




IDHUS Institute
Policy Paper
PP-002

Building Institutional Memory for Reuse

Copyright 2026 © IDHUS Institute - idhus.org

Editor and Director: David González

All rights reserved.

This publication forms part of the constitutional editorial architecture of the IDHUS Institute and belongs to the Institutional Cognition Policy Paper Series.

No part of this publication may be reproduced, distributed, translated or adapted without prior written permission from the Institute, except where permitted under applicable copyright legislation or for legitimate academic quotation with appropriate citation.

Index

Executive Summary	4
I. The Problem of Stored but Unused Knowledge	6
II. Where Institutional Memory Breaks	9
III. When Memory Failure Becomes an Institutional Problem	13
IV. Designing Institutional Memory for Reuse	17
V. From Intervention to Institutional Function	22
VI. Responsible Memory: Challenge, Renewal and Judgement	27
VII. Evaluating Capability Gain.....	32
VIII. Implications, Limits and Capability Direction	38
Conclusion — Capability Direction	43
References.....	44

Executive Summary

Public institutions can retain extensive knowledge while remaining unable to use it when later decisions, operations or institutional processes require it. Reports, evaluations, decisions, lessons and records may survive in repositories or organisational systems without becoming reliable inputs to future action. **Retention is therefore a necessary condition of institutional memory, but it is not institutional memory capability.**

This policy paper, PP-002, addresses this gap between stored knowledge and functional reuse. It defines institutional memory as a capability through which materially relevant retained knowledge can be recovered when needed, reconstructed with sufficient context, brought into appropriate institutional workflows and subjected to informed judgement about its present applicability. The policy problem is consequently not how institutions can preserve the greatest possible amount of knowledge, but how they can make relevant accumulated knowledge reliably reusable without creating disproportionate burden or allowing the past to acquire inappropriate authority over present decisions.

The analysis identifies four principal points at which the route from retention to reuse can fail. **Retention without discovery** occurs when knowledge survives but intended users cannot reliably find it. **Retrieval without meaning** occurs when material can be recovered but insufficient context prevents later users from interpreting it responsibly. **Understanding without invocation** occurs when relevant knowledge can be found and understood but does not enter the decisions or processes in which it could matter. **Invocation without appropriate use** occurs when prior knowledge enters present work but is followed, ignored or weighted without adequate judgement of its relevance, validity or authority.

These failures can generate duplicated reconstruction, repeated avoidable error, discontinuity in institutional reasoning, coordination and implementation weaknesses, loss of preparedness and reduced ability to explain how institutional decisions evolved. Their significance, however, varies by context. **Not every memory failure warrants formal intervention.** Intervention is justified where the failure creates material consequences or significant exposure, its mechanism is sufficiently understood and a proportionate response offers plausible value greater than the institutional burden it creates.

Where that threshold is crossed, intervention should begin with the failed function rather than a preferred organisational or technological solution. Problems of discovery call for stronger findability; problems of interpretation for sufficient contextual preservation; failures of invocation for proportionate routes through which relevant prior knowledge can re-enter institutional workflows; and inappropriate use for mechanisms that preserve challenge, renewal and independent applicability judgement. Ownership and stewardship may be required across these functions to keep memory arrangements viable over time.

The resulting mechanisms form a **design space rather than a mandatory institutional package**. Retrieval architecture, contextual arrangements, human and social memory, workflow integration, stewardship and renewal mechanisms should be configured according to the diagnosed failure, institutional context and readiness. Institutions should therefore **transfer the function before transferring the form** and begin with the minimum sufficient mechanism set capable of addressing the material problem. Technology, including increasingly capable retrieval and navigation tools, can support these functions but should not be treated as a substitute for contextual understanding or legitimate institutional judgement.

Implementation must also distinguish the presence of an intervention from the existence of capability. Adoption does not establish deployment; deployment does not establish mechanism activation; and an activated mechanism does not by itself demonstrate capability gain. Institutions should verify that the mechanism actually performs its intended function and that this function becomes sufficiently dependable across the cases for which it was designed. Institutionalisation

should preserve the required function rather than make the original intervention form permanently resistant to change.

Stronger memory also requires safeguards against the opposite of institutional forgetting: excessive dependence on the past. **Retrievability is not the same as relevance, validity or authority.** Previous knowledge should remain challengeable, its changing status should be visible where material, and historical trace should be distinguishable from current operational authority. Successful reuse can therefore include applying previous knowledge, adapting or combining it with newer knowledge, qualifying it, rejecting it with reason or superseding it. Responsible non-use can be evidence of strong institutional memory where prior knowledge has been recovered, understood and deliberately judged not to apply.

Evaluation should consequently distinguish **activity, function, capability and outcome.** Repository use, completed records or workflow interactions can show activity. Mechanism performance can show that an intended function operates. Capability gain requires evidence that responsible reuse has become more reliable across relevant cases and over time. Wider effects — such as reduced reconstruction, stronger decision continuity or improved preparedness — may provide additional evidence but require proportionate causal interpretation.

Capability gain must also be assessed against burden and negative effects. Documentation, maintenance, governance, workflow requirements, technology and increased information exposure all consume institutional capacity and can generate cognitive load, procedural compliance, maintenance debt or precedent bias. Institutions should therefore remain able to continue, adapt, broaden, de-escalate, replace or retire memory mechanisms as evidence changes. **More memory is not inherently better memory.**

The evidence supports this capability-oriented and conditional approach more strongly than any universal institutional model. Specific intervention forms will depend upon institutional setting, materiality, readiness and the characteristics of the knowledge involved. The framework should therefore guide diagnosis, mechanism selection, implementation and evaluation rather than be treated as a standardised maturity model or universally validated measurement instrument.

Institutional memory is not the ability to preserve the greatest possible amount of past knowledge. It is the ability to recover the right prior knowledge when relevant, reconstruct enough context to judge it, bring it into present institutional work and decide responsibly whether to apply, adapt, combine, qualify, reject or supersede it.

David González
Director
IDHUS Institute

I. The Problem of Stored but Unused Knowledge

Public institutions accumulate knowledge continuously. Policies generate analysis, programmes produce evaluations, operational processes create records, regulatory decisions establish reasoning, crises generate lessons, and successive generations of officials leave behind reports, datasets, guidance and institutional precedent. Much of this material survives in document-management systems, shared drives, archives, case-management platforms, specialised databases, knowledge repositories and personal working environments. Yet the continued existence of this knowledge does not mean that the institution can use it when a later decision, programme or operational problem makes it relevant.

This distinction is the starting point for institutional memory. An institution may preserve extensive evidence of what it has previously known and done while remaining functionally unable to recover that knowledge at the point of need. A team can commission analysis substantially similar to earlier work because no one knows that the previous analysis exists. A programme can encounter an implementation problem already experienced elsewhere because the relevant lesson remains confined to another unit. A decision-maker can retrieve an earlier decision but be unable to reconstruct why it was made. An evaluation can identify lessons that remain technically available yet never influence the design of subsequent work. In each case, knowledge has survived while its capacity to inform future institutional action has not.

The problem is therefore not primarily one of preservation. **It is one of reuse.** The relevant institutional question is not simply whether knowledge has been retained, but whether a reliable route exists through which materially relevant prior knowledge can re-enter future institutional action. That route requires more than storage or technical access. Knowledge must be recoverable, interpretable, recognisable as potentially relevant, capable of entering the institutional process where it could matter, and subject to informed judgement about what role it should play.

Retention is not institutional memory capability

Retention is a necessary condition for reuse: knowledge that has disappeared cannot normally be recovered later. But retention alone says little about an institution's practical ability to remember. A report may remain intact for decades while becoming functionally invisible because future users do not know that it exists, cannot find it using the language of their present problem, lack access to the system in which it is stored or depend upon colleagues with unusual knowledge of the organisation's information environment (Walsh & Ungson, 1991; Stein, 1995).

Increasing the volume retained does not necessarily correct these weaknesses. It can even make selection more difficult if additional material enters an environment in which future users already struggle to distinguish what is relevant. Repository size, archive completeness and document volume are therefore weak measures of institutional memory capability. They describe what the institution possesses, not what it can reliably recover and use.

The distinction matters because institutional knowledge often survives organisationally without remaining functionally available. A document can exist in a formal system while disappearing from practical institutional awareness. A lesson can remain recorded while becoming disconnected from the people and processes for which it would later be relevant. A previous rationale can remain accessible while losing the contextual information required to understand it. **Preservation protects the possibility of future memory; it does not guarantee that memory will function.**

Access is not reuse

The same problem persists even when retained knowledge can be technically accessed. Finding a document is not equivalent to recovering usable institutional knowledge from it. Future users may need to understand why it was produced, which problem it addressed, what evidence informed it, which assumptions shaped the analysis, what constraints affected the decision and whether later developments altered its standing. Without sufficient context, technically accessible knowledge can remain difficult to interpret or easy to misapply.

Nor does interpretation guarantee institutional influence. Relevant prior knowledge may be both recoverable and understandable yet remain outside the decision, programme or operational process in which it could matter. If its use depends entirely upon an individual remembering to search, knowing where to look and recognising that an earlier institutional experience is relevant, the institution still lacks a dependable route from accumulated knowledge to present work.

Institutional memory therefore requires more than access. It requires a sequence of functions through which prior knowledge can move from passive retention towards informed consideration in later institutional action.

Institutional knowledge must be able to travel through time

This requirement becomes particularly important because the people, structures and systems surrounding institutional knowledge change. Officials move between roles or leave organisations. Teams are reorganised. Responsibilities migrate between departments. Information systems are replaced. Terminology changes. Programmes close and later initiatives confront similar problems under different names. What appears easily recoverable to the people who created a body of knowledge may become difficult to locate or interpret several years later.

Institutions often compensate through human and social memory. Experienced officials know whom to ask, where older material is kept, which previous initiative resembles the present problem or why an apparently authoritative document should be treated cautiously. These forms of memory are valuable and should not be treated as defects simply because they are informal. They become fragile, however, when the institution's ability to recover important knowledge depends excessively on the continued presence, availability or personal networks of particular individuals (Olivera, 2000).

The institutional challenge is therefore not to replace human memory with repositories or technology. It is to ensure that materially important knowledge can survive changes in personnel, structure and systems without requiring future actors to reconstruct the institution's past from accidental personal continuity. Institutional knowledge must be capable of travelling through time in forms that remain sufficiently recoverable and intelligible to later users.

Memory becomes valuable when it can re-enter action

The functional threshold for institutional memory is reached when retained knowledge can do more than survive. Materially relevant prior knowledge must be capable of being recovered when needed, interpreted with sufficient context, recognised as relevant to the present problem and brought into the institutional work in which it can influence judgement.

Even then, memory does not determine what the institution should do. Prior knowledge may remain useful, require adaptation or have ceased to apply. The function of institutional memory is to make relevant accumulated knowledge available to present judgement under conditions in which its significance can be assessed. Reuse therefore means informed consideration, not automatic reproduction of previous decisions or practices.

This distinction separates institutional memory from a model in which the past becomes instruction. A capable institution does not merely reproduce what it remembers. It creates a reliable informational connection between previous institutional experience and present

institutional cognition while preserving the capacity to reach a different conclusion where circumstances justify it.

The problem is institutional, not merely informational

Weak institutional memory is often visible through information problems: fragmented repositories, inconsistent metadata, difficult search, missing documentation or inaccessible records. These conditions matter, but the capability problem extends beyond them. Knowledge that can be found may still lack context. Knowledge that can be understood may still fail to enter the relevant workflow. Knowledge that enters a workflow may be given inappropriate weight. The weakness can therefore lie in information architecture, organisational routines, ownership, governance, professional practice or the relationships between them (Stark et al., 2026; Kankanhalli et al., 2005).

Technology can strengthen some parts of this pathway. Search, classification, cross-repository navigation and automated surfacing can make retained knowledge easier to recover. But improved retrieval does not by itself determine whether what has been found is sufficiently contextualised, presently relevant or appropriate to the institutional decision being made. The object of reform is consequently not the technology through which memory operates, but the institutional capability that the technology or other mechanisms may help to support.

This broader perspective also prevents institutional memory from becoming a generic knowledge-management ambition. The objective is not to make every piece of retained information permanently prominent or reusable. Some knowledge has little future material value. Some can reasonably be reconstructed if required. Some deserves historical preservation without continuing operational prominence. The relevant concern is whether failures to recover and use particular prior knowledge create consequences significant enough to matter institutionally.

From accumulated knowledge to reusable institutional memory

Institutional memory can therefore be understood as a repeatable capacity to reconnect materially relevant prior knowledge with present institutional work. Retention provides the informational base, but functional memory requires a route through which that knowledge can be recovered, interpreted, recognised as relevant and brought into circumstances where authorised institutional actors can judge what significance it should have.

This framing changes the practical question facing institutional leaders. Instead of asking how much knowledge the organisation stores, or whether it possesses a sufficiently sophisticated repository, they can ask whether knowledge that may materially affect future action can actually reach that action when required. The answer will not always justify intervention. Nor will every failure require the same response. But it establishes the capability problem that must be diagnosed before any response can be designed.

The next question is therefore more precise: **if retained knowledge exists but still fails to become usable institutional memory, where does the route from retention to future institutional action actually break?**

II. Where Institutional Memory Breaks

Once institutional memory is understood as a route from retained knowledge to future institutional action, the problem becomes easier to diagnose. Knowledge can survive and still fail at several different points before it becomes a usable input to present judgement. It may remain undiscoverable despite being retained. It may be found but no longer carry enough context to be interpreted responsibly. It may be both findable and understandable yet fail to enter the institutional process in which it could matter. Or it may enter that process and receive inappropriate weight because retrieval is confused with continuing relevance or authority.

These are different institutional failures. They can coexist, and weakness at one point can amplify weakness at another, but they should not be treated as interchangeable. A new repository will not necessarily correct missing context. Better contextual documentation will not ensure that prior knowledge enters future workflows. Stronger workflow integration can become counterproductive if what is surfaced is obsolete or poorly understood. Effective intervention therefore depends first upon identifying where the materially relevant break occurs.

Retention without discovery

The first failure arises when relevant knowledge has been retained but intended future users cannot reliably discover it. **A document can remain perfectly preserved and yet be functionally absent from the institution that produced it.** It may exist in a repository that users do not know to search, be described using terminology that no longer corresponds to the way later problems are framed, remain inaccessible across organisational boundaries or depend upon informal knowledge of where particular material is kept (Kankanhalli et al., 2005; Markus, 2001).

This is not simply a question of whether a search function exists. Discovery depends upon the route between a future institutional need and the knowledge that could help address it. Users may know that earlier work probably exists but not where it is located. They may search one system while relevant material sits in another. They may retrieve documents matching familiar keywords while missing more relevant material described differently. Formal access may exist while practical navigation remains dependent on experienced colleagues who understand the institution's information landscape.

Such dependence on informal navigation is not inherently problematic. Colleagues often provide context and judgement that formal systems cannot reproduce. The weakness appears when materially relevant knowledge can be recovered reliably only by people who already possess unusual knowledge of where to look or whom to ask. Personnel changes, organisational restructuring and the passage of time then increase the probability that retained knowledge becomes functionally invisible.

The capability issue is therefore **findability rather than storage volume**. The relevant question is whether intended reusers have a sufficiently reliable route to potentially relevant prior knowledge when a future need arises. Increasing retention without improving that route can enlarge the informational environment while leaving the underlying memory failure unchanged.

Retrieval without meaning

Discovery does not complete the transition from storage to usable memory. A second failure occurs when users can retrieve prior knowledge but cannot reconstruct enough of the circumstances surrounding it to understand what it means or judge whether it applies.

Institutional documents are produced for particular purposes. Analyses respond to specific questions. Decisions are made under particular legal, political, financial, evidential and

operational conditions. Programme lessons emerge from particular implementation environments. Guidance may reflect assumptions that were reasonable when it was issued but later changed. When the contextual information required to reconstruct those conditions is lost, the document may remain readable while its institutional meaning becomes uncertain.

The relevant context varies according to the knowledge involved. Future users may need to know why a document was produced, what decision it supported, which evidence informed it, which assumptions shaped the analysis, who had responsibility for the judgement, what constraints applied or whether subsequent developments changed its standing. The objective is not to preserve every possible contextual detail. Excessive documentation can itself create burden and make useful information harder to distinguish from background material. The requirement is **minimum sufficient context**: enough information for a plausible future user to reconstruct what the retained knowledge represents and to make an informed judgement about its possible applicability (Markus, 2001; UK Resilience Academy, 2025).

This failure is particularly important because improved retrieval can expose it rather than solve it. A stronger search system may make more historical material visible while simultaneously revealing how little information remains about why that material was produced or how it should be interpreted. Retrieval technology cannot reconstruct context that the institution never preserved. Where meaning has become detached from the retained object, making that object easier to find does not necessarily make institutional memory stronger.

Understanding without invocation

A third failure appears when relevant knowledge can be found and understood but still does not enter the later institutional work in which it could influence action. The knowledge is available, yet its reuse depends upon someone independently remembering that earlier experience may exist, deciding to search for it and doing so while the present decision remains open to influence.


This problem is common where memory mechanisms are predominantly passive. Reports, evaluations, lessons and previous decisions may be stored competently and remain understandable to anyone who retrieves them, but no connection exists between those resources and recurring institutional processes. A programme team can design a new intervention without being prompted to consider earlier implementation evidence. A policy review can proceed without relevant previous analysis entering the process. An operational unit can confront a recurring problem while lessons from another part of the organisation remain outside its normal workflow (U.S. Government Accountability Office, 2002).

Timing matters. Prior knowledge that is discovered after a decision has effectively been made may contribute to retrospective learning without having influenced the judgement for which it was most relevant. Similarly, a general instruction to “consult previous lessons” may create little practical effect if no role, trigger or checkpoint connects that expectation to the point at which the knowledge can still alter the work.

The capability problem is therefore not simply availability but **workflow re-entry**. Relevant prior knowledge needs a proportionate route into the institutional processes where it can matter. That does not imply that every workflow should contain mandatory searches or that all prior knowledge should be surfaced automatically. Excessive invocation can create noise, delay and procedural compliance. The objective is to reduce unreasonable dependence on accidental recollection where the consequence of missing relevant prior knowledge is material.

Figure 1. Where Institutional Memory Breaks

Four failure points in the transition from retained knowledge to responsible institutional reuse.



Invocation without appropriate use

Connecting prior knowledge to institutional work introduces a fourth and more subtle failure. An institution can become better at finding, interpreting and invoking accumulated knowledge while becoming worse at judging what significance that knowledge deserves. Earlier decisions can harden into precedent even where circumstances have changed. Guidance can retain operational prominence after its evidential or policy basis has weakened. Familiar interpretations can narrow the range of alternatives considered. A memory arrangement designed to prevent forgetting can therefore create a different problem: an institution that remembers too rigidly (Kluge & Gronau, 2018; Njiru et al., 2025).

The fact that prior knowledge has been recovered does not establish that it should determine present action. Its relevance may depend upon whether the present case sufficiently resembles the circumstances in which the knowledge was produced. Its standing may have changed as evidence, law, policy objectives or operating conditions evolved. Even knowledge that remains relevant can require adaptation rather than direct application.

This creates an essential boundary for institutional memory. **Reuse is not compliance with the past.** A capable institution may recover an earlier analysis and decide that the present circumstances justify a different conclusion. It may identify a precedent in order to understand why it no longer applies, combine earlier experience with newer evidence or determine that previous guidance has been superseded. Informed rejection or qualification can therefore represent successful institutional memory rather than failure.

The capability problem at this stage is one of **appropriate use and independent judgement**. Prior knowledge must be capable of influencing present action without becoming authoritative merely because it is documented, familiar or easy to retrieve. The more effectively an institution brings its past into current work, the more important it becomes that users retain practical capacity to question what has been surfaced and determine what role, if any, it should play. The fuller distinction between retrievability, relevance, validity and authority will become important later in the paper when the safeguards required for responsible memory are considered.

Diagnosing the break before prescribing the response

The four failures form a connected sequence, but institutions should not assume that every memory problem contains all four or that each stage needs strengthening simultaneously. A discovery problem does not automatically require extensive contextual redesign. A workflow problem can persist even where retrieval systems perform well. Stronger workflow integration can be counterproductive if the knowledge being surfaced lacks sufficient context or has become

obsolete. Conversely, sophisticated challenge mechanisms cannot compensate for materially relevant knowledge that intended users cannot locate.

Observable symptoms can therefore be misleading. Repeated work, inconsistent decisions, failure to use previous lessons or heavy dependence on individual recollection may arise from different weaknesses in the route between retention and responsible reuse. More than one failure may also be present, while weaknesses in ownership, stewardship or knowledge quality can cut across the sequence. Intervention based on the visible symptom alone risks strengthening a function that was not responsible for the material problem.

The first task is consequently diagnostic. Institutions need to establish whether materially relevant knowledge cannot be found; can be found but not adequately understood; can be understood but does not enter the appropriate workflow; or enters that workflow without sufficiently independent judgement about its present significance. This diagnosis should precede preferences for a new repository, documentation requirement, workflow procedure or technological layer.

Identifying the break, however, does not yet establish that institutional intervention is warranted. Some memory failures create little more than manageable friction; others generate repeated costs, weaken continuity or expose the institution to significant future consequences. **The next question is therefore not yet how to correct the failure, but when a diagnosed memory failure becomes material enough to justify an institutional response.**

III. When Memory Failure Becomes an Institutional Problem

Not every difficulty in recovering or reusing prior knowledge constitutes an institutional problem requiring intervention. Large organisations inevitably contain informational friction. Documents take time to locate, terminology changes, responsibilities move between units, and some knowledge becomes difficult to reconstruct because it was created for a purpose that no longer exists. Attempting to eliminate every such friction would be neither realistic nor desirable. Institutional memory itself has costs: knowledge must be captured, organised, maintained, contextualised, surfaced and periodically reviewed, while the people expected to use it must devote attention to material produced in the past rather than exclusively to the problem immediately before them.

The relevant policy question is therefore not whether an institution can make all retained knowledge permanently accessible and reusable. It is whether a particular failure of retrieval or reuse is consequential enough, recurrent enough or associated with sufficient future exposure that a proportionate institutional response is justified. Memory failure can range from minor inconvenience to significant capability weakness, and intervention should reflect that difference.

The cost of reconstructing what the institution already knows

The most visible consequence of weak reuse is often unnecessary reconstruction. When previous analysis, experience or decisions cannot be recovered, institutional actors may repeat searches, reconstruct chronologies, rediscover earlier constraints, commission analysis substantially similar to work already undertaken or rebuild an understanding of why previous choices were made. Some repetition is legitimate. Circumstances change, new evidence emerges and earlier analysis may need to be reconsidered. The problem arises when reconstruction occurs not because fresh judgement is required, but because the institution cannot effectively recover what it has already learned or established.

Repeated reconstruction consumes more than time. It absorbs institutional attention that could otherwise be used to analyse what is genuinely new, test assumptions, engage stakeholders, improve implementation or address emerging problems. The effect can become cumulative in organisations undertaking recurring policy, regulatory, programme or operational activities. Individual episodes may appear too small to attract senior attention while their aggregate burden remains substantial. In this sense, memory failure can operate as a **hidden tax on institutional capacity**: resources are repeatedly spent recovering an informational position the institution had previously reached but could not retain in sufficiently reusable form (U.S. Government Accountability Office, 2002; UK Government, 2022).

That does not mean every instance of duplication warrants a formal response. In some settings, reconstructing knowledge occasionally may be cheaper than maintaining mechanisms designed to preserve its future usability. The institutional concern arises when the expected cost or consequence of repeated reconstruction becomes material relative to the burden of a proportionate memory response.

Repeated error and decision discontinuity

The consequences become more significant when memory failure prevents institutions from recognising problems they have encountered before. Implementation weaknesses can recur because earlier lessons are not recovered. Known constraints can be rediscovered only after resources have already been committed. Risks identified in one cycle can disappear from practical institutional awareness before the next. In such circumstances, weak memory does not simply create duplicated work; it can reduce the institution's ability to learn cumulatively from its own experience.

A related problem is decision discontinuity. Changes in personnel, leadership, organisational structure or programme ownership can interrupt the transmission of reasoning that connected earlier choices to their circumstances. Later actors may know what was decided without knowing why, or may encounter only the formal outcome after the alternatives, assumptions and trade-offs that shaped it have disappeared from view. The result can be unnecessary reversal, unexplained inconsistency or repeated debate over questions the institution has already considered.

Continuity should not be confused with preservation of the previous decision. New officials may legitimately reach different conclusions, and changing evidence, law, political priorities or institutional conditions can justify a new direction. The value of memory lies in allowing that change to occur **with awareness of what preceded it**, rather than forcing each generation of institutional actors to begin from an artificially blank informational position. Decision continuity therefore means continuity of informed reasoning, not automatic continuity of outcome.

Coordination and implementation consequences

Memory failure can also become material across organisational boundaries. Knowledge generated in one unit may remain invisible to another facing a similar implementation problem. Lessons from a pilot may not reach the programme responsible for subsequent scale-up. Operational experience may remain separated from policy design. A regulatory or administrative interpretation developed in one part of an institution may be reconstructed independently elsewhere.

These failures can produce inconsistent practice, repeated implementation problems and avoidable coordination costs even when each organisational unit retains its own records effectively. The relevant weakness lies in the institution's inability to reconnect knowledge with later work across the boundaries through which responsibility has moved or similar problems have reappeared.

This does not imply that stronger institutional memory requires centralising all knowledge. Distributed institutions may have good reasons to retain specialised repositories, professional communities and locally governed information. The capability requirement is narrower: where knowledge generated in one part of the organisation is materially relevant to subsequent work elsewhere, institutional arrangements should not make its recovery depend unreasonably on chance, personal acquaintance or retrospective discovery.

Preparedness and low-frequency knowledge

Frequency alone is an inadequate measure of the institutional value of memory. Some knowledge may be required rarely but become highly consequential when the relevant situation occurs. Crisis procedures, experience from unusual regulatory cases, responses to infrastructure failure, lessons from emergency procurement or knowledge generated during previous high-impact disruptions may remain unused for long periods while retaining considerable potential value (UK Covid-19 Inquiry, 2024).

This creates a particular challenge for institutional memory because low observed use can be mistaken for low relevance. Repository views, search frequency or reuse counts may indicate how often knowledge enters ordinary work, but they cannot establish how important that knowledge would be under conditions that occur infrequently. A memory arrangement designed only around frequently accessed material may therefore perform well during routine operations while leaving the institution poorly prepared for exceptional but consequential situations.

The materiality question should consequently include **exposure**, not only recurrence. Where the consequence of being unable to recover particular knowledge would be substantial, the expected value of maintaining that knowledge in reusable form may remain high even if direct use is rare. This does not justify maximal preservation of everything that might conceivably become useful. It means that institutional leaders should not infer low memory value simply from low historical frequency.

Accountability and the reconstruction of institutional reasoning

A further consequence arises where institutions need to explain not only what they did but why. Formal records may establish that a decision occurred, yet provide insufficient basis for reconstructing the reasoning, evidence, assumptions or alternatives that informed it. This can weaken internal review, external accountability and the institution's own ability to understand how previous judgement developed (Williams, 2020).

Institutional memory therefore has an explanatory dimension. Retaining sufficient context around materially important decisions can allow later actors to distinguish between a conclusion that remains applicable, one that depended upon conditions that have changed and one that should be reconsidered in light of new evidence. The same capacity can make legitimate departure from precedent easier to explain because the institution can identify what has changed rather than treating the earlier decision as either binding or irrelevant.

The objective is not comprehensive documentation of every deliberation. Excessive capture can impose substantial burden and may itself make significant reasoning harder to recover. The relevant question remains proportional: which decisions, experiences or knowledge domains create enough institutional consequence or accountability exposure that future reconstruction of their reasoning has material value?

Materiality depends on consequence, recurrence and exposure

Taken together, these consequences show why memory failure should not be judged through a single measure. A weakness may be material because it occurs repeatedly, because each occurrence imposes significant cost, because failure could produce severe consequences even if rare, or because the institution lacks reasonable alternatives for reconstructing the required knowledge when needed.

Materiality therefore depends upon the interaction between the knowledge involved, the institutional process in which it may later matter, the consequence of its absence, the probability or recurrence of future need, and the cost of compensating when memory fails. Different combinations can produce different intervention thresholds. A frequently repeated but low-cost retrieval problem may justify a simple local correction. A rare but high-consequence memory failure may justify stronger preparedness arrangements. Other weaknesses may remain tolerable because informal practices already compensate adequately or because formal correction would create more burden than value.

This is why institutional memory should not be strengthened uniformly. Some knowledge deserves more durable findability, context, stewardship or workflow connection than other knowledge. Some processes justify explicit memory mechanisms while others can continue to rely on lighter arrangements. **The objective is not maximum memory, but sufficient memory capability where failure materially affects institutional action.**

Establishing the intervention threshold

A diagnosed memory failure becomes a defensible candidate for institutional intervention when three conditions are present.

First, the failure is **materially consequential or creates sufficiently significant exposure**. The institution should be able to identify what is at stake if relevant prior knowledge cannot be recovered or responsibly reused, rather than assuming that retrieval difficulty is intrinsically important.

Second, the **mechanism producing the failure is understood sufficiently well** to distinguish the actual capability weakness from adjacent problems. Repeated reinvention may reflect poor discovery, but it may instead result from knowledge never having been documented adequately. Apparent retrieval failure may actually be contextual failure. Relevant knowledge may already be findable and understandable but remain disconnected from workflow. Stronger retrieval may even create new risks where obsolete material is surfaced without adequate

challengeability. Diagnosis must therefore be precise enough to identify the function that is genuinely missing.

Third, there must be a **plausible proportionate response whose expected institutional value exceeds the burden of introducing and maintaining it**. A real capability weakness does not automatically justify formal intervention. Existing informal practices may be adequate, cases may be too heterogeneous for a standard mechanism, correction may cost more than occasional reconstruction, or the institution may lack the conditions required to sustain the proposed response. Intervention can also be justified for a bounded process or knowledge domain without implying organisation-wide reform.

These conditions establish a deliberately demanding threshold between recognising a weakness and prescribing a solution:

A material reuse failure should justify institutional intervention only where its consequences or exposure warrant action, its mechanism is sufficiently understood, and a proportionate response offers plausible value greater than the burden it creates.

The practical starting point is therefore assessment rather than presumption. Institutions should identify the reuse failure, determine why it matters, establish which part of the retention-to-reuse pathway is responsible and judge whether correction is likely to create sufficient institutional value. This preserves an important discipline: institutional memory should be strengthened because doing so improves the institution's capacity to act, not because additional memory infrastructure is assumed to be desirable in itself.

Once that threshold has been crossed, the problem changes. The institution no longer needs to ask whether a memory weakness warrants action. It can ask a more specific design question: **which institutional function must be strengthened to correct the diagnosed failure, and what is the least burdensome arrangement capable of providing it?**

IV. Designing Institutional Memory for Reuse

Once a material memory failure has been identified and the case for intervention established, the central design question is not which knowledge-management system an institution should acquire or which repository model it should adopt. It is **which institutional function is failing and what proportionate combination of mechanisms could restore it**. If materially relevant retained knowledge cannot reliably be found, the institution needs stronger findability. If knowledge can be retrieved but cannot be adequately interpreted, it needs stronger contextual interpretability. If relevant knowledge remains outside the work in which it could matter, it needs a more reliable route back into institutional workflows. If prior knowledge enters those workflows but receives inappropriate weight, it needs stronger challenge, renewal and applicability judgement. Across these situations, ownership, stewardship and evaluation may also need strengthening to keep the resulting capability viable over time.

This function-first approach prevents institutional memory reform from beginning with a preferred instrument. Search platforms, knowledge portals, metadata arrangements, decision records, workflow checkpoints and automated retrieval capabilities can all contribute to institutional memory, but none is intrinsically a memory solution. **The value of any such mechanism depends upon whether it addresses the failure that actually matters under the conditions in which the institution operates**. The same mechanism can strengthen capability in one setting, add little value in another and create unnecessary burden in a third.

Intervention design should therefore proceed from **diagnosed failure** → **required function** → **appropriate mechanism**. The practical consequence is a selective rather than comprehensive approach to reform. Institutions do not need to build an organisation-wide memory system before they can strengthen a particular memory capability, and they should not assume that every material weakness requires intervention across the entire retention-to-reuse pathway.

Strengthen findability where discovery is the binding failure

Where materially relevant knowledge exists but intended users cannot reliably discover it, the intervention objective is to improve the route between a future institutional need and the knowledge already retained. Depending on the cause of failure, this may involve clearer knowledge-location structures, more usable classification or metadata, improved search and navigation, better visibility across repositories, appropriate access arrangements or reduced dependence on individuals who happen to know where material is kept.

The design target is not simply faster search. It is greater reliability that potentially relevant prior knowledge can be recovered by intended reusers without requiring unusual familiarity with the institution's information environment. The appropriate mechanism therefore depends upon what prevents discovery. Greater search sophistication inside one repository will add little if users do not know which repository contains the relevant material. Better metadata may matter more where future users are unlikely to employ the terminology used by the original authors. Cross-repository visibility may be important where organisational boundaries fragment access. Human and social memory can remain valuable where experienced colleagues provide navigation and contextual knowledge that formal systems do not yet reproduce (Kankanhalli et al., 2005; U.S. Government Accountability Office, 2002).

Technology can materially strengthen this function. Improved indexing, federated search, semantic retrieval and other forms of automated navigation can increase the probability that relevant retained knowledge becomes visible. But technological sophistication does not remove the underlying conditions of successful retrieval. Content quality, permissions, classification,

interoperability and continuing maintenance still shape what users can recover. The intervention objective remains reliable findability rather than adoption of a particular retrieval technology.

Preserve sufficient context where interpretation is the problem

Where knowledge can be recovered but later users cannot reconstruct what it means or judge its possible applicability, intervention should strengthen contextual interpretability. The design question becomes what information must accompany materially relevant knowledge so that plausible future users can understand the conditions under which it was produced and used.

That context may include purpose, provenance, decision rationale, relevant assumptions, temporal conditions, evidence used, institutional constraints or information about subsequent changes in status. The appropriate combination will vary. A major regulatory decision may require a more durable record of rationale and evidential basis than a routine operational note. Knowledge expected to remain relevant across personnel transitions may justify stronger contextual preservation than material whose future use is unlikely or whose meaning is obvious from the record itself.

The governing principle should therefore be **minimum sufficient context rather than maximal capture**. Attempting to preserve every contextual detail increases documentation and maintenance burden and can make important information harder to distinguish from background material. Too little context makes later interpretation unreliable; too much can make the memory arrangement unsustainable. The design task is to preserve enough for responsible future judgement in proportion to the materiality of the knowledge involved (Markus, 2001; UK Resilience Academy, 2025).

This principle also establishes a limit to what retrieval technology can accomplish. Better search can expose more retained material, and automated tools may help users navigate or summarise it, but they cannot reliably reconstruct institutional context that was never captured or has become detached from the underlying knowledge. Where contextual loss is the binding failure, the intervention must address the conditions under which meaning survives, not merely the speed with which documents can be retrieved.

Connect prior knowledge to work where it can matter

Where relevant knowledge can be found and understood but remains outside subsequent institutional work, the design objective is workflow re-entry. The institution needs a proportionate mechanism through which relevant prior knowledge can reach the decision, programme, review or operational process while judgement remains actionable.

Possible arrangements vary widely. A recurring process may include a checkpoint for relevant previous decisions or lessons. Responsibility for surfacing prior knowledge may be attached to an existing role. Particular triggers may prompt consultation of previous evidence when defined conditions arise. In other settings, communities of practice or structured handovers may provide a more appropriate route than formal workflow integration.

The objective is not to require a search before every decision or to place historical material automatically in front of every user. Such arrangements can generate noise, delay and procedural compliance without improving judgement. Workflow integration should be selective and materiality-sensitive. It is justified where reliance on individual recollection creates a significant probability that relevant knowledge will otherwise remain absent from work in which it could affect the result.

Invocation also remains distinct from obligation. Bringing prior knowledge into a workflow creates an opportunity for informed consideration; it does not establish that the knowledge should be followed. This boundary becomes increasingly important as institutions make memory mechanisms more proactive (U.S. Government Accountability Office, 2002).

Establish proportionate ownership, stewardship and renewal

Retrieval, context and workflow mechanisms deteriorate if no one remains responsible for the functions they are expected to perform. Search structures become outdated, contextual records lose meaning, workflow triggers cease to correspond to actual practice and obsolete material can retain prominence long after its institutional standing has changed.

Material memory interventions therefore require proportionate ownership and stewardship. This does not necessarily mean creating a new knowledge-management unit or assigning a dedicated owner to every item of retained knowledge. Ownership can reside within existing policy, operational, records, programme or professional responsibilities. What matters is that responsibility for maintaining the relevant memory function is identifiable rather than assumed to exist somewhere within the organisation.

Stewardship includes more than preservation. It can involve maintaining findability, protecting necessary context, reviewing status, identifying supersession, adjusting workflow connections and deciding when material no longer warrants the same level of prominence or maintenance. In this sense, institutional memory requires the capacity not only to retain knowledge but also to govern how that knowledge changes over time (Zyngier & Burstein, 2012).

The required intensity should follow materiality. High-consequence knowledge may justify clearer ownership and more deliberate renewal, while lower-value material may reasonably receive lighter treatment. Uniform stewardship across all retained knowledge would recreate the same problem as maximal capture: institutional capacity would be consumed maintaining memory irrespective of whether future reuse justifies the burden.

Preserve challengeability and independent judgement


As retrieval and workflow mechanisms become stronger, a further design requirement becomes more important: prior knowledge must remain challengeable. Greater visibility can increase the practical influence of earlier decisions, established guidance and well-documented interpretations. Without safeguards, an intervention designed to reduce institutional forgetting can unintentionally strengthen precedent dependence.

Memory architecture should therefore preserve the distinction between making knowledge available and granting it authority. Users need sufficient information to understand the status of what they retrieve, recognise where material has been superseded or qualified and question its applicability where present circumstances differ. Historical trace may need to remain available even when operational authority changes, allowing the institution to understand previous reasoning without presenting it as current instruction.

The objective is not to weaken memory by encouraging routine disregard of previous knowledge. It is to ensure that stronger memory improves the informational foundation of present judgement rather than substituting for that judgement. The fuller responsible-memory architecture is developed in Section VI; at the design stage, the essential requirement is that challenge and renewal remain possible wherever stronger retrieval or invocation could otherwise give the past excessive authority.

Figure 2. From Diagnosed Failure to Institutional Intervention

Function-first selection of institutional-memory mechanisms.



Configure interventions rather than copy solutions

These functions can be supported through different combinations of mechanisms. Retrieval architecture, context-preservation arrangements, human and social memory, workflow integration, governance and stewardship, and renewal or challenge mechanisms together form a **design space**, not a mandatory package (Argote & Guo, 2016; Boh, 2008; Olivera, 2000).

This distinction matters because institutional forms do not transfer mechanically. A retrieval arrangement that performs well in a large central administration may be unnecessarily complex for a smaller agency. A formal decision-recording process appropriate to a regulated high-consequence domain may create excessive burden in a rapidly changing operational setting. A technology-intensive solution may perform poorly where underlying information quality or maintenance capacity is weak, while a relatively simple human or procedural mechanism may provide the required function more reliably.

The appropriate transfer principle is therefore:

Transfer the function before transferring the form.

Institutions can learn from mechanisms used elsewhere, but they should first identify the capability function those mechanisms perform and then determine whether the same function is required under their own conditions. The institutional form should follow that assessment rather than be copied because it represents apparent best practice.

This approach also prevents organisational size or technological sophistication from being mistaken for memory capability. A large institution may possess advanced systems and still lack dependable routes from retained knowledge to action. A smaller institution may rely on simpler arrangements that perform the required functions effectively. Capability should be judged by what the institution can reliably do, not by how elaborate its memory infrastructure appears.

Design to readiness and use the minimum sufficient mechanism set

A materially justified intervention can still fail if the conditions required for its operation are absent. Before introducing a mechanism, institutions should therefore establish whether the minimum dependencies necessary for it to function are present. At a minimum, the relevant capability should have identifiable ownership, the knowledge upon which the intervention depends should be of sufficient quality, enough capacity should exist to maintain the arrangement, and there should be a viable route through which the resulting memory can enter institutional work. Other dependencies — including access, interoperability, skills, organisational acceptance or technical infrastructure — will depend upon the intervention (Rashman et al., 2009).

This assessment should not become a generic maturity exercise. The relevant question is narrower: **does this institution, in this domain, currently possess the conditions required for this mechanism to operate as intended?** Where the answer is no, the appropriate first step may be enabling work rather than deployment. Sophisticated retrieval over poorly governed information may increase noise. Workflow checkpoints without functional ownership may create compliance activity. Extensive contextual requirements without maintenance capacity may produce records that quickly become incomplete or obsolete.

Intervention should consequently begin with the **minimum sufficient mechanism set** capable of addressing the diagnosed material failure under existing institutional conditions. This is not a prescription for permanently minimal ambition. It is a discipline against unnecessary complexity. Mechanisms can be expanded, combined or strengthened where evidence shows that the initial configuration is insufficient; they can also be simplified or retired where burden exceeds contribution.

Taken together, the intervention logic produces a conditional portfolio rather than a universal reform programme. Institutions should strengthen findability where discovery is the binding constraint, preserve sufficient context where interpretation is impaired, create proportionate workflow connections where invocation fails, establish ownership and renewal where memory quality deteriorates, and preserve challengeability where stronger access could distort judgement. The configuration should follow materiality, readiness and context rather than an assumed model of what a mature memory architecture should look like.

This is the point at which institutional memory design moves beyond storage without becoming either a technology programme or an implementation toolkit. The intervention is not the repository, search engine, contextual record, workflow trigger or stewardship role considered in isolation. It is the institutional arrangement through which one or more mechanisms restore a diagnosed function in the pathway from retained knowledge to responsible reuse.

Yet even a well-designed arrangement remains only a proposition about how capability might improve. **Deployment does not establish that the intended mechanism is operating, and an operating mechanism does not by itself establish that institutional capability has strengthened.** The next question is therefore how a designed intervention becomes a dependable institutional function rather than merely an adopted initiative.

V. From Intervention to Institutional Function

A well-designed intervention does not become institutional memory capability simply because it has been approved, procured, configured or launched. A retrieval system can be deployed without materially improving discovery. A contextual-recording requirement can be introduced while users continue to lack the information needed to interpret earlier knowledge. A workflow checkpoint can exist formally while relevant prior knowledge still fails to enter decisions. A stewardship role can be assigned without meaningful renewal taking place. The existence of an intervention therefore demonstrates only that an institutional response has been adopted. It does not establish that the mechanism on which that response depends is actually functioning.

Implementation must consequently preserve a series of distinctions:

Adoption is not deployment. Deployment is not mechanism activation. Mechanism activation is not repeated institutional function. Repeated institutional function is not, by itself, capability gain (UK Resilience Academy, 2025; U.S. Government Accountability Office, 2002).

Each transition requires different evidence. Adoption establishes that an intervention has been authorised. Deployment establishes that the arrangement exists in operational practice. Mechanism activation establishes that it is performing the function for which it was introduced. Repeated function establishes that this performance is sufficiently dependable to form part of institutional operation rather than appearing only in isolated cases. Whether that repeated function ultimately makes the institution more capable is a subsequent evaluative question.

This distinction matters because implementation programmes often measure what is easiest to observe. Systems go live, templates are completed, training is delivered, responsibilities are assigned and workflow steps are introduced. These outputs may be necessary, but they are weak evidence of institutional memory if the mechanism expected to connect retained knowledge with later action has not activated. The implementation problem is therefore not simply how to deploy an intervention. It is how to make the intended institutional function observable and dependable.

Prepare the conditions required for the mechanism to work

Implementation begins before deployment. The design stage should already have established whether the intervention has identifiable ownership, sufficiently usable knowledge, maintenance capacity and a viable route into institutional work. Implementation must convert those dependencies from design assumptions into operational conditions.

A retrieval mechanism, for example, depends upon the underlying knowledge being sufficiently accessible and structured for the intended discovery route. A contextual mechanism depends upon actors being able to capture the information required for later interpretation without imposing unsustainable burden. Workflow integration depends upon a real decision or operational process into which prior knowledge can enter while it remains actionable. Stewardship and renewal depend upon someone possessing both responsibility and practical authority to maintain status, correct weaknesses and respond to change.

Where these dependencies are absent, deployment should not be treated as the default next step. Enabling work may be required first. Poor-quality material may need remediation. Ownership may need clarification. Access barriers may need resolution. A proposed workflow connection may need redesign because the intended decision point occurs too late. The relevant implementation question is not whether the institution is ready for institutional memory in general, but whether the conditions required for **this mechanism, in this setting**, are sufficiently present.

This preserves the readiness discipline established in Section IV without turning implementation into a maturity programme. Readiness is a practical test of whether the causal assumptions embedded in the intervention can operate.

Assign functional ownership

Institutional memory mechanisms require ownership that corresponds to the function being performed. Formal responsibility for a system or repository is not necessarily the same as responsibility for whether relevant knowledge can be found, interpreted, brought into work or appropriately renewed.

Functional ownership asks who is responsible for maintaining the connection between intervention and intended effect. For a retrieval mechanism, this may include responsibility for whether intended users can actually recover relevant material rather than merely whether the platform remains technically available. For contextual preservation, ownership concerns whether future users receive sufficient meaning rather than whether a required field has been completed. For workflow integration, responsibility concerns whether relevant prior knowledge reaches the appropriate point in the process rather than whether a procedural step appears in guidance.

Ownership need not be concentrated in a single role. Technical, records, policy, operational and professional responsibilities may legitimately be distributed. What matters is that the function does not disappear between them. Where several actors contribute to a memory mechanism, the institution should be able to identify who can detect functional failure, who can act upon it and who has authority to adapt the arrangement when its assumptions no longer hold.

This is particularly important over time. Institutional memory interventions often outlive the project teams that introduce them. Without durable functional ownership, mechanisms can remain formally present while gradually ceasing to perform the purpose for which they were created.

Configure before scaling

Implementation should normally begin within a bounded context where the mechanism can be configured and tested against the diagnosed failure. The objective is not to prove that a generic model works, but to determine whether the proposed arrangement performs the required function under actual institutional conditions.

Configuration may involve adjusting retrieval structures to the terminology and search behaviour of intended users, calibrating contextual requirements to the knowledge involved, locating workflow triggers at points where they can still influence action or defining how stewardship distinguishes current, qualified and superseded material. Human and social mechanisms may require clarification of roles, participation or handover practices. Technological mechanisms may require adjustment of permissions, data sources, ranking behaviour or interfaces (UK Resilience Academy, 2025).

This stage is important because mechanisms that appear plausible at design level can behave differently in practice. A workflow prompt may occur too frequently and become routine noise. A contextual template may demand information that users cannot reasonably provide. Search may expose large volumes of marginally relevant material while failing to recover the knowledge associated with the diagnosed problem. A renewal process may exist but lack authority to change how outdated material is presented.

These are not necessarily reasons to abandon the intervention. They are evidence that the mechanism needs configuration. Early implementation should therefore create enough reversibility to permit adjustment before the institutional form becomes difficult to change.

Test mechanism activation

The central implementation question is:

Is the mechanism actually doing what it was designed to do?

This is the mechanism activation test.

For interventions addressing findability, activation means that intended users can in practice recover the relevant retained knowledge through the route the intervention was designed to provide. The existence of indexing, search functionality or a knowledge portal does not establish activation if materially relevant knowledge remains effectively undiscoverable.

For contextual interpretability, activation means that retrieved knowledge carries enough information for intended users to reconstruct its meaning and relevant conditions. A completed decision record or contextual field demonstrates activity; activation requires the resulting information to perform the interpretive function for which it was captured.

For workflow re-entry, activation means that relevant prior knowledge reaches the appropriate institutional process while it can still affect judgement. The presence of a formal checkpoint is insufficient if the material arrives too late, is routinely bypassed or is surfaced without a meaningful connection to the work.

For stewardship, challenge and renewal, activation means that the institution can genuinely maintain status, identify changed conditions, qualify or supersede problematic material and preserve practical routes for challenge. A governance arrangement that exists on paper but cannot alter how prior knowledge is treated has not activated the intended mechanism.

These tests should be proportionate to the intervention. A bounded local mechanism does not require an elaborate evaluation architecture simply to determine whether it functions. Direct observation, structured user testing, case review, process evidence or other suitable methods may be sufficient. What matters is that deployment is not used as a proxy for function.

From activation to repeated institutional function

A mechanism that works once has demonstrated possibility, not dependable institutional function. Institutional memory requires a degree of repeatability because the capability problem exists precisely when relevant knowledge needs to survive changes in time, personnel, circumstances or organisational attention (UK Resilience Academy, 2025).

Repeatability does not imply identical performance in every case. Institutional work varies, and some memory mechanisms apply only under defined conditions. Nor does it require large numbers of observations where relevant situations occur infrequently. The requirement is that the mechanism becomes sufficiently dependable across the class of cases for which it was designed that its operation does not rely primarily on exceptional effort, unusual individual knowledge or continued attention from the implementation team.

This distinction is especially important during pilots. A pilot can show that a retrieval route works when closely supported, that contextual information can be captured when participants receive intensive guidance or that a workflow mechanism can function while the project team actively monitors compliance. The implementation question is whether the function persists when those exceptional supports are reduced and ordinary institutional conditions resume.

Evidence of repeated function may therefore include consistency across relevant cases, successful operation after personnel changes, continued use without intensive project support, appropriate handling of exceptions and the ability to detect and correct degradation. The exact evidence should reflect the mechanism and the frequency of the institutional process involved.

The threshold is not permanence. It is **dependability sufficient for the function to become institutional rather than episodic**.

Institutionalise the function, not the project

Once a mechanism has activated and shown sufficient repeatability, the institution can decide how the function should be sustained. This is the point at which implementation may move towards institutionalisation, but institutionalisation should not be confused with making the original intervention permanent (Zyngier & Burstein, 2012; UK Resilience Academy, 2025).

The objective is to preserve the required function. The form through which that function is delivered may legitimately change. A temporary project role may become part of an existing operational responsibility. A manual retrieval process may later be supported by better technology. A formal workflow checkpoint may be simplified once the relevant behaviour becomes embedded in ordinary practice. A pilot governance arrangement may be integrated into an existing review structure.

This follows directly from the function-first principle. If institutionalisation focuses on preserving the intervention itself, organisations can accumulate procedures, systems and governance arrangements long after those forms cease to be the best way of performing the required function. Institutional memory then creates its own maintenance burden.

Institutionalisation should therefore remain **reversible and adaptable**. The institution should be able to modify the mechanism when knowledge environments, workflows, technology or organisational responsibilities change. It should also be able to simplify or retire an arrangement if the original failure diminishes, another mechanism can perform the function more effectively or the burden of the intervention becomes disproportionate.

Permanent expansion is not the measure of implementation success. A mechanism succeeds when the institution can sustain the required function appropriately, not when the mechanism becomes difficult to remove.

Use implementation evidence to adapt the mechanism

Implementation should generate evidence capable of changing implementation. Where a mechanism fails to activate, the institution needs to determine whether the problem lies in configuration, missing dependencies, incorrect diagnosis or the intervention logic itself. Where activation occurs inconsistently, the relevant question is what prevents dependable operation across the cases for which the mechanism was intended.

This creates an implementation loop:

prepare → configure → test → deploy → verify activation → observe repeated function → adapt.

The sequence should not be interpreted as mechanically linear. Testing may expose missing dependencies and return the intervention to preparation. Deployment may reveal that configuration works in one unit but not another. Repeated operation may expose maintenance requirements that were invisible during design. The value of the sequence is analytical: it prevents the institution from treating implementation milestones as evidence that the intended function exists.

Adaptation can therefore include simplifying a mechanism, changing ownership, improving the underlying knowledge, altering a workflow trigger, revising contextual requirements, modifying technology or narrowing the intervention to cases where it provides greater value. Expansion is only one possible response. The objective remains reliable institutional function rather than growth of the intervention.

The boundary between function and capability

Successful implementation establishes something important but deliberately limited. It can show that the institution has moved beyond adopting an intervention and that the mechanism performs its intended function with sufficient repeatability to form part of ordinary institutional operation. That is a necessary step towards stronger institutional memory.

It is not yet sufficient evidence of capability gain.

A retrieval mechanism may reliably improve discovery while the knowledge recovered remains difficult to interpret. A workflow mechanism may consistently surface prior knowledge while users give that knowledge excessive authority. A contextual intervention may function well but impose a maintenance burden greater than the institutional problem justified. Even where

several mechanisms work as designed, wider improvement should not be inferred automatically from their operation.

This boundary protects the causal logic of the paper:

Deployment establishes presence. Activation establishes function. Repeatability establishes dependable institutional function. Capability evaluation must establish whether those functions make responsible reuse more reliable and whether the resulting gain remains proportionate to its burden.

The final question belongs to evaluation, not implementation. Before reaching it, however, one further condition requires closer examination. A memory mechanism can function exactly as intended — knowledge can become easier to find, understand and bring into institutional work — while still weakening institutional judgement if the past acquires excessive authority merely because it has become more visible and easier to invoke.

A functioning memory mechanism therefore needs more than operational reliability. It must preserve the institution's ability to challenge, renew and judge what it remembers.

VI. Responsible Memory: Challenge, Renewal and Judgement

A stronger institutional memory is not necessarily a better institutional memory. Institutions can fail because materially relevant knowledge disappears from practical use, forcing later actors to reconstruct what was previously known or repeat avoidable mistakes. But they can also fail in the opposite direction: prior knowledge can become so visible, persistent or routinely invoked that earlier decisions, assumptions and interpretations acquire more authority than present circumstances justify.

The first failure produces unnecessary forgetting. The second produces uncritical continuity. **Responsible institutional memory must protect against both.**

This becomes increasingly important as memory mechanisms improve. Better retrieval makes previous knowledge easier to recover. Stronger contextual records make earlier reasoning easier to reconstruct. Workflow integration increases the probability that prior knowledge reaches present decisions. Stewardship can make accumulated knowledge more durable across organisational change. Each improvement can strengthen institutional capability, but each can also increase the practical influence of the past. The better an institution becomes at remembering, the more important it becomes that remembering remains compatible with independent judgement.

Institutional memory should therefore be designed not merely to make prior knowledge available, but to make it **available under conditions in which its present significance can be judged.**

Retrievability is not relevance, validity or authority

The central distinction is between **retrievability, relevance, validity and authority** (Kluge & Gronau, 2018; Njiru et al., 2025).

Retrievability asks whether prior knowledge can be recovered. Relevance asks whether it bears materially upon the present problem. Validity asks whether the reasoning, evidence or conclusion remains sufficiently sound under current conditions. Authority asks what institutional weight that knowledge should carry in the present decision or process.

These questions are connected but not interchangeable. The fact that a previous analysis can be recovered does not establish that it concerns the present situation. A relevant analysis may depend upon assumptions that no longer hold. A conclusion that remains analytically valid may still lack authority over a new decision because responsibility, law, policy objectives or institutional circumstances have changed. Conversely, older knowledge can remain materially important even where its formal status has changed, because it helps explain how the institution reached its current position or reveals experience that present actors should consider.

Collapsing these distinctions creates a particular risk as retrieval becomes easier. Material that is well documented, strongly indexed or repeatedly surfaced can acquire practical prominence simply because it is more visible than competing knowledge. Earlier institutional positions can then appear more authoritative than they actually are. What begins as an improvement in discoverability can become a source of retrieval bias or precedent rigidity.

Memory architecture should therefore make status and provenance sufficiently visible to support judgement without pretending that metadata can resolve judgement itself. Users may need to know when material was produced, for what purpose, under which conditions, whether it has subsequently been reviewed or superseded and what institutional standing it currently holds. These signals can inform applicability judgement. They cannot determine it automatically.

Reuse is not compliance

This distinction changes what successful reuse means. Institutional memory is not successful merely when retrieved knowledge is followed. A capable institution may recover a previous decision and apply it because the relevant circumstances remain sufficiently similar. But it may also conclude that the earlier decision requires adaptation, that elements of several previous experiences should be combined, that an earlier conclusion remains useful only with qualification, or that changed circumstances justify rejection or supersession.

Responsible reuse therefore includes the practical capacity to:

apply, adapt, combine, qualify, reject or supersede prior knowledge according to its present relevance, validity and appropriate authority.

None of these responses is intrinsically superior. Direct application is not evidence of stronger memory than reasoned rejection. The relevant capability lies in ensuring that materially relevant prior knowledge enters present judgement and that the institution can decide responsibly what role it should play.

This point is particularly important for evaluation. If an institution measures memory success by how frequently retrieved knowledge is adopted, it can create incentives to reproduce the past rather than consider it. A decision that recovers an earlier analysis, reconstructs its context and deliberately departs from it because material conditions have changed may demonstrate stronger institutional memory than a decision that follows precedent without examining whether it still applies.

Responsible non-use is therefore compatible with successful institutional memory.

What matters is not whether prior knowledge determines the outcome, but whether relevant institutional experience was available to informed judgement.

Prior knowledge must remain challengeable

For this form of reuse to operate, challenge cannot remain merely theoretical. Institutional actors need practical routes through which the status, interpretation or applicability of retained knowledge can be questioned.

Challengeability may require different arrangements depending upon the knowledge involved. Users may need to flag material whose assumptions appear obsolete, identify conflict with newer evidence, question whether a precedent remains applicable or request review of guidance whose institutional status is unclear. In some settings, challenge can occur through ordinary professional judgement. In others, formally governed material may require a defined process through which status can be reviewed or supersession recorded (Njiru et al., 2025).

The purpose is not to encourage constant reopening of settled questions. Stable institutional knowledge has value precisely because actors should not need to reconstruct every conclusion continuously. Challengeability instead protects the boundary between useful continuity and inappropriate persistence. It ensures that stability remains conditional upon the knowledge continuing to deserve the role it occupies.

This is also why stronger memory should not be equated with greater prominence of historical material. Some retained knowledge should remain discoverable for accountability or institutional learning while no longer appearing as current operational guidance. Other material may deserve active surfacing because its continued relevance is high. A responsible architecture can distinguish historical trace from present authority rather than forcing institutions to choose between deletion and continued operational prominence.

Renewal is part of remembering

Institutional knowledge changes even when the underlying document does not. Evidence develops, law changes, programmes evolve, responsibilities move, terminology shifts and new institutional experience modifies how earlier conclusions should be understood. A memory

architecture that preserves knowledge without preserving the possibility of changing its status can therefore make outdated knowledge increasingly easy to retrieve without making it easier to interpret responsibly.

Renewal is the institutional function through which remembered knowledge can be reclassified, qualified, linked to newer material, marked as superseded or otherwise repositioned as circumstances change. The intensity of renewal should remain proportionate. Not every retained item requires continuous review. Material with high operational prominence, significant decision consequences or a substantial risk of becoming misleading may justify more deliberate renewal than knowledge retained primarily for historical or occasional reference.

This creates an important distinction between **preserving the record and preserving its authority**. Historical trace may remain valuable after a decision, guidance document or interpretation ceases to govern current action. Removing outdated material entirely can weaken accountability and institutional learning. Leaving it undifferentiated from current knowledge can mislead future users. Responsible memory therefore needs mechanisms capable of preserving what happened while signalling how its present standing has changed.

Renewal also prevents institutional memory from becoming a one-way accumulation process. Stronger memory capability does not mean continuously adding material to an ever-growing body of equally prominent knowledge. It includes the capacity to reduce prominence, qualify, connect, supersede and, where appropriate, cease maintaining material whose future value no longer justifies the burden (Kluge & Gronau, 2018).

Technology can surface knowledge; it cannot assign legitimate authority

Technology can materially strengthen institutional memory. Search and retrieval tools can make dispersed knowledge easier to discover. Classification and navigation can improve the connection between present questions and earlier material. Automated systems can help surface potentially relevant knowledge at points in institutional workflows where manual discovery would be unreliable.

These capabilities may become increasingly important as institutional information environments grow in scale and complexity. They do not, however, remove the distinctions on which responsible memory depends.

A system can identify similarity without establishing institutional relevance. It can rank material without determining whether the most prominent result remains valid. It can surface a previous decision without deciding what authority that decision should possess over a new one. It can summarise retained material without resolving whether omitted context changes its significance. Even where automated systems become increasingly capable of assisting these judgements, responsibility for legitimate institutional applicability cannot simply be inferred from technical confidence.

This boundary is particularly important because automation can amplify the practical consequences of information architecture. If retrieval systems systematically privilege well-documented, frequently referenced or easily indexed knowledge, they can narrow what institutional actors see. Historical prominence can reinforce itself: knowledge surfaced more often is used more often, and knowledge used more often may become easier to surface again. Without attention to provenance, diversity of relevant sources and challengeability, improved retrieval can therefore create new forms of institutional path dependence (Ruscheimer & Hondrich, 2024).

Technology should consequently strengthen the informational conditions for judgement without acquiring authority it cannot legitimately exercise. **Its role is to improve the institution's capacity to know what it has previously known, not to decide what the institution must continue to believe or do.**

Preserve judgement without returning to individual dependence

Protecting independent judgement does not mean returning institutional memory to informal individual recollection. The solution to precedent rigidity is not to weaken retrieval until officials once again depend upon personal networks or memory to discover what the institution previously knew.

The objective is instead to distribute the appropriate functions between institutional memory architecture and authorised institutional actors. The architecture should perform functions that individuals should not repeatedly have to reconstruct: making materially relevant prior knowledge discoverable, preserving sufficient provenance and context, connecting knowledge to appropriate workflows, signalling changes in status and providing routes for challenge. Institutional actors should then exercise the contextual, professional and legitimate judgement required to determine what that knowledge means for present action.

This relationship is complementary. Professional expertise remains necessary because applicability often depends upon features of the current situation that retained knowledge alone cannot resolve. But expertise should not become an excuse for weak institutional memory. Experienced officials should not have to reconstruct the organisation's past from personal networks when materially relevant knowledge could reasonably be made available through institutional mechanisms.

The aim is therefore neither automated institutional remembering nor wholly individual interpretation. It is a stronger shared informational foundation upon which authorised actors can exercise informed judgement.

Responsible memory is memory that can change

The ultimate safeguard against becoming trapped by institutional memory is not forgetting. It is the capacity to change what remembered knowledge means for present action.

A responsible institution can recover its past without treating that past as fixed instruction. It can preserve historical trace while distinguishing current authority. It can retain previous reasoning while recognising changed assumptions. It can make accumulated experience accessible while allowing new evidence, legitimate new priorities and changed institutional conditions to alter its direction.

This capacity for renewal turns memory from an archive of previous action into an input to institutional cognition. Remembering strengthens capability when it enlarges the relevant informational basis available to judgement. It becomes constraining when accumulated knowledge narrows that basis by making previous choices appear inevitable.

The institutional objective should therefore be neither systematic deference to precedent nor continuous reinvention, but **informed continuity**: the ability to know what the institution has previously learned and done, understand enough of why it did so, and decide deliberately what should follow.

Under this standard, relevant knowledge should be recoverable, sufficient context should permit interpretation, status and provenance should help users understand its present standing, and prior knowledge should remain practically challengeable. Historical trace should be preserved where it serves accountability or learning. Technology should strengthen access without substituting for legitimate judgement. And the institution should retain the practical capacity to apply, adapt, combine, qualify, reject or supersede what it remembers.

Under those conditions, stronger institutional memory does not make an institution more dependent upon its past. It makes the institution better able to use its past as evidence while remaining capable of changing course.

The remaining question is therefore evaluative. **If institutional memory must combine reliable retrieval and reuse with challengeability, renewal and independent judgement, how**

can an institution determine whether these arrangements have actually produced capability gain rather than merely more memory-related activity?

VII. Evaluating Capability Gain

A functioning memory mechanism is necessary for capability gain, but it is not sufficient evidence of it. Section V established the implementation question: whether an intervention has activated its intended mechanism and become a dependable institutional function. Evaluation must now ask a different question: **has the institution become more reliable at recovering and responsibly using materially relevant prior knowledge because those functions operate?**

This distinction matters because institutional interventions generate visible activity long before they demonstrate capability. Searches increase, repositories receive more visits, contextual records are completed, workflow prompts are triggered and governance processes meet. Such evidence can show that an intervention is being used. It cannot establish by itself that institutional memory has strengthened.

Evaluation should therefore preserve four analytical levels:

Activity → Function → Capability → Outcome

Each level answers a different question. Activity asks what people and systems are doing. Function asks whether the mechanism performs the task for which it was designed. Capability asks whether the institution has become more dependable at the relevant form of responsible reuse across the cases for which the intervention matters. Outcome asks whether wider institutional effects — such as reduced reconstruction, stronger decision continuity or improved preparedness — have plausibly followed.

Confusing these levels creates two opposite risks. Institutions may declare success too early because implementation activity is visible, or they may demand evidence of broad organisational outcomes before a more bounded capability improvement can reasonably be demonstrated. Evaluation should instead establish the strongest conclusion justified at each level.

Activity shows use, not capability

Activity evidence is often useful. Search behaviour can show whether a retrieval mechanism is being used. Completion of contextual records can indicate whether a new documentation practice has entered ordinary work. Workflow interactions can reveal whether a memory checkpoint is being encountered. Participation in stewardship or review processes can indicate that governance arrangements are active.

But activity remains evidence of behaviour around an intervention, not of the institutional capability the intervention was intended to strengthen. A repository can receive more searches while relevant knowledge remains difficult to find. A contextual template can be completed without giving future users enough information to reconstruct meaning. A workflow prompt can be acknowledged without materially relevant prior knowledge influencing consideration. A review process can operate regularly while obsolete knowledge retains inappropriate prominence (U.S. Government Accountability Office, 2002; UK Resilience Academy, 2025).

This is why implementation evidence must not be asked to carry an evaluative conclusion it cannot support. Activity can help explain whether an intervention is being used and where further investigation is needed. It cannot substitute for evidence that the required institutional function operates.

Function establishes mechanism performance

Function-level evidence asks whether the mechanisms identified in Sections IV and V are performing their intended roles. The detailed activation tests belong to implementation; evaluation needs only to preserve their result as a distinct evidential level.

For findability, the question is whether intended users can recover materially relevant retained knowledge with greater reliability. For contextual interpretability, it is whether they can reconstruct sufficient meaning to understand what the knowledge represents and the conditions under which it was produced. For workflow re-entry, it is whether relevant prior knowledge reaches institutional work while it can still affect judgement. For stewardship and renewal, it is whether status can be maintained, problematic material can be challenged and changed conditions can alter how prior knowledge is treated.

These tests establish that the mechanism works. They provide the causal bridge between intervention and capability, but **a functioning mechanism is still not the same as a more capable institution**. Capability requires the relevant functions to combine into a sufficiently dependable improvement in responsible reuse.

Capability asks whether responsible reuse has become more reliable

Capability evaluation moves from the mechanism to the institution. The central question is whether, across relevant cases and over time, the organisation has become more reliable at recovering materially relevant prior knowledge, reconstructing sufficient context, bringing that knowledge into appropriate workflows and using it through informed applicability judgement.

Five dimensions help define this capability.

Retrieval capability concerns whether relevant retained knowledge can be recovered when needed with greater reliability.

Contextual interpretability concerns whether later users can reconstruct enough meaning to understand what that knowledge represents and the conditions under which it was produced.

Workflow re-entry concerns whether relevant prior knowledge reaches institutional work at points where it can still influence judgement.

Appropriate utilisation concerns whether institutional actors can determine responsibly what role retrieved knowledge should play rather than treating retrieval as authority.

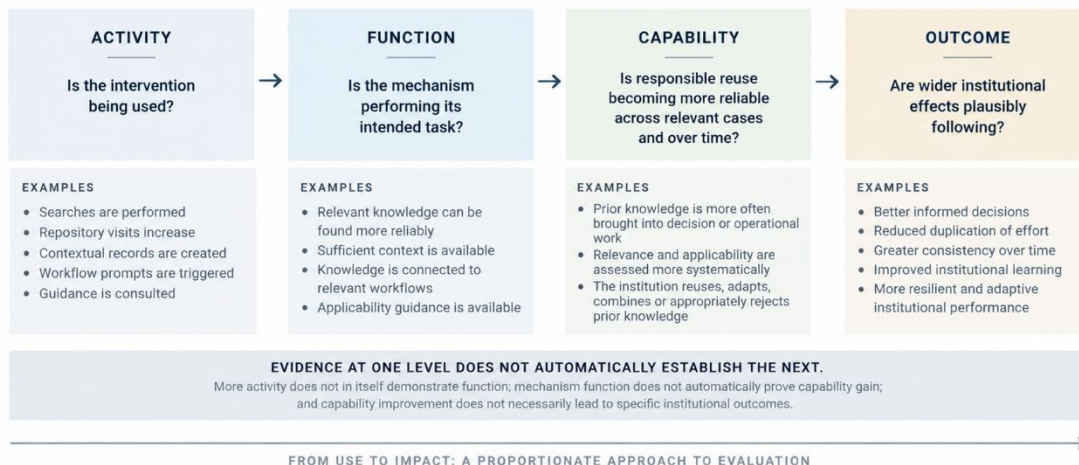
Durability and governance concern whether the institution can maintain, challenge and adapt the arrangements through which these functions operate without unreasonable dependence on particular individuals or disproportionate institutional burden.

These dimensions define the capability developed in PP-002. They should **not** be treated as a universal scorecard requiring identical indicators, weights or thresholds across institutions. Evaluation should follow the material failure diagnosed and the intervention actually implemented. A bounded retrieval intervention may reasonably concentrate on findability and subsequent use in a particular workflow. A broader intervention addressing context, workflow and stewardship will require correspondingly broader evidence.

The purpose of the framework is therefore to identify what needs to be demonstrated, not to impose a standardised measurement instrument where the evidence does not justify one.

Figure 3. From Activity to Outcome

Evaluating institutional memory: distinguishing activity, function, capability and outcome.



Capability requires repeatability

A single successful case demonstrates that responsible reuse is possible. Capability requires greater reliability across the relevant class of cases (UK Resilience Academy, 2025).

The degree and form of repetition will differ by institutional context. High-frequency processes may generate enough observations for patterns to become visible relatively quickly. Low-frequency, high-consequence situations may provide few opportunities for direct repeated observation. In such cases, preparedness exercises, simulations, structured review, historical case testing or other forms of evidence may be needed to assess whether the required memory function is likely to remain available when the relevant situation occurs.

Repeatability should not be interpreted as uniformity. Different cases can require different responses to prior knowledge, and a capable institution should not be expected to produce identical reuse outcomes. What should become more dependable is the **route through which materially relevant prior knowledge can be recovered, understood, considered and judged**.

This distinction prevents capability from being confused with standardisation. Institutional memory is stronger when the institution can repeatedly bring relevant accumulated knowledge into present cognition, not when it repeatedly reaches the same decision.

Responsible non-use can demonstrate capability

One of the most important consequences of evaluating capability rather than reuse volume is that successful memory does not require previous knowledge to be followed.

A decision may recover relevant prior knowledge, reconstruct the conditions under which it was produced and then determine that the knowledge should not govern the present case. The earlier analysis may depend upon assumptions that have changed. New evidence may justify another conclusion. A previous practice may have produced lessons precisely because it should not be repeated. An earlier decision may remain historically important while no longer possessing current authority.

In such cases, **reasoned non-use can demonstrate stronger institutional memory than automatic reuse**.

Evaluation should therefore ask not only whether retrieved knowledge was used in the sense of being adopted, but whether it was **considered adequately and subjected to responsible applicability judgement**. Evidence may show that prior knowledge was applied, adapted, combined with newer evidence, qualified, rejected with rationale or superseded. The relevant success condition is not conformity with the past but informed treatment of it.

This has practical consequences for measurement. Indicators based primarily on reuse rates can create the wrong institutional incentive. If teams are rewarded for demonstrating that previous knowledge influenced decisions, they may become reluctant to reject material that no longer applies. Evaluation should instead recognise responsible departure as a legitimate outcome where the institution can show that materially relevant prior knowledge was recovered, understood and considered before that departure occurred.

Responsible non-use is therefore not an exception to the capability model. It is evidence that memory remains subordinate to judgement.

Wider outcomes matter, but attribution should remain proportionate

Stronger institutional memory is valuable partly because it may contribute to wider institutional effects. Reduced duplicated reconstruction, fewer repeated avoidable errors, stronger continuity of reasoning, improved coordination, better preparedness and stronger explanatory accountability are all plausible consequences of improved memory capability (Rashman et al., 2009).

These outcomes can strengthen the case that an intervention is producing institutional value. They should not, however, be treated as automatic proof that memory caused them.

Institutional outcomes usually have multiple causes. Reduced duplication may reflect organisational restructuring as well as improved retrieval. Greater decision continuity may result from personnel stability, clearer policy direction or stronger governance alongside better memory. Preparedness may improve through training, resources and leadership as well as retained knowledge. Evaluation should therefore make causal claims in proportion to the evidence available.

The evidential sequence matters:

activity can support a claim of use; function can support a claim of mechanism performance; repeated responsible reuse can support a claim of capability gain; wider outcomes require additional evidence before stronger causal attribution is justified.

This discipline allows institutions to recognise meaningful capability improvement without overstating what an intervention has proven.

Capability gain must justify its burden

An intervention can work exactly as designed and still fail the institutional net-value test.

Memory mechanisms consume capacity. Knowledge may need to be documented, classified, maintained and reviewed. Contextual records take time to create. Workflow integration can add procedural steps. Stewardship requires responsibility and attention. Technology creates acquisition, integration and maintenance costs. Greater availability of prior material can increase cognitive load if users must distinguish relevant knowledge from informational noise. Governance and consultation can consume resources that might otherwise support current analysis or delivery (OECD, 2025; U.S. Government Accountability Office, 2002).

These burdens do not make memory intervention undesirable. They make proportionality part of capability evaluation.

The relevant question is not simply whether responsible reuse has improved, but whether the improvement is valuable enough to justify the institutional capacity required to sustain it. This relationship may change over time. Initial implementation costs may fall as practices stabilise. Maintenance burden may rise as retained material accumulates. New technology may reduce some retrieval costs while creating additional governance requirements. The materiality of the original failure may itself change.

A functioning intervention should therefore remain open to the possibility that its configuration is too extensive, too weak or no longer appropriate.

Memory capability must not consume more institutional capacity than the failure it is intended to correct reasonably justifies.

There is no universal numerical threshold for this judgement. Institutions need evidence appropriate to the scale, consequence and context of the intervention. The purpose of the principle is to prevent successful mechanism operation from being treated as sufficient justification for indefinite continuation.

Evaluate negative effects as well as intended gains

Evaluation should also examine whether stronger memory produces effects that undermine the capability it was intended to create.

Improved retrieval can increase exposure to obsolete material. Strong workflow integration can turn useful prompts into procedural compliance. Greater visibility of precedent can strengthen bias towards established interpretations. Extensive documentation can generate maintenance debt. Large volumes of surfaced knowledge can increase cognitive load. Automated retrieval can privilege particular sources or forms of knowledge because they are easier to index or rank (Njiru et al., 2025; Ruschemeier & Hondrich, 2024).

These effects are especially important because they may coexist with successful mechanism activation. A search system can improve findability while increasing retrieval bias. A workflow mechanism can ensure that prior knowledge is considered while making departure from precedent more difficult. A documentation practice can preserve context while imposing a burden that discourages substantive analysis elsewhere.

Evaluation should therefore ask not merely whether the intended mechanism works, but whether the resulting institutional configuration remains consistent with responsible memory. Capability gain is weakened where improvements in one part of the retention-to-reuse pathway systematically degrade another.

Evaluation should change what happens next

Evaluation becomes institutionally useful when its findings can alter the intervention. A system that measures performance but cannot modify what follows has established monitoring, not adaptive governance.

Several responses should remain available.

Where the mechanism functions, responsible reuse becomes more reliable and burden remains proportionate, the institution may **continue** the intervention.

Where the mechanism is broadly sound but configuration, ownership or safeguards need improvement, it may **adapt** it.

Where evidence shows that the mechanism should cover additional high-materiality cases, it may **escalate or broaden** the intervention.

Where burden exceeds value or the original materiality was overestimated, it may **de-escalate**.

Where another mechanism could perform the required function more effectively, it may **replace** the existing arrangement.

Where the original problem has diminished, the intervention no longer adds sufficient value or negative effects outweigh benefits, it may **retire** it.

None of these decisions should be treated as inherently superior. Scaling is not evidence of greater institutional maturity, and retirement is not necessarily evidence of failure. An intervention that solves a bounded problem and can later be simplified may have succeeded precisely because the institution no longer requires the original arrangement. De-escalation can demonstrate better calibration rather than declining commitment.

The evaluation loop should therefore remain open:

verify function → assess capability → examine outcomes, burden and negative effects → continue, adapt, escalate, de-escalate, replace or retire → evaluate again.

This is adaptive governance applied to institutional memory.

The standard is reliable responsible reuse

The final evaluative test is not whether the institution stores more knowledge, generates more memory-related activity or successfully launches a particular intervention. It is whether the institution becomes **more reliably capable of recovering materially relevant prior knowledge, reconstructing sufficient context, bringing that knowledge into appropriate institutional work and subjecting it to informed applicability judgement across the cases in which the capability matters.**

That improvement should be demonstrated at the capability level rather than inferred from deployment. It should be sufficiently repeatable to represent an institutional function rather than an isolated success. It should preserve challengeability and independent judgement. And its institutional value should remain proportionate to the burden required to sustain it.

Wider outcomes can provide additional evidence where they are observable and plausibly connected, but they should not replace direct evidence of capability. Nor should responsible departure from previous knowledge be counted as failure simply because the retrieved knowledge was not adopted.

Evaluated in this way, institutional memory becomes a capability that can be governed rather than an infrastructure whose expansion is assumed to be beneficial. The institution can distinguish activity from function, function from capability and capability from wider outcome; determine whether responsible reuse has genuinely become more reliable; and adapt the mechanisms through which that capability is sustained.

The analysis can therefore move from evaluation to its broader implications. **If institutional memory should be selective, proportionate, challengeable and judged by reliable responsible reuse rather than by the scale of memory infrastructure, what does this imply for institutional leaders — and how far can the framework developed here reasonably be generalised?**

VIII. Implications, Limits and Capability Direction

The analysis developed in this paper changes the practical question institutional leaders should ask about memory. The relevant issue is not whether an organisation possesses enough repositories, records or knowledge-management infrastructure. Nor is it whether every part of its accumulated knowledge can be made permanently accessible. The more useful question is whether **materially relevant prior knowledge can reliably re-enter institutional action when it matters, under conditions that preserve sufficient context and independent judgement.**

This framing has several implications for institutional design. It also establishes limits on how far the framework developed here should be generalised. Institutional memory operates across highly different organisational, legal, professional and technological environments, and the evidence does not support a universal model of what every institution should build. The capability direction can nevertheless be stated more confidently than any single implementation form: institutions should become better at reconnecting relevant accumulated knowledge with present work without turning that accumulated knowledge into an automatic constraint on future action.

Diagnose memory functionally

The first implication is that institutional memory problems should be diagnosed through the function that fails rather than through the apparent absence of a particular solution.

An organisation that repeatedly reconstructs earlier work may need stronger findability, but it may instead suffer from poor retention, insufficient contextualisation or fragmentation across organisational boundaries. A repository that receives little use may reflect weak search and navigation, but it may also contain knowledge with limited future material value. A programme that fails to learn from previous implementation may need better access to lessons, but the more important weakness may be that those lessons never enter subsequent design processes.

Starting from function changes the intervention logic. It directs attention towards the route between retained knowledge and future institutional action and asks where that route becomes unreliable. This makes it possible to distinguish problems requiring retrieval, context, workflow, stewardship or challenge mechanisms from those for which formal intervention is unnecessary.

For institutional leaders, the practical discipline is therefore to resist solution-first diagnosis. **The absence of a knowledge portal, AI retrieval layer, decision log or formal memory process is not itself evidence of a capability gap.** The gap lies in an institutional function that matters and cannot be performed with sufficient reliability.

Intervene selectively and proportionately

The second implication is that stronger institutional memory should not be pursued uniformly.

Knowledge differs in future value, consequence, recurrence, reconstructability and accountability significance. Institutional processes also differ in their exposure to memory failure. Some depend heavily on accumulated experience; others legitimately require substantial new analysis each time. Some failures occur frequently but create modest cost. Others arise rarely while carrying significant consequences.

Intervention should follow those differences. Institutions may reasonably invest more in the findability, context, stewardship or workflow connection of high-materiality knowledge while allowing lower-value material to remain under lighter arrangements. They may use formal mechanisms in one domain and rely more heavily on professional or social memory in another.

They may decide that occasional reconstruction is less burdensome than maintaining an extensive reuse mechanism.

This selectivity is not a weakness in institutional memory. It is part of responsible capability design. A system that attempts to preserve, contextualise, surface and govern everything with equal intensity can consume substantial institutional capacity while making important knowledge harder to distinguish from informational noise.

The practical objective is therefore **minimum sufficient institutional memory for the material problem**, with the possibility of adaptation as evidence develops.

Evaluate function and capability rather than project completion

The third implication is that institutional memory programmes should not be judged primarily through implementation milestones.

Systems can be deployed, processes introduced, records completed and responsibilities assigned without the intended mechanism functioning. Conversely, a relatively modest intervention can materially strengthen institutional memory if it reliably restores the function that was previously failing.

Institutional leaders therefore need evidence at different levels. They should know whether the intervention is being used, whether its mechanism functions, whether responsible reuse has become more dependable and, where wider effects are claimed, whether the evidence supports attributing those effects to stronger memory. Keeping these levels distinct prevents deployment from being mistaken for capability.

It also changes the meaning of scale. Expansion is not necessarily evidence of success. A bounded intervention that corrects a material failure with manageable burden may represent stronger institutional design than an organisation-wide programme whose activity is extensive but whose contribution to responsible reuse remains unclear.

The same principle applies to institutionalisation. The objective is to preserve the required function, not the original project form. Mechanisms should remain capable of adaptation, simplification, replacement or retirement when evidence shows that another configuration would perform the function more effectively.

Treat responsible judgement as part of memory capability

The fourth implication is that institutions should not evaluate memory through reuse volume alone.

Better institutional memory can increase the probability that previous knowledge influences later action, but influence is not the same as compliance. Earlier evidence may remain valid, require adaptation, conflict with newer knowledge or no longer apply. Previous decisions may provide important context without possessing continuing authority. Lessons can matter precisely because they explain what should not be repeated.

A capable institution therefore needs to recover its past **and remain able to disagree with it**.

This requirement becomes more important as retrieval mechanisms become stronger. When prior knowledge is easier to discover and more routinely surfaced, the risk of excessive institutional influence can increase alongside the benefit of reduced forgetting. Challengeability, renewal, visible status and independent applicability judgement are consequently not secondary safeguards around institutional memory. They are part of the capability itself.

For leaders, this implies that a memory intervention should be questioned not only when relevant knowledge remains absent, but also when the intervention makes established knowledge too difficult to challenge. Institutional forgetting and precedent rigidity are opposite failures of the same capability.

Technology should strengthen access without acquiring authority

Technological change can materially alter the conditions under which institutional memory operates. Improved search, semantic retrieval, automated classification and cross-repository navigation can reduce some of the barriers that historically made accumulated knowledge difficult to recover. Systems may increasingly be able to surface relevant previous decisions, analyses or lessons within ordinary workflows rather than waiting for users to search manually.

These developments can strengthen memory capability, particularly where discovery and invocation are the binding failures. They also make the governance principles developed in this paper more important.

More powerful retrieval does not remove the need for sufficient context. Greater automation does not establish the present validity of what is surfaced. Ranking does not confer institutional authority. Easier access can increase exposure to obsolete material or informational noise. Automated systems may privilege sources that are better structured, more frequently referenced or easier to index, thereby influencing what institutional actors encounter.

The appropriate institutional response is neither technological avoidance nor technological presumption. Technology should be assessed according to the function it performs, the conditions required for that function to operate and the safeguards necessary to preserve responsible judgement.

As technological capability changes, the appropriate configuration of institutional memory may change with it. The **function-first principle** therefore provides a more durable basis for institutional adaptation than commitment to any particular technological form.

Transfer the function before transferring the form

The same principle applies across institutional contexts.

Public institutions vary in size, mandate, legal environment, organisational structure, professional composition, information infrastructure and administrative tradition. An arrangement that performs well in one organisation cannot be assumed to produce the same result elsewhere simply because its formal design can be reproduced (Rashman et al., 2009).

Transfer should begin by identifying the function an intervention performs. Institutions can then determine whether the same function is materially weak in their own setting, whether the proposed mechanism depends upon conditions they possess and whether another arrangement could perform that function with lower burden or greater reliability.

Transfer the function before transferring the form.

This principle allows institutions to learn from one another without converting observed practice into universal prescription. It also places organisational readiness in the correct role. Readiness is not a ranking of institutional maturity; it is an assessment of whether the dependencies required for a particular mechanism are sufficiently present for that mechanism to work.

The same intervention may therefore be appropriate in one domain, premature in another and unnecessary in a third.

The evidence supports a capability framework, not a universal institutional model

The framework developed in PP-002 should be interpreted within the limits of the available evidence.

The evidence base supports the distinction between retaining knowledge and possessing a functioning capacity to recover and reuse it. It also supports the practical importance of discoverability, contextual information, workflow connection, organisational stewardship and the risks associated with outdated or inappropriately authoritative knowledge. Across these areas,

the evidence provides a sufficiently coherent basis for the capability architecture and the conditional recommendations developed in the paper.

The evidence is less decisive about complete intervention packages. Comparative causal evidence showing that a particular end-to-end institutional memory model produces superior results across different public institutions remains limited. Evidence is also geographically and institutionally concentrated in parts of the underlying literature and practice base, reducing confidence in universal transfer of particular organisational forms (Rashman et al., 2009).

The retention-to-reuse capability chain developed here should therefore not be treated as a universally validated measurement instrument. It is an analytical architecture for identifying where institutional memory can fail and what functions may need strengthening. Institutions will need indicators appropriate to their own material failures, processes and intervention designs.

Nor does the evidence support a universal numerical threshold for acceptable burden, sufficient repeatability or intervention materiality. These judgements remain context-dependent. The framework can discipline them by clarifying what should be considered, but it cannot replace institutional judgement with a single metric.

Outcome attribution requires similar restraint. Stronger memory may plausibly reduce duplicated reconstruction, repeated error or decision discontinuity, but those outcomes are influenced by many other institutional conditions. Claims about wider effects should therefore remain proportionate to the evidence available.

These limitations narrow the strength of the prescriptions without undermining the capability direction. PP-002 does not establish that every institution should adopt the same memory architecture. It establishes a structured way of diagnosing material reuse failure, selecting mechanisms conditionally, testing whether they function and evaluating whether responsible reuse becomes more reliable at acceptable burden.

Capability development should remain adaptive

Institutional memory should not be treated as a one-time infrastructure project. The knowledge environment changes, organisational responsibilities move, workflows evolve, technology develops and the future value of retained knowledge changes over time. Mechanisms that perform well under one configuration may later become burdensome, insufficient or misaligned with the problem they were designed to correct.

Capability development should therefore remain adaptive. Institutions should be able to strengthen a mechanism where evidence shows persistent weakness, broaden it where additional material cases justify coverage, simplify it where burden exceeds value, replace it where another mechanism performs the function better and retire it where the original problem no longer warrants intervention.

This adaptive orientation is particularly important because memory infrastructure has a tendency to accumulate. New repositories, documentation requirements, governance processes and technological layers can be added more easily than old ones are removed. Without deliberate adaptation, institutional memory reform can therefore generate maintenance debt and procedural complexity even while attempting to reduce the costs of institutional forgetting (OECD, 2025; Kluge & Gronau, 2018).

The objective is not permanent growth of memory infrastructure. It is continued alignment between **material institutional need and the minimum sufficient mechanisms required to support responsible reuse**.

The capability direction

The direction emerging from the analysis is consequently neither maximal retention nor maximal reuse. Institutions do not need to remember everything, and they should not seek to reproduce everything they remember.

They need a selective and dependable capacity through which knowledge that may materially affect future action can remain recoverable across time and organisational change, carry enough context to be interpreted, enter the work in which it can matter and remain subject to judgement about its present significance. That capability must be maintainable at a burden proportionate to the failure it addresses and adaptable as institutional conditions change.

Under this model, institutional memory becomes less a question of preserving the past than of **governing the relationship between past and present**. The institution is not attempting to make all accumulated knowledge continuously prominent. It is trying to ensure that knowledge which deserves to matter later can matter later, while preserving the ability to determine that some of what it remembers should be adapted, challenged or left behind.

The limits of the evidence counsel against a universal institutional model. They do not remove the practical direction established by the analysis: **reliable, selective and responsible reuse is a more defensible objective than either comprehensive retention or automatic continuity with the past**.

The final question is therefore no longer how institutional memory should be diagnosed, designed, implemented or evaluated. It is what this capability ultimately changes about the institution itself — what it means for an organisation not merely to preserve its accumulated knowledge, but genuinely to remember.

Conclusion — Capability Direction

Public institutions can preserve extensive traces of what they have previously known and done while remaining unable to use that accumulated knowledge when it becomes relevant again. Reports survive, decisions remain recorded, evaluations are archived and institutional experience leaves documentary traces, yet later actors can still be forced to reconstruct earlier reasoning, rediscover previous constraints or confront problems the organisation has already encountered. **An institution can therefore retain its past without being able to remember it institutionally.**

The paradox is resolved once institutional memory is understood as a capability rather than a stock of information. Knowledge must be able to travel from retention into present institutional cognition. What matters is not simply that prior knowledge exists somewhere within the organisation, but that materially relevant knowledge can become available to later actors with enough context to understand what it represents and through routes that allow it to enter the work in which it can matter.

This movement across time must remain compatible with judgement. The past becomes useful when it enlarges the informational basis from which present actors can reason; it becomes constraining when retrieval is mistaken for authority or continuity becomes an obligation to reproduce earlier choices. **The purpose of memory is not to make the institution obey its past. It is to ensure that the past can inform present judgement without replacing it.**

A capable institution must therefore be able both to remember and to change. It should not needlessly reconstruct what it has already learned, but neither should it treat accumulated experience as fixed instruction. It should be able to recognise what remains relevant, understand what has changed and depart from previous knowledge when present evidence or circumstances justify doing so. In this sense, institutional memory supports informed continuity rather than precedent dependence: it allows change to occur with knowledge of what came before rather than through unnecessary institutional forgetting.

This capability does not eliminate the need for new analysis, professional expertise or institutional adaptation. It improves the informational conditions under which they occur. Officials can reach new conclusions with greater awareness of earlier reasoning. Programmes can change direction without unnecessarily rediscovering what previous implementation revealed. Institutions can depart from precedent while understanding why that precedent existed and what has changed since it was established.

The capability direction is therefore neither maximal retention nor maximal reuse. **Institutional memory is the selective and dependable capacity to recover the right prior knowledge when it matters, reconstruct enough context to understand it, bring it into present institutional work and decide responsibly whether to apply, adapt, combine, qualify, reject or supersede it.**

An institution that can do this repeatedly is not simply better at preserving information. It is better able to connect accumulated experience with present cognition while retaining the freedom to learn, reconsider and change course.

That is the difference between storing institutional knowledge and possessing institutional memory as a functioning capability.

References

- Argote, L., & Guo, J. (2016).** Routines and transactive memory systems: Creating, coordinating, retaining, and transferring knowledge in organizations. *Research in Organizational Behavior*, 36, 65–84. <https://doi.org/10.1016/j.riob.2016.10.002>
- Boh, W. F. (2008).** Reuse of knowledge assets from repositories: A mixed methods study. *Information & Management*, 45(6), 365–375. <https://doi.org/10.1016/j.im.2008.06.001>
- Kankanhalli, A., Tan, B. C. Y., & Wei, K.-K. (2005).** Understanding seeking from electronic knowledge repositories: An empirical study. *Journal of the American Society for Information Science and Technology*, 56(11), 1156–1166. <https://doi.org/10.1002/asi.20219>
- Kluge, A., & Gronau, N. (2018).** Intentional forgetting in organizations: The importance of eliminating retrieval cues for implementing new routines. *Frontiers in Psychology*, 9, 51. <https://doi.org/10.3389/fpsyg.2018.00051>
- Markus, M. L. (2001).** Toward a theory of knowledge reuse: Types of knowledge reuse situations and factors in reuse success. *Journal of Management Information Systems*, 18(1), 57–93.
- Njiru, D. K., Mugo, D. M., Musyoka, F. M., & Murikah, W. (2025).** How knowledge management systems trap dynamic organisations into outdated practices: A systematic review. *Scientific African*, 29, e02946. <https://doi.org/10.1016/j.sciaf.2025.e02946>
- OECD. (2025).** Everybody cares about using education research sometimes: Perspectives of knowledge intermediaries. *Educational Research and Innovation*. OECD Publishing. <https://doi.org/10.1787/5ef88972-en>
- Olivera, F. (2000).** Memory systems in organizations: An empirical investigation of mechanisms for knowledge collection, storage and access. *Journal of Management Studies*, 37(6), 811–832. <https://doi.org/10.1111/1467-6486.00205>
- Rashman, L., Withers, E., & Hartley, J. (2009).** Organizational learning and knowledge in public service organizations: A systematic review of the literature. *International Journal of Management Reviews*, 11(4), 463–494. <https://doi.org/10.1111/j.1468-2370.2009.00257.x>
- Ruscheimer, H., & Hondrich, L. J. (2024).** Automation bias in public administration – an interdisciplinary perspective from law and psychology. *Government Information Quarterly*, 41(3), 101953. <https://doi.org/10.1016/j.giq.2024.101953>
- Stark, A., Gajurel, H., Corbett, J., Grube, D. C., Lovell, H., Scott, R., et al. (2026).** *Memory and institutional amnesia in government*. Oxford University Press.
- Stein, E. W. (1995).** Organization memory: Review of concepts and recommendations for management. *International Journal of Information Management*, 15(1), 17–32. [https://doi.org/10.1016/0268-4012\(94\)00003-C](https://doi.org/10.1016/0268-4012(94)00003-C)
- UK Covid-19 Inquiry. (2024).** Module 1 report: The resilience and preparedness of the United Kingdom. UK Covid-19 Inquiry.
- UK Government. (2022).** *Levelling up the United Kingdom*. HM Government.
- UK Resilience Academy. (2025).** *Lessons management best practice guidance*. UK Government.
- U.S. Government Accountability Office. (2002).** *NASA: Better mechanisms needed for sharing lessons learned* (GAO-02-195).
- Walsh, J. P., & Ungson, G. R. (1991).** Organizational memory. *Academy of Management Review*, 16(1), 57–91. <https://doi.org/10.5465/AMR.1991.4278992>
- Williams, W. (2020).** *Windrush lessons learned review*. Home Office.

Zyngier, S., & Burstein, F. (2012). Knowledge management governance: The road to continuous benefits realization. *Journal of Information Technology*, 27(2), 140–155. <https://doi.org/10.1057/jit.2011.31>

