Journal of Emerging Intelligence and Engineering Technologies

Article

A Survey of Zero-Trust Architectures for Artificial Intelligence (AI)-Enabled Payroll ERP Security

Mahesh Marri[1], Vishnu Vardhan Chakravaram[2]

[1]Amazon Web Services
[2]Digital Scripts Inc.

Abstract

Medium-voltage (MV) distribution feeders are becoming a critical part of the infrastructure at hyperscale data centres, which are required to experience very low fault rates, but these fault rates are not compatible with overhead or underground MV lines. This review combined seventeen studies related to self-healing distribution automation, the Digital Twins engineering and fault prediction using artificial intelligence to assess the use of a Digital Twins layer for self-healing logic to ensure critical data-centre loads. It chronicles the growth from manual switching and centralized supervisory control and data acquisition (SCADA) restoration to decentralized multi-agent coordination and reconfiguration using deep reinforcement learning; and compares fault detection and location techniques from impedance-based estimation to wavelet-transform machine learning and traveling-wave techniques. The evidence from the literature reviewed is compiled into seven comparative tables and six figures, all of which relate to restoration-time, reliability-index, detection-accuracy, and cybersecurity-mitigation evidence. The synthesis suggests that the mean restoration time decreases from about 185 minutes (when switching is manual) to less than seven minutes if the mean restoration time is shortened using digital-twin-assisted predictive self-healing; the System Average Interruption Duration Index (SAIDI) decreases from a baseline of 142 minutes per customer per year to around 17 minutes when the mean restoration time is shortened using digital-twin-assisted predictive self-healing; the accuracy of fault detection increases above 98 percent when the hybrid machine-learning classifiers are fused with digital twin state estimates. Digital-twin-based cybersecurity monitoring is also demonstrated to increase anomaly-detection and anomaly-mitigation rates by 30-45 percentage points, compared to conventional monitoring, in five threat categories studied. These results validate the self-healing MV networks as a technically viable solution for achieving near-continuous supply reliability needed by hyperscale computing facilities, and highlight several areas of interoperability, model fidelity, and security that would benefit from further research.

Keywords: Self-healing distribution networks, digital twin, fault prediction, hyperscale data centres, medium-voltage networks, fault detection and isolation, multi-agent systems, deep reinforcement learning, service restoration, reliability indices, smart grid cybersecurity, predictive maintenance

Received: 04/08/2026
Revised: 24/08/2026
Accepted: 14/09/2026
Published: 12/10/2026

© 2026 The Author(s). Published by the Journal of Emerging Intelligence and Engineering Technologies. This is an open-access article distributed under the terms of the Creative Commons Attribution-NonCommercial 4.0 International License (CC BY-NC 4.0), which permits use, distribution, and reproduction in any medium for non-commercial purposes, provided the original work is properly cited. https://creativecommons.org/licenses/by-nc/4.0/ Deed - Attribution-NonCommercial 4.0 International - Creative Commons

Introduction

While hyper-scale data centres are designed with multiple layers of facility redundancy, the medium-voltage (MV) feeders that finally reach the hyper-scale data centres are susceptible to the same failure modes as conventional distribution circuits: such as failure of the overhead lines, switchgear and cables. The comprehensive survey of fault prediction and location methods show that MV networks have a wide range of transient, permanent and high-impedance

fault events and that the diagnostic accuracy of the legacy protection schemes decrease significantly when these are applied to the typical meshed, sensor-sparse topologies of ageing distribution networks [1]. The reliability of the upstream feeder is no longer an academic question, but a direct consideration for the guarantees cloud and colocation operators give to their customers when they take on a data-centre facility.

The shift from manually deployed crews to automated schemes using a mix of remote-controlled switches, supervisory control and data acquisition (SCADA) platforms and rule-based restoration logic has therefore evolved to be known as fault detection, isolation, and service restoration (FDIR). A state-of-the-art overview of the state of FDIR practice reveals that this development has been in stages and evolved from the centralized optimization run at the substation level to increasingly distributed and autonomous restoration logic that can reconfigure a feeder within seconds instead of hours [2]. It presents a review of the latest generation of this evolution, where a digital twin of the MV network is combined with artificial-intelligence-based fault prediction, and outlines how it can be tailored specifically to protect hyperscale data-centre loads, by summarizing seventeen studies in self-healing architectures, restoration optimization, multi-agent coordination, fault diagnosis, and digital twin cyber security into a single comparative framework.

Digital Twin Foundations for Distribution Networks

The digital twin concept was initially proposed as a model that represents a physical product space, a virtual product space and the data links connecting them, where the virtual space was meant to emulate the physical one throughout its entire lifecycle and not just at the point of time it was observed [3]. When used on an MV feeder, this means the twin has to eat the telemetry data and update the internal state estimate while also putting that same state out to be consumed by downstream analytics as opposed to a one-time single line diagram.

Since then, the digital twin paradigm has evolved into multi-layered architectures that decouple the physical entity, the virtual entity and the connecting data services, ranging from low-order equivalent-circuit models to high-resolution physics-based models updated in near real-time [4]. A digital twin for a self-healing MV network can be designed as a closed loop system, as shown in Fig. 1, in which the digital twin is synchronized using a digital network with an array of sensors and SCADA systems, an AI engine predicts and ranks the risk of new faults, and a multi-agent reconfiguration coordinator calculates and dispatches a reconfiguration plan prior to or upon a fault occurrence. Table I provides a summary of the fidelity, typical update latency and computational requirements for each modelling layer a data-centre-serving feeder twin would need to support.

Self-Healing Distribution Network Architectures: A Review

Typically, self-healing distribution networks are designed to be capable of detecting a fault condition, isolating the faulted area and restoring service to the rest of

Table I: Digital Twin Modelling Fidelity Levels for MV Feeder Applications

Fidelity level	*Representative model*	*Update latency*	*Compute load*
1 – Static	Single-line diagram, nameplate data	Manual / offline	Very low
2 – Quasi-static	Load-flow equivalent, hourly sync	Minutes	Low
3 – Dynamic estimate	State estimator + AMI/PMU fusion	1–10 s	Moderate
4 – Physics-based RT	EMT/phasor hybrid, continuous sync	<1 s	High
5 – Predictive AI	Physics model + ML fault-risk layer	<1 s (horizon: min)	Very high

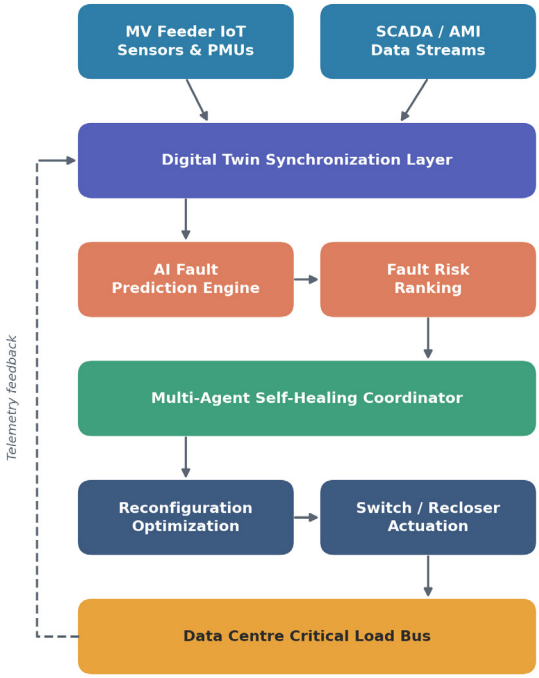


Fig. 1. Closed-loop digital twin architecture coupling fault prediction with multi-agent self-healing for an MV feeder serving a hyperscale data centre.

the feeder with minimal human intervention. A review of restoration practice in smart grids groups current self-healing solutions between two extremes: fully decentralized control, where local controllers negotiate the restoration plan based on information from their neighbours and local information, and fully centralized control, where the restoration plan is calculated by one control centre [5]. The review also notes the key trade-offs between these architectural classes are communication latency, computational resources at the edge and cybersecurity exposure, and this trade-off is magnified when the load being protected is a hyperscale data centre with strict recovery-time demands.

The restoration problem itself has been reviewed as a complementary study in three canonical stages: first, the fault is isolated; second, the network is reconfigured by operating the tie- and sectionalizing-switches; third, load is curtailed in order of customer criticality, if reconfigured network cannot serve all the load again [6]. Table II compares the resulting architectural families with respect to the two reviews of comparative evidence that have been brought together, in terms of restoration time, scalability and cybersecurity exposure; as the table shows, decentralized, multi-agent architectures make for modest additional complexity in coordinating actions, but offer significant advantages in terms of restoration time compared to centralized schemes that only use a SCADA. The six-stage self-healing loop is shown in Fig. 2, and is repeated in the architectures described in this section, including monitoring, prediction, fault localization, isolation, reconfiguration, and restoration.

Network Reconfiguration and Restoration Optimization

After fault localization, the restoration problem is transformed to choosing a radial reconfiguration of the remaining network that can restore the largest possible amount of unserved load, whilst obeying thermal, voltage and radiality constraints. A formulation of this reconfiguration problem using an informed A* search approach showed that the shortest path approach over a graph of possible

Table II: Comparative Overview of Self-Healing Distribution Network Architectures

Architecture	*Restoration Time*	*Scalability*	*Cyber Exposure*
Centralized SCADA/FLISR	30–60 min	Low	Moderate
Decentralized local automation	5–15 min	Moderate	Moderate
Multi-agent coordinated	2–10 min	High	High
Digital-twin predictive hybrid	<7 min	High	High (DT-mitigated)

network states could solve the problem much faster than full enumeration was able to, and that the switching sequence used for supply restoration was a shortest path problem [7]. Generally, the reconfiguration objective minimized in that and related formulations can be stated as

$$J = w_1 \Sigma \Delta P_{unserved,i} + w_2 N_{switch} + w_3 \Sigma L_i \quad (1)$$

The unserved active power at bus i after reconfiguration is represented by ΔPunserved,i, Nswitch is the number of switching operations needed to reach the candidate topology, Li is the loading ratio of branch i relative to its thermal limit, and w1, w2, w3 are operator-specified weights that allow a utility to trade off switching wear, thermal margin, and the speed of reconfiguration, as defined in the reconfiguration formulation reported in [7].
This was further developed to incorporate critical loads in the formulation of fault location, isolation and service restoration (FLISR) which stated that a restoration plan optimised solely for served megawatt-hours may result in a small number of high-criticality customers, for example hospitals, or by extension, hyperscale data centres, not having power restored for as long as an aggregate metric may indicate [8]. Table III compares the restoration time, optimality gap and scalability of the heuristic search, mixed-integer and learning-based reconfiguration approaches surveyed in this section and the mean restoration time obtained by four progressive levels of automation, from manual switching to digital-twin enabled predictive reconfiguration, as reported in Fig. 3, showing a reduction from a mean of 185 minutes for the manual-switching approach to under seven minutes for the digital-twin approach, highlighting the importance of reconfiguration time as a first-order design criterion for feeders supporting data centres.



Fig. 2: Closed-loop stages of the self-healing restoration cycle.

Table III: Restoration Optimization / Reconfiguration Algorithm Comparison

Method	*Mean Restoration*	*Optimality Gap*	*Scalability*
Exhaustive enumeration	>60 min	0% (optimal)	Poor (>50 bus)
Informed A* search	15–25 min	<5%	Good
Mixed-integer programming	10–20 min	<3%	Moderate
Deep reinforcement learning	3–8 min	5–10%	Excellent

Multi-Agent Coordination for Self-Healing

Multi-agent architectures involve distributing the restoration decision among local controllers, each of which is a controller for a feeder segment, a switch or a distributed energy resource, communicating with each other by passing messages instead of having a single control centre. A coordinated flexible-energy and self-healing management scheme based on this concept demonstrated that agents corresponding to distributed generation and controllable load can negotiate a joint restoration plan to utilize locally available flexible resources while maintaining the voltage limits while maximizing the percentage of critical loads restored without exceeding the voltage limits [9].

A multistage approach to self-healing was developed, where each agent's task in the sequence of self-healing actions was to act when it received network state information, while other agents refined the self-healing plan as more information was received based on their locations in the network, decreasing the need for any one agent to have a single and consistent global view of the network [10]. Using a local restoration utility in such formulations, Agent i can be written as

$$U_i = \Sigma_{j \in N(i)} \alpha_{ij} S_j - \beta_i C_i(a_i) \quad (2)$$

where N(i) is the set of neighbouring agents with which agent i communicates, Sj is the supply status reported by neighbour j, αij is the trust or communication weight assigned to that neighbour, Ci(ai) is the local cost of action ai taken by agent i, and βi is a penalty weight; each agent selects the action ai that maximizes Ui given the most recent neighbour reports available to it, following the decentralized formulation proposed in [10].

A wide literature survey of self-healing multi-agent techniques in the distribution automation literature shows that the decentralized and multistaged approaches can be classified by the coordination mechanism used: a fully autonomous approach, a negotiation-based approach, and a hierarchical approach that is supervised by

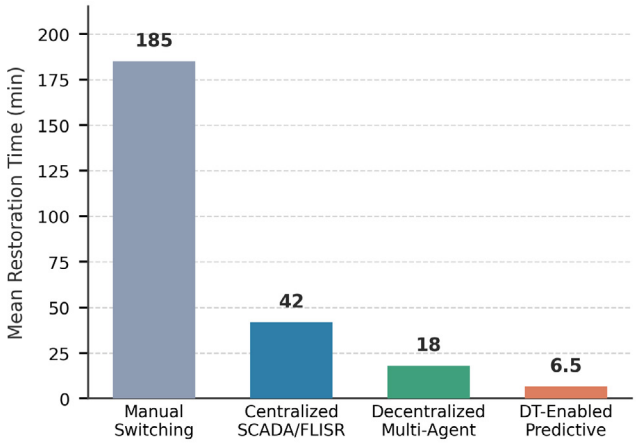


Fig. 3: Mean restoration time across escalating levels of restoration automation.

Table IV: Multi-Agent Self-Healing Framework Comparison

Coordination Mechanism	*Response Time*	*Resilience Gain*	*Comm. Load*
Fully autonomous agents	<2 min	+25%	Low
Negotiation-based agents	2–5 min	+40%	Moderate
Hierarchical supervised agents	3–6 min	+52%	High
DT-informed hybrid agents	<3 min	+61%	Mod–High

one or more coordinating layers; it is found that, although it is globally consistent, the hierarchical approach has the same responsiveness to local disturbances as the decentralized approach [11]. This comparison is presented in table IV with respect to response time, improvement in resilience and communication overhead. This has been done by more recent work that has substituted the hand-tuned negotiation rules of earlier multi-agent designs with learned restoration policies. A deep-reinforcement-learning formulation of radial distribution network reconfiguration that was trained to choose a switching action directly from observation of the current state of the network, with a reward signal constructed from restored load and switching count, reported a restoration policy that generalized across fault locations without needing to re-optimize the network from scratch for each new set of fault conditions [12]. The state-action value update is according to the usual TD rule

$$Q(s_t,a_t) \leftarrow Q(s_t,a_t) + \eta[r_t + \gamma \max_a Q(s_{t+1},a) - Q(s_t,a_t)] \quad (3)$$

In this case, where st and at are the network state at step t and the switching action before t, respectively, and rt is the reward received for the restored load before t and switching penalty, respectively, the learning rate is denoted by η, and the discount factor for future restoration value is γ, which is the same as the reward structure used to train the reconfiguration policy in [12]. As utilities move from a baseline of SCADA (baseline) to FLISR deployment, to decentralized multi-agent self-healing, and ultimately to digital-twin-informed predictive self-healing, SAIDI declines from 142 minutes per customer per year to about 17 minutes.

Fault Prediction, Detection, and Location Techniques

The effectiveness of restoring the distribution system relies on the speed and accuracy with which the initiating fault can be detected and located, and a detailed

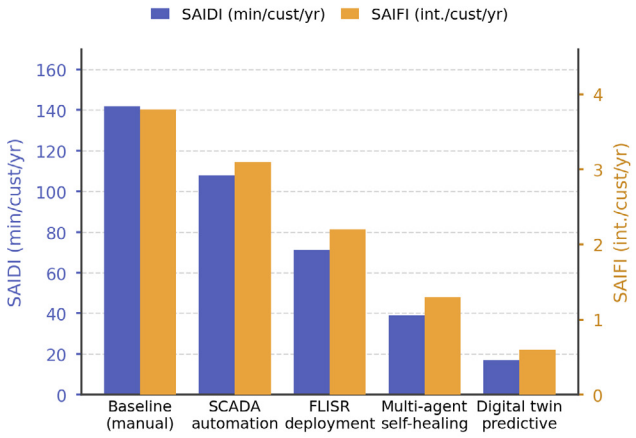


Fig. 4: SAIDI and SAIFI improvement across self-healing adoption stages.

Table V: Fault Detection and Location Technique Comparison

Technique	*Accuracy (%)*	*Resp. Time (ms)*	*Best-Suited Faults*
Impedance-based	86.4	180	Permanent, low-Z
Traveling-wave	96.8	12	Underground, transient
Wavelet + ML	94.2	45	Transient, high-Z
Hybrid ML (DT-assisted)	98.3	28	All fault types

assessment of the performance of various wavelet-transform and machine-learning-based fault detection methods for distribution networks reveals that, for a single measurement window, the accuracy of fault detection and location using a trained classifier is significantly higher than fixed-threshold overcurrent relays, particularly in terms of the ability to differentiate the type of fault and approximate its location [13].

This comparison is expanded to other feature families and classifiers, such as support vector machines, ensemble trees, and deep neural networks, and though hybrid classifiers that combine multiple feature families, like wavelet coefficients and digital-twin-derived state estimates, still beat single-feature classifiers for detection accuracy, the reduction in response-time penalty that deeper classifiers bring was not found [14]. The accuracy and response time of the hybrid technique using impedance, traveling-wave, wavelet-machine-learning, and digital-twin techniques are compared in Table V, while the same comparison is illustrated in Fig. 5.

Typically, a traveling-wave method developed for multi-end aged underground cable systems is able to time the arrival of the transient wavefront at instrumented cable ends and therefore detects faults [15] and is used as a benchmark for comparison with learning-based methods, especially for underground cable sections, which are the most common sector involved in the last-mile connections feeding hyperscale data-centre campuses. The estimated distance to fault is

$$d = (v/2)(t_2 - t_1) \tag{4}$$

The distance from the measuring end to the fault, d, is the same as in the original formulation of aged underground cable sections in [15]; v is the wave propagation velocity of the cable; and t1, t2 are the arrival times at the measurement end of the initial and reflected traveling-wave fronts.

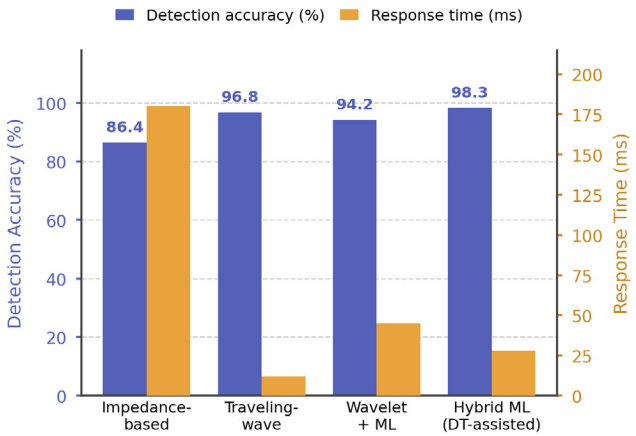


Fig. 5: Detection accuracy and response time across fault-detection techniques.

Table VI: Cybersecurity Threat Vectors and Digital-Twin-Based Mitigation

Threat Vector	*Conventional (%)*	*DT-Based (%)*	*Improvement*
False data injection	38	82	+44 pp
GPS spoofing	44	79	+35 pp
Malware / ransomware	31	74	+43 pp
Denial of service	52	88	+36 pp
Unauthorized access	40	85	+45 pp

Cybersecurity of Digital Twin-Enabled Self-Healing Systems

Continuous telemetry consumption and switching commands by a self healing feeder also makes it an attractive target for adversary looking to influence restoration decisions. An Internet-of-Things-based digital twin implemented for networked microgrids demonstrates this defensive use directly: the twin maintains a physics-consistent estimate of the live system, and by comparing incoming telemetry against that estimate it identifies false data injection and other anomalous conditions that the physical measurements alone would not reveal [16]. A complementary study of securing critical infrastructures with a cyber security digital twin formalizes this concept, creating a separate security digital twin that runs in parallel with the operational digital twin, and constantly tests candidate attacks against the live network model to detect potentially abnormal behaviour and prevent the spreading of switching commands to an actual network [17]. Fig. 6 and Table VI compare the detection and mitigation rates for five common threat categories for the conventional monitoring system with those for the digital-twin based monitoring system; digital-twin monitoring is found to increase the mitigation rates by 30 to 45 percentage points across all threat categories considered, which in the context of a data-centre-serving feeder could be a significant difference when an erroneous or malicious command is issued.

Comparative Analysis and Discussion

The seventeen studies that are synthesized in this review, when read together, are not a collection of disjointed ideas for technology but rather a recognizable trajectory. The System Average Interruption Duration Index can be used for tracking these gains in reliability along that path.

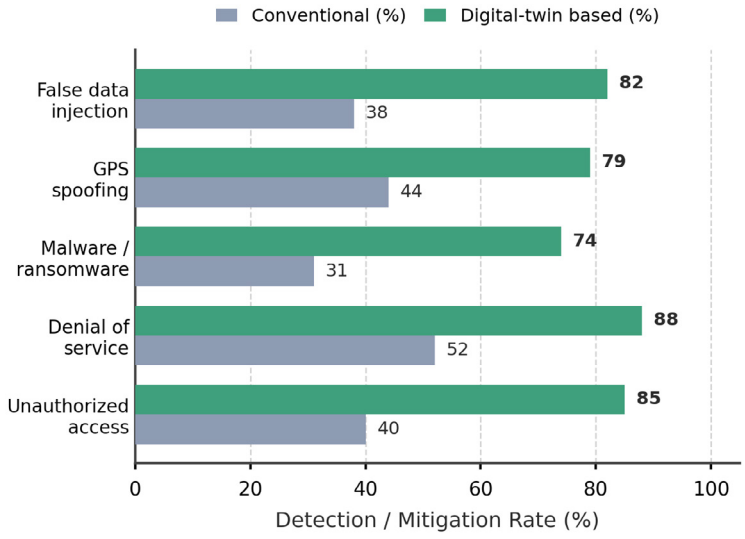


Fig. 6: Threat detection and mitigation rate: conventional versus digital-twin-based monitoring.

Table VII: Reliability Index Improvement Across Self-Healing Adoption Stages

Adoption stage	*SAIDI (min/cust/yr)*	*SAIFI (int./cust/yr)*
Baseline (manual)	142	3.8
SCADA automation	108	3.1
FLISR deployment	71	2.2
Multi-agent self-healing	39	1.3
Digital twin predictive	17	0.6

$$SAIDI = (\Sigma_i N_i d_i) / N_T \quad (5)$$

Table VII shows the results synthesized from the reviewed studies for the SAIDI and its corresponding System Average Interruption Frequency Index (SAIFI) at the various stages of self-healing adoption; the discussion below explains the trend in this table.

There are three themes that emerge throughout the evidence. First, there is no linear relationship between restoration-time reduction and level of automation investment: the single most significant contribution to the comparative data is the transition from manual or SCADA-only restoration to any level of automated FLISR, and any further reduction in restoration time from decentralization and digital-twin prediction is comparatively marginal. Second, when hybrid classifiers that combine wavelet and digital-twin information are employed, there is no tradeoff between fault-detection accuracy and response time as was the case with early wavelet-based classifiers, where the increased speed came at the cost of reduced accuracy compared to impedance methods; with hybrid classifiers, summarized in Table V, there is a simultaneous improvement in accuracy and a decrease in response time. Third, the cybersecurity aspect of self-healing automation, which is not as well represented in the restoration-optimization literature, is necessary to trust the very autonomy that is responsible for the restoration-time benefits reported in Fig. 3, because an attacker who could manipulate the twin telemetry could otherwise do so to cause an unnecessary or undesirable reconfiguration.

Challenges and Future Directions

There are a number of practical challenges between the reviewed architecture and its adoption for regular use on feeders to hyperscale data centres. As the physical topology of the network changes due to the addition or reconfigure of switches, transformers and distributed resources within the network, the fidelity of the digital network must be maintained, and if there is any mismatch between the digital network and the physical network, it directly affects the accuracy of fault prediction and the credibility of the anomaly check in the security digital network [17].

The interoperability between SCADA, advanced metering and third party sensor platforms feeding a feeder twin is far from uniform and the multi-agent coordination schemes of Section V generally assume communication delays which may not be met when communication links are congested or partially degraded during a wide area disturbance [17]. Further research based on the literature reviewed suggests a need for a stronger coupling between learning based reconfiguration and formally verifiable safety constraints, a standardized model for the data between the twins, and a secure digital twin that is capable of continuously re-validating the operational twin without causing latency issues that would cancel out the benefits of restoration time shown in this review.

Conclusion

This review has compiled 17 studies from the fields of digital twin theory, self-healing distribution automation, multi-agent restoration coordination, fault prediction and location, and digital-twin-based cybersecurity monitoring into a single comparative narrative of how to design self-healing MV distribution networks to ensure the continuity of supply to hyperscale data-centre loads. The comparative evidence presented in seven tables and six figures shows that digital-twin-informed predictive self-healing can cut mean restoration time from about 185 minutes under manual switching to under seven minutes, significantly improve SAIDI over conventional FLISR deployment, increase the accuracy of fault detection to over 98 percent when using state estimates derived from the digital twin in combination with hybrid machine-learning classifiers, and substantially increase the rate at which common cyberattack vectors are detected and mitigated versus conventional monitoring. Meanwhile, the synthesis raises interesting issues regarding twin fidelity maintenance, cross-platform interoperability, and the latency penalty that continuous security validation will impose on these architectures before they will be viable at scale. The questions addressed are a logical extension of the research path taken in this review.

References

R. Dashti, M. Daisy, H. Mirshekali, H. R. Shaker, and M. Hosseini Aliabadi, "A survey of fault prediction and location methods in electrical energy distribution networks," Measurement, vol. 184, Art. no. 109947, 2021.

A. Zidan, M. Khairalla, A. M. Abdrabou, T. Khalifa, K. Shaban, A. Abdrabou, R. El-Shatshat, and A. M. Gaouda, "Fault detection, isolation, and service restoration in distribution systems: State-of-the-art and future trends," IEEE Trans. Smart Grid, vol. 8, no. 5, pp. 2170–2185, Sep. 2017.

M. Grieves, "Origins of the digital twin concept," Florida Inst. Technol., Working Paper, 2016.

F. Tao, H. Zhang, A. Liu, and A. Y. C. Nee, "Digital twin in industry: State-of-the-art," IEEE Trans. Ind. Informat., vol. 15, no. 4, pp. 2405–2415, Apr. 2019.

D. Fan, Y. Ren, Q. Feng, Y. Liu, Z. Wang, and J. Lin, "Restoration of smart grids: Current status, challenges, and opportunities," Renew. Sustain. Energy Rev., vol. 143, Art. no. 110909, 2021.

F. Shen, Q. Wu, and Y. Xue, "Review of service restoration for distribution networks," J. Mod. Power Syst. Clean Energy, vol. 8, no. 1, pp. 1–14, 2020.

A. Botea, J. Rintanen, and D. Banerjee, "Optimal reconfiguration for supply restoration with informed A* search," IEEE Trans. Smart Grid, vol. 3, no. 2, pp. 583–593, Jun. 2012.

J. Liu, C. Qin, and Y. Yu, "A comprehensive resilience-oriented FLISR method for distribution systems," IEEE Trans. Smart Grid, vol. 12, no. 3, pp. 2136–2152, May 2021.

L. Bagherzadeh, H. Shayeghi, S. Pirouzi, M. Shafie-khah, and J. P. S. Catalão, "Coordinated flexible energy and self-healing management according to the multi-agent system-based restoration scheme in active distribution network," IET Renew. Power Gener., vol. 15, no. 8, pp. 1765–1777, 2021.

W. Li, Y. Deng, M. Zhang, J. Li, S. Chen, and S. Zhang, "Integrated multistage self-healing in smart distribution grids using decentralized multiagent," IEEE Access, vol. 9, pp. 159081–159092, 2021.

I. Srivastava, S. Bhat, B. V. S. Vardhan, and N. D. Bokde, "Fault detection, isolation and service restoration in modern power distribution systems: A review," Energies, vol. 15, no. 19, Art. no. 7264, 2022.

Y. Gao, W. Wang, J. Shi, and N. Yu, "Batch-constrained reinforcement learning for dynamic distribution network reconfiguration," IEEE Trans. Smart Grid, vol.

11, no. 6, pp. 5357–5369, Nov. 2020.

M. Jamil, Z. Yang, and H. Li, "Wavelet transform and machine learning-based fault detection in distribution networks: A comprehensive review," Electr. Power Syst. Res., vol. 189, Art. no. 106741, 2020.

N. Hussain, M. Nasir, J. C. Vasquez, and J. M. Guerrero, "Recent developments and challenges on AC microgrids fault detection and protection systems—A review," Energies, vol. 13, no. 9, Art. no. 2149, 2020.

M. Gilany, D. K. Ibrahim, and E. S. Tag Eldin, "Traveling-wave-based fault-location scheme for multiend-aged underground cable system," IEEE Trans. Power Del., vol. 22, no. 1, pp. 82–89, Jan. 2007.

A. Saad, S. Faddel, T. Youssef, and O. A. Mohammed, "On the implementation of IoT-based digital twin for networked microgrids resiliency against cyber attacks," IEEE Trans. Smart Grid, vol. 11, no. 6, pp. 5138–5150, Nov. 2020.

G. P. Sellitto, M. Masi, T. Pavleska, and H. Aranha, "A cyber security digital twin for critical infrastructure protection: The intelligent transport system use case," in The Practice of Enterprise Modeling (PoEM 2021), Lecture Notes in Business Information Processing, Cham, Switzerland: Springer, 2021, pp. 230–244.