Addressable Memory Networks

Governed Institutional Memory and the Path Toward Knowledge Resolution

Public Manuscript v3.0

Addressable Memory Network Research Group

July 28, 2026

Abstract

Modern organizations have become exceptionally good at storing information, increasingly good at retrieving it, and still surprisingly poor at preserving what they know. Documents, databases, emails, reports, contracts, logs, and AI-generated artifacts may all be searchable, yet the organization often lacks a durable way to determine what is currently believed, why it is believed, where it came from, whether it has changed, whether it has been superseded, and whether it is safe to act upon.

This paper introduces Addressable Memory Networks (AMN) as a framework for governed institutional memory. AMN is not proposed as a replacement for databases, search engines, retrieval-augmented generation systems, or knowledge graphs. Instead, it describes a complementary memory layer in which organizational knowledge can be represented as addressable, traceable, revisable, and governable knowledge objects.

The central argument is that AI-augmented organizations require more than retrieval at question time. They require memory that can preserve provenance, lifecycle state, revision history, uncertainty, stewardship, and operational status over time. We position AMN in relation to retrieval-augmented generation, knowledge graphs, agent memory systems, and provenance frameworks, and we describe a conceptual architecture for transforming heterogeneous sources into governed memory.

The paper further argues that governed institutional memory is a prerequisite for a broader possibility: knowledge resolution. If memory objects become persistently addressable and governed, systems may eventually move beyond locating documents toward resolving what is known, what supports it, what contradicts it, and how it has evolved. While the present work remains focused on institutional settings, the same principles suggest a future in which independently governed memory systems may become interoperable without requiring a single global repository or universal authority.

1 Introduction

Organizations have solved many problems of information storage. They have not solved institutional memory.

A modern enterprise may contain millions of documents, contracts, invoices, reports, tickets, meeting notes, spreadsheets, source-code repositories, audit logs, customer records, scientific publications, project decisions, and machine-generated artifacts. These materials are increasingly indexed, embedded, summarized, and made available to large language models (LLMs). They can often be searched within seconds.

Yet when an organization asks a higher-level question, the weakness becomes visible:

> What do we currently know? Why do we believe it? Which evidence supports it? Has it been superseded? Which version should an agent or employee act upon?

These are not merely retrieval questions. They are memory-governance questions.

Search systems find documents. Databases retrieve records. Knowledge graphs represent entities and relationships. Retrieval-augmented generation (RAG) systems assemble relevant context for a prompt. These capabilities are useful and necessary, but none of them, by themselves, defines what an organization should remember as operational knowledge, how that memory changes, who or what may rely on it, and how it should be audited over time.

The gap matters because AI systems are becoming operational actors. When models and agents summarize reports, draft contracts, reconcile invoices, plan engineering work, answer regulatory questions, or execute workflows, the organization needs more than plausible answers. It needs a reliable memory substrate that preserves the lineage and status of the knowledge being used.

This paper introduces Addressable Memory Networks (AMN) as a framework for governed institutional memory. The core idea is simple: important knowledge should not remain trapped as an incidental fragment inside a document, prompt, embedding, or generated answer. It should become an addressable memory object with identity, provenance, lifecycle state, revision history, and governance context.

AMN is not presented as a universal truth system. Nor is it a claim that all knowledge can be centralized. Its purpose is narrower and more practical: to help bounded organizations preserve, revise, inspect, and act upon knowledge without repeatedly reconstructing it from raw artifacts.

At the same time, the implications are broader. Once memory becomes governed and addressable inside an institution, the possibility of interoperability naturally emerges. Independent organizations may eventually need ways to reference governed knowledge across boundaries. In that sense, AMN begins as a framework for institutional memory, but it also points toward a longer-term question: whether memory itself can become interoperable.

1.1 Contributions

This paper makes the following contributions:

1. **Governed institutional memory.** We define the institutional memory gap and argue that modern AI-augmented organizations require a memory layer distinct from storage, retrieval, and generation.
2. **Addressable knowledge objects.** We introduce the idea of persistently referenceable knowledge objects that carry provenance, lifecycle state, revision history, and governance context.
3. **A conceptual AMN architecture.** We describe a layered architecture for transforming heterogeneous sources into governed memory through interpretation, validation, admission, and use.
4. **Memory lifecycle and stewardship.** We frame institutional knowledge as something that moves through a lifecycle rather than something merely stored once.
5. **Knowledge resolution.** We introduce knowledge resolution as a future-facing capability: resolving what is currently understood, why it is believed, what supports it, and what challenges it.
6. **Interoperable memory networks.** We outline how governed memory could eventually support interoperability between independently governed memory systems while preserving autonomy.

2 The Problem of Institutional Memory

Most enterprise systems are optimized around storage, workflow, or retrieval. They answer questions such as:

- Where is the file?
- Which record matches this query?
- Which document contains this phrase?
- Which database row contains this value?
- Which text chunk is semantically similar to this prompt?

These questions matter, but they do not capture the full problem of organizational knowledge. Institutions also need to answer:

- What do we currently believe?
- Why do we believe it?
- Which source supports it?
- Which version is current?
- What has been revised or retired?
- What evidence contradicts it?
- Who or what is allowed to act on it?
- Can this memory be audited after the fact?

These questions require more than retrieval. They require governed memory.

2.1 Information Fragmentation

Organizational knowledge is scattered across many systems. Contracts may live in one repository, invoices in another, emails in a third, decisions in chat tools, and operational data in databases. A human can sometimes reconstruct the picture by searching across these systems, but the result is slow, inconsistent, and often unrepeatable.

This fragmentation is not only technical. It is institutional. Different teams maintain different local truths. Legal, finance, operations, research, engineering, and customer support may each hold partial knowledge about the same situation.

2.2 Provenance Decay

Over time, organizations often remember conclusions but lose the evidence chain behind them. A spreadsheet cell, policy note, model output, or executive summary may survive long after the original source has been forgotten.

This produces an obvious risk: the organization may act on knowledge whose origin, confidence, or scope is no longer understood. In regulated or high-stakes environments, this is not a minor inconvenience. It is an operational and legal exposure.

2.3 Memory Drift

Knowledge changes. Contracts are amended. Payments are made. scientific evidence evolves. Engineering decisions are superseded. Regulations change. Internal policies are revised. AI agents produce new artifacts that may depend on earlier assumptions.

Without lifecycle management, old information remains retrievable and may appear just as convincing as current information. Retrieval alone cannot reliably distinguish active memory from retired memory.

2.4 The AI Amplification Problem

LLMs make the institutional memory gap more visible. They can summarize, transform, and reason over retrieved material, but they can also produce confident answers from stale, contradictory, or weakly supported context.

As AI systems become agents that plan and act, the problem becomes sharper. An agent does not merely need relevant information. It needs governed memory: information whose status, source, authority, and operational permission are inspectable.

2.5 The Confident Miss

A distinct and underappreciated risk arises whenever a system narrows a candidate set before returning a result: through classification, routing, ranking, or access scoping. Narrowing is necessary for any system operating at scale, but an incorrectly narrowed operation does not fail visibly. It returns a plausible, well-formed answer, drawn confidently from the wrong subset of the underlying knowledge, with no signal that anything went wrong. We term this failure mode the **confident miss**: an answer that is wrong not because no answer was found, but because the system stopped looking in the wrong place and reported success anyway.

The confident miss is distinguishable in kind from an absent result. A system that returns nothing, or that reports low confidence, exposes its own incompleteness and invites correction. A confident miss does not. It is, in this sense, the governance-layer counterpart to hallucination in generative retrieval: hallucination is a generation-time failure to ground an answer in evidence; the confident miss is a narrowing-time failure to search the evidence that was actually relevant, while presenting the result with the same authority as a correct one.

The risk is general rather than specific to any one mechanism. Any process that can *exclude* candidates before returning a result is a candidate site for this failure; processes that only ever add candidates cannot produce it, by construction. This has a direct design consequence: narrowing operations should be treated as a distinct risk surface requiring their own instrumentation, independent of how well-designed the narrowing criteria appear on paper (Section 10).

3 Related Systems and Their Limits

AMN is adjacent to several existing technical traditions. It is not intended to replace them. Rather, it occupies a different layer.

3.1 Databases

Databases are excellent at storing structured records and supporting transactional consistency. They are less suited to representing the lifecycle of interpreted knowledge across heterogeneous sources. A database may contain a contract row, a payment record, or an audit entry. It does not automatically define what the organization currently knows as governed memory.

3.2 Search Engines

Search engines locate documents or passages. They are indispensable for discovery, but they do not decide whether a statement is supported, current, superseded, contradicted, or approved for action.

3.3 Retrieval-Augmented Generation

Retrieval-Augmented Generation [1] retrieves relevant passages and supplies them to a generative model. RAG improves model grounding, but it is primarily a query-time architecture. It assembles context for an answer. It does not, by itself, create a durable governed memory object with lifecycle state and institutional status.

3.4 Graph-Based Retrieval

Graph-based retrieval approaches, including GraphRAG [2], organize corpora into graph structures to improve sensemaking across large document collections. These approaches are valuable, especially for broad questions over complex corpora. However, retrieval over a graph is still not the same as governing organizational memory.

3.5 Knowledge Graphs

Knowledge graphs represent entities and relationships [7, 8]. They are powerful tools for structured knowledge representation and reasoning. AMN can use graph structures, but it is not reducible to a graph. The distinguishing concern is the governance of memory objects over time: provenance, lifecycle, revision, contradiction, and permission to act.

3.6 Agent Memory Systems

Agent memory systems show that long-lived agents benefit from persistent memory, reflection, and retrieval [4, 3]. AMN extends the concern from individual or conversational agent memory to institutional memory shared by many humans, applications, and agents under explicit governance constraints.

3.7 Provenance Frameworks

The FAIR principles [5] and provenance models such as W3C PROV-O [6] provide important foundations for traceability and reuse. AMN is aligned with these principles but applies them to the formation and lifecycle of operational memory.

Table 1: AMN in relation to adjacent systems.

System	Primary function	Typical unit	Limit for governed memory
Database	Store and query records	Row, table, transaction	Does not by itself govern interpreted knowledge across sources
Search engine	Locate documents or passages	Document, page, passage	Finds material but does not establish lifecycle or trust status
RAG	Retrieve context for generation	Chunk, passage, prompt context	Often ephemeral and answer-centered rather than memory-centered
Knowledge graph	Represent entities and relations	Entity, edge, relation	May lack provenance, revision, stewardship, and action status
Agent memory	Preserve agent experience	Observation, reflection, memory note	Often local to an agent rather than institutional and auditable
AMN	Govern institutional memory	Governed knowledge object	Requires domain setup, validation rules, and stewardship

4 Addressable Memory Networks

An Addressable Memory Network is a framework for converting heterogeneous source material into governed memory objects that can be persistently referenced, traced to evidence, revised over time, and used by humans, applications, and AI agents.

The term *addressable* means that a memory object has a stable reference within an organizational knowledge environment. It should not exist only as a phrase inside a report, a paragraph inside a contract, an embedding vector, or a generated answer. It should be possible to refer to it, inspect it, revise it, retire it, and audit its history.

4.1 From Information to Governed Memory

AMN treats memory formation as a process:

1. Source material enters the system.
2. Candidate knowledge is interpreted from the source.
3. Evidence is attached.
4. Validation and governance rules are applied.
5. Accepted objects enter governed memory.

6. Objects are revised, challenged, or retired as new information arrives.

The crucial distinction is between interpretation and admission. AI systems may assist in proposing memory candidates, but the act of admitting something into governed memory requires validation.

4.2 Governed Knowledge Objects

The central unit of AMN is the governed knowledge object.

A governed knowledge object is a persistently referenceable unit of institutional knowledge whose source, status, lifecycle, and governance context remain inspectable over time. It may represent a conclusion, obligation, policy interpretation, measurement, finding, decision, definition, or other unit of knowledge.

This paper intentionally avoids implementation-specific schemas. At the public conceptual level, a governed knowledge object should minimally support:

- a stable identity;
- source provenance;
- evidence references;
- lifecycle state;
- revision history;
- uncertainty or contradiction markers;
- stewardship or authorization context;
- operational use constraints;
- a placement or classification, expressed as a confidence-weighted assignment rather than a single fixed category.

On this last point, classification deserves a brief note of its own. It is tempting to treat a knowledge object's category, domain, or topical placement as a fixed attribute assigned once at admission time. Subsequent experience with governed memory systems suggests this is too rigid an assumption: a nontrivial share of real knowledge objects legitimately belong to more than one classificatory context at once, and forcing a single assignment discards information rather than simplifying it. AMN therefore treats placement as a governed, revisable, and potentially plural judgment, distinct from the object's identity, which remains singular and stable throughout.

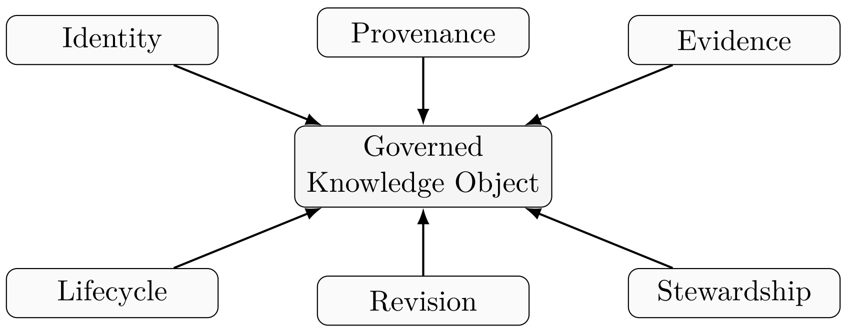

Figure 1: Conceptual components of a governed knowledge object.

5 A Conceptual AMN Architecture

AMN can be described as a layered architecture. The layers are conceptual rather than prescriptive; implementations may vary by domain and organization.

1. **Source layer.** Documents, databases, logs, emails, reports, repositories, workflows, and other raw artifacts.
2. **Interpretation layer.** Parsers, extractors, models, and human processes that propose candidate knowledge objects.
3. **Governance layer.** Validation, evidence binding, contradiction checks, policy constraints, and admission control.
4. **Memory layer.** Persistent governed knowledge objects, lifecycle states, revision history, and lineage.
5. **Use layer.** Search, dashboards, audit tools, APIs, applications, and AI agents that query and act upon governed memory.

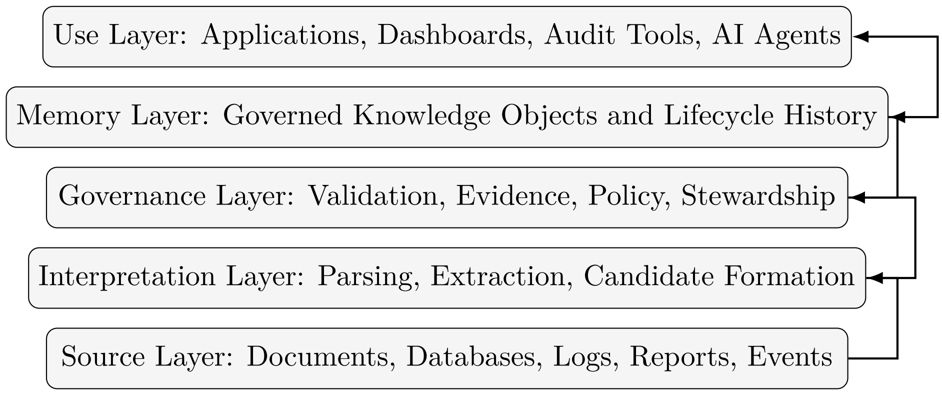

Figure 2: Conceptual layered architecture for Addressable Memory Networks.

5.1 Admission Rather Than Extraction Alone

Many systems stop at extraction. They parse a document, extract fields or entities, and make those available to downstream tools. AMN treats extraction as only the beginning. A candidate must still be checked for source support, conflict, lifecycle implications, and permitted use before becoming operational memory.

This separation is important because generative systems are useful but fallible. They may propose candidates, but they should not automatically canonize them.

5.2 Memory-Time Governance

Retrieval-time systems decide what to retrieve for a given question. AMN introduces the complementary idea of memory-time governance: deciding what should enter memory, what status it should hold, and how it should evolve before any particular question is asked.

This reduces reliance on repeated reconstruction. When a future query arrives, the system does not need to rediscover the entire evidence chain from scratch. It can resolve against governed memory.

6 Memory Lifecycle and Stewardship

Institutional memory should not be treated as a static store. Knowledge objects move through lifecycle states as they are observed, interpreted, supported, accepted, revised, challenged, or retired.

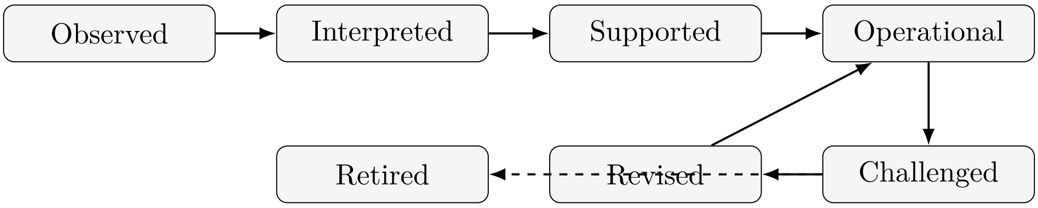

Figure 3: Illustrative lifecycle for governed institutional memory.

6.1 Stewardship

Stewardship refers to the organizational responsibility for maintaining the quality and status of memory. It need not imply a single central authority. Different domains may require different stewards, rules, and review patterns.

A legal memory object may require different admission rules than an engineering memory object. A medical conclusion may require different evidence handling than a customer-support policy. AMN therefore treats stewardship as domain-specific.

Subsequent validation of the framework at personal scale (Section 11) did not surface any failure attributable to the governance or stewardship layer itself; the corrections that emerged were confined to the addressing and placement mechanisms discussed above, which is consistent with treating stewardship and governance as a separately sound layer.

7 Enterprise and Institutional Scenarios

The purpose of AMN becomes clearer in concrete institutional settings.

7.1 Contracts and Obligations

A contract may contain renewal dates, penalties, payment schedules, service obligations, and conditional clauses. A search system can locate the contract. A language model can summarize it. But the organization may still lack a governed memory object representing the current obligation, its source clause, its version, and its operational status.

In an AMN-style system, the obligation can become an addressable memory object. It can be linked to evidence, updated when an amendment arrives, flagged when a payment event changes its state, and audited later.

7.2 Invoices and Financial Operations

An invoice is not only a document. It may imply a payable obligation, a disputed charge, a contractual dependency, a tax record, or an approval workflow. AMN can represent these as governed memory objects rather than leaving each system to reconstruct meaning from the raw invoice.

7.3 Research and Technical Knowledge

Research organizations often accumulate findings that evolve over time. A result may be supported by one report, contradicted by another, and revised by later evidence. AMN can preserve not only the finding but also its support, dispute status, and revision history.

7.4 AI-Agent Workspaces

Autonomous agents generate plans, code, summaries, decisions, and intermediate artifacts. Without governed memory, these artifacts remain scattered across logs and chats. With AMN, significant decisions and outputs can become memory objects linked to review, evidence, and later revisions.

8 Knowledge Resolution

The concept of knowledge resolution is central to the longer-term significance of AMN.

A database retrieves records. A search engine finds documents. DNS resolves names to network locations. AMN may eventually allow systems to resolve knowledge: to move from a question or reference to the governed knowledge object, its evidence, its status, its history, and its competing interpretations.

Knowledge resolution does not mean producing a universal truth. It means answering a more structured set of questions:

- What is currently known or accepted in this context?
- What supports it?
- What contradicts it?
- What changed since the previous version?
- Which steward or process admitted it?
- Is it approved for operational use?

Table 2: Resolution functions across systems.

System	Resolution target	Typical answer
DNS	Name to network location	Where the named resource can be reached
Database	Query to records	Which records satisfy the query
Search engine	Query to documents	Which documents are relevant
RAG system	Prompt to context	Which passages may support generation
AMN	Reference to governed knowledge	What is known, why, with what status, and under which governance context

8.1 Why Resolution Is Different From Search

Search is concerned with locating material. Resolution is concerned with interpreting governed status. A search result may include old, new, contradicted, duplicate, unofficial, and speculative material side by side. Knowledge resolution attempts to expose the state of the knowledge itself.

This distinction is especially important for AI agents. An agent should not merely retrieve plausible context. It should be able to determine whether the information is approved, superseded, disputed, or insufficient for action.

A further clarification is warranted here. Assigning a governed knowledge object a stable, structured identity is necessary but not sufficient for resolution. Identity answers the question of what is meant; resolution answers the separate question of how a system, given an incoming query expressed in the requester's own language, locates the object that identity picks out. A structured addressing scheme does not by itself supply this second capability: a query is typically a paraphrase of the knowledge it is seeking, not a literal statement of that knowledge's assigned coordinates, and a resolver that expects the two to align closely will fail silently rather than informatively. Practical AMN implementations therefore require a resolution mechanism grounded in the meaning of the query and the content of candidate objects, with the governed address supplying status, authority, and lifecycle context once a candidate has been located, rather than serving as the lookup key itself.

9 Toward Interoperable Memory Networks

Once memory becomes governed and addressable within an institution, a natural question follows: can independently governed memory systems become interoperable?

This paper does not propose a global database. It does not propose a universal source of truth. It proposes a more modest possibility: if institutions adopt compatible conventions for identifying, governing, and resolving memory objects, they may eventually be able to reference knowledge across organizational boundaries while preserving local control.

9.1 From Isolated Memory to Shared Reference

Organizations already exchange documents, reports, API payloads, and database exports. What they rarely exchange is governed memory itself.

A future memory ecosystem could allow one institution to reference a governed knowledge object maintained by another institution, inspect its provenance and status, and decide how much trust to assign within its own context.

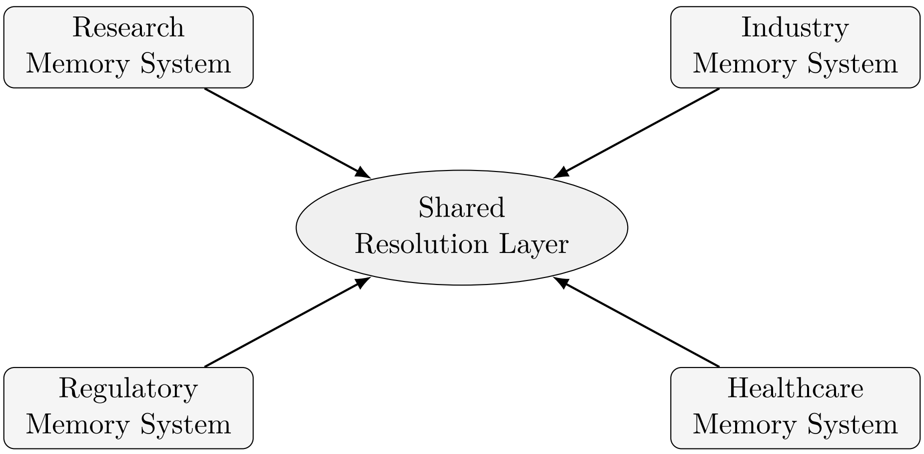

Figure 4: Conceptual interoperability between independently governed memory systems.

9.2 Open Participation and Earned Trust

A durable knowledge ecosystem must balance openness and stability.

If participation is too restricted, new knowledge cannot enter. If modification is unrestricted, trust collapses. A plausible model is open proposal with governed promotion:

> Participation may be open. Operational trust must be earned.

In such a model, anyone may propose new knowledge, corrections, or challenges. Not everyone may promote those proposals into operationally trusted memory. Promotion depends on domain-specific stewardship, validation, review, and evidence.

9.3 No Universal Truth Engine

The objective of interoperable memory is not to force universal agreement. Many domains contain legitimate disagreement, incomplete evidence, jurisdictional variation, and changing interpretations.

A mature memory network should make disagreement visible, not erase it. Different institutions may assign different status to the same knowledge object. The useful function is to show the evidence, lineage, and stewardship behind those differences.

10 Security, Privacy, and Access Control

Governed memory is sensitive infrastructure. AMN-style systems must treat privacy, permissioning, and auditability as core design concerns rather than external additions.

Key requirements include:

- **Source access boundaries.** Not all users or agents may access all source material.
- **Memory object permissions.** Some objects may be visible only to specific roles or systems.
- **State-transition controls.** Only authorized actors or processes should alter operational memory status.
- **Audit trails.** Admissions, revisions, challenges, and retirements should be inspectable. Audit trails are also the primary countermeasure against the confident miss (Section 2.5): a narrowing operation that records what it excluded, and why, converts a silent failure into an inspectable one.
- **Separation of visibility and action.** A system may be allowed to read a memory object without being allowed to act on it.

11 Evaluation Methodology

An AMN implementation should be evaluated not only by answer accuracy but by memory-governance behavior.

Since the framework was first outlined, an independent implementation has been evaluated against roughly seventeen thousand real memory items at personal scale. That evaluation confirmed the identity, provenance, and governance layers described above without qualification, and produced the two corrections to the addressing and placement mechanisms discussed in Sections 2.5 and 4.2. The evaluation was conducted at the scale of a single authority and a single corpus; validation at organizational scale, with contested authority and multiple stewards, remains open.

Table 3: Evaluation dimensions for governed institutional memory.

Dimension	Question	Example measure
Traceability	Can each operational answer identify supporting evidence?	Share of answers with inspectable source trail
Persistence	Does memory survive across sessions and workflows?	Recall consistency over time
Revision handling	Can superseded information be distinguished from current memory?	Accuracy on temporal update tasks
Contradiction handling	Are conflicting claims surfaced rather than hidden?	Detection and escalation rate
Governance compliance	Are admission and state transitions controlled?	Policy violation rate
Actionability	Can memory support operational workflows?	Task completion under governed memory
Auditability	Can humans inspect how memory changed?	Audit success rate

12 Limitations and Open Challenges

AMN is a framework, not a solved implementation. Several challenges remain. The evaluation described in Section 11 narrowed two of these (resolution mechanism design, placement rigidity) into the specific corrections already incorporated above; the challenges below remain open.

12.1 Domain Setup Cost

Governed memory requires domain-specific definitions, evidence expectations, lifecycle rules, and stewardship processes. This setup is costly and cannot be eliminated entirely.

12.2 Ambiguous Sources

Low-quality documents, informal notes, contradictory records, and incomplete metadata can make memory formation difficult. AMN should not pretend certainty where the source material is weak.

12.3 Governance Burden

If every memory transition requires manual review, the system becomes slow. If too much is automated, trust may degrade. Practical systems need calibrated thresholds and escalation paths.

12.4 Institutional Incentives

Memory governance changes organizational behavior. Teams must agree on stewardship, responsibility, and review processes. The technical architecture cannot solve those incentives alone.

12.5 Interoperability Standards

Interoperable memory would require conventions for identity, provenance, lifecycle states, access control, and trust exchange. These conventions would need to emerge gradually and remain domain-sensitive.

13 Discussion

The rise of generative AI has made the institutional memory gap harder to ignore. LLMs can generate fluent answers, but fluency is not memory. Retrieval can supply context, but context is not governance. Knowledge graphs can model relationships, but relationships are not necessarily operational memory.

AMN proposes a shift in emphasis. The question is not only how to retrieve relevant material for a prompt. The question is how an organization forms, preserves, revises, audits, and uses knowledge over time.

The paper's central claim is therefore modest but important: organizations need governed memory. Addressability, provenance, lifecycle management, stewardship, and auditability are not secondary features. They are the conditions under which AI systems can safely participate in institutional work.

The broader implication is that governed memory may become a foundation for knowledge resolution. If memory objects can be referenced and governed within one institution, then future systems may learn to reference and reason across independently governed memory networks. That possibility should be approached cautiously. It requires standards, incentives, privacy models, and trust mechanisms that do not yet exist at scale.

Still, the direction is worth naming. The long-term significance of AMN may not lie in building larger repositories. It may lie in making memory itself more structured, governed, and eventually interoperable.

14 Conclusion

The future of AI-augmented organizations depends not only on what models can generate, but on what institutions can reliably remember.

Modern systems store information, retrieve documents, and generate answers, but they often fail to preserve governed memory: knowledge whose identity, provenance, lifecycle, revision history, and operational status remain visible over time.

Addressable Memory Networks are proposed as a framework for this missing layer. AMN treats important knowledge as addressable, governable memory objects rather than transient text fragments or isolated records. It separates interpretation from admission, emphasizes evidence and lifecycle, and supports both human and machine use.

The immediate value of AMN is practical: better institutional memory, stronger auditability, safer AI-agent operation, and clearer handling of revisions and contradictions.

The longer-term implication is larger. Once memory becomes governed and addressable, knowledge resolution becomes possible. And once knowledge resolution becomes possible, independently governed memory systems may eventually interoperate without surrendering autonomy.

The long-term significance of governed memory may therefore not lie in constructing larger repositories of information, but in enabling institutions to reference, understand, and reason about knowledge together.

References

[1] Lewis, P., et al. "Retrieval-Augmented Generation for Knowledge-Intensive NLP Tasks." *Advances in Neural Information Processing Systems*, 2020.

[2] Edge, D., et al. "From Local to Global: A Graph RAG Approach to Query-Focused Summarization." *arXiv preprint arXiv:2404.16130*, 2024.

[3] Packer, C., et al. "MemGPT: Towards LLMs as Operating Systems." *arXiv preprint arXiv:2310.08560*, 2023.

[4] Park, J. S., et al. "Generative Agents: Interactive Simulacra of Human Behavior." *Proceedings of the ACM Web Conference*, 2023.

[5] Wilkinson, M. D., et al. "The FAIR Guiding Principles for Scientific Data Management and Stewardship." *Scientific Data*, 2016.

[6] Lebo, T., et al. "PROV-O: The PROV Ontology." *W3C Recommendation*, 2013.

[7] Ji, S., et al. "A Survey on Knowledge Graphs: Representation, Acquisition and Applications." *IEEE Transactions on Neural Networks and Learning Systems*, 2021.

[8] Jiang, X., et al. "On the Evolution of Knowledge Graphs: A Survey and Perspective." *arXiv preprint arXiv:2310.04835*, 2023.

[9] Jarnac, L., et al. "Uncertainty Management in the Construction of Knowledge Graphs: A Survey." *arXiv preprint arXiv:2405.16929*, 2024.

A Public Disclosure Boundary

This public manuscript intentionally avoids implementation-specific details. It does not disclose proprietary addressing schemes, admission algorithms, validation topology, scoring functions, domain-specific schemas, orchestration mechanics, or internal system thresholds. The framework is presented at the level required to define a research category, motivate evaluation, and support discussion while preserving implementation flexibility and commercial confidentiality.

B Suggested Future Work

Future work may include:

- reference implementations for bounded domains;
- benchmark tasks for governed memory;
- formal models for lifecycle state transitions;
- interoperability conventions for memory object references;
- human-in-the-loop governance studies;
- evaluation of AMN-style systems in AI-agent workflows.