

Memory Disaggregation for AI Infrastructure

Cameron Reed - Corresponding author

Abstract

This systems review examines memory disaggregation for artificial-intelligence infrastructure. The organizing question is when pooled and remote memory improves utilization without unacceptable latency, interference, or recovery complexity. Ten related scholarly sources are synthesized through a decision-centered framework spanning problem definition, mechanism, measurement, evaluation, implementation, and governance. The review does not invent experiments, pooled estimates, or unreported quantitative results. It instead evaluates the strength and transferability of the available evidence, with particular attention to comparing capacity gains without modeling workload locality and failure domains. The resulting framework links technical or empirical performance to explicit use conditions and identifies tests that should precede wider adoption in large-scale AI training and inference platforms.

Keywords: memory disaggregation, CXL, AI infrastructure, performance, resource pooling

Synthesis lens	Principal question
Mechanism	the chain connecting evidence inputs, memory disaggregation for artificial-intelligence infrastructure, intermediate outputs, and downstream decisions
Evaluation	construct validity, comparative performance, uncertainty, transfer across settings, and decision utility
Primary risk	comparing capacity gains without modeling workload locality and failure domains

1. Introduction

Research on memory disaggregation for artificial-intelligence infrastructure sits at the boundary between a promising component and a consequential decision. The core question is when pooled and remote memory improves utilization without unacceptable latency, interference, or recovery complexity. Answering it requires more than assembling favorable results. It requires a chain of evidence that makes the problem boundary, mechanism, measurement, comparator, uncertainty, and accountable decision visible to the reader.

This article presents a structured narrative review of ten related publications. Sources were selected for topical relevance, traceable publication metadata, and complementary coverage of technical, empirical, and governance questions. No meta-analytic pooling is attempted because the included studies differ in designs, populations, system boundaries, and outcomes. The purpose is decision-oriented synthesis: to identify what each source can support, where transfer remains uncertain, and which evidence would justify the next stage of use.

The central risk is comparing capacity gains without modeling workload locality and failure domains. A top-line metric or broad policy label can appear decisive while concealing the conditions that produced it. Accordingly, the synthesis separates observation from interpretation and implementation from validation. This separation is especially important when the same system creates benefits for one user or institution while transferring cost, risk, or administrative burden to another.

2. Evidence Base and Problem Boundary

A defensible review begins by defining the unit of analysis. For the present topic, the relevant boundary includes inputs, institutional or technical context, the mechanism producing an output, the person or system acting on that output, and the downstream outcome. Evidence falls outside the boundary when it measures only an intermediate proxy yet is used to claim an end-to-end benefit.

Aguilera et al. (2023) addresses "Memory disaggregation: why now and what are the challenges." In this synthesis, that work is used as a source on problem definition within memory disaggregation for artificial-intelligence infrastructure. The connection is deliberately bounded: the cited study informs the evidence map, but it does not by itself establish readiness for large-scale AI training and inference platforms. That distinction keeps the review close to the published record and prevents an adjacent result from being converted into a broader causal claim.

Maruf and Chowdhury (2023) addresses "Memory Disaggregation: Advances and Open Challenges." In this synthesis, that work is used as a source on conceptual boundary within memory disaggregation for artificial-intelligence infrastructure. The connection is deliberately bounded: the cited study informs the evidence map, but it does not by itself establish readiness for large-scale AI training and inference platforms. That distinction keeps the review close to the published record and prevents an adjacent result from being converted into a broader causal claim.

Wang et al. (2022) addresses "The case for distributed shared-memory databases with RDMA-enabled memory disaggregation." In this synthesis, that work is used as a source on measurement within memory disaggregation for artificial-intelligence infrastructure. The connection is deliberately bounded: the cited study informs the evidence map, but it does not by itself establish readiness for large-scale AI training and inference platforms. That distinction keeps the review close to the published record and prevents an adjacent result from being converted into a broader causal claim.

Together, these works provide complementary entry points, but their conclusions should not be flattened into a single ranking. Differences in data generation, setting, temporal horizon, and outcome definition may be substantively meaningful. A review should therefore preserve those differences and state where analogy, rather than direct replication, connects one domain to another.

3. Mechanism and System Design

The operative mechanism is the chain connecting evidence inputs, memory disaggregation for artificial-intelligence infrastructure, intermediate outputs, and downstream decisions. This formulation discourages component-only reasoning. A model, platform, policy, assay, or infrastructure service acquires practical meaning only when its interfaces and use conditions are specified. The evidence chain should describe what information is available, what transformation occurs, which assumptions make that transformation credible, and who is authorized to act.

Wang et al. (2024) addresses "Rcmp: Reconstructing RDMA-Based Memory Disaggregation via CXL." In this synthesis, that work is used as a source on system mechanism within memory disaggregation for artificial-intelligence infrastructure. The connection is deliberately bounded: the cited study informs the evidence map, but it does not by itself establish readiness for large-scale AI training and inference platforms. That distinction keeps the review close to the published record and prevents an adjacent result from being converted into a broader causal claim.

Xia et al. (2020) addresses "Non-intrusive load disaggregation based on composite deep long short-term memory network." In this synthesis, that work is used as a source on representation choice within memory disaggregation for artificial-intelligence infrastructure. The connection is deliberately bounded: the cited study informs the evidence map, but it does not by itself establish readiness for large-scale AI training and inference platforms. That distinction keeps the review close to the published record and prevents an adjacent result from being converted into a broader causal claim.

Kwon et al. (2023) addresses "Failure Tolerant Training With Persistent Memory Disaggregation Over CXL." In this synthesis, that work is used as a source on failure analysis within memory disaggregation for artificial-intelligence infrastructure. The connection is deliberately bounded: the cited study informs the evidence map, but it does not by itself establish readiness for large-scale AI training and inference platforms. That distinction keeps the review close to the published record and prevents an adjacent result from being converted into a broader causal claim.

These studies also show why modular claims require interface tests. A component may perform well while the complete pathway fails because inputs drift, users interpret outputs differently, recovery semantics are ambiguous, or institutional incentives change. Design documentation should therefore include not only the intended pathway but also foreseeable deviations, degraded modes, and conditions under which the system should stop or defer.

4. Evaluation and Uncertainty

Evaluation should center on construct validity, comparative performance, uncertainty, transfer across settings, and decision utility. Average performance remains useful, but it should be accompanied by distributions, error categories, relevant subgroups or operating regimes, and uncertainty at the point where a decision changes. Comparators must isolate the value of the proposed addition rather than compare a mature intervention with an intentionally weak baseline.

Heo et al. (2021) addresses "Hardware-assisted Trusted Memory Disaggregation for Secure Far Memory." In this synthesis, that work is

used as a source on comparative evaluation within memory disaggregation for artificial-intelligence infrastructure. The connection is deliberately bounded: the cited study informs the evidence map, but it does not by itself establish readiness for large-scale AI training and inference platforms. That distinction keeps the review close to the published record and prevents an adjacent result from being converted into a broader causal claim.

Quek et al. (2019) addresses "Load Disaggregation Using One-Directional Convolutional Stacked Long Short-Term Memory Recurrent Neural Network." In this synthesis, that work is used as a source on uncertainty within memory disaggregation for artificial-intelligence infrastructure. The connection is deliberately bounded: the cited study informs the evidence map, but it does not by itself establish readiness for large-scale AI training and inference platforms. That distinction keeps the review close to the published record and prevents an adjacent result from being converted into a broader causal claim.

External validation should introduce meaningful shifts in setting, time, population, workload, market structure, or environmental conditions. A nominally independent sample can still be too similar to test transfer. Conversely, a failed transfer is scientifically informative when the changed conditions are measured and the mechanism of failure is examined rather than hidden in an aggregate score.

5. Translation and Governance

Translation to large-scale AI training and inference platforms depends on more than technical readiness. Decision rights, documentation, maintenance, monitoring, appeal, and recovery are part of the intervention. The governance requirement is traceable evidence, named responsibility, version control, monitoring, appeal, and explicit revalidation. Each control should be connected to a named failure mode and should identify an owner, evidence source, review interval, and escalation path.

Mehra and Coughlin (2022) addresses "Taming Memory With Disaggregation." In this synthesis, that work is used as a source on implementation within memory disaggregation for artificial-intelligence infrastructure. The connection is deliberately bounded: the cited study informs the evidence map, but it does not by itself establish readiness for large-scale AI training and inference platforms. That distinction keeps the review close to the published record and prevents an adjacent result from being converted into a broader causal claim.

Yu et al. (2025) addresses "Rethinking Virtual Machines Live Migration for Memory Disaggregation." In this synthesis, that work is used as a source on governance within memory disaggregation for artificial-intelligence infrastructure. The connection is deliberately bounded: the cited study informs the evidence map, but it does not by itself establish readiness for large-scale AI training and inference platforms. That distinction keeps the review close to the published record and prevents an adjacent result from being converted into a broader causal claim.

A credible implementation plan also defines sunset and revalidation conditions. New data, software updates, institutional restructuring, population change, or shifts in incentives can invalidate previously reasonable assumptions. Monitoring should therefore test the validity of the evidence chain, not merely whether a service remains online or a headline metric remains stable.

6. Research Agenda

Future studies should preregister or otherwise predefine the intended decision, principal outcomes, exclusions, and analysis plan when feasible. They should report missingness, sensitivity analyses, negative or null findings, and the cost of errors to differently situated stakeholders. Reproducible artifacts should include sufficient metadata to reconstruct data provenance, model or analytical versions, parameter choices, and the sequence from evidence to conclusion.

Cross-disciplinary work needs a shared object of explanation rather than a collection of adjacent vocabularies. For memory disaggregation for artificial-intelligence infrastructure, that shared object is the complete decision pathway. Technical, clinical, social, environmental, or economic expertise should each change the design or interpretation of the study. When evidence remains incomplete, the useful output is a bounded hypothesis, a transparent uncertainty statement, or a request for additional measurement.

7. Conclusion

The literature supports a disciplined approach to memory disaggregation for artificial-intelligence infrastructure: define the decision, preserve system boundaries, test consequential failure modes, connect uncertainty to action, and assign accountable control. The strongest conclusion is not that one method is universally superior. It is that credible use in large-scale AI training and inference platforms requires an auditable link from source evidence to design choice, evaluation result, and governed decision.

Declarations

Ethics approval: Not applicable. This article synthesizes previously published literature and reports no new research involving humans, animals, human tissue, or identifiable sensitive data.

Consent to participate and consent for publication: Not applicable.

Funding: No external funding is reported for this commissioned synthesis.

Competing interests: The authoring group declares no competing interests for this commissioned synthesis.

AI-assisted writing: Generative AI supported drafting, source organization, and language editing. Accountable authors and editors must verify every claim, citation, and disclosure before publication.

Data, code, and materials availability: No new dataset or code was generated. All evidence is identified in the reference list.

Licence: This manuscript is prepared for publication under the Creative Commons Attribution 4.0 International licence (CC BY 4.0).

References

- Aguilera, M. K., Amaro, E., Amit, N., Hunhoff, E., Yelam, A., & Zellweger, G. (2023). Memory disaggregation: why now and what are the challenges. *ACM SIGOPS Operating Systems Review*, 57(1), 38-46. <https://doi.org/10.1145/3606557.3606563>
- Heo, T., Kang, S., Lee, S., Hwang, S., & Huh, J. (2021). Hardware-assisted Trusted Memory Disaggregation for Secure Far Memory. *arXiv (Cornell University)*. <https://doi.org/10.48550/arxiv.2108.11507>
- Kwon, M., Jang, J., Choi, H., Lee, S., & Jung, M. (2023). Failure Tolerant Training With Persistent Memory Disaggregation Over CXL. *IEEE Micro*, 43(2), 66-75. <https://doi.org/10.1109/mm.2023.3237548>
- Maruf, H. A., & Chowdhury, M. (2023). Memory Disaggregation: Advances and Open Challenges. *ACM SIGOPS Operating Systems Review*, 57(1), 29-37. <https://doi.org/10.1145/3606557.3606562>
- Mehra, P., & Coughlin, T. (2022). Taming Memory With Disaggregation. *Computer*, 55(9), 94-98. <https://doi.org/10.1109/mc.2022.3187847>

- Quek, Y. T., Woo, W. L., & Logenthiran, T. (2019). Load Disaggregation Using One-Directional Convolutional Stacked Long Short-Term Memory Recurrent Neural Network. *IEEE Systems Journal*, 14(1), 1395-1404. <https://doi.org/10.1109/jsyst.2019.2919668>
- Wang, R., Wang, J., Idreos, S., Özsu, M. T., & Aref, W. G. (2022). The case for distributed shared-memory databases with RDMA-enabled memory disaggregation. *Proceedings of the VLDB Endowment*, 16(1), 15-22. <https://doi.org/10.14778/3561261.3561263>
- Wang, Z., Guo, Y., Lu, K., Wan, J., Wang, D., Yao, T., & Wu, H. (2024). Rcmp: Reconstructing RDMA-Based Memory Disaggregation via CXL. *ACM Transactions on Architecture and Code Optimization*, 21(1), 1-26. <https://doi.org/10.1145/3634916>
- Xia, M., Liu, W. A., Wang, K., Song, W., Chen, C., & Li, Y. (2020). Non-intrusive load disaggregation based on composite deep long short-term memory network. *Expert Systems with Applications*, 160, 113669. <https://doi.org/10.1016/j.eswa.2020.113669>
- Yu, X., Jia, X., Zhang, J., Wang, Y., Yu, S., & Qi, Z. (2025). Rethinking Virtual Machines Live Migration for Memory Disaggregation. *IEEE Transactions on Parallel and Distributed Systems*, 36(11), 2310-2324. <https://doi.org/10.1109/tpds.2025.3597149>