

DHUS Institute
Policy Paper
PP-001

**Preventing
Institutional
Knowledge Loss**

Copyright © 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
1. The Institutional Problem	6
2. Why the Problem Persists	8
3. What the Evidence Shows	10
4. Intervention Options	13
5. Recommendations.....	15
6. From Recommendation to Institutional Action.....	17
7. How to Know Whether Capability Has Improved.....	20
8. Conclusion	22
References	23

Executive Summary

Public institutions depend on specialised knowledge held by individuals. This is neither unusual nor inherently problematic: expertise, professional judgement and differentiated responsibilities are essential to complex institutional work. A continuity risk emerges, however, when a consequential institutional function depends so heavily on knowledge held by one person, or by a very small number of people, that the institution cannot reproduce the required capability once those individuals become unavailable. Retirement, resignation, reassignment, rotation, restructuring or unexpected absence can then produce what appears to be sudden knowledge loss even though the underlying vulnerability may have existed for years.

Personnel transition often exposes institutional dependency rather than creating it. While an experienced employee remains available, the organisation can compensate for weaknesses in its own knowledge architecture simply by consulting that person. Formal procedures may appear sufficient because the employee supplies missing context; difficult cases remain manageable because accumulated judgement is still accessible; and relationships continue to function because the individual through whom they were built remains present. The institution may therefore appear capable while part of that capability remains dependent on continued access to particular people.

The evidence reviewed for this Policy Paper shows that turnover-related knowledge loss is a genuine but heterogeneous organisational phenomenon. Not every departure produces meaningful loss, and not every form of lost knowledge has significant institutional consequences. Whether a transition weakens capability depends on the importance of the affected function, the concentration and substitutability of relevant knowledge, the availability of alternative pathways, and the difficulty of reproducing or reacquiring what the institution still requires (Galan, 2023a). The appropriate unit of analysis is therefore not the departing employee alone, but the relationship between the person, the knowledge and the institutional function that depends upon it.

Different dependencies also create different reproduction burdens. Some knowledge can be represented effectively through records and documentation, while other forms depend more heavily on context, experience, practice or relationships. Continuity may therefore require representation, structured social transfer, practice-based reproduction, distributed capability or, where proportionate, a credible route to reacquisition. The objective is not universal duplication of expertise, but sufficient institutional reproducibility.

PP-001 consequently recommends that public institutions govern material person-dependent knowledge as a continuity risk. This principle translates into four connected areas of action. Institutions should see the risk before transition by identifying consequential functions that depend excessively on particular knowledge holders and prioritising those dependencies according to criticality and vulnerability. They should design the right continuity response by matching mechanisms to the reproduction burden and creating sufficient alternative pathways. They should make continuity actionable by acting early enough for the selected pathway to work and assigning clear institutional ownership. Finally, they should test whether capability survived rather than treating completion of transfer activity as evidence of success.

Implementation can follow a common logic while remaining proportionate to context: own the risk, identify the dependency, decide whether it matters enough to act, determine what must survive, build the lightest credible alternative pathway, activate it while the necessary conditions still exist, and test the result after transition. Some dependencies will justify little more than accessible representation and confirmation that another route exists. Others may require structured transfer, practice or additional holders. A smaller number of high-consequence dependencies may justify multiple pathways or deliberate redundancy. Sudden departures require a corresponding shift from preventive reproduction toward recovery through records, colleagues, external expertise, temporary workarounds or redesign.

Completion of transfer activity is not evidence that continuity has been achieved. The relevant test is whether an alternative institutional pathway actually functions after transition and whether the resulting continuity remains proportionate to its organisational cost and effects. Capability gain should therefore be distinguished from implementation activity and from the mere existence of a planned mechanism.

This proportionality is essential because the objective of knowledge continuity is not institutional stasis. Mobility can introduce new expertise, inherited practices may become obsolete, and some knowledge can be reconstructed more efficiently than it can be preserved. Public institutions therefore need minimum sufficient continuity: preserving or reproducing enough capability-relevant knowledge to keep consequential functions viable without attempting to retain everything the institution currently knows.

The central test is straightforward. When a consequential knowledge holder becomes unavailable, can the institution still do what it materially needs to do with sufficient understanding, reliability and independence? If it can, and the resources required to create that continuity were proportionate to the institutional risk, capability has improved. If transfer activities have been completed but the function remains dependent on privileged access to the former holder, the underlying vulnerability remains.

David González
Director
IDHUS Institute

1. The Institutional Problem

Institutions depend on people. Expertise accumulates through years of practice, repeated exposure to difficult cases, professional judgement, relationships, familiarity with procedures and an often imperfectly documented understanding of why particular decisions, routines or arrangements developed in the way they did. None of this is inherently problematic. Specialisation is indispensable to complex institutional work, and organisations would function poorly if every member were expected to possess the same knowledge or if every form of expertise had to be duplicated. The institutional problem begins somewhere else: when a consequential function depends so heavily on knowledge held by one person, or by a very small group of people, that the institution cannot continue to perform that function with comparable understanding or reliability once those individuals are no longer available. At that point, what appears to be organisational capability may in fact depend partly on continued access to particular people.

Personnel transitions make these dependencies visible. Retirement is the most obvious example, but it is only one among many. Employees resign, rotate, move internally, take extended leave, change responsibilities, become unavailable unexpectedly or leave as organisations restructure. A systematic review of 91 empirical studies confirms that knowledge loss associated with organisational member turnover is a recognisable phenomenon across a wide range of organisational and sectoral contexts while also showing that its effects are heterogeneous: not every departure produces meaningful knowledge loss, and not every loss produces the same consequences (Galan, 2023a). This distinction matters because the institutional problem is not turnover itself. Organisations should expect movement, renewal and specialisation. What matters is whether a transition removes access to knowledge that the institution still requires and cannot readily obtain elsewhere.

The difficulty is that much of this dependency can remain invisible while the knowledge holder is present. A function may appear stable because an experienced employee knows whom to contact, how to interpret ambiguous cases, why an exception was introduced years earlier, which warning signs deserve attention or how to adapt a formal procedure when circumstances do not match the standard case. Colleagues may rely repeatedly on that person without treating the reliance as a continuity risk because, under normal conditions, the knowledge remains readily accessible. The weakness becomes apparent only after access disappears.

The problem extends beyond explicit information. Institutional knowledge may include context, judgement, experience and relational understanding that cannot be assumed to survive merely because formal records remain available. Documentation can reproduce some dependencies effectively, but its adequacy depends on whether the knowledge required for competent use can itself be represented. Research in public enterprises has similarly linked turnover and retirement with tacit-knowledge loss while identifying organisational conditions as relevant to whether knowledge remains available (Phaladi & Ngulube, 2024).

The consequences are also not always immediate. A transition can appear successful because routine work continues, only for a more serious gap to emerge later when an unusual case, crisis, dispute, technical failure or complex decision requires knowledge that was previously concentrated in the departed employee. This delayed visibility makes institutional knowledge loss particularly difficult to govern. Recent evidence from the United States federal government illustrates the practical relevance of the issue. In 2026, the U.S. Government Accountability Office reported that the Office of Personnel Management had experienced a substantial workforce reduction and that 57 per cent of those who left had eleven or more years of service, characterising the pattern as a significant loss of institutional knowledge and linking it with risks to skills and operational capacity (U.S. Government Accountability Office, 2026).

Yet the opposite error would be to conclude that institutions should preserve everything. Not all knowledge is critical, not every experienced employee represents a continuity vulnerability, and not every loss deserves an organised preservation effort. Some knowledge can be replaced easily, reacquired externally, reconstructed at acceptable cost or allowed to disappear because the function itself has changed. New personnel may also bring better methods, different expertise and perspectives that an institution needs. Knowledge continuity should therefore not mean institutional stasis. The relevant question is how to recognise the smaller subset of dependencies whose disappearance would materially weaken an institutional function and whose preservation or reproduction is proportionate to the consequences of losing them.

This distinction is especially important for public institutions. Governments, public agencies and other mission-driven organisations often combine long-lived mandates with staff mobility, specialised expertise, formal accountability requirements and operating environments in which institutional context accumulates over periods much longer than an individual posting or career stage. In such environments, continuity cannot depend solely on the expectation that knowledgeable individuals will remain accessible indefinitely. Nor can it be treated purely as a human-resources problem concerned with replacing departing employees. The capability question is different: can the institution continue to know and do what it materially needs to know and do after the people through whom that knowledge was previously accessible have moved on?

PP-001 addresses that narrower institutional problem. It is concerned with preventing avoidable loss of consequential knowledge through personnel transition by reducing excessive dependence on particular knowledge holders. It is not a general framework for knowledge management, records retention or organisational learning, and it does not address the broader question of how preserved institutional memory should later be retrieved and reused across future decisions. Its focus is continuity at the point of transition: whether critical knowledge, context, judgement and relationships survive sufficiently for the institution to continue functioning once the original holder is no longer available.

2. Why the Problem Persists

Institutional knowledge loss is difficult to prevent partly because the vulnerability that produces it is often hidden by normal organisational functioning. As long as experienced people remain available, institutions can continue to make decisions, solve unusual problems, interpret ambiguous procedures, maintain relationships and recover the reasoning behind earlier choices without needing to distinguish carefully between what the organisation itself can reproduce and what particular individuals make possible through their continued presence. Successful day-to-day performance does not necessarily demonstrate that the underlying knowledge has become an institutional capability. It may demonstrate only that the institution still has access to the people in whom important parts of that capability remain concentrated.

This creates a characteristic form of hidden dependency. Over time, specialised employees accumulate more than factual information about their work. They learn which elements of a formal process matter most, how different rules interact, which exceptions have proved significant, where previous attempts failed, which informal relationships enable coordination, how particular stakeholders respond and what signals indicate that an apparently routine situation requires closer attention. Some of this knowledge becomes embedded in documents, systems, routines and other people; some does not. The institution can therefore compensate for weaknesses in its own knowledge architecture simply by asking the person who knows.

The resulting vulnerability is not created suddenly when the employee leaves. Personnel transition is often the event that exposes a reproduction failure that already existed. A departure removes or weakens the dominant pathway through which the institution previously accessed the relevant knowledge, and only then does the difference between possessing expertise and possessing institutional capability become visible. A procedure that appeared adequate may prove unable to guide non-routine cases; a successor may know what tasks need to be performed without understanding the reasoning that shaped them; or a technically competent replacement may still require repeated assistance from the former holder because contextual judgement was never reproduced elsewhere.

Knowledge is also not equally reproducible. Explicit information can often be represented with considerable fidelity, while capability grounded in context, experience, pattern recognition, judgement or relationships may require forms of reproduction that depend more heavily on interaction or practice. This does not make documentation intrinsically inadequate; it means that continuity depends on whether the mechanism corresponds to what the institution actually needs to reproduce. A document, mentoring arrangement, period of shadowing or formal handover can therefore be useful without necessarily being sufficient, particularly when the relevant capability is demonstrated only through independent performance.

Timing compounds the problem because many institutions respond to knowledge continuity primarily when departure becomes imminent. The feasible mechanism set narrows as the transition approaches. Where knowledge can be represented readily, a short transfer window may be sufficient; where reproduction depends on repeated interaction, practice or relationship transfer, recognising the dependency only at the point of departure may already be too late for the preferred pathway.

Responsibility can be equally ambiguous. Knowledge continuity often sits across human resources, operational management, records or knowledge functions, individual employees and successors. When responsibility is distributed in this way, substantial continuity activity can occur without anyone clearly owning the institutional outcome. The departing employee may consequently carry responsibility for a risk that properly belongs to the institution.

There is also a problem of proportionality. Institutions cannot and should not eliminate all person-dependency. Complex organisations need specialists, professional judgement and differentiated expertise, and attempting to duplicate every capability would impose substantial costs while potentially weakening the benefits of specialisation. Nor should every piece of accumulated knowledge be preserved indefinitely. Institutions therefore face a problem of discrimination as well as preservation. They need to distinguish valuable specialisation from excessive dependency and consequential knowledge from knowledge that can safely disappear, be reconstructed or be reacquired. Seniority, years of service or proximity to retirement cannot establish criticality on their own; what matters is the relationship between the affected function, knowledge concentration and substitutability, reproduction or reacquisition difficulty and the consequence if institutional access disappears.

Finally, organisations often verify activity more readily than capability. It is easy to establish that documents were uploaded, handover meetings occurred or overlap was completed; it is harder to know whether the institution can later interpret an unusual case, reconstruct an earlier rationale or perform a specialised function independently. The absence of immediate disruption is therefore not necessarily evidence that continuity succeeded.

Taken together, these dynamics explain why institutional knowledge loss cannot be addressed through better handovers alone. The underlying problem is a mismatch between the knowledge on which consequential functions depend and the pathways through which the institution can reproduce that knowledge when individual availability changes. Personnel transition exposes the mismatch; unsuitable mechanisms can leave it intact; late recognition can narrow the feasible response; diffuse ownership can delay action; and activity-based measures can conceal continued dependency. The question is therefore not simply whether knowledge was transferred, but whether the institution has become sufficiently capable of reproducing what the function still requires.

3. What the Evidence Shows

The evidence on institutional knowledge loss does not support a simple proposition in which employee departure automatically produces organisational decline, nor does it justify treating every concentration of expertise as a problem requiring intervention. What emerges instead is a conditional picture in which the consequences of transition depend on the knowledge involved, the function it supports, the availability of alternative sources, the difficulty of reproduction or reacquisition and the organisational conditions surrounding the transition. The evidence therefore supports treating person-dependent knowledge as a potential continuity risk while rejecting indiscriminate preservation as a general institutional response.

Knowledge loss through personnel transition is real, but its consequences are heterogeneous

Research across organisational settings provides substantial support for the proposition that employee turnover can result in the loss of valuable knowledge while also showing considerable variation in antecedents and consequences (Galan, 2023a). The relevant causal sequence is therefore not departure followed automatically by institutional damage. The same personnel event can be inconsequential where knowledge is substitutable or readily recoverable and materially disruptive where a consequential function depends upon concentrated knowledge for which no credible alternative pathway exists. Seniority or tenure may indicate accumulated experience, but neither establishes criticality on its own.

Public-sector evidence reinforces this distinction. Long-lived functions, specialised expertise and recurring mobility can make continuity particularly salient, but the appropriate unit of analysis remains the relationship between the person, the knowledge and the institutional function that depends upon it. Work on critical knowledge in public administration similarly suggests that risk cannot be inferred mechanically from retirement proximity, age or job title alone (Arimany-Serrat, Antentas-Peraile, & Tarrats-Pons, 2024).

Whether knowledge survives depends partly on the organisation around the individual

A second finding is that knowledge continuity cannot be explained exclusively by characteristics of the departing employee. Knowledge may already be distributed across teams, embedded in routines, represented in records or accessible through several people. Alternatively, formally similar organisations can allow expertise to remain concentrated because work is repeatedly routed through one specialist, responsibilities are highly individualised or opportunities for reproduction are limited. Empirical work on retiring employees also indicates that the usefulness of transferred knowledge depends materially on organisational conditions rather than on transfer behaviour alone (Biron, Turgeman-Lupo, & Zaid-Dominik, 2024).

Knowledge continuity is therefore influenced not only by what an individual knows but by the organisational environment through which that knowledge is represented, shared, practised and made accessible to others. This does not imply that every specialist capability should be deliberately distributed. It establishes the narrower point that person-dependency is partly an organisational condition and can, where justified, be modified through organisational action.

Different forms of knowledge require different continuity responses

The evidence also provides strong reasons to reject the idea that knowledge transfer is a uniform process. Explicit and relatively stable information can often be represented with considerable fidelity, while capability grounded in experience, judgement, repeated practice or relationships may be less separable from the conditions through which it is used. This does not mean that tacit knowledge cannot be reproduced or that documentation is intrinsically weak. It means that the reproducibility of a dependency depends partly on what must be represented, interpreted, practised, socially transmitted or reconstructed.

The evidence consequently does not support treating any single transfer instrument as universally sufficient. The preventive literature identifies multiple intervention families while also showing that comparative effectiveness remains limited and context-dependent (Galan, 2023b). A mechanism that reproduces one dimension of knowledge effectively may leave another dimension unresolved.

Continuity should be proportionate to institutional consequence

The evidence and counter-evidence converge on another conclusion: preserving knowledge has costs, and those costs matter. Redundancy consumes resources; overlap can be expensive; extensive documentation can become difficult to maintain; duplicated expertise may reduce the efficiency gained through specialisation; and preserving outdated practices can inhibit adaptation. Mobility can also introduce new experience and circulate knowledge between organisations. Institutional practice in high-consequence sectors similarly uses risk-based approaches to decide what knowledge warrants preservation effort rather than assuming that everything should be retained (International Atomic Energy Agency, 2012).

These considerations mean that the existence of person-dependency does not itself establish a case for intervention. The institutional question is whether the expected consequence and recovery burden of losing access justify the cost of creating an alternative pathway beforehand. A proportionate continuity approach must therefore permit not only preservation and reproduction, but also rational non-intervention, later reacquisition and deliberate acceptance of some knowledge loss.

Transfer activity is not evidence that institutional capability survived

Much of what organisations can readily observe during personnel transition consists of activities and outputs: documents completed, meetings held, successors appointed, training delivered, files transferred or overlap concluded. These measures can establish that an intervention occurred, but they cannot by themselves establish that the underlying dependency changed. Where the continuity risk arises because a consequential function depends excessively on one individual, the relevant outcome is whether the institution subsequently possesses another credible route through which the required knowledge can be accessed and applied. Administrative completion is therefore an insufficient proxy for continuity.

What the evidence does not establish

These findings provide a coherent foundation for institutional action but also establish important limits. The evidence does not identify a universally optimal knowledge-retention programme, does not support a single mechanism suitable for every dependency, does not establish that all tacit knowledge should be made explicit, does not show that redundancy is always preferable to specialisation, does not establish that continuity investment necessarily produces a positive institutional return and does not justify a universal numerical threshold at which a person-dependent capability becomes critical.

These limitations determine the form that responsible institutional action can take. The evidence supports conditional rather than universal intervention: the appropriate response depends on the function at risk, the characteristics and concentration of the relevant knowledge, available alternatives, reproduction or reacquisition difficulty, the consequences of loss and the burden of continuity action. It therefore points away from both passive reliance on normal staffing processes and comprehensive preservation and toward a risk-based approach in which continuity effort is directed toward material dependencies and judged ultimately by whether consequential institutional capability remains viable.

4. Intervention Options

If institutional knowledge loss results from excessive dependence on particular individuals, the purpose of intervention is not simply to extract as much information as possible before those individuals leave. It is to create another credible route through which the institution can continue to access, interpret, reproduce or, where appropriate, reacquire the knowledge required for a consequential function. The relevant choice is therefore among different ways of maintaining sufficient institutional access to capability-relevant knowledge, not between comprehensive transfer and unavoidable loss.

Representation

One family of interventions works through representation. Documentation, process descriptions, decision records, technical notes, structured knowledge elicitation and other durable representations can reduce dependency by separating useful knowledge from the continuing availability of its original holder. Representation is strongest where knowledge can be articulated with sufficient fidelity and where future users can interpret what has been preserved. More documentation does not necessarily create more continuity when competent performance depends materially on judgement, practice or relationships.

Social transfer

A second family works through social transfer. Structured handovers, mentoring, guided conversations, interviews and relationship introductions can surface reasoning, contextual knowledge and assumptions that might otherwise remain implicit. Social transfer can improve understanding without necessarily creating independent capability. It is therefore most useful when it contributes to an alternative pathway rather than being treated as evidence of continuity in itself.

Practice-based reproduction

Where knowledge depends heavily on experience, practice-based reproduction becomes more important. Shadowing, guided practice, temporary role overlap, rotation, simulation and progressive assumption of responsibility allow recipients to encounter situations in which knowledge is actually used. The value of these mechanisms lies in moving from description toward performance. Practice-based reproduction requires time and sufficiently representative opportunities; even prolonged overlap may fail to expose a successor to rare situations in which the most consequential judgement becomes necessary.

Distributed reproduction

A fourth family reduces vulnerability through distributed reproduction. Rather than moving knowledge from one indispensable person to another, institutions can develop additional holders,

cross-train colleagues, share responsibility for critical processes or create multiple pathways through which a function can be sustained. Distributed reproduction differs from simple succession because it seeks to reduce concentration rather than move the same dependency from one indispensable individual to another. Its resilience benefits must be weighed against the cost of maintaining additional capability, particularly for highly specialised or infrequently used knowledge.

Recovery or reacquisition

Institutions can also choose recovery or reacquisition where maintaining continuous internal access would be disproportionate. Knowledge that is inexpensive to reconstruct, available through external expertise, likely to become obsolete or associated with a low-consequence function may rationally be allowed to disappear, provided a credible route exists to restore sufficient capability if required. Knowledge can have institutional value without justifying permanent internal preservation.

These families are complementary rather than mutually exclusive. Complex dependencies may require representation for explicit information, social transfer for context, practice for judgement and distribution where concentration itself is material. Other dependencies may require only a concise record or a credible recovery route. The appropriate combination depends on what must survive and on the consequence of losing access.

Intervention should therefore not begin with a preferred tool. A repository, handover template, mentoring programme, succession process or knowledge-mapping exercise is useful only insofar as it activates the mechanism the particular dependency requires. The correct sequence is to understand the dependency first and select the intervention second.

The intervention families should not be read as a hierarchy. Representation is not inherently inferior to practice-based reproduction, distribution is not automatically preferable to succession and recovery is not evidence of continuity failure. A mechanism is appropriate insofar as it provides sufficient reproduction for the dependency and consequence involved.

Timing further constrains what remains feasible. Planned transitions may permit repeated interaction, practice and gradual relationship transfer; short-notice departures may require prioritisation and acceptance of incomplete reproduction; sudden unavailability removes direct transfer altogether and shifts the institution toward reconstruction, external expertise, temporary workarounds or redesign.

The intervention space should therefore be understood as a set of proportionate ways to alter institutional dependency rather than as a catalogue of preferred knowledge-management practices. The institutional task is to understand the dependency first and then select the lightest credible pathway capable of providing minimum sufficient continuity. “Lightest” does not mean cheapest in isolation; it means the least burdensome pathway that can credibly maintain or restore the level of capability the institutional consequence requires.

5. Recommendations

Preventing institutional knowledge loss requires institutions to distinguish ordinary specialisation from material dependency. The objective is neither to preserve as much knowledge as possible nor to create a transfer process for every personnel transition, but to recognise where consequential functions depend excessively on particular knowledge holders and to intervene where the likely consequences of losing that access justify the burden of reproduction. Institutions should therefore govern material person-dependent knowledge as a continuity risk, creating proportionate alternative pathways through which consequential capability can survive personnel transition.

This changes the unit of attention from the characteristics of the employee to the vulnerability of the institutional function. Continuity effort should follow the consequence of losing access, not hierarchy, tenure or the volume of knowledge an individual possesses.

See the Risk Before the Transition

Institutions should identify material knowledge dependencies before personnel transitions expose them through failure. This does not require an exhaustive inventory of organisational knowledge. A proportionate approach begins with consequential functions and asks where continued performance depends heavily on knowledge accessible through only one person or a very small number of people.

The diagnostic question should therefore be functional rather than biographical: if this knowledge holder became unavailable tomorrow, what would the institution struggle to understand, decide or do? The purpose is not to label individuals as indispensable, but to make the underlying dependency visible and determine whether another credible route already exists.

Once identified, dependencies should be prioritised according to institutional consequence and vulnerability. Relevant considerations include the importance of the affected function, concentration and substitutability of the required knowledge, transition exposure, reproduction or reacquisition difficulty and the consequences if access disappears. These considerations should support judgement rather than a universal numerical score. A low-consequence dependency involving readily recoverable knowledge may justify no intervention, while difficult-to-reproduce knowledge supporting a consequential function may warrant a substantially stronger response.

Design the Right Continuity Response

Where intervention is justified, institutions should match continuity mechanisms to the reproduction burden of the capability that must survive. The relevant question is not simply what information should be captured, but what another person, team, representation or external pathway would need in order to support the function after the original holder becomes unavailable. Mechanisms should be selected because they address the dependency, not because they are familiar organisational tools.

Institutions should create sufficient alternative pathways rather than assume that continuity requires complete duplication of an outgoing employee's expertise. Different components of capability may remain accessible through different routes, provided those routes collectively allow the consequential function to continue. Where consequence warrants,

institutions should also consider whether succession merely transfers the same vulnerability to a new single holder. Additional pathways may be justified in high-consequence cases, but universal redundancy would be disproportionate. The objective is sufficient institutional reproducibility, not maximum duplication.

Make Continuity Actionable

Institutions should activate continuity measures early enough for the selected reproduction mechanism to work. The continuity window should therefore follow the requirements of the knowledge and pathway involved rather than a standard administrative handover timetable. Early recognition need not trigger immediate intensive transfer; it may simply preserve the conditions required for stronger action if transition becomes probable.

Continuity also requires identifiable institutional ownership. Departing employees can contribute to reproduction, but they should not carry sole responsibility for deciding what is institutionally consequential or whether continuity has been achieved. An owner of the affected function should be capable of authorising proportionate action, coordinating responsibility and accepting any residual risk that remains. This does not require a new knowledge-continuity office. Existing managerial, operational, programme, HR or knowledge functions can perform these roles according to institutional context. What matters is functional clarity rather than organisational form.

Test Whether Capability Survived

Institutions should verify continuity after transition rather than treat completion of transfer activity as evidence of success. The relevant question is whether another credible pathway now allows the consequential function to operate without excessive dependence on the former holder.

Where verification reveals remaining gaps, institutions should distinguish between dependency requiring further action, residual risk that can be managed through a credible recovery pathway and vulnerability that can reasonably be accepted. The important requirement is that remaining dependency should be visible and owned rather than disappearing from attention when the personnel process formally ends. Post-transition experience should also inform future continuity decisions, particularly where the selected mechanism proved inadequate or merely relocated dependency.

Taken together, these recommendations create a deliberately selective model of institutional knowledge continuity: see material dependency before it becomes a transition failure, decide whether the consequence justifies intervention, determine what capability must survive, create a proportionate alternative pathway, act while the necessary reproduction conditions remain available, assign responsibility for the continuity outcome and verify afterwards that sufficient capability remains. The objective is not to eliminate specialisation or preserve every accumulated form of knowledge, but to reduce consequential person-dependency only to the extent required to keep important institutional functions viable.

6. From Recommendation to Institutional Action

Translating knowledge continuity into institutional practice does not require every organisation to create a new programme, governance body or administrative process. The necessary functions can often be incorporated into existing arrangements around operational risk, workforce planning, succession, programme management, records, professional development or organisational change. What matters is whether the institution can move reliably from recognising a consequential dependency to creating an adequate alternative pathway and determining afterwards whether the relevant capability survived. A common continuity logic can therefore operate through different institutional forms.

Institutions will begin from different conditions. Some may discover dependencies only when personnel announce their departure; others may know informally which specialists are difficult to replace; some may already operate handover or succession practices without testing functional continuity; and others may already identify material dependencies but need stronger mechanism selection or verification. Implementation should begin from existing capability rather than require every institution to construct the same organisational model.

The first operational question is who owns the continuity outcome. The affected function requires an identifiable actor capable of judging consequence, authorising action and accepting residual risk. This may be a programme director, operational manager, process owner or another appropriate institutional actor; it need not be a knowledge-management specialist. What matters is that continuity is not simultaneously everyone's concern and no one's responsibility.

The institution should then establish where consequential functioning depends excessively on particular knowledge holders. This does not require mapping all organisational knowledge. Analysis can begin from critical functions, ordinary risk or workforce discussions, known transitions, periods of unexpected absence or near-misses that reveal weaker distribution than previously assumed. The objective is sufficient visibility to support a decision, not exhaustive knowledge capture.

Once visible, the dependency should be subjected to a proportionality judgement. Relevant considerations include functional consequence, concentration, available substitutes, transition exposure, reproduction or reacquisition difficulty and the burden of intervention. These factors should inform judgement rather than operate as a mechanical score. Legitimate outcomes include no action, light intervention, a more substantial continuity response or reliance on a credible recovery pathway.

Where action is justified, the institution should determine what must actually survive. The question is not what the individual knows in general, but what another person, team, representation or external pathway would need in order for the consequential function to remain viable. The analysis should be detailed enough to establish the reproduction burden, but no more elaborate than the decision requires.

The institution can then build the lightest credible alternative pathway capable of meeting that burden. One or more of the intervention families described earlier may be used, but combination is not an objective in itself. A more elaborate intervention is not inherently a better one. The appropriate design is the least burdensome arrangement that credibly maintains or restores the level of capability justified by the institutional consequence.

Activation should occur while the selected pathway remains feasible. Where reproduction requires repeated practice, progressive responsibility or relationship development, action must begin while those opportunities still exist; where knowledge is readily representable, a shorter

window may be sufficient. Institutions can recognise a dependency well in advance without immediately launching an intensive intervention, activating stronger measures only when justified but before the necessary reproduction conditions disappear.

Implementation also requires sufficient readiness: an owner capable of acting, enough understanding of the dependency, feasible access to recipients or alternative pathways and the information or infrastructure required by the chosen mechanism. Readiness should be treated proportionately. Weaknesses can sometimes be corrected during implementation rather than becoming a reason for inaction.

Continuity arrangements should be reconsidered when material conditions change. A dependency may become more consequential if alternative holders disappear, or less consequential if technology, organisational redesign or external availability makes the knowledge easier to reproduce. Continuity mechanisms should therefore not be maintained simply because they were once created.

Personnel transition exposes the planned pathway to reality. At that point the institution should know what alternative routes are expected to remain and what uncertainty still exists. Continuity begins to become institutional when the required knowledge can be accessed and used through a credible pathway that no longer depends exclusively on the original holder.

Not every transition permits preventive reproduction. Sudden resignation, illness, unexpected reassignment or other immediate unavailability can remove the possibility of direct transfer. The institution should then move into recovery mode, reconstructing what it can from records, colleagues, previous decisions and communications, external expertise or other available sources, while using temporary workarounds or redesign where necessary. Failure to complete preventive transfer changes the mechanism set; it does not eliminate the possibility of continuity action.

Complete reproduction will not always be feasible or proportionate. Institutions may rationally retain only part of a capability, rely on reacquisition for rare needs or accept temporary reduced performance. For material dependencies, however, residual vulnerability should be explicit: the institution should know what remains dependent, why the gap persists, what consequences could follow, what recovery route exists and who owns the decision to accept the remaining risk.

Implementation may consequently vary from light to standard or heightened intensity. A light response may require only accessible representation, an alternative route and basic verification; a standard response may combine structured analysis, transfer and explicit verification; and a heightened response may justify multiple pathways, more extensive practice or stronger redundancy where consequence warrants. These are illustrative differences of intensity, not maturity levels, mandatory tiers or numerical thresholds.

The same proportionality should govern scale. A continuity practice that works for one critical function should not automatically become an institution-wide administrative requirement. Expansion is justified where the underlying dependency pattern recurs and where coordination benefits exceed the additional burden; otherwise continuity can remain embedded locally within existing processes.

The pathway could later support an Institutional Knowledge Continuity Assessment examining dimensions such as functional criticality, concentration, substitutability, transition exposure, reproduction difficulty and recovery burden. Such an assessment should support judgement rather than create a universal risk score, and it is not required for institutions to apply the continuity logic developed here.

The implementation challenge is therefore less about constructing a new administrative system than about introducing a disciplined sequence of institutional decisions: own the risk, identify the dependency, decide whether it matters enough to act, determine what must survive, build the lightest credible alternative pathway, activate it while the necessary conditions remain

available and test the result after transition. Institutions can perform these functions through different structures and instruments, and most dependencies should require only limited intervention. What remains constant is the causal logic: continuity action is justified where individual dependency creates material institutional vulnerability and succeeds only where sufficient institutional capability remains after individual availability changes.

7. How to Know Whether Capability Has Improved

Knowledge-continuity initiatives are relatively easy to measure when success is defined in terms of activity. Institutions can record completed transfer plans, documents produced, mentoring participation, periods of overlap or roles reviewed before transition. Such measures can establish whether an intervention was implemented, but they answer only the first evaluation question. A meaningful assessment must distinguish between implementation, activation of the intended mechanism, improvement in institutional capability and the broader net effect of the intervention.

At the implementation level, the institution should establish whether the planned continuity response actually occurred. Relevant representation should have been produced and made accessible, intended opportunities for practice or relationship transfer should have taken place, and any planned redistribution of responsibility should have occurred. These observations matter because an intervention that was never implemented cannot produce its intended effect. Yet implementation evidence remains evidence of activity, not evidence that the underlying dependency was resolved.

The second level concerns mechanism activation. Here the question is whether the intervention has actually created another credible route to the knowledge required by the function. Another employee may now be able to perform the work, a team may possess sufficient distributed understanding, contextual information may be recoverable without the original holder, a relationship may continue through another institutional actor or an external source may provide credible reacquisition. These are stronger indicators because they show that the condition creating the continuity risk has begun to change.

Mechanism activation is nevertheless not the final outcome. An alternative pathway can appear credible before transition and still prove inadequate after the original holder becomes unavailable. The decisive test is therefore capability gain: can the institution continue to perform the consequential function at the required level? Can relevant knowledge be found, understood and applied; can important exceptions and non-routine situations be handled; can critical relationships continue; and can these things occur without repeated default reliance on the former holder?

Verification timing should reflect the function being evaluated. Gaps in frequently used knowledge may become visible quickly, while capabilities required only during unusual events, annual cycles, complex disputes or infrequent technical problems may take much longer to test. Institutions should therefore avoid universal verification deadlines. Until the reproduced capability has encountered a meaningful test, the most accurate conclusion may simply be that transfer has been completed but capability remains unverified.

Continuity need not be treated as binary. Transition may reveal material capability failure, partial continuity with consequential gaps, sufficient continuity accompanied by consciously accepted residual dependency, or strong continuity through functioning alternative pathways. Sufficient continuity does not mean reproducing everything an individual knew; it means retaining enough capability to perform what materially matters.

Evaluation should also determine whether dependency was reduced or merely relocated. Successful transfer to a successor may restore immediate continuity while making that successor the new exclusive holder of the same knowledge. This may be entirely proportionate for a low-consequence function, but it should not automatically be interpreted as increased resilience where the original concentration was itself considered material.

Remaining gaps should generate an institutional decision rather than disappear when the transition process ends. Further practice, representation, another holder, reconstruction, external expertise or redesign may be appropriate, but further intervention may also be disproportionate. The important distinction is between accepted residual risk and invisible residual risk: an institution that understands the remaining vulnerability, its consequence and its recovery route is in a fundamentally different position from one that assumes continuity because administrative transition has concluded.

Genuine capability improvement does not automatically imply a positive net institutional effect. Reducing dependency consumes resources; additional holders can weaken specialisation efficiencies; prolonged overlap can impose disproportionate costs; extensive continuity requirements can divert effort from current work; and preservation can reinforce practices that should instead evolve. Continuity should therefore be judged not only by whether capability survived, but by whether the gain was proportionate to the institutional burden required to produce it.

This broader perspective also protects knowledge continuity from becoming an objective of preservation for its own sake. Institutions need renewal as well as continuity. New employees may question inherited assumptions, introduce expertise from elsewhere and replace routines that no longer serve their purpose. The objective is not to ensure that an institution knows tomorrow exactly what it knew yesterday. It is to prevent avoidable loss of knowledge whose disappearance would materially weaken the institution while allowing knowledge that is obsolete, replaceable or no longer valuable to change or disappear. Continuity without adaptation can become stagnation; adaptation without sufficient continuity can become institutional amnesia.

In practice, evaluation can therefore ask whether material dependencies remain visible, whether credible alternatives actually function, whether consequential work continues under routine and non-routine conditions, whether excessive dependency has fallen rather than moved, whether residual gaps are understood and owned, and whether the continuity achieved justified its institutional burden. These questions provide a stronger basis for assessing progress than the volume of captured knowledge or the number of completed transfer activities.

Evaluation should also improve subsequent continuity decisions. Failed pathways can reveal underestimated reproduction difficulty; unnecessarily elaborate interventions can expose overestimated consequence; repeated reliance on former holders can reveal hidden dependency; and successful light-touch interventions can demonstrate where limited formalisation is sufficient.

The ultimate test is straightforward even if measuring it is not always easy: when a consequential knowledge holder becomes unavailable, can the institution still do what it materially needs to do, with sufficient understanding, reliability and independence? If the answer is yes, and the resources required to achieve that continuity were proportionate to the risk, institutional capability has improved. If the institution has merely completed the prescribed activities while remaining dependent on the same individual knowledge pathway, it has not. This distinction between administrative completion and functioning capability is what turns knowledge continuity from a transfer exercise into an institutional capability.

8. Conclusion

Institutional knowledge loss is most visible when experienced people leave, but the underlying vulnerability usually begins much earlier. An institution can continue to perform successfully while important parts of its capability remain concentrated in particular individuals because their continued availability compensates for weaknesses in institutional reproducibility. Personnel transition removes that compensation and reveals what the organisation could not independently recover, interpret or perform. Preventing knowledge loss should therefore begin by identifying where consequential functions depend excessively on individual access and determining whether that dependency is important enough to justify intervention.

The objective is not to capture everything people know, eliminate specialisation or create comprehensive redundancy. Institutions benefit from expertise, mobility and renewal, and some knowledge can disappear without meaningful consequence or be reacquired more efficiently when required. What institutions need is minimum sufficient continuity: enough alternative access to capability-relevant knowledge to prevent a foreseeable personnel transition from unnecessarily weakening a consequential function.

Achieving that proportionality requires institutions to recognise material dependencies before transition exposes them, determine what must survive, select mechanisms appropriate to the reproduction burden, act while those mechanisms remain feasible, assign responsibility for the continuity outcome and verify afterwards whether sufficient capability remains. Documentation, mentoring, overlap and succession can all contribute, but none should be mistaken for continuity itself.

An institution does not become more capable simply by transferring knowledge; it becomes more capable when consequential knowledge is no longer excessively dependent on the continued availability of particular individuals. That outcome may be achieved through another capable holder, distributed expertise, durable representations, functioning relationships, external reacquisition or a combination of pathways. It does not require perfect replication of everything the original holder knew. It requires sufficient institutional reproducibility for the relevant function to continue with an acceptable level of understanding, reliability and residual risk.

Knowledge continuity should also remain compatible with institutional adaptation. Preserving obsolete routines, creating unnecessary duplication or imposing extensive transfer requirements on low-consequence roles would substitute one form of institutional weakness for another. The purpose is not to make organisations resistant to change, but to ensure that change does not unintentionally destroy capabilities they still need. Continuity and renewal are therefore not opposing objectives: effective institutions must preserve what remains consequential while allowing knowledge, practices and expertise to evolve where change creates greater value.

For public institutions in particular, this is ultimately a question of whether knowledge resides only in the people who currently perform institutional work or becomes, where consequence requires it, sufficiently reproducible beyond them. The institutional capability to be built is not the ability to prevent knowledge from ever being lost, but the ability to ensure that the knowledge which materially sustains institutional functioning can survive the movement of the people who currently hold it.

References

Arimany-Serrat, N., Antentas-Peraile, M., & Tarrats-Pons, E. (2024). Protocol for identifying and retaining critical knowledge in a public health administration. *Systems*, 12(11), 505. <https://doi.org/10.3390/systems12110505>

Beazley, H., Boenisch, J., & Harden, D. (2003). Knowledge continuity: The new management function. *Journal of Organizational Excellence*, 22(3), 65–81. <https://doi.org/10.1002/npr.10080>

Biron, M., Turgeman-Lupo, K., & Zaid-Dominik, O. (2024). Contextualizing the usefulness of knowledge received from retiring employees: Leader behaviour and organisational culture. *Knowledge Management Research & Practice*, 22(6), 620–631. <https://doi.org/10.1080/14778238.2023.2297060>

Galan, N. (2023a). Knowledge loss induced by organizational member turnover: A review of empirical literature, synthesis and future research directions (Part I). *The Learning Organization*, 30(2), 117–136. <https://doi.org/10.1108/TLO-09-2022-0107>

Galan, N. (2023b). Knowledge loss induced by organizational member turnover: A review of empirical literature, synthesis and future research directions (Part II). *The Learning Organization*, 30(2), 137–161. <https://doi.org/10.1108/TLO-09-2022-0108>

International Atomic Energy Agency. (2012). Knowledge management for nuclear research and development organizations (IAEA-TECDOC-1675). Vienna: IAEA.

Phaladi, M. P., & Ngulube, P. (2024). Understanding tacit knowledge loss in public enterprises of South Africa. *SA Journal of Human Resource Management*, 22, a2229. <https://doi.org/10.4102/sajhrm.v22i0.2229>

U.S. Government Accountability Office. (2026). GAO Overview: Recent federal workforce changes at OPM (GAO-26-108916). Washington, DC: U.S. Government Accountability Office.



IDHUS
Institute for the Development
of Human Societies