold **insurance that responds to failure of that outcome**.

Nothing in that specification names a supplier, and it should not. If another supplier can prove the outcome, warrant it and insure it, they have met the standard and the buyer is protected. That is the point. Naming a product in a specification also invites challenge and usually has to carry "or equivalent" anyway, whereas naming an outcome does not.

9.3 A guarantee must survive reconciliation

A guarantee cannot exist in isolation. It must remain valid when all relevant documents are read together: certificates, contracts, software licence terms, service terms and liability clauses. If a certificate asserts irretrievability but the governing licence denies responsibility for effectiveness, the guarantee does not survive reconciliation, and the apparent assurance is overridden by the contractual reality.

9.4 Control and responsibility must align

Sufficient guarantees exist only where responsibility aligns with control. The party guaranteeing the outcome must control method selection, enforcement of the correct method, prevention of unsafe alternatives, execution conditions and validation of results. Guarantees issued by parties who do not control those factors are necessarily conditional, even where that conditionality is not stated.

9.5 A guarantee must be device specific

Modern storage does not allow generic guarantees. A sufficient guarantee must take account of storage type, interface, device capabilities, firmware behaviour and media specific risks. A statement that destruction was successful, without reference to device specific behaviour, is a general assertion rather than a guarantee of irretrievability.

9.6 A guarantee must be meaningful when tested

The value of a guarantee is not assessed when everything goes to plan. It is assessed when something goes wrong: when data is later recovered, when a regulator investigates, when an auditor challenges assumptions, when a breach notification is required. If it evaporates under those conditions, it was never sufficient.

9.7 The four questions

The standard can be tested with four questions. Put them to any supplier, and ask for the answers in writing.

ASK	WHAT A GOOD ANSWER LOOKS LIKE
How does your software select the correct destruction method for each individual drive?	The device is interrogated and the method determined per drive, not one method applied across a batch.
Does your certificate warrant that the data on that specific device is irretrievable, or only that a process completed?	A warranty of the outcome on the certificate, and the same warranty in the governing licence.
When the correct method is refused or unsupported, do you fall back to an overwrite and still certify?	Fallback is not permitted and the record says so. A failure is recorded as a failure.
What do you pay if data is later recovered from a device you certified?	A stated liability position with the clause number, and insurance that responds to the destruction having failed.

A supplier who can answer all four on paper is giving outcome evidence. One who cannot is giving a process receipt. The test applies to the author of this paper as much as to anyone else.

9.8 Why sufficient guarantees are rare

They are rare because they require a vendor or service provider to accept responsibility for an outcome that is technically difficult and commercially risky. Avoiding that responsibility is understandable. Normalising ambiguity is not.

Ask one question. If data believed to be destroyed later resurfaces, **who is contractually and explicitly responsible for that failure?** If the honest answer is unclear, disputed, or defaults to the controller, sufficient guarantees do not exist.

Why certification does not equal a guarantee

Necessary, valuable, and not what many people think it is.

10.1 What certification is for

Certification and standards define good practice, establish baseline capability, provide consistency across suppliers, enable comparability in procurement and support audit and governance. Without them the industry would be fragmented and far harder to regulate. Certification is valuable.

10.2 What certification actually confirms

Certification confirms that a product or process has been assessed against defined criteria under defined conditions: that a tool is capable of performing certain functions, that a method aligns with a published standard, that testing has been conducted in controlled scenarios, and that documentation meets prescribed requirements.

What it does not do is warrant that every real world execution will produce an irretrievable outcome. Certification validates capability. It does not accept responsibility.

10.3 Certification does not override contracts

If a software licence states that the product is supplied as is, disclaims effectiveness, or excludes liability for data loss, certification does not negate those clauses. The contractual position still governs responsibility. Certification may demonstrate that a tool can perform a task. It does not change who is responsible if that task fails.

10.4 Why it is mistaken for assurance

Certification logos are prominent, certificates reference recognised standards, membership frameworks encourage the use of certified tools, and buyers assume certification implies risk transfer. Over time certification becomes shorthand for safety. That interpretation is understandable and incorrect.

Certification bodies do not operate the destruction process, select methods for live devices, control execution conditions, warrant irretrievability or accept liability for failure. Their role is to assess conformance, not to underwrite outcomes. Treating certification as a guarantee creates an assurance gap that no certification body intends to fill.

10.5 The effect of encouraged certification

Where a membership body or framework requires or encourages the use of certified tools, an additional risk emerges. Certification may be read not only as capability validation but as compliance sufficiency, leading organisations to believe that using a certified product is by itself enough to transfer risk.

Certification answers "can this be done". **It does not answer "who is responsible if it goes wrong".**

A certified tool that disclaims irretrievability creates the same accountability challenge as an uncertified one. Where irretrievability matters, certification must be combined with explicit outcome responsibility. Without that, certification is context rather than protection.

The Data Safe Solutions position

Stated plainly, and open to the same test as everyone else.

Everything to this point is supplier agnostic. This section is not, and it is separated for that reason. Data Safe Solutions publishes this paper, and states its own position rather than leaving the reader to infer it.

11.1 A different starting assumption

Most destruction solutions begin from the assumption that responsibility must be limited. This one begins from the assumption that if a certificate is issued stating that destruction was successful, the party issuing it should be willing to stand behind that statement. That is not an incremental improvement. It is a different view of responsibility.

11.2 The loop at Section 5, implemented

Section 5 sets out what outcome evidence requires, stated so that it can be put to any supplier. This is how it is implemented here, and it is the whole of the answer rather than a summary of it.

Each device is interrogated in place and the method determined from what that physical unit reports at the interface, not from the model on the label or the batch it arrived in. Verification data is written to randomly selected locations and confirmed present before anything is erased, establishing the baseline the outcome will be tested against. If that baseline cannot be established, the device fails there.

The determined method is then executed. Where it does not complete, the device is not dropped to a weaker method that would pass; it is failed and diverted to physical destruction. There is nothing to fall back to, because the method was chosen to destroy the data rather than to produce a certificate.

After execution, recovery of the seeded data is attempted from the same locations. That is the outcome test, and it is independent of anything the drive reports about itself. If any of it survives, the device fails at a second and distinct gate, regardless of what the command returned. What is recorded is the state of those locations before and after, the result, the device state, and a sanitisation signature binding the record to that physical unit, so the outcome can be re-verified later rather than taken on trust.

A certificate is issued only for a device that has passed. No human makes a method decision at any point in that chain.

No single stage in that loop is novel, and this paper does not claim otherwise. The integrated chain is the subject of published UK patent application **GB2638704A**, with a related United States filing.

11.3 Certificate, system and contract are aligned

The misalignment identified throughout this paper is the gap between what a certificate asserts and what the governing terms allow. Here there is no such gap. The contractual deliverable is the **Data Destruction Outcome**, achieved only where a Successful Execution Event is recorded.

- **Clause 5.1.** Irretrievability is warranted for each successful execution.
- **Clause 17.1.** Liability for that outcome is uncapped, and is not qualified by a disclaimer elsewhere in the agreement.
- Where destruction does not complete successfully, no warranty is given, no charge is made, and the failure is recorded as a failure.
- Device level evidential records are produced for each item. Certificates are evidential records, not process receipts.

11.4 The warranty is insured

A warranty is only as good as the party standing behind it. The outcome carries bound insurance of **£10 million Technology Professional Indemnity and £10 million Public and Products Liability**, each on an each and every claim basis, with **the standard inefficacy exclusion expressly removed by endorsement**.

That endorsement is the material term rather than the figure. Insurance policies routinely exclude claims arising from a product failing to perform its intended function, which in this industry is the only event that counts. With that exclusion expressly removed, the cover responds to the destruction itself having failed. The policy therefore sits alongside the warranty rather than around it.

11.5 Responsibility sits where control exists

Responsibility is accepted because the controlling factors are controlled: device identification, method determination, enforcement of execution, the two fail closed gates, and the empirical test of the result. That alignment is what allows an outcome warranty to exist without contradiction, and it is the test set out at 9.4 applied to ourselves. A warranty can be given for a tested outcome in a way it cannot be given for an assumed one.

11.6 What this means in an audit

When assurance is tested, clarity matters more than volume of documentation. The certificate answers the decisive question directly. Was the data irretrievable, and who took responsibility for ensuring that it was. There is no need to infer intent from logs or to reconcile disclaimers, because the responsibility is explicit.

Independently certified to ADISA Product Assurance Level 5. Patent application GB2638704A published.

11.7 This is not a claim that others are wrong

We do not claim that other vendors or service providers are negligent or non compliant. The industry evolved in response to genuine technical complexity and commercial pressure, and disclaiming an outcome that cannot be controlled is a rational decision. We simply chose not to normalise the ambiguity that resulted.

The four questions at 9.7 apply to us. Put them to your current supplier, then put the same four to us, and compare the answers side by side.

The decision facing boards and data controllers

Once the assurance gap is visible, it cannot be unseen.

12.1 This is no longer a technical debate

The issue is not which tool is faster, cheaper or more widely deployed, and it is not a debate about standards alignment or certification logos. It is a governance decision about where responsibility sits when data is believed to be destroyed. Once the distinction between process evidence and outcome responsibility is understood, technical preference becomes secondary.

12.2 The decision cannot be delegated away

Data destruction is often treated as an operational detail. It is a board level risk decision. Under UK GDPR accountability remains with the controller regardless of how many suppliers, processors or certifications are involved, so the decision cannot be delegated to tooling, procurement frameworks or service providers without understanding the consequences.

12.3 The two defensible positions

Option one

Accept that the certificates relied upon today certify that a process was run, not that data is irretrievable, and that residual risk remains with the organisation.

Option two

Require arrangements where irretrievability is explicitly warranted, responsibility is aligned with control, and the certificate issued survives reconciliation against the governing terms without contradiction.

There is no third position that removes the accountability gap. Continuing existing arrangements without reassessment is not neutral. It is an active decision to retain ambiguity. If data later resurfaces, the organisation will not be judged on what it believed at the time, but on what it could reasonably have understood and addressed.

12.4 What this costs the organisations that got it right

It is worth recording what the alternative has cost. 241 of the bodies in this study moved to physical destruction because they could not satisfy themselves that software would render data unrecoverable, and a number recorded that assessment in their data protection impact assessments.

They were right about what they read. Every product whose terms could be examined disclaims the destruction outcome, and a controller reviewing those licences and concluding that software could not

give it confidence was reading the contracts correctly. But they paid for that certainty in environmental impact, in residual asset value, in reduced reuse and digital inclusion, and in higher direct cost.

Six bodies in the same dataset reached the identical conclusion about disclaimed software and then required the outcome to be warranted instead. They hold what the 241 went to such lengths for, without destroying a working machine. That is the choice this paper is really about.

12.5 What regulators and auditors will ask

When assurance is tested the questions are simple. Was the data irretrievable. Who took responsibility for ensuring that it was. On what basis that responsibility was accepted. Documents that do not answer those questions directly will be reconciled until they do, and where responsibility is unclear, accountability defaults to the controller.

12.6 The final question

The assumptions that underpinned legacy destruction models no longer hold. Storage has changed, risk has increased and regulatory expectations have sharpened. What was once reasonable reliance has become unjustified comfort.

This paper does not tell any organisation what decision to make. It exists to ensure the decision is made consciously. Before relying on any Certificate of Data Destruction, ask one question.

If this data is later recovered, **who stands behind the statement that it was irretrievable?** If the answer is unclear, the risk remains.

Data destruction is binary. Either responsibility is accepted, or it is retained. Certificates do not change that reality. Only outcome responsibility does.

Sources and method

Legislation

UK General Data Protection Regulation, Articles 5(1)(f), 5(2), 24, 28(1) and 32, in particular Article 32(1) (d) on regularly testing, assessing and evaluating the effectiveness of technical and organisational measures. Data Protection Act 2018.

Regulatory correspondence

Information Commissioner's Office, case reference IC-471076-R5B2, correspondence dated 21 January 2026 and 28 January 2026. All quotations are verbatim, including the original's punctuation. Data Safe Solutions Ltd made a submission to the Commissioner under Article 5(2) arising from this correspondence.

The Freedom of Information programme

1,036 requests issued to UK public bodies under the Freedom of Information Act 2000 from January 2026. The requests were made by **Chris Littlewood in a personal capacity, as a private individual**, not by Data Safe Solutions Ltd, and were submitted through whatdotheyknow.com under his own name, where every request and every response is published in full and permanently. 684 classified responses form the analysed set. Data Safe Solutions analyses that public record; it did not make the requests. Figures are the cut current at the date of publication, the programme continues, and later cuts supersede this one. The 1,036 figure is requests issued, not responses received.

Named exhibits: Scottish Environment Protection Agency response F0201023; Environment Agency responses FOI2026/07163 of 12 March 2026 and FOI2026/12919 of 7 April 2026. Both are on the public record at whatdotheyknow.com.

Patent

Published UK patent application GB2638704A, with a related United States filing. The published application is a public document and the claims are open to inspection.

Technical standards

NIST SP 800-88 Revision 2, September 2025, which supersedes Revision 1 and directs compliance with IEEE 2883, NSA specifications or an organisationally approved standard. IEEE 2883-2022 and IEEE 2883.1-2025. NCSC guidance on secure sanitisation and disposal of storage media. The three NIST sanitisation methods are Clear, Purge and Destroy.

On the limits of overwrite on flash media, the citable experimental source is Wei, Grupp, Spada and Swanson, *Reliably Erasing Data from Flash-Based Solid State Drives*, USENIX FAST'11.

Vendor terms

The characterisation of licence positions in this paper is drawn from the published end user licence agreements and terms of the erasure software vendors whose terms could be located, each read in full

and retrieved on 9 September 2026, with copies held on file. No vendor is named in this paper. The vendor by vendor record, with verbatim quotations and clause references, is maintained separately and is available to counsel on request.

Statement on scope

This paper makes no allegation of unlawful conduct by any vendor, service provider, certification body or public authority. Publishing terms that disclaim an outcome is lawful. No accreditation scheme operator has been consulted on, or has endorsed, this document. The finding of this paper is the gap between what those terms say and what purchasers believe they are buying.

Data Safe Solutions Ltd, 2nd Floor, Port of Liverpool Building, Mann Island, Liverpool L3 1BY. Telephone 0151 440 3200. hello@datasafe360.com. Company registration 12245066. VAT registration 425206721.