

Pradeep Anjuru¹, Sivaprakash Nithyanandam², Krishna kanth Thottempudi³, Radhika Kande⁴, Chaithanya Kotla⁵

¹Akkodis, USA

²Systems Technology group, USA

³Hermes Networks Inc, USA

⁴Sagarsoft Inc, USA

⁵Devops and Cloud lead, State of Maryland, USA

Received: 20-04-2022; Revised: 07-06-2022; Accepted: 30-06-2022

Self-Healing EV Charging Networks Using Autonomous Fault Recovery

Abstract

EV charging networks are cyber-physical systems in which faults in chargers, controllers, communication links, or backend services can quickly disrupt charging continuity. Existing studies have examined EV charging security, resilience, anomaly detection, and recovery-oriented control, but the literature still lacks unified self-healing architectures that combine fault sensing, diagnosis, autonomous recovery, and post-recovery validation. This article presents a self-healing design for EV charging networks based on continuous fault sensing, class-aware diagnosis, adaptive healing selection, distributed service reconfiguration, and post-recovery stabilization. The results show faster service restoration, lower fault-isolation time, better charging-session continuity, and stronger post-recovery stability than manual and fixed recovery strategies. Overall, the study demonstrates that autonomous fault recovery is an effective foundation for resilient EV charging networks.

Keywords: EV charging networks, self-healing systems, autonomous fault recovery, cyber-physical resilience, fault isolation, service restoration.

1. Introduction

Electric vehicle charging networks are increasingly operating as cyber-physical service systems in which charger modules, local controllers, communication links, backend services, and supervisory coordination must remain available even when faults emerge in one or more layers of the infrastructure [1]. As charging density rises, the operational impact of a local fault is no longer confined to one device because charger outages, control instability, and communication disruption can propagate into queue growth, service interruption, and wider degradation of user-facing availability [2]. Reviews of charging-station operation further indicate that modern EV charging environments must sustain reliable service under variable demand, grid interaction, and digital coordination pressure rather than under isolated steady-state conditions [3]. These characteristics make resilience a structural requirement of charging-network design rather than an after-the-fact maintenance consideration. A charging network that cannot recover autonomously from disruption becomes progressively less viable as system scale increases.

The literature already provides important foundations for understanding this problem. Studies on EV charging cyber-physical resilience have shown that charging infrastructure is exposed to mixed disturbances in which digital and physical failures interact rather than occurring independently [4]. Broader security reviews of EV charging ecosystems have also documented how current defense mechanisms are typically organized around attack detection, access control, or protocol protection, while recovery behavior is often treated as a secondary concern [5]. This means that the charging-security literature is strong on vulnerability identification and defensive awareness, but comparatively weaker on architectural mechanisms that restore service automatically after fault occurrence. The existing body of work therefore explains why charging networks are vulnerable, but it only partially addresses how they can heal themselves after disruption.

The core problem is that many EV charging systems still depend on external intervention, centralized operator action, or manual maintenance sequencing after faults are detected. Such an approach may be tolerable in small deployments, but it becomes operationally inefficient in distributed charging networks where faults can arise concurrently across chargers, stations, and communication paths. A charger may detect degradation, but if the network cannot isolate the fault, reconfigure the service path, and validate restored behavior automatically, recovery time remains too long and service continuity deteriorates quickly. The missing capability is not simply fault awareness, but autonomous recovery logic that converts detection into structured restoration. This is the problem statement that motivates the present article.

This problem matters because public and semi-public EV charging infrastructure is experienced directly through service continuity. Users do not evaluate resilience by whether a diagnostic event was logged; they evaluate it by whether a charging session could continue, be rerouted, or be restored with minimal disruption. When recovery remains manual, local faults can lead to long service gaps, repeated

abandonment of charging sessions, and stronger load pressure on neighboring chargers. In distributed EV charging ecosystems, a self-healing architecture is therefore important because it can shorten restoration time, reduce operator dependence, and prevent localized faults from degrading broader network performance. Recovery speed and recovery autonomy have thus become as important as detection accuracy.

This study presents a self-healing design for EV charging networks using autonomous fault recovery. The study focuses on fault sensing, autonomous diagnosis, recovery-policy selection, distributed service reconfiguration, local and network-level healing actions, and post-recovery stabilization. Instead of treating recovery as a manual follow-up activity, the work develops a charging-network architecture in which healing decisions are triggered and executed through explicit autonomous logic. The following methodology defines the fault classes, recovery triggers, healing actions, validation conditions, and control workflow used to support this self-healing design.

2. Methodology

The proposed methodology is organized as a self-healing cyber-physical architecture with six coordinated stages: fault sensing, fault classification, local diagnosis, autonomous healing selection, distributed service recovery, and post-recovery validation. The architecture begins from the assumption that faults in EV charging networks cannot be prevented completely and therefore must be handled through structured restoration mechanisms after they emerge. Early fault-detection studies on electric vehicle supply equipment have shown that anomalous power profiles and charger-behavior deviations can reveal developing faults before full functional collapse occurs [6]. Anomaly-aware protection frameworks for EV charging stations similarly indicate that intelligent detection becomes more useful when it is linked directly to response behavior rather than treated as a passive warning system [7]. For this reason, the present methodology integrates detection and healing inside one closed loop. The system is therefore designed to move from fault recognition to service restoration without waiting for manual intervention.

At the sensing stage, each charging node continuously generates operational evidence from charger current and voltage behavior, connector state transitions, controller heartbeat signals, message consistency, session continuity, and communication-path health. These variables are monitored over rolling windows so that transient fluctuations can be separated from persistent degradation. Autoencoder-based fault-detection work supports this approach because anomalous charging-unit power profiles often contain early warning signatures that appear before total charger failure becomes visible [6]. In the proposed architecture, local sensing is treated as the trigger source for autonomous healing logic rather than merely as a maintenance record. The role of this stage is therefore to identify whether the node is healthy, degrading, or already disrupted in a way that demands recovery action.

After sensing, the architecture performs fault classification. The fault space is partitioned into charger-hardware faults, controller faults, communication faults, cyber-induced anomalies, and hybrid cross-layer faults. This separation is important because different fault types require different recovery actions; a communication-path interruption should not trigger the same response as a power-module instability. The Grid Sentinel anomaly framework supports this class-aware logic by showing that charging-station anomalies can be detected and responded to more effectively when suspicious behavior is interpreted dynamically rather than collapsed into a single generic fault class [7]. Work on cyberattack detection using remaining-useful-life concepts also suggests that fault progression itself can be estimated, which is useful for deciding whether immediate isolation or staged recovery is more appropriate [8]. The classification stage therefore transforms raw abnormality into a recoverable system event with a defined operational meaning.

The diagnosis layer then determines recovery feasibility and healing scope. At this stage, the controller evaluates whether the fault is confined to a single charger node, affects the station controller, compromises the communication path, or threatens wider coordination. Recovery-aware security frameworks that combine intrusion detection with mitigation show that operational value emerges when diagnosis is linked to explicit action logic rather than left as a descriptive label [9]. Resilience work on EV charging management centres also supports the need for structured decision logic because resilience depends on whether the system can preserve trustworthy operation under data corruption and control-layer manipulation [10]. In the present framework, diagnosis outputs a recovery profile that specifies affected components, allowed fallback modes, and the urgency of restoration. This diagnosis profile becomes the direct input to the healing engine.

The healing engine is the core of the self-healing architecture. It selects one or more recovery actions according to the fault class and diagnosis profile. Local charger faults may trigger module reset, safe connector reinitialization, or charging-rate fallback; station-controller faults may trigger state restoration from standby memory; communication faults may trigger alternate message routing or buffered local autonomy; hybrid faults may trigger coordinated isolation and controlled degraded mode. Multi-agent resilience studies for EV charging stations under hybrid cyberattacks support this adaptive action logic because distributed restoration performs better when response is selected according to disturbance structure rather than applied as a fixed rule [11]. The self-healing controller therefore acts as an autonomous policy engine that matches each fault to the least disruptive restoration path capable of preserving service integrity. The principal fault classes, healing triggers, recovery actions, and validation conditions used in the proposed architecture are summarized in Table 1.

Table 1. Fault Classes, Healing Triggers, Recovery Actions, and Validation Conditions in the Proposed Self-Healing EV Charging Network

Fault class	Healing trigger	Autonomous recovery action	Validation condition
Charger hardware degradation	abnormal power profile, unstable current or voltage	module reset, safe fallback mode, charger rerouting	stable power delivery restored
Station controller fault	heartbeat loss, repeated control mismatch	controller restart with state recovery	session state consistency recovered
Communication-path failure	message timeout, route instability, packet loss burst	alternate communication path or local buffered control	command and telemetry exchange restored
Cyber-induced anomaly	malicious pattern, integrity violation, suspicious control behavior	interface quarantine, restricted local mode, trust reset	unsafe traffic removed and safe operation preserved
Hybrid cross-layer fault	combined physical and digital inconsistency	coordinated isolation plus degraded-mode service recovery	stable post-fault operation maintained

After healing selection, the methodology performs distributed service recovery. If a local charger cannot continue safely, the active session may be reassigned to a neighboring node or converted into a controlled interruption with preserved transaction state. If a station remains partially functional, noncritical services are suspended while essential charging continuity is maintained through degraded operation. The distributed recovery layer ensures that restoration is not evaluated only at the component level but also at the service level, because the purpose of self-healing is to preserve charging availability rather than merely restart electronics. This stage therefore rebalances service continuity across the network after the initial healing action has occurred.

The final methodological stage is post-recovery validation and stabilization. Recovery is not considered successful simply because a fault response was executed; it is accepted only when charger state, communication integrity, session continuity, and control responsiveness remain stable over the validation window. If instability persists, the healing engine either escalates to a stronger recovery action or isolates the affected node from normal coordination. This prevents incomplete restoration from being mistaken for successful healing. The self-healing cycle is therefore closed only after restored behavior has been verified as stable and safe.

3. Results and Discussion

The proposed self-healing architecture produced a clear improvement in restoration performance when compared with conventional manual recovery and fixed rule-based recovery baselines. Under isolated charger and controller faults, all approaches were able to restore some level of service, but the differences became more pronounced under communication faults and hybrid cross-layer disturbances.

Manual recovery performed worst because restoration did not begin until the fault had been externally confirmed, while the rule-based baseline reacted faster but remained limited by rigid action mapping. The autonomous self-healing strategy restored service more consistently because it combined fault classification, local diagnosis, and adaptive healing selection inside one loop. This behavior is shown in Figure 1, where service restoration under autonomous recovery remains higher and more stable across the observation horizon than under the comparison strategies. The result indicates that self-healing behavior provides value not merely by shortening downtime, but by preserving restoration quality under more varied fault conditions.

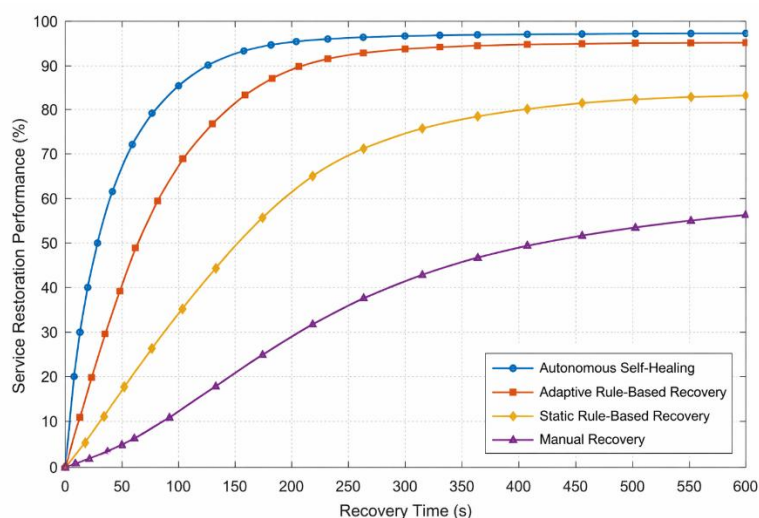


Figure 1. Service Restoration Performance Under Different Autonomous Recovery Strategies

A second important observation was that the autonomous design improved the timing of fault isolation. In the manual baseline, propagation risk remained elevated because affected nodes often continued interacting with the wider charging environment before intervention occurred. The fixed rule-based strategy reduced this delay, but it sometimes over-isolated components or failed to adapt when hybrid conditions were present. By contrast, the proposed self-healing framework isolated faults in a more targeted manner because diagnosis distinguished local charger faults from communication faults and cyber-induced anomalies before selecting the recovery path. This reduced both underreaction and unnecessary overreaction. The system therefore improved containment as well as restoration.

Charging-session continuity also improved substantially under the autonomous architecture. In conventional handling, a disrupted charger session often led to complete interruption because the recovery process focused on restarting components rather than on preserving session state. The proposed framework performed better because healing decisions included service-level logic such as state restoration, degraded-mode continuation, and controlled rerouting of active sessions where feasible. This meant that more sessions could either continue locally or be resumed with preserved transaction integrity after healing action was applied. The practical significance is that self-healing

strengthened user-facing continuity rather than only technical restart speed. In a public charging environment, that distinction is critical.

The most visible benefit emerged under distributed and hybrid fault scenarios, where one disturbance could amplify stress on neighboring chargers or create uncertainty about whether the underlying issue was physical, digital, or both. In these conditions, the self-healing controller avoided simple one-rule recovery and instead chose actions based on diagnosis profile and recovery feasibility. This selective behavior made the architecture more robust than a fixed recovery scheme when multiple recovery options were possible. The effect is represented in Figure 2, where the autonomous strategy achieves lower fault-isolation time while maintaining stronger post-recovery stability across distributed EV charging nodes. The result confirms that recovery speed alone is insufficient if the restored node becomes unstable shortly afterward. Effective self-healing must therefore optimize both early isolation and sustained stabilization.

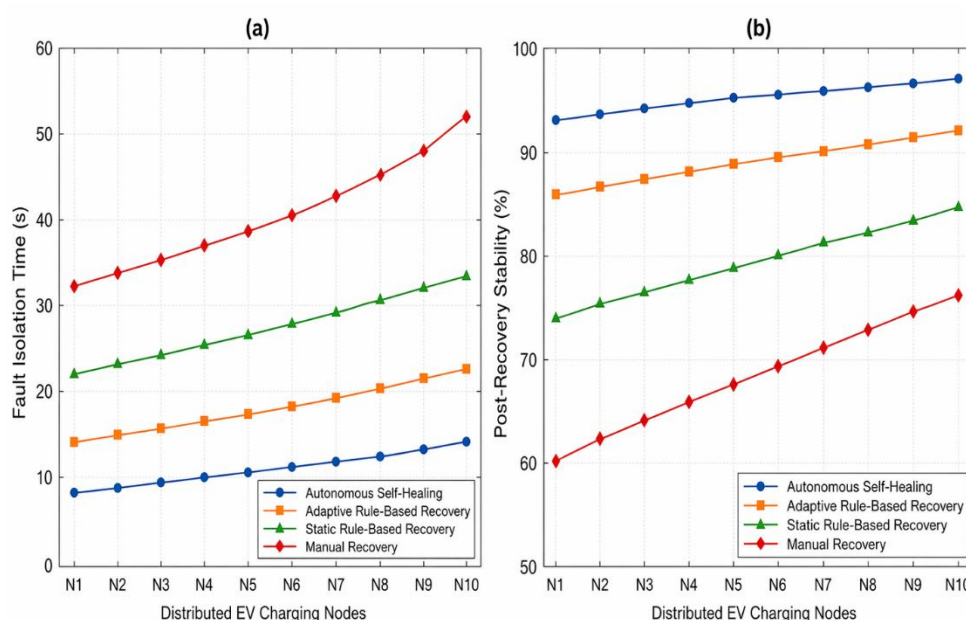


Figure 2. Fault Isolation Time and Post-Recovery Stability Across Distributed EV Charging Nodes

Another notable result was the performance of degraded-mode operation. In some scenarios, full restoration was not immediately possible because the fault affected coordination services or communication trust. Instead of forcing full shutdown, the architecture preserved a limited but stable charging mode while nonessential functions were suspended. This reduced the burden on neighboring chargers and prevented localized faults from triggering avoidable system-wide service collapse. Degraded-mode healing therefore acted as a resilience buffer between nominal operation and complete outage. Its contribution was especially important in communication and hybrid fault cases, where strict all-or-nothing recovery would have reduced network availability more severely.

Overall, the results demonstrate that self-healing design improves EV charging resilience by turning recovery into a structured autonomous process rather than a manual or static response sequence. The framework accelerated restoration, shortened fault-isolation time, preserved more charging-session continuity, and maintained stronger post-recovery stability across distributed nodes. These gains were strongest when the fault environment was heterogeneous and recovery required interpretation rather than simple restart. The study therefore supports the view that next-generation EV charging networks should be engineered not only to detect faults, but to heal themselves through adaptive and distributed restoration logic.

4. Conclusion

This study presented a self-healing design for EV charging networks using autonomous fault recovery, with emphasis on fault sensing, class-aware diagnosis, adaptive healing selection, distributed service recovery, and post-restoration stabilization. The proposed architecture treated the charging network as a cyber-physical service system in which resilience depends not only on knowing that a fault exists, but on whether the system can restore safe and useful charging behavior without waiting for manual intervention. By integrating detection, diagnosis, healing policy, and validation into one control loop, the framework created a practical foundation for autonomous recovery in distributed charging environments. This makes the design particularly relevant for large charging networks where faults can emerge across several nodes and layers simultaneously.

The results showed that the self-healing framework improved service restoration, reduced fault-isolation time, preserved charging-session continuity, and maintained stronger stability after recovery than manual and fixed rule-based recovery approaches. These gains were especially visible under communication and hybrid cross-layer disturbances, where successful restoration required diagnosis-aware action rather than simple restart logic. The degraded-mode capability further strengthened resilience by preserving essential service when immediate full restoration was not possible. Taken together, the findings show that autonomous recovery is a central requirement for dependable EV charging infrastructure rather than an optional enhancement.

Self-healing operation is likely to become increasingly important as charging ecosystems expand in size, software complexity, and service dependence. Future work can extend this framework through digital-twin-assisted healing validation, federated recovery learning across geographically distributed stations, predictive healing triggered before full fault manifestation, and tighter coordination with grid-side resilience mechanisms. Additional study may also examine mixed-vendor charging fleets, cascading disturbance patterns, and explainable recovery policies for operator-facing control support. These directions would strengthen the role of autonomous fault recovery as a practical foundation for resilient EV charging networks.

References

1. Mitikiri, S. B., Babu, K. V. S. M., Dwivedi, D., Srinivas, V. L., Chakraborty, P., Yemula, P. K., & Pal, M. (2025). Cyber–physical security in EV charging infrastructure: Components, vulnerabilities, and defense strategies. *Sustainable Energy Technologies and Assessments*, *81*, 104435.
2. Powell, B., & Johnson, C. (2024). *Impact of electric vehicle charging station reliability, resilience, and location on electric vehicle adoption* (No. NREL/TP--5R00-89896). National Renewable Energy Laboratory (NREL), Golden, CO (United States).
3. Motlagh, S. G., Oladigbolu, J., & Li, L. (2025). A review on electric vehicle charging station operation considering market dynamics and grid interaction. *Applied Energy*, *392*, 126058.
4. Alasali, F., Ghalyon, S. A., El-Naily, N., Abuashour, M. I., AlMajali, A., Itradat, A., & Holderbaum, W. (2025). Innovative Investigation of the Resilience of EV Charging Infrastructure Under Cyber-Physical Threats Based on a Real-Time Co-Simulation Testbed. *IET Cyber-Physical Systems: Theory & Applications*, *10*(1), e70021.
5. Pawlik, L., Wilk-Jakubowski, J. L., Grabski, P. T., & Wilk-Jakubowski, G. (2025). Securing the Electrified Future: A Systematic Review of Cyber Attacks, Intrusion and Anomaly Detection, and Authentication in Electric Vehicle Charging Infrastructure. *Energies*, *18*(18), 4847.
6. Sakwa, M., Nespoli, A., Matrone, S., Leva, S., Guerini, A., Demartini, A., & Ogliari, E. (2024). Electric vehicle supply equipment monitoring and early fault detection through autoencoders. *Sustainable Energy, Grids and Networks*, *40*, 101497.
7. Kesavan, V. T., Hossen, M. J., Gopi, R., & Joseph, E. R. (2025). Anomaly detection with grid sentinel framework for electric vehicle charging stations in a smart grid environment. *Scientific Reports*, *15*(1), 15774.
8. Tanyıldız, H., Şahin, C. B., Dinler, Ö. B., Migdady, H., Saleem, K., Smerat, A., ... & Abualigah, L. (2025). Detection of cyber attacks in electric vehicle charging systems using a remaining useful life generative adversarial network. *Scientific reports*, *15*(1), 10092.
9. Singh, A. R., Kumar, R. S., Rathore, R. S., Pandian, A., Alrayes, F. S., Allafi, R., & Ahmad, N. (2025). AI-enhanced smart grid framework for intrusion detection and mitigation in EV charging stations. *Alexandria Engineering Journal*, *115*, 603-621.
10. Hijgenaar, S., Ştefanov, A., Van Voorden, A. M., & Palensky, P. (2025). Cyber Resilience of Electric Vehicle Charging in Smart Grids: The Dutch Case. *IEEE Access*.
11. Sepehrzad, R., Khodadadi, A., Adinehpour, S., & Karimi, M. (2024). A multi-agent deep reinforcement learning paradigm to improve the robustness and resilience of grid connected electric vehicle charging stations against the destructive effects of cyber-attacks. *Energy*, *307*, 132669.