Federated and Continual Learning for Autonomous Self-Healing Datacenter Networks

Dheeraj Ramasahayam

Abstract—Most research on self-healing datacenter networks still assumes that learning occurs inside a single administrative domain. That assumption is increasingly mismatched to operational reality. Large outages, rare failure modes, routing anomalies, and security incidents are distributed across cloud regions, colocation environments, and enterprise operators that cannot exchange raw telemetry because of privacy, contractual, and security constraints. This paper proposes a federated and continual learning framework for autonomous self-healing datacenter networks in which multiple organizations train local models, share only protected model updates, and continually adapt to site-specific drift over time. The design extends a six-paper research program from temporal failure prediction, drift-adaptive intrusion detection, topology-grounded root cause analysis (RCA), recovery validation, and explainable RCA toward cross-organization collaboration. The proposed architecture combines cross-silo federated optimization, secure aggregation, site-local drift detection, continual model refresh, and local explainability. Rather than claiming a completed multi-organization benchmark that does not yet exist publicly, the paper makes three concrete contributions: a systems architecture for collaborative failure intelligence without raw-data exchange, a continual adaptation protocol that couples drift triggers to federated model updates, and a reproducible evaluation plan comparing local-only, centralized, federated, and federated-plus-continual regimes. The central claim is that collaborative self-healing intelligence can be scaled across organizations without exposing sensitive network telemetry.

Index Terms—datacenter networks, self-healing systems, federated learning, continual learning, concept drift, privacy-preserving machine learning, root cause analysis

I. Introduction

Datacenter networks now underpin cloud computing, artificial-intelligence training, financial infrastructure, healthcare systems, logistics, and digital public services. Failures in one environment are rarely unique: routing instability, telemetry drift, misconfiguration cascades, and attack patterns often recur across organizations with slightly different signatures. Yet operational learning remains fragmented because the strongest incident corpora are private. A single operator rarely observes enough severe failures to build broadly robust self-healing intelligence, while multiple operators cannot simply pool packet traces, telemetry histories, and topology snapshots into one centralized training set.

This gap matters because the author's prior papers already establish five building blocks of a self-healing stack: temporal failure prediction [1], drift-adaptive intrusion detection [2], topology-grounded RCA and recovery validation [3], explainable RCA [4], and a unified closed-loop self-healing architecture [5]. The remaining systems question is no longer whether one site can learn from its own data. It is whether multiple organizations can share failure intelligence without sharing the underlying sensitive data.

Federated learning (FL) is a natural systems answer because it supports collaborative model training over decentralized data [6], [7]. However, plain FL is insufficient for self-healing networks. Datacenter telemetry is strongly non-IID across organizations, failures are rare, labels are delayed, topologies evolve, and deployed models encounter continuous drift. A practical design therefore also needs continual learning, drift-aware triggering, and privacy controls such as secure aggregation [8]–[10].

This paper proposes Federated Continual Self-Healing (FCSH), a cross-silo architecture for collaborative failure intelligence across independent datacenter operators. Each site trains locally on its own telemetry, flow, topology, and incident data; each site keeps explainability artifacts and raw traces local; and only protected model updates and limited aggregate statistics are shared with the federation coordinator. Drift detectors trigger site-local continual adaptation between federation rounds, which keeps models responsive to changing conditions without forcing continuous raw-data exchange.

The paper makes four contributions:

- A cross-organization architecture that extends autonomous self-healing datacenter networks from a single-site closed loop to a federated, privacy-preserving learning system.
- A continual adaptation protocol that couples site-local drift detection with regularized incremental updates, reducing catastrophic forgetting while preserving the value of globally learned patterns.
- A privacy and trust design that combines secure aggregation, local explainability, and organization-specific deployment boundaries.
- A defensible evaluation protocol for Paper 6 that compares local-only, centralized, federated, and federated-plus-continual baselines without overstating the existence of a public multi-organization benchmark.

II. Positioning in the Six-Paper Program

Table I makes the portfolio role of this paper explicit. Paper 6 is not another metric-tuning exercise. Its purpose is to scale the already-developed self-healing stack across organizational boundaries.

Dheeraj Ramasahayam is an Independent Researcher based in Durham, North Carolina, USA.

TABLE I
Position of Paper 6 in the six-paper datacenter-network research sequence.

Paper	Primary contribution	Role
1	Temporal attention-guided failure prediction	Prediction
2	Drift-adaptive intrusion detection	Adaptation
3	ClosRCA-Bench with recovery validation	Evaluation
4	Unified autonomous self-healing architecture	Integration
5	Explainable RCA for datacenter networks	Trust
6	Federated and continual learning across organizations	Scale + collaboration

The thesis progression is cumulative. Paper 1 asks whether failures can be forecast early; Paper 2 asks how models remain useful under drift; Paper 3 establishes a reproducible topology-grounded RCA and recovery benchmark; Paper 4 organizes the modules into a closed loop; Paper 5 exposes operator-facing explanations; and Paper 6 asks how that closed loop can learn from many organizations without violating data sovereignty. This final step is important because rare, high-severity network incidents are precisely the events least likely to be numerous within one site and most likely to benefit from shared intelligence.

III. Problem Statement

Consider a federation of S organizations, each operating one or more datacenter environments. Organization s observes a local stream of telemetry windows, flow records, topology snapshots, incident tickets, RCA labels, and recovery outcomes:

$$\mathcal{D}_s = \{(\mathbf{x}_t, \mathbf{f}_t, G_t, y_t, r_t)\}_{t=1}^{T_s}. \tag{1}$$

Here $\mathbf{x}_t$ denotes device telemetry, $\mathbf{f}_t$ denotes traffic or flow features, G_t is a topology graph, y_t contains task labels such as failure, attack, cause, and target annotations when available, and r_t represents recovery or post-action outcomes.

The goal is to learn a global self-healing model Θ that improves generalization to rare and unseen failure patterns across sites, while satisfying three constraints:

- C1: raw telemetry, flows, and incident traces never leave site s.
- C2: the model remains adaptive under local drift over time.
- C3: operators retain local explanations and execution control.

This setting is closer to cross-silo federated learning than mobile-device FL because each client is an organization with substantial local compute and stricter governance boundaries [7]. It also differs from standard centralized training because non-IID site distributions, asynchronous participation, and delayed labels are expected rather than exceptional.

IV. Related Work

Predictive maintenance, intrusion detection, and RCA have each advanced substantially in isolation. Learned sequence models improve early warning on telemetry streams [11], [12]. Drift-aware IDS pipelines improve deployment-time robustness relative to static detectors [9], [13]. Graph-based RCA better matches the topology-dependent propagation structure of infrastructure incidents [14]–[17]. The author's earlier self-healing manuscript shows how these modules can be integrated into one closed-loop controller [5].

Federated learning addresses a different bottleneck: how multiple parties collaborate without centralizing raw data [6], [7]. The standard federated averaging protocol is a reasonable baseline for cross-silo organizations, especially when communication rounds are infrequent and each site can perform several local optimization epochs [6]. Secure aggregation strengthens the privacy story by hiding individual model updates from the coordinator and exposing only an aggregate update [8].

Continual learning is required because self-healing models operate on non-stationary environments. Concept drift is well established in networked systems [9]. Regularization-based continual-learning methods such as Elastic Weight Consolidation (EWC) and distillation-based approaches such as Learning without Forgetting (LwF) are useful because they adapt models without requiring complete replay of historical raw data [10], [18], [19]. The gap in the literature is the intersection of these ideas for operational self-healing datacenter networks: cross-organization collaboration, drift-aware continual updates, topology-grounded RCA, and operator-facing explainability inside one architecture.

V. Federated Continual Self-Healing Architecture

The proposed system maintains a modular global parameter set

$$\Theta = \{\theta_{\text{pred}}, \theta_{\text{ids}}, \theta_{\text{rca}}, \theta_{\text{xai}}\}, \tag{2}$$

covering temporal failure prediction, intrusion and anomaly detection, topology-aware RCA, and explanation heads. Each participating organization maintains a local copy Θ_s and performs training on its own incident corpus.

A. Local Training

At federation round r, site s initializes from the current global model $\Theta^{(r)}$ and optimizes a local objective

$$\mathcal{L}_s = \mathcal{L}_{\text{task}} + \lambda_{\text{cl}}\mathcal{L}_{\text{cont}} + \lambda_{\text{reg}}\|\Theta_s - \Theta^{(r)}\|_2^2. \tag{3}$$

The task loss aggregates the site's active modules, for example prediction, IDS, RCA, and recovery-calibration heads. The regularization term keeps local updates from diverging too aggressively from the global model under strongly non-IID data.

B. Federated Aggregation

After local optimization, sites transmit encrypted or otherwise protected updates to the coordinator. The coordinator computes a weighted aggregate:

$$\Theta^{(r+1)} = \sum_{s=1}^{S} \frac{n_s}{\sum_{j=1}^{S} n_j} \Theta_s^{(r+1)}, \tag{4}$$

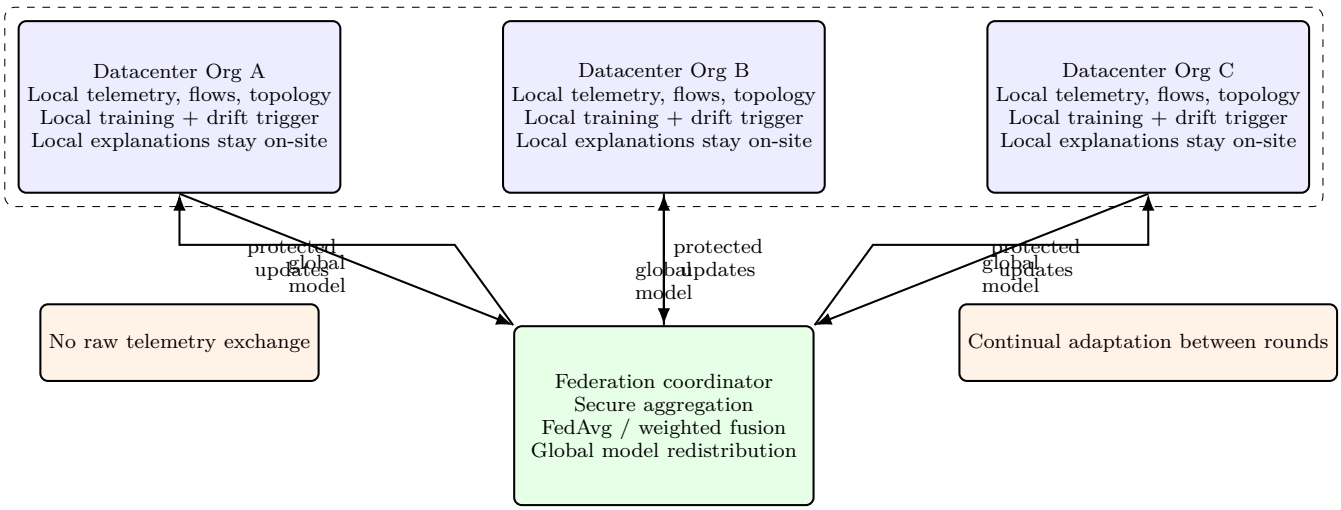

Fig. 1. Federated Continual Self-Healing architecture. Each organization trains locally, shares only protected model updates, receives a refreshed global model, and continues site-local drift-aware adaptation between rounds.

where n_s is the effective number of site-local training examples or windows. This is the standard FedAvg baseline [6]. For Paper 6, FedAvg is the correct starting point because the novelty is the self-healing application and its continual adaptation loop, not a claim that a brand-new federated optimizer is required from the outset.

C. Drift-Triggered Continual Adaptation

Plain periodic aggregation is too slow for live operations when site distributions shift. Each organization therefore runs a drift detector on local predictive residuals, alert streams, or feature distributions. When ADWIN or a comparable detector signals change [9], [20], the site enters a continual-learning phase before the next federation round.

One practical local continual objective is

$$\mathcal{L}_{\text{cont}} = \sum_i F_i(\theta_i - \theta_i^\star)^2 + \mu \|\phi_{\Theta_s}(\mathbf{z}) - \phi_{\Theta_{\text{old}}}(\mathbf{z})\|_2^2, \quad (5)$$

where the first term is an EWC-style consolidation penalty using parameter-importance scores F_i [18], and the second term is a distillation term that preserves useful behavior from the previous local model on a retained buffer or pseudo-replay set [19]. This combination is appropriate for self-healing systems because it balances two needs: rapid local plasticity under new failure signatures and retention of globally useful rare-event knowledge.

D. Local Explainability and Deployment Control

Explainability remains local by design. Sites generate operator-facing evidence from their own telemetry windows, topology graphs, and RCA outputs. This includes suspicious-node rankings, salient features, and propagation traces inherited from the earlier explainable RCA work [4]. The federation may optionally share aggregated explanation priors, such as anonymized feature-importance histograms or cause-family confusion summaries, but it does not transmit raw incident narratives, topology screenshots, or switch-level operator reports. This preserves both privacy and operational trust boundaries.

VI. Threat Model and Privacy Assumptions

The proposed design assumes an honest-but-curious federation coordinator. Organizations are willing to contribute to collaborative model improvement but do not trust the coordinator or one another with raw telemetry, packet captures, or incident details. The architecture therefore makes the following assumptions explicit:

- Raw telemetry, traffic logs, graph windows, and recovery traces remain local.
- Model updates are protected with secure aggregation so the coordinator observes only an aggregate update, not an individual site's gradient or weights [8].
- Deployment decisions remain local. The global model recommends; the site decides whether to alert, escalate, or execute recovery.
- Optional differential-privacy noise can be added in high-sensitivity settings, but the base Paper 6 design treats secure aggregation as the minimum viable protection layer.

This is not a complete defense against every privacy attack on collaborative learning. Update inversion, poisoning, and Byzantine participation are real threats in FL systems [7]. For that reason, the paper is careful to frame privacy claims at the architectural level: the system avoids raw-data centralization and reduces direct exposure of site-specific telemetry, but it does not claim perfect cryptographic immunity under all adversarial models.

VII. Evaluation Plan

No public benchmark currently contains synchronized multi-organization telemetry, flow records, topology snapshots, RCA labels, explainability targets, and recovery outcomes across several datacenter operators. Paper 6

should therefore be evaluated through a partitioned and staged protocol rather than by pretending such an artifact already exists.

A. Federation Setup

The experimental plan simulates 3–5 organizations with intentionally non-IID data distributions. Site heterogeneity should reflect realistic operational differences:

- different hardware families or routing environments for failure prediction,
- different traffic and attack mixes for drift-adaptive IDS,
- different topology slices or incident families for RCA,
- different explanation and recovery histories for operator-facing calibration.

Existing artifacts from the earlier papers provide the starting point for that partitioning: telemetry prediction benchmarks from Paper 1, drift-oriented intrusion benchmarks from Paper 2, topology-grounded RCA windows from Paper 3, and local explainability templates from Paper 5. The system-level handoff logic from Paper 4 supplies the unified deployment frame.

B. Baselines

Table II defines the minimum comparison set.

C. Metrics

The evaluation should report both task metrics and operational metrics:

- failure-prediction warning F1, lead time, and cross-site generalization;
- IDS weighted F1, post-drift recovery, detection delay, and latency;
- RCA cause F1, target-device F1, hidden-target accuracy, and temporal delay;
- explanation compactness and operator-effort proxies;
- recovery-action success, unsafe-action blocking, and service-restoration proxies;
- federated communication cost per round and total adaptation overhead.

The core hypothesis is not that federated training will beat centralized pooled training. Centralized training remains the best-case access regime if privacy is ignored. The meaningful hypothesis is narrower and operationally stronger: the federated-plus-continual regime should close a substantial portion of the gap between local-only and centralized training, while preserving data sovereignty and improving robustness to site-local drift.

D. Primary Claims to Test

Paper 6 should explicitly test the following claims:

1) Federated training improves generalization to unseen local failure patterns relative to isolated local training.
2) Continual site-local adaptation improves post-drift performance relative to periodic federated aggregation alone.
3) The privacy-preserving regime avoids raw telemetry exchange while retaining most of the practical value of collaborative learning.
4) Local explainability can be preserved under federated training because incident evidence remains on-site even when model knowledge is shared globally.

VIII. Deployment Scenario

The target deployment is a federation of cloud providers, enterprise datacenters, colocation operators, or multi-region business units that each maintain their own network operations center. Each site exports telemetry and flow features into its existing monitoring stack, trains locally on a scheduled cadence, and participates in a federation round daily or weekly depending on incident volume and policy. Sites encountering acute drift can run emergency local continual updates without waiting for the next global round.

The operational benefit is straightforward. One organization may observe a new BGP instability signature, another may observe a workload-induced congestion cascade, and a third may observe an attack-induced topology anomaly. None wants to hand over raw telemetry or customer-sensitive traces. All benefit if the shared model family learns from the others' experiences. In one sentence, the architecture enables collaborative failure intelligence across organizations without exposing sensitive network telemetry.

IX. Discussion

Three design choices matter most. First, cross-silo FL is the right collaboration model because the clients are organizations, not phones. Second, continual learning is not optional; it is necessary because self-healing models operate against moving operational targets. Third, explainability must remain local because trust in network operations depends on reviewable site-specific evidence, not only on higher global accuracy.

This paper also has clear limits. It does not yet claim a native public multi-organization benchmark. It does not claim perfect privacy under every adversarial setting. It does not argue that FedAvg is the final optimizer for all deployment regimes. Those are acceptable limits for Paper 6 because the contribution is conceptual expansion with a defensible architecture and evaluation protocol, not unsupported completionism. The stronger mistake would be to fabricate a cross-organization dataset or report results that have not been measured.

X. Conclusion

This paper defines the sixth step in a datacenter-network research program: scaling autonomous self-healing intelligence across organizations. The proposed Federated Continual Self-Healing architecture extends

TABLE II
Proposed evaluation regimes for Paper 6.

Regime	Training visibility	Adaptation mode	Purpose
Local only	Each site trains only on its own data	Optional local refresh	Lower bound for isolated organizations.
Centralized pooled	All sites' data pooled into one training set	Offline retraining	Unrealistic upper bound for data access.
Federated	Sites share protected model updates only	Periodic federation rounds	Tests whether collaboration helps without raw-data sharing.
Federated + continual (proposed)	Protected updates only	Drift-triggered local adaptation plus federation rounds	Tests whether privacy-preserving collaboration remains robust under evolving site conditions.

prior work on prediction, adaptation, RCA, explainability, and unified orchestration into a privacy-preserving collaborative learning system. Organizations train locally, share only protected updates, adapt continuously under drift, and keep explanations and execution authority on-site. The key research claim is both practical and nationally relevant: datacenter operators should be able to collaborate on failure intelligence without centralizing sensitive telemetry. That is the right scope for Paper 6 because it turns a strong single-site self-healing story into a scalable multi-organization systems agenda.

References

[1] D. Ramasahayam, "Temporal attention-guided sequence learning for zero-shot generalization in datacenter network failures," 2026, unpublished manuscript on early-warning failure prediction.

[2] ——, "Drift-adaptive intrusion detection for enterprise networks," 2026, unpublished manuscript on drift-aware intrusion detection.

[3] ——, "Closrca-bench: An open topology-grounded benchmark and counterfactual recovery framework for self-healing datacenter networks," 2026, unpublished manuscript on topology-aware RCA and recovery validation.

[4] ——, "Explainable ai for root cause analysis in large-scale datacenter networks," 2026, unpublished manuscript on operator-facing RCA explanations.

[5] ——, "Autonomous self-healing datacenter networks: A unified ai system for prediction, detection, root cause analysis, and recovery," 2026, unpublished manuscript on the integrated self-healing architecture.

[6] B. McMahan, E. Moore, D. Ramage, S. Hampson, and B. A. y Arcas, "Communication-efficient learning of deep networks from decentralized data," in Proceedings of the 20th International Conference on Artificial Intelligence and Statistics, ser. Proceedings of Machine Learning Research, vol. 54, 2017, pp. 1273–1282. [Online]. Available: https://proceedings.mlr.press/v54/mcmahan17a.html

[7] P. Kairouz, H. B. McMahan, B. Avent, A. Bellet, M. Bennis, A. N. Bhagoji, K. Bonawitz, Z. Charles, G. Cormode, R. Cummings, R. G. L. D'Oliveira, H. Eichner, S. E. Rouayheb, D. Evans, J. Gardner, Z. Garrett, A. Gascón, B. Ghazi, P. B. Gibbons, M. Gruteser, Z. Harchaoui, C. He, L. He, Z. Huo, B. Hutchinson, J. Hsu, M. Jaggi, G. Joshi, M. Khodak, J. Konečný, A. Korolova, F. Koushanfar, S. Koyejo, T. Lepoint, Y. Liu, P. Mittal, M. Mohri, O. Saeedi, A. T. Suresh, S. Wu, F. X. Yu, H. Yu, S. Zhao et al., "Advances and open problems in federated learning," Foundations and Trends in Machine Learning, vol. 14, no. 1–2, pp. 1–210, 2021.

[8] K. A. Bonawitz, V. Ivanov, B. Kreuter, A. Marcedone, H. B. McMahan, S. Patel, D. Ramage, A. Segal, and K. Seth, "Practical secure aggregation for privacy-preserving machine learning," in Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security, 2017, pp. 1175–1191.

[9] J. Gama, I. Žliobaitė, A. Bifet, M. Pechenizkiy, and A. Bouchachia, "A survey on concept drift adaptation," ACM Computing Surveys, vol. 46, no. 4, pp. 44:1–44:37, 2014.

[10] G. I. Parisi, R. Kemker, J. L. Part, C. Kanan, and S. Wermter, "Continual lifelong learning with neural networks: A review," Neural Networks, vol. 113, pp. 54–71, 2019.

[11] X. Chen, Y. Zhao, H. Liu, and J. Wang, "Predictive maintenance in datacenters: A machine learning approach," IEEE Transactions on Network and Service Management, vol. 18, no. 1, pp. 60–73, 2021.

[12] A. Putina, D. Rossi, A. Bifet, S. Barth, D. Pletcher, C. Precup, and P. Nivaggioli, "Telemetry-based stream-learning of BGP anomalies," in Proceedings of the ACM SIGCOMM 2018 Workshop on Big Data Analytics and Machine Learning for Data Communication Networks, 2018, pp. 15–21.

[13] M. Roesch, "Snort: Lightweight intrusion detection for networks," in Proceedings of the 13th USENIX Conference on System Administration, 1999, pp. 229–238.

[14] A. Roy, H. Zeng, J. Bagga, and A. C. Snoeren, "Passive realtime datacenter fault detection and localization," in 14th USENIX Symposium on Networked Systems Design and Implementation, 2017, pp. 595–612.

[15] L. Li, X. Zhang, X. Zhao, H. Zhang, Y. Kang, P. Zhao, B. Qiao, S. He, P. Lee, J. Sun, F. Gao, L. Yang, Q. Lin, S. Rajmohan, Z. Xu, and D. Zhang, "Fighting the fog of war: Automated incident detection for cloud systems," in 2021 USENIX Annual Technical Conference, 2021, pp. 131–146.

[16] Z. Yu, Q. Ouyang, C. Pei, X. Wang, W. Chen, L. Su, H. Jiang, X. Wang, J. Li, and D. Pei, "Causality enhanced graph representation learning for alert-based root cause analysis," in 2024 IEEE 24th International Symposium on Cluster, Cloud and Internet Computing, 2024, pp. 77–86.

[17] T. N. Kipf and M. Welling, "Semi-supervised classification with graph convolutional networks," in International Conference on Learning Representations, 2017. [Online]. Available: https://openreview.net/forum?id=SJU4ayYgl

[18] J. Kirkpatrick, R. Pascanu, N. Rabinowitz, J. Veness, G. Desjardins, A. A. Rusu, K. Milan, J. Quan, T. Ramalho, A. Grabska-Barwinska, D. Hassabis, C. Clopath, D. Kumaran, and R. Hadsell, "Overcoming catastrophic forgetting in neural networks," Proceedings of the National Academy of Sciences, vol. 114, no. 13, pp. 3521–3526, 2017.

[19] Z. Li and D. Hoiem, "Learning without forgetting," in Computer Vision – ECCV 2016, 2016, pp. 614–629.

[20] A. Bifet and R. Gavaldà, "Learning from time-changing data with adaptive windowing," in Proceedings of the 2007 SIAM International Conference on Data Mining, 2007, pp. 443–448.